

THE CHINESE UNIVERSITY OF HONG KONG,

IERG 5254: MULTIUSER INFORMATION THEORY

Lecture Notes

Chandra Nair

*Scribes: Devon Ding, Chaorui (Claire) Yao, Ken Lau, Ling Chih Wei,
Mehrdad Tahernia, Sida Liu*

Spring 2027, revised October 2026

Contents

1	Information measures and properties	8
1.1	Heuristics	8
1.2	Axiomatic formulation of entropy	11
1.3	Properties of entropy	14
1.4	Properties of Kullback–Leibler divergence	17
1.5	Properties of differential entropy	22
1.6	Maximum entropy distributions	26
1.7	Properties of mutual information	27
1.8	Properties of Fisher information	31
1.9	Entropy power inequality	37
1.9.1	Conditional EPI	43
1.10	Further inequalities	46
1.11	Gaussian optimality via rotation: general idea	53
1.11.1	Gaussian maximises differential entropy under a covariance constraint: alternate proof	54
1.11.2	Alternate proof of the EPI via Gaussian extremality (advanced, optional)	56
1.11.3	Geometric interpretation of the EPI	66
1.12	Perturbation method	67
1.13	Sub-additivity: the key principle behind outer bounds and converses	69
1.13.1	The abstract template	70
1.13.2	Example: the single-letterisation of $I(X^n; Y^n)$	72
1.13.3	Example: the functional of the perturbation proof of the EPI	72
1.13.4	Roadmap	73
2	Point-to-point information theory	74
2.1	Typicality	74
2.2	Channel coding	80
2.3	Examples	88
2.3.1	Binary symmetric channel	88
2.3.2	Binary erasure channel	89
2.3.3	Product channel	90
2.3.4	Direct sum channel	91
2.4	Source coding	93
3	Multiple access channel	96
3.1	Capacity region	97

3.2	Cardinality bound	98
3.3	Converse	102
3.4	Achievability	105
3.5	Time-sharing	108
3.6	Extension to multiuser multiple access channel	110
3.7	Capacity of the Gaussian multiple access channel	115
3.8	Simple implementation on the AWGN MAC	118
3.9	Weighted sum-rate of the Gaussian MAC	119
3.10	Bounds for the determinant of the sum of Hermitian matrices	120
3.11	Multi-letter characterization of the capacity region	122
4	Broadcast channel	125
4.1	Properties of the capacity region	126
4.2	Time-division inner bound	127
4.3	Degraded broadcast channel	128
4.4	Superposition coding	129
4.5	Upper concave envelope and sub-additivity	135
4.5.1	Sub-additivity of the envelope in a degraded channel	136
4.6	AWGN degraded broadcast channel	138
4.7	Degraded broadcast channel with K receivers	139
4.8	Less noisy and more capable broadcast channels	140
4.9	UV outer bound	149
4.10	UVW outer bound	153
4.11	Marton's inner bound	158
4.11.1	Simplification of Marton's region	164
4.12	Deterministic and semi-deterministic broadcast channels	165
4.13	Marton's sum rate	167
4.13.1	Cardinality bound via the perturbation method	168
4.14	Functional representation and multi-letter Marton	170
4.15	Equivalence of average- and maximal-error capacity regions	171
4.16	Gaussian vector broadcast channel	171
5	Interference channel	173
5.1	Model and Definitions	173
5.2	Some Simple Inner and Outer Bounds	174
5.2.1	Time-division inner bound	174
5.2.2	Treating interference as noise (TIN)	175
5.2.3	Simultaneous-decoding inner bound	175
5.3	Strong and Very Strong Interference	176
5.3.1	Very strong interference	176
5.3.2	Strong interference	177
5.4	Han-Kobayashi Inner Bound	180
5.5	The Scalar Gaussian Interference Channel	182
5.5.1	Strong interference	182

5.6	The Injective Deterministic Interference Channel	185
5.7	Outer Bounds	187
5.7.1	A genie-based outer bound	187
5.7.2	The auxiliary-receiver principle	191
5.7.3	The Etkin–Ordentlich outer bound	192
5.8	Very Weak Interference and the Noisy-Interference Sum Capacity	193
5.9	Mixed Interference	195
5.10	Z-Interference and Binary Deterministic Examples	196
5.10.1	Clean Z-interference	196
5.10.2	The binary skew-symmetric Z-IC	197
5.10.3	The binary deterministic IC (El Gamal–Costa model)	198
6	Channels with state	199
6.1	Setting and code definitions	199
6.2	The compound channel	200
6.3	Channels that change over time: the averaging argument	202
6.3.1	State known to neither party	202
6.3.2	State known at the receiver only (CSIR)	203
6.3.3	State known at the transmitter and receiver (CSIT + CSIR)	204
6.4	State known causally at the transmitter: Shannon strategies	205
6.5	State known non-causally at the transmitter: Gelfand–Pinsker	207
6.6	Costa’s dirty-paper coding	212
6.7	Extensions and remarks	214
7	Distributed source coding	216
7.1	The distributed source coding problem	217
7.2	Slepian–Wolf: lossless coding of two correlated sources	218
7.3	The one-helper problem: source coding with side information at the decoder	221
7.4	One source, two helpers	224
7.5	The Gray–Wyner system	225
7.5.1	The supporting hyperplane and the corner points	227
7.6	The Körner–Marton modulo-two sum problem	229
7.6.1	Linear-code achievability	229
7.6.2	The Ahlswede–Han inner bound	230
7.6.3	An outer bound via the Csiszár sum identity	230
7.7	Hypercontractivity: the sharp contraction calculus	233
7.7.1	Reverse hypercontractivity and the Bonami–Beckner / log-Sobolev state- ments	235
7.8	Closing: the sub-additivity / contraction calculus	237
7.9	Exercises	237
8	Relay Channels	239
8.1	The relay channel: setting and code definition	239
8.2	The cut-set outer bound	241

8.3	Decode-and-forward	244
8.4	Compress-and-forward	246
8.5	Remarks	249
8.6	Exercises	250
9	Appendix	252
9.1	Diffusion Equation, Heat Flow and Fisher Information	252
9.2	Sanov's inequality in discrete spaces: Types and Typical Sequences	255
9.2.1	Types	255
9.2.2	Maximal couplings	262
9.3	A Min-Max Theorem	264
9.4	d-separation: proof of the graphical independence rule	266

Notation

The following notation is used throughout these notes.

$\triangleq$	equal by definition
$\Longleftrightarrow$	if and only if
$\bigcirc$	end of theorem, definition, remark, ...
$\blacksquare$	end of proof
$[1 : n]$	the set $\{1, 2, \dots, n\}$; for a random variable Q uniform on $[1 : n]$, X_Q denotes the random variable X_i with $i = Q$
$\log$	$\log = \log_2$ unless explicitly overridden; $\ln$ is the natural logarithm and $\exp(t) = e^t$ (base-2 exponentials are written 2^t explicitly)
RV	random variable
X	random variable (upper-case letters); x denotes a realisation (lower-case letters)
$\mathcal{X}, \mathcal{Y}$	alphabets of X, Y (calligraphic letters)
$x \in \mathcal{X}$	x is an element of $\mathcal{X}$
$\mathbf{X}, \mathbf{Y}, \dots$	bold upper-case letters denote random vectors in $\mathbb{R}^d$; bold lower-case letters (e.g. $\mathbf{x}$) denote deterministic vectors
X^n	the vector-valued RV $(X_1, X_2, \dots, X_n)$; $\mathbf{x}^n$ denotes the sequence of realisations
X_ℓ^n	the block $(X_{\ell 1}, \dots, X_{\ell n})$ transmitted by user ℓ (do not confuse with a suffix, which we avoid)
$X^{n \setminus i}$	the vector-valued RV with components of X^n excluding X_i
X^{i-1}	the vector-valued RV $(X_1, \dots, X_{i-1})$ (similarly Y^{i-1})
$\Pr(A)$	probability of an event A
$p_X(x) = p(x)$	the pmf value $\Pr(X = x)$; the subscript is omitted when there is no ambiguity
P_X	the distribution (probability measure) of X ; for finite alphabets this is a stochastic row vector of length $ \mathcal{X} $ with entries $p_X(x)$
$p_{Y X}(y x) = p(y x)$	conditional pmf $\Pr(Y = y X = x)$
$W : \mathcal{X} \rightarrow \mathcal{Y}$	stochastic matrix with rows indexed by elements of $\mathcal{X}$ and columns indexed by elements of $\mathcal{Y}$; $W(\cdot x)$ is a pmf on $\mathcal{Y}$ for every $x \in \mathcal{X}$
$W^{\otimes n}$	the memoryless channel W used n times, i.e. $W^{\otimes n}(y^n x^n) = \prod_{i=1}^n W(y_i x_i)$
$E[g(X)]$	expectation of $g(X)$ with respect to the distribution of X
$H(X) \triangleq H(P_X)$	(Shannon) entropy of a discrete RV X

$H_b(x) = H_2(x)$	binary entropy function $-x \log x - (1-x) \log(1-x)$ (both symbols are used interchangeably)
$H(X Y)$	conditional entropy $\sum_y p_Y(y) H(X Y=y)$
$h(\mathbf{X}) \triangleq h(f_{\mathbf{X}})$	differential entropy of a continuous random vector $\mathbf{X}$ with density $f_{\mathbf{X}}$
$I(X; Y) \triangleq I(P_X, W_{Y X})$	mutual information; $I(X; Y Z)$ denotes conditional mutual information
$D(P_X \ P_Y)$	Kullback–Leibler divergence of P_X and P_Y
$J(\mathbf{X})$	Fisher information of a random vector (see the Fisher information section of Chapter 1 for the convention)
$\pi(\cdot x^n)$	type (empirical distribution) of the sequence x^n
$\mathcal{T}_\epsilon^{(n)}$	typical set $\{x^n : \pi(x x^n) - p_X(x) \leq \epsilon p_X(x) \text{ for all } x\}$
$X \rightarrow Y \rightarrow Z$	X, Y, Z form a Markov chain, i.e. $p(z x, y) = p(z y)$ (equivalently $I(X; Z Y) = 0$)
$X \perp\!\!\!\perp Y$	X and Y are independent; $X \perp\!\!\!\perp Y Z$ denotes conditional independence given Z
$\mathcal{N}(\mu, K)$	Gaussian distribution with mean μ and covariance K
$p_e^{(n)}$	probability of error of a length- n code
$K \succeq 0$	K is a positive semidefinite matrix; $A \succeq B$ means $A - B \succeq 0$
$\mathfrak{C}_X[f]$	upper concave envelope of the function f with respect to the distribution of X
sup, inf	supremum, infimum
sub-additivity	a functional f on joint laws is <i>sub-additive</i> if $f^{(2)}(p_{X_1 X_2}) \leq f(p_{X_1}) + f(p_{X_2})$ for every admissible joint law; see the end of Chapter 1 for the role of sub-additivity in outer bounds and converses

1 Information measures and properties

Disclaimer: Unless stated explicitly, random variables in these notes are assumed to take values in a finite set. The main emphasis of this course is on the key ideas and concepts that underlie results in network information theory. Most of these ideas can be understood using finitely supported random variables, without getting bogged down in technicalities that arise in more general settings. An interested reader should be able to transfer these ideas to more general spaces while being mindful of possible pitfalls. Some of these pitfalls are highlighted along the way.

1.1 Heuristics

Entropy measures the uncertainty, or randomness, of a random variable. In other words, the entropy of a random variable quantifies the amount of information it carries, or equivalently, the amount of information revealed when its value is observed.

Entropy was first introduced as a parameter of thermodynamic systems. Boltzmann realised that entropy can be interpreted, in modern terms, as the amount of information lost in the abstraction from microstates to macrostates. Shannon was the first to study entropy as the amount of information needed to specify the realisation of a general random variable.

Consider a physical system consisting of a large number n of particles, each of which is in some state in $\mathcal{X}$. We distinguish:

- **Microstate:** The states $x_1, \dots, x_n \in \mathcal{X}$ of the constituent particles.
- **Macrostate:** The empirical distribution

$$p(x) \triangleq \frac{1}{n} |\{i \in [1 : n] : x_i = x\}|$$

of the states of the constituent particles.

The physical system is completely described by its microstate, which consists of microscopic parameters such as the positions and velocities of the particles. In practice, however, we typically observe only macroscopic quantities such as pressure, volume, and temperature, which depend only on the macrostate of the system.

We now count the number of microstates consistent with a given macrostate whose empirical distribution is $p(x)$. Recall Stirling's approximation for factorials:

$$n! \sim \sqrt{2\pi n} \left(\frac{n}{e}\right)^n.$$

Write $S_p \triangleq \{x \in \mathcal{X} : p(x) > 0\}$ for the support of p and $m \triangleq |S_p|$. Then the number of microstates corresponding to the empirical distribution $p(x)$ is

$$\begin{aligned} \frac{n!}{\prod_{x \in S_p} (np(x))!} &\sim \frac{\sqrt{2\pi n} \left(\frac{n}{e}\right)^n}{\prod_{x \in S_p} \sqrt{2\pi np(x)} \left(\frac{np(x)}{e}\right)^{np(x)}} \\ &= (2\pi n)^{\frac{1-m}{2}} \left(\prod_{x \in S_p} p(x)\right)^{-\frac{1}{2}} \cdot 2^{-n \sum_{x \in S_p} p(x) \log p(x)} \\ &= 2^{-n \sum_{x \in S_p} p(x) \log p(x) + o(n)} \\ &= 2^{nH(X) + o(n)}, \end{aligned}$$

where we have set $H(X) \triangleq -\sum_{x \in \mathcal{X}} p(x) \log p(x)$ (terms with $p(x) = 0$ contribute zero). The Stirling prefactor is restricted to the support S_p : for $x \notin S_p$ one has $(np(x))! = 0! = 1$, so those factors simply do not appear in the denominator, while the formal substitution $\sqrt{2\pi np(x)}$ with $p(x) = 0$ would be invalid. The ratio asymptotics displayed here is valid for a fixed type p with positive support (all $np(x)$ on the support tending to infinity); for arbitrary varying types the uniform polynomial type-count bound of the appendix replaces this nonuniform Stirling argument. The quantity $H(X)$ is called the entropy of the distribution $p(x)$.

This leads to the following interpretation: entropy is a measure of uncertainty. A macrostate with entropy $H(X)$ corresponds to approximately $2^{nH(X)}$ indistinguishable microstates. Thus, observing a macrostate with higher entropy leaves us more uncertain about the underlying microstate than observing a macrostate with lower entropy.

We now explain why information entropy and thermodynamic entropy coincide in this setting. To avoid base-conversion constants, throughout this derivation we use natural logarithms and write $H_{\text{nat}}(p) = -\sum_x p(x) \ln p(x)$; the conversion to bits is $H_2 = H_{\text{nat}}/\ln 2$. Suppose that a single-particle state x has energy $E(x)$, and that the internal energy *per particle*, that is, the average energy of a particle, is

$$\begin{aligned} U &\triangleq \mathbb{E}[E(X)] \\ &= \sum_{x \in \mathcal{X}} p(x) E(x). \end{aligned}$$

(Here x is a single-particle state, not a complete system microstate $(x_1, \dots, x_n)$: the total energy of the n -particle system is nU , and the leading term of the total thermodynamic entropy is $n k_B H_{\text{nat}}(p)$.) Statistical mechanics postulates that, in thermal equilibrium, all microstates of the whole system with the same total energy are equally likely. For a physical system with internal energy per particle U in thermal equilibrium, an observer is most likely to see a macrostate that contains the largest number of compatible microstates, that is, a macrostate corresponding to the maximiser of

$$\max_{\substack{p(x) \\ \mathbb{E}[E(X)] = U}} H_{\text{nat}}(X).$$

To solve this, consider

$$\begin{aligned}
& \max_{\substack{p_i \geq 0 \\ \sum_i p_i = 1 \\ \sum_i p_i E_i = U}} - \sum_i p_i \ln p_i \\
&= \sup_{p_i \geq 0} \inf_{\lambda, \mu} \left(- \sum_i p_i \ln p_i + \lambda \left(\sum_i p_i - 1 \right) + \mu \left(\sum_i p_i E_i - U \right) \right) \\
&\leq \inf_{\lambda, \mu} \sup_{p_i \geq 0} \left(- \sum_i p_i \ln p_i + \lambda \left(\sum_i p_i - 1 \right) + \mu \left(\sum_i p_i E_i - U \right) \right) \\
&= \inf_{\lambda, \mu} \left(\sup_{p_i \geq 0} \sum_i (-p_i \ln p_i + p_i(\lambda + \mu E_i)) - \lambda - \mu U \right) \\
&= \inf_{\lambda, \mu} \left(\sum_i e^{-1+\lambda+\mu E_i} - \lambda - \mu U \right) \quad (\text{maximised by } p_i = e^{-1+\lambda+\mu E_i}) \\
&= \inf_{\mu} \left(\ln \sum_i e^{\mu E_i} - \mu U \right) \quad (\text{minimised by } \lambda = 1 - \ln \sum_i e^{\mu E_i}) \\
&= \ln \sum_i e^{\mu^* E_i} - \mu^* U \\
&= - \sum_i p_i^* \ln p_i^*,
\end{aligned}$$

where μ^* is the unique root of $\mu \mapsto \frac{\sum_i E_i e^{\mu E_i}}{\sum_i e^{\mu E_i}} - U$, and $p_i^* = \frac{e^{\mu^* E_i}}{\sum_j e^{\mu^* E_j}}$. Therefore,

$$\max_{\substack{p(x) \\ E[E(X)] = U}} H_{\text{nat}}(X)$$

is attained by

$$p(x) = \frac{e^{-\beta E(x)}}{\sum_{\tilde{x} \in \mathcal{X}} e^{-\beta E(\tilde{x})}},$$

where $\beta = -\mu^*$ satisfies $E[E(X)] = U$. This distribution is called the *Boltzmann distribution*. It is customary to write $\beta = \frac{1}{k_B T}$, where k_B is the Boltzmann constant, T is the temperature, and β is sometimes called the inverse temperature or coldness. The derivation above assumes that the energies E_i are not all equal and that the target U satisfies $\min_i E_i < U < \max_i E_i$; under these hypotheses the minimiser μ^* (equivalently the finite Gibbs parameter β) exists and is unique, because the derivative of $\mu \mapsto \ln \sum_i e^{\mu E_i} - \mu U$ is the tilted mean energy *minus* U , namely $\mu \mapsto \frac{\sum_i E_i e^{\mu E_i}}{\sum_i e^{\mu E_i}} - U$, which is strictly increasing in μ (its derivative is the variance of E under the tilted law, positive because the energies are not all equal) and hence has a unique root. The endpoint cases $U = \min_i E_i$ or $U = \max_i E_i$ are obtained as limiting Gibbs laws, uniform on the corresponding extremal energy level, and values of U outside $[\min_i E_i, \max_i E_i]$ are infeasible.

Under an equality mean-energy constraint β may have either sign (an inequality constraint $E[E(X)] \leq U$ would force $\beta \geq 0$).

One can verify that

$$\frac{dH_{\text{nat}}(X)}{d\beta} = \beta \frac{dU}{d\beta},$$

which in bits reads $\frac{dH_2(X)}{d\beta} = \frac{\beta}{\ln 2} \frac{dU}{d\beta}$. In terms of physical quantities,

$$\frac{dS}{dU} = \frac{1}{T},$$

where $S \triangleq k_B H_{\text{nat}}(X) = k_B \ln 2 H_2(X)$ is the thermodynamic entropy per particle. In other words, we have recovered the classical relation between entropy S , internal energy U and temperature T .

1.2 Axiomatic formulation of entropy

Let $\vec{p} = (p_1, \dots, p_n)$ be a probability vector. We seek a function $H(\vec{p})$ satisfying:

1. **Nonnegativity:** $H(\vec{p}) \geq 0$ for every probability vector $\vec{p}$.
2. **Continuity:** H is continuous in $\vec{p}$ on each probability simplex.
3. **Normalisation:** $H([\frac{1}{2}, \frac{1}{2}]) = 1$.
4. **Grouping:**

$$H([p_1, p_2, \dots, p_n]) = H([p_1 + p_2, p_3, \dots, p_n]) + (p_1 + p_2) H\left(\left[\frac{p_1}{p_1 + p_2}, \frac{p_2}{p_1 + p_2}\right]\right),$$

where the weighted term on the right is declared to be zero when $p_1 + p_2 = 0$ (so that its H -argument is not evaluated).

5. **Symmetry:**

$$H([p_1, \dots, p_n]) = H([p_{\sigma(1)}, \dots, p_{\sigma(n)}])$$

for any permutation σ of $[1 : n]$.

These are axioms of the usual Faddeev type, with nonnegativity stated explicitly; the binary normalisation fixes the base of the logarithm to 2. Nonnegativity is a property of each candidate entropy, whereas the difference of two candidate entropies can change sign; the functional equation derived below controls that signed difference.

One can verify that the Shannon entropy

$$H([\vec{p}]) = - \sum_i p_i \log p_i$$

satisfies these axioms. We now give a simple argument that any such H is unique.

Define the binary entropy function

$$H_b(x) \triangleq H([x, 1-x]), \quad x \in [0, 1].$$

By repeated use of the grouping axiom, H on general probability vectors is uniquely determined by H_b , so it suffices to characterise H_b .

Lemma 1.1. $H_b(x) = H_b(1-x)$ for all $x \in [0, 1]$. ○

Proof. By symmetry, $H([x, 1-x]) = H([1-x, x])$, i.e. $H_b(x) = H_b(1-x)$. ■

Lemma 1.2. $H_b(0) = H_b(1) = 0$. ○

Proof. Apply the grouping axiom to $[\frac{1}{2}, \frac{1}{2}, 0]$ in two ways:

$$H([\frac{1}{2}, \frac{1}{2}, 0]) = H([1, 0]) + H([\frac{1}{2}, \frac{1}{2}])$$

(by grouping the first two coordinates) and also

$$H([\frac{1}{2}, \frac{1}{2}, 0]) = H([\frac{1}{2}, \frac{1}{2}]) + \frac{1}{2} H([1, 0])$$

(by grouping the last two coordinates). Subtracting gives

$$H([1, 0]) = \frac{1}{2} H([1, 0]),$$

so $H([1, 0]) = 0$. Thus $H_b(0) = H([0, 1]) = H([1, 0]) = 0$, and by symmetry also $H_b(1) = 0$. ■

Next, consider the ternary probability vector

$$\vec{p} = [xy, x-xy, 1-x], \quad x, y \in [0, 1].$$

Applying the grouping axiom in two different ways yields a functional identity for H_b .

Grouping the first two coordinates first gives

$$H(\vec{p}) = H([x, 1-x]) + x H([y, 1-y]) = H_b(x) + x H_b(y).$$

On the other hand, grouping the last two first gives

$$\begin{aligned} H(\vec{p}) &= H([1-xy, xy]) + (1-xy) H\left(\left[\frac{x-xy}{1-xy}, \frac{1-x}{1-xy}\right]\right) \\ &= H_b(1-xy) + (1-xy) H_b\left(\frac{x-xy}{1-xy}\right). \end{aligned}$$

Equating the two expressions for $H(\vec{p})$, and using symmetry of H_b , we obtain for all $x, y \in [0, 1]$ with $xy < 1$:

$$H_b(x) + x H_b(y) = H_b(1-xy) + (1-xy) H_b\left(\frac{x-xy}{1-xy}\right). \quad (1.1)$$

The restriction $xy < 1$ is only to keep the argument $(x - xy)/(1 - xy)$ well defined: the single value $x = 1$ is not a singular case unless also $y = 1$ (at $x = 1, y < 1$ the equation reduces to the symmetry identity $H_b(y) = H_b(1 - y)$), and all specialisations used below are regular.

We now show that the binary entropy function solving (1.1), together with symmetry and normalisation, is unique.

Lemma 1.3. (Zero exclusion for the binary functional equation) Let $g : [0, 1] \rightarrow \mathbb{R}$ be continuous, with $g(0) = g(1) = 0$ and $g(x) = g(1 - x)$ for all $x \in [0, 1]$. Suppose

$$g(x) + xg(y) = g(1 - xy) + (1 - xy)g\left(\frac{x - xy}{1 - xy}\right)$$

for all $x, y \in [0, 1]$ with $xy < 1$. Then either $g \equiv 0$ on $[0, 1]$, or g is strictly positive on $(0, 1)$, or g is strictly negative on $(0, 1)$. In particular, a single interior zero of g — or a sequence of positive zeros accumulating at 0 — forces $g \equiv 0$. $\circ$

Proof. By symmetry of g , for $0 < x < 1$ the functional equation is equivalent to

$$g(x) + xg(y) = g(xy) + (1 - xy)g\left(\frac{x(1 - y)}{1 - xy}\right), \quad xy < 1.$$

Set $C = \int_0^1 g(y) dy$ and integrate both sides with respect to $y \in [0, 1]$. The first term on the right integrates to $x^{-1} \int_0^x g(u) du$. For the second, substitute

$$u = \frac{x(1 - y)}{1 - xy}, \quad y = \frac{x - u}{x(1 - u)}, \quad 1 - xy = \frac{1 - x}{1 - u}, \quad dy = -\frac{1 - x}{x(1 - u)^2} du,$$

so that as y runs over $[0, 1]$, u runs over $[x, 0]$. Thus

$$xg(x) + x^2C = \int_0^x g(u) du + (1 - x)^2 \int_0^x \frac{g(u)}{(1 - u)^3} du. \quad (1.2)$$

The right-hand side of (1.2) is continuously differentiable in $x \in (0, 1)$, so g is continuously differentiable there. Writing $B(x) = \int_0^x \frac{g(u)}{(1 - u)^3} du$ and differentiating (1.2) gives, after cancellation,

$$xg'(x) + 2xC = -2(1 - x)B(x) + \frac{g(x)}{1 - x}. \quad (1.3)$$

This identity in turn makes g twice continuously differentiable on $(0, 1)$. Differentiating (1.3) again yields

$$xg''(x) + g'(x) + 2C = 2B(x) + \frac{g'(x)}{1 - x} - \frac{g(x)}{(1 - x)^2}.$$

From (1.3),

$$2B(x) = \frac{g(x)}{(1 - x)^2} - \frac{xg'(x)}{1 - x} - \frac{2xC}{1 - x}.$$

Substituting this into the previous display and cancelling gives

$$g''(x) = -\frac{2C}{x(1-x)}, \quad x \in (0, 1).$$

If $C > 0$, then g is strictly concave on $(0, 1)$ and its zero endpoint values imply $g > 0$ on $(0, 1)$. If $C < 0$, g is strictly convex there and $g < 0$ on $(0, 1)$. If $C = 0$, then g is affine on $(0, 1)$ and continuity at the endpoints forces $g \equiv 0$. The final assertion follows because a single interior zero, or zeros accumulating at 0, rules out the two strict-sign alternatives and hence forces $g \equiv 0$. ■

Remark 1.1. Integrating $g''(x) = -2C/(x(1-x))$ explicitly gives $g(x) = 2C[-x \ln x - (1-x) \ln(1-x)] + Ax + B$; continuity and the endpoint conditions force $A = B = 0$. This explicit classification is not needed for the zero-exclusion statement. ○

We now conclude uniqueness. Let $H^{(1)}$ and $H^{(2)}$ both satisfy the axioms above, and set

$$f_b(x) \triangleq H^{(1)}([x, 1-x]) - H^{(2)}([x, 1-x]).$$

Then f_b is continuous, symmetric, and satisfies $f_b(0) = f_b(1) = 0$ (the grouping argument that gave $H_b(0) = H_b(1) = 0$ applies verbatim to each of $H^{(1)}$ and $H^{(2)}$). Both $H^{(1)}$ and $H^{(2)}$ obey the functional equation (1.1), so their difference f_b satisfies all the hypotheses of Lemma 1.3; the normalisation axiom gives the interior zero $f_b(\frac{1}{2}) = 0$, which forces $f_b \equiv 0$ on $[0, 1]$. Thus the two binary entropy functions coincide, and by repeated use of the grouping axiom the two functionals $H^{(1)}$ and $H^{(2)}$ agree on every finite probability vector.

Therefore, the Shannon entropy is the unique (up to the chosen normalisation) functional satisfying nonnegativity, continuity, symmetry, and the grouping axiom.

1.3 Properties of entropy

Definition 1.4. (Entropy) The entropy of a random variable X on $\mathcal{X}$ with distribution $p(x)$ is

$$H(X) \triangleq - \sum_{x \in \mathcal{X}} p(x) \log p(x).$$

○

Remark 1.2. Unless stated otherwise, the binary logarithm $\log_2$ is used in the definition of entropy, and entropy is then measured in *bits*. If the natural logarithm $\ln$ is used instead, entropy is measured in *nats*. One bit equals $\ln 2$ nats. The base of the logarithm will be clear from context. ○

Remark 1.3. We adopt the convention that $x \log x = 0$ when $x = 0$. ○

Theorem 1.5. (Jensen's inequality) Let X be a (possibly infinitely supported) random vector on $\mathbb{R}^d$. Let f be a real-valued convex function on some convex subset of $\mathbb{R}^d$. Suppose that $E\|X\| < \infty$, that X lies in the domain of f almost surely, and that $E[X]$ lies in the interior of the domain of f . Then

$$E[f(X)] \geq f(E[X]).$$

Moreover, if f is strictly convex, then equality holds if and only if X is almost surely constant. $\circ$

Proof. By convexity, f has a subgradient at $E[X]$, i.e. there exists some $a \in \mathbb{R}^d$ such that $f(x) - f(E[X]) \geq \langle a, x - E[X] \rangle$ for all x . Then

$$\begin{aligned} E[f(X)] &\geq E[f(E[X]) + \langle a, X - E[X] \rangle] \\ &= f(E[X]). \end{aligned}$$

The hypothesis $E\|X\| < \infty$ makes the linear term integrable, and the inequality is understood in the extended sense: it permits $E[f(X)] = +\infty$, while $f(E[X])$ is finite because $E[X]$ lies in the interior of the domain of f . Equality holds if and only if the subgradient inequality is tight almost surely, i.e. $f(X) - f(E[X]) = \langle a, X - E[X] \rangle$ almost surely. For a strictly convex f , a supporting hyperplane touches the graph at a single point, so this forces $X = E[X]$ almost surely, i.e. X is almost surely constant. $\blacksquare$

Proposition 1.6. $H(X) \geq 0$, with equality if and only if X is almost surely constant. $\circ$

Proof. From $0 \leq p(x) \leq 1$ we have $-p(x) \log p(x) \geq 0$. Hence

$$H(X) = - \sum_{x \in \mathcal{X}} p(x) \log p(x) \geq 0.$$

Equality holds if and only if $-p(x) \log p(x) = 0$ for all $x \in \mathcal{X}$, i.e. $p(x) = 0$ or $p(x) = 1$ for all $x \in \mathcal{X}$, so X is almost surely constant. $\blacksquare$

Proposition 1.7. Assume $\mathcal{X}$ is finite. Then, $H(X) \leq \log |\mathcal{X}|$, with equality if and only if X is uniformly distributed on $\mathcal{X}$. $\circ$

Proof. Let $S = \{x \in \mathcal{X} : p(x) > 0\}$ be the support of p . Applying Jensen's inequality on S ,

$$\begin{aligned} H(X) &= - \sum_{x \in S} p(x) \log p(x) \\ &= \sum_{x \in S} p(x) \log \frac{1}{p(x)} \\ &\leq \log \left(\sum_{x \in S} p(x) \frac{1}{p(x)} \right) \quad (\text{Jensen's inequality}) \end{aligned}$$

$$= \log |S| \leq \log |\mathcal{X}|.$$

Equality in the final inequality $|S| \leq |\mathcal{X}|$ requires $S = \mathcal{X}$, and equality in Jensen's step requires $\frac{1}{p(x)}$ to be constant in $x \in S$, i.e. $p(x) = \frac{1}{|\mathcal{X}|}$ for all x , so X is uniform on $\mathcal{X}$. ■

Definition 1.8. (Conditional entropy) Let X be a random variable on $\mathcal{X}$ and Y a random variable on $\mathcal{Y}$. The conditional entropy of X given Y is

$$\begin{aligned} H(X|Y) &\triangleq \sum_{y \in \mathcal{Y}} p(y) H(X|Y=y) \\ &= - \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) \log \frac{p(x, y)}{p(y)}, \end{aligned}$$

where for each y with $p(y) > 0$, $H(X|Y=y)$ is the entropy of the conditional distribution of X given $Y=y$, and the term $p(y)H(X|Y=y)$ contributes zero when $p(y) = 0$. In the double sum, the summand with $p(x, y) = 0$ is read as zero by the convention $x \log x = 0$ at $x = 0$. ○

Proposition 1.9. $H(X) \geq H(X|Y)$, with equality if and only if X and Y are independent. ○

Proof. Consider

$$\begin{aligned} H(X) - H(X|Y) &= - \sum_{x \in \mathcal{X}} p(x) \log p(x) + \sum_{x, y} p(x, y) \log \frac{p(x, y)}{p(y)} \\ &= - \sum_{x, y} p(x, y) \log p(x) + \sum_{x, y} p(x, y) \log \frac{p(x, y)}{p(y)} \\ &= \sum_{x, y} p(x) p(y) \frac{p(x, y)}{p(x) p(y)} \log \frac{p(x, y)}{p(x) p(y)} \\ &\geq \left(\sum_{x, y} p(x) p(y) \frac{p(x, y)}{p(x) p(y)} \right) \log \left(\sum_{x, y} p(x) p(y) \frac{p(x, y)}{p(x) p(y)} \right) \\ &= 0, \end{aligned}$$

where Jensen's inequality is applied to the convex function $u \mapsto u \log u$. Equality holds if and only if the ratio $\frac{p(x, y)}{p(x) p(y)}$ is constant on the support of the product measure, i.e. $\frac{p(x, y)}{p(x) p(y)} = c$ for all (x, y) with $p(x) p(y) > 0$; summing both sides of $p(x, y) = c p(x) p(y)$ over all (x, y) gives $1 = c \cdot 1$, so $c = 1$ and $p(x, y) = p(x) p(y)$ for all x, y , i.e. X and Y are independent. (Constancy merely on the joint support $\{(x, y) : p(x, y) > 0\}$ would not suffice: for $X = Y$ uniform on $\{0, 1\}$ the ratio equals 2 on the joint support, yet X and Y are not independent.) ■

Proposition 1.10. (Chain rule) $H(X, Y) = H(X) + H(Y|X)$. ○

Proof. We have

$$H(X, Y) = - \sum_{x, y} p(x, y) \log p(x, y)$$

$$\begin{aligned}
&= - \sum_{x,y} p(x,y) \log p(x) - \sum_{x,y} p(x,y) \log \frac{p(x,y)}{p(x)} \\
&= - \sum_x p(x) \log p(x) - \sum_{x,y} p(x,y) \log \frac{p(x,y)}{p(x)} \\
&= H(X) + H(Y|X).
\end{aligned}$$

■

Proposition 1.11. (Concavity of entropy) $H(X)$ is a concave function of the distribution $p(x)$. ○

Proof. For $\lambda \in [0, 1]$ let $p = \lambda p_1 + (1 - \lambda)p_2$, and let $Q \in \{1, 2\}$ be a random variable with $\Pr(Q = 1) = \lambda$. Let X be a random variable whose conditional law given $Q = i$ is p_i ; then X has marginal law p . Since conditioning reduces entropy,

$$H(X) \geq H(X|Q) = \lambda H(X|Q = 1) + (1 - \lambda)H(X|Q = 2) = \lambda H(p_1) + (1 - \lambda)H(p_2),$$

so H is concave in its argument. ■

1.4 Properties of Kullback–Leibler divergence

Definition 1.12. (KL divergence) Let p, q be (possibly infinitely supported) distributions on a common measurable space. If p is absolutely continuous with respect to q , i.e. for any event E , $\int_E dq = 0$ implies $\int_E dp = 0$, the Kullback–Leibler (KL) divergence, also called the relative entropy, of p with respect to q is

$$D(p||q) \triangleq \int \log \frac{dp}{dq} dp,$$

where $\frac{dp}{dq}$ is the Radon–Nikodym derivative of p with respect to q . If p is not absolutely continuous with respect to q , we define $D(p||q) \triangleq +\infty$. ○

Remark 1.4. In the finite discrete case where $p(x)$ and $q(x)$ are distributions on a finite set $\mathcal{X}$, absolute continuity of p with respect to q reduces to the condition that $q(x) = 0$ implies $p(x) = 0$. In this case,

$$D(p(x)||q(x)) = \sum_{x \in \mathcal{X}} p(x) \log \frac{p(x)}{q(x)},$$

with the convention $p(x) \log \frac{p(x)}{q(x)} = 0$ when $p(x) = 0$ and $+\infty$ when $q(x) = 0 < p(x)$. Unlike Shannon entropy on a finite alphabet — which does not have a canonical generalisation to continuous spaces, since its continuous analogue (differential entropy, defined below) depends on a reference measure and is not coordinate-invariant — discrete entropy itself extends immediately

to countable alphabets, possibly with value $+\infty$, and KL divergence extends naturally to general distributions, with absolute continuity (or $D = +\infty$) as the only requirement. KL divergence is related to entropy as follows: if $p(x)$ is a distribution on a finite set $\mathcal{X}$ and $\text{Unif}(\mathcal{X})$ is the uniform distribution on $\mathcal{X}$, then

$$D(p \parallel \text{Unif}(\mathcal{X})) = \log |\mathcal{X}| - H(X).$$

Thus $D(p \parallel q)$ should be interpreted as a measure of how far p deviates from q , rather than as an amount of “information” in the same sense as entropy. $\circ$

Proposition 1.13. $D(p \parallel q) \geq 0$, with equality if and only if $p = q$ almost everywhere. $\circ$

Proof. If p is not absolutely continuous with respect to q , the conclusion is immediate since $D(p \parallel q) = +\infty$. Otherwise $p \ll q$, and we compute

$$\begin{aligned} D(p \parallel q) &= \int \log \frac{dp}{dq} dp \\ &= \int \frac{dp}{dq} \log \frac{dp}{dq} dq \\ &\geq \left(\int \frac{dp}{dq} dq \right) \log \left(\int \frac{dp}{dq} dq \right) \quad (\text{Jensen's inequality}) \\ &= \left(\int dp \right) \log \left(\int dp \right) \\ &= 1 \cdot \log 1 \\ &= 0. \end{aligned}$$

Since $x \mapsto x \log x$ is strictly convex, equality holds if and only if $\frac{dp}{dq} = 1$ q -almost surely, i.e. $p = q$ almost everywhere. $\blacksquare$

Lemma 1.14. (Data-processing inequality for KL divergence) Let $W_{Y|X}$ be a linear operator that maps distributions on $\mathcal{X}$ to distributions on $\mathcal{Y}$ (assume both alphabets are finite). Given two distributions P_X and Q_X on $\mathcal{X}$, let WP and WQ denote their induced distributions on $\mathcal{Y}$. Then

$$D(WP \parallel WQ) \leq D(P \parallel Q).$$

$\circ$

Proof. Consider the joint distributions $P_{XY} = P_X W_{Y|X}$ and $Q_{XY} = Q_X W_{Y|X}$. We can also express these as $P_{XY} = (WP)_Y W_{X|Y}^{(a)}$ and $Q_{XY} = (WQ)_Y W_{X|Y}^{(b)}$ for some conditional distributions $W_{X|Y}^{(a)}$ and $W_{X|Y}^{(b)}$. Then

$$D(P_X \parallel Q_X) = \sum_x P_X(x) \log \frac{P_X(x)}{Q_X(x)}$$

$$\begin{aligned}
&= \sum_{x,y} P_X(x) W_{Y|X}(y|x) \log \frac{P_X(x) W_{Y|X}(y|x)}{Q_X(x) W_{Y|X}(y|x)} \\
&= \sum_{x,y} (WP)_Y(y) W_{X|Y}^{(a)}(x|y) \log \frac{(WP)_Y(y) W_{X|Y}^{(a)}(x|y)}{(WQ)_Y(y) W_{X|Y}^{(b)}(x|y)} \\
&= D(WP \| WQ) + \sum_y (WP)_Y(y) D(W_{X|Y=y}^{(a)} \| W_{X|Y=y}^{(b)}) \\
&\geq D(WP \| WQ),
\end{aligned}$$

by non-negativity of KL divergence. ■

Now, we state Pinsker's inequality, which relates the KL divergence and total variation distance between two distributions. Before stating the general case, we prove it for the special case where the random variables have a Bernoulli distribution. The general case can be reduced to this two-point case using the data-processing inequality.

Proposition 1.15. (Two-point Pinsker's inequality) For any $0 \leq p \leq 1$ and $0 \leq q \leq 1$,

$$D_2(p \| q) \geq \frac{2}{\ln 2} (p - q)^2,$$

where

$$D_2(p \| q) \triangleq p \log \frac{p}{q} + (1 - p) \log \frac{1 - p}{1 - q}$$

is the KL divergence between the Bernoulli distributions with parameters p and q . ○

Proof. Define

$$f(p, q) \triangleq D_2(p \| q) - \frac{2}{\ln 2} (p - q)^2.$$

Then

$$\begin{aligned}
\frac{\partial f(p, q)}{\partial q} &= \frac{1}{\ln 2} \left(-\frac{p}{q} + \frac{1 - p}{1 - q} + 4(p - q) \right) \\
&= -\frac{p - q}{\ln 2} \left(\frac{1}{q(1 - q)} - 4 \right) \\
&= -\frac{p - q}{\ln 2} \cdot \frac{(1 - 2q)^2}{q(1 - q)},
\end{aligned}$$

so

$$\frac{\partial f(p, q)}{\partial q} \begin{cases} \leq 0, & \text{if } q < p, \\ \geq 0, & \text{if } q > p, \\ = 0, & \text{if } q = p, \end{cases}$$

for $q \in (0, 1)$; note that the derivative also vanishes at $q = \frac{1}{2}$ (where $(1 - 2q)^2 = 0$) even when $q \neq p$. Thus $q \mapsto f(p, q)$ is non-increasing on $[0, p]$ and non-decreasing on $[p, 1]$, so $f(p, q)$ is

minimised at $q = p$, and $f(p, q) \geq f(p, p) = 0$. The cases $p \in \{0, 1\}$ follow by continuity (limits from the interior), under the standard extended conventions for D_2 at boundary parameters.

Alternate proof: Let $p, q \in (0, 1)$, and write

$$p = \frac{1+x}{2}, \quad q = \frac{1+y}{2},$$

so that $x, y \in (-1, 1)$. Then

$$D_2\left(\frac{1+x}{2} \parallel \frac{1+y}{2}\right) = \frac{1}{\ln 2} \left(\frac{1+x}{2} \ln \frac{1+x}{1+y} + \frac{1-x}{2} \ln \frac{1-x}{1-y} \right).$$

Using the Taylor series

$$\ln \frac{1+u}{1+v} = \sum_{k \geq 1} (-1)^{k+1} \frac{u^k - v^k}{k}, \quad |u|, |v| < 1,$$

we obtain

$$\begin{aligned} D_2\left(\frac{1+x}{2} \parallel \frac{1+y}{2}\right) &= \frac{1}{\ln 2} \sum_{k \geq 1} \left(\frac{x^{2k}}{2k-1} - \frac{x^{2k}}{2k} - x \frac{y^{2k-1}}{2k-1} + \frac{y^{2k}}{2k} \right) \\ &= \frac{1}{\ln 2} \sum_{k \geq 1} \frac{x^{2k} + (2k-1)y^{2k} - 2kxy^{2k-1}}{2k(2k-1)}. \end{aligned}$$

For each $k \geq 1$, observe that

$$x^{2k} + (2k-1)y^{2k} - 2kxy^{2k-1} \geq 0$$

by convexity of $t \mapsto t^{2k}$: its tangent at y is a supporting line, $x^{2k} \geq y^{2k} + 2ky^{2k-1}(x-y)$ for all $x, y \in \mathbb{R}$, and rearranging gives exactly the displayed inequality (no AM–GM on signed variables is invoked). Hence every summand is non-negative. In particular, the series is bounded below by its first term ($k = 1$):

$$D_2\left(\frac{1+x}{2} \parallel \frac{1+y}{2}\right) \geq \frac{1}{\ln 2} \cdot \frac{(x-y)^2}{2} = \frac{(x-y)^2}{2 \ln 2}.$$

Since $x - y = 2(p - q)$, we have

$$\frac{(x-y)^2}{2 \ln 2} = \frac{4(p-q)^2}{2 \ln 2} = \frac{2}{\ln 2} (p-q)^2,$$

and therefore

$$D_2(p \parallel q) \geq \frac{2}{\ln 2} (p-q)^2.$$

■

Proposition 1.16. (Pinsker's inequality) Let p, q be (possibly infinitely supported) distributions on a common measurable space. Then

$$D(p\|q) \geq \frac{2}{\ln 2} \delta(p, q)^2,$$

where

$$\delta(p, q) \triangleq \sup_E \left| \int_E dp - \int_E dq \right|$$

is the total variation distance between p and q . ○

Proof. If p is not absolutely continuous with respect to q , the left-hand side is $+\infty$ and there is nothing to prove. So assume $p \ll q$ and write $r = dp/dq$. For a measurable set E , conditional Jensen's inequality for the convex function $u \mapsto u \log u$, applied on E and on E^c separately, gives

$$D(p\|q) = \int r \log r dq \geq p(E) \log \frac{p(E)}{q(E)} + p(E^c) \log \frac{p(E^c)}{q(E^c)},$$

with the usual extended conventions for the zero-probability terms. The right-hand side is exactly the binary divergence between the Bernoulli laws with parameters $p(E)$ and $q(E)$, so the two-point Pinsker inequality gives

$$D(p\|q) \geq \frac{2}{\ln 2} (p(E) - q(E))^2$$

for every measurable set E . Taking the supremum over E yields

$$D(p\|q) \geq \frac{2}{\ln 2} \delta(p, q)^2,$$

as claimed. (This conditional Jensen step is the binary data-processing inequality for the test E , proved here directly because Lemma 1.14 is stated only for finite alphabets; the general measurable-kernel version of the data-processing inequality admits an analogous Radon–Nikodym proof, and we use it only through this binary case.) ■

Proposition 1.17. (Joint convexity of relative entropy) For probability mass functions p_1, p_2, q_1, q_2 on a finite set X and $0 \leq \lambda \leq 1$, let

$$p := \lambda p_1 + (1 - \lambda) p_2, \quad q := \lambda q_1 + (1 - \lambda) q_2.$$

Then

$$D(p\|q) \leq \lambda D(p_1\|q_1) + (1 - \lambda) D(p_2\|q_2),$$

i.e. $D(\cdot\|\cdot)$ is jointly convex in (p, q) . ○

Proof. Recall that

$$D(p\|q) = \sum_{x \in X} q(x) f\left(\frac{p(x)}{q(x)}\right), \quad f(t) := t \log t,$$

so $D(p\|q)$ is the f -divergence associated with the convex function f . Assume first that $p_1, p_2, q_1, q_2 > 0$ pointwise. The general case follows by approximation: regularise each of the four distributions by mixing it with a common full-support distribution, say $p_i^{(\varepsilon)} := (1 - \varepsilon)p_i + \varepsilon \text{Unif}(X)$ and similarly for q_i , apply the strictly positive case to the regularised quadruple (for which the perspective $q f(p/q)$ is finite and continuous), and let $\varepsilon \downarrow 0$; along this specific regularisation the divergences converge under the standard extended conventions ($q(x) f(p(x)/q(x)) = 0$ when $p(x) = q(x) = 0$, and $+\infty$ when $q(x) = 0 < p(x)$). (The perspective $q f(p/q)$ is not finite and continuous at $(p, q) = (0, 0)$ along arbitrary paths, which is why the approximation is routed through a full-support regularisation.) For each x , define

$$\alpha_1(x) := \lambda \frac{q_1(x)}{q(x)}, \quad \alpha_2(x) := (1 - \lambda) \frac{q_2(x)}{q(x)}.$$

Then $\alpha_1(x), \alpha_2(x) \geq 0$ and $\alpha_1(x) + \alpha_2(x) = 1$, and

$$\frac{p(x)}{q(x)} = \frac{\lambda p_1(x) + (1 - \lambda)p_2(x)}{\lambda q_1(x) + (1 - \lambda)q_2(x)} = \alpha_1(x) \frac{p_1(x)}{q_1(x)} + \alpha_2(x) \frac{p_2(x)}{q_2(x)}.$$

By convexity of f ,

$$f\left(\frac{p(x)}{q(x)}\right) \leq \alpha_1(x) f\left(\frac{p_1(x)}{q_1(x)}\right) + \alpha_2(x) f\left(\frac{p_2(x)}{q_2(x)}\right).$$

Multiplying by $q(x)$ and summing over x yields

$$\begin{aligned} D(p\|q) &= \sum_x q(x) f\left(\frac{p(x)}{q(x)}\right) \\ &\leq \sum_x q(x) \left[\alpha_1(x) f\left(\frac{p_1(x)}{q_1(x)}\right) + \alpha_2(x) f\left(\frac{p_2(x)}{q_2(x)}\right) \right] \\ &= \lambda \sum_x q_1(x) f\left(\frac{p_1(x)}{q_1(x)}\right) + (1 - \lambda) \sum_x q_2(x) f\left(\frac{p_2(x)}{q_2(x)}\right) \\ &= \lambda D(p_1\|q_1) + (1 - \lambda) D(p_2\|q_2), \end{aligned}$$

which proves the claim. ■

1.5 Properties of differential entropy

Definition 1.18. (Differential entropy) Let $\mathbf{X}$ be a real random vector with density $f_{\mathbf{X}}$. The differential entropy of $\mathbf{X}$ is defined as

$$h(\mathbf{X}) \triangleq - \int f_{\mathbf{X}}(\mathbf{x}) \log f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x}.$$

○

Remark 1.5. Writing the density as $f = \frac{dP_X}{dm}$ with respect to Lebesgue measure m , the formal relative entropy of P_X to m is $\int f \log f \, dm$; differential entropy is its *negative*, $h(X) = -\int f \log f \, dm$, although m is not a probability measure (in particular, this formal “relative entropy” is not nonnegative). Differential entropy depends on the reference measure: it is not invariant under non-volume-preserving coordinate changes, it can be negative, and it may fail to exist for some distributions. ○

Proposition 1.19. $h(X + x_0) = h(X)$ for any deterministic vector x_0 . ○

Proof. The density of $X + x_0$ is $f_{X+x_0}(y) = f_X(y - x_0)$. Thus

$$\begin{aligned} h(X + x_0) &= - \int f_{X+x_0}(y) \log f_{X+x_0}(y) \, dy \\ &= - \int f_X(y - x_0) \log f_X(y - x_0) \, dy \\ &= - \int f_X(x) \log f_X(x) \, dx && (\text{with } x \triangleq y - x_0) \\ &= h(X). \end{aligned}$$

■

Proposition 1.20. $h(AX) = h(X) + \log |\det A|$ for any invertible real square matrix A . ○

Proof. The density of AX is $f_{AX}(y) = \frac{1}{|\det A|} f_X(A^{-1}y)$. Thus

$$\begin{aligned} h(AX) &= - \int f_{AX}(y) \log f_{AX}(y) \, dy \\ &= - \int \frac{1}{|\det A|} f_X(A^{-1}y) \log \left(\frac{1}{|\det A|} f_X(A^{-1}y) \right) \, dy \\ &= - \int \frac{1}{|\det A|} f_X(x) \log \left(\frac{1}{|\det A|} f_X(x) \right) |\det A| \, dx && (\text{with } x \triangleq A^{-1}y) \\ &= - \int f_X(x) \log f_X(x) \, dx + (\log |\det A|) \int f_X(x) \, dx \\ &= h(X) + \log |\det A|. \end{aligned}$$

The absolute determinant is used even when $\det A < 0$, since the Jacobian of the change of variables $y \mapsto A^{-1}y$ is $|\det A^{-1}| = |\det A|^{-1}$. ■

Definition 1.21. (Conditional differential entropy) Let (X, Y) be a real random vector with joint density $f_{X,Y}$. The conditional differential entropy of X given Y is

$$h(X|Y) \triangleq \int f_Y(y) h(X|Y = y) \, dy$$

$$= - \iint f_{X,Y}(\mathbf{x}, y) \log \frac{f_{X,Y}(\mathbf{x}, y)}{f_Y(y)} d\mathbf{x} dy,$$

where $h(X|Y = y)$ is the differential entropy of the conditional distribution of X given $Y = y$. $\circ$

Proposition 1.22. Let (X, Y) have a joint density and suppose $h(X)$ and $h(X|Y)$ are finite. Then $h(X) \geq h(X|Y)$, with equality if and only if X and Y are independent. $\circ$

Proof. Consider

$$\begin{aligned} h(X) - h(X|Y) &= - \int f_X(\mathbf{x}) \log f_X(\mathbf{x}) d\mathbf{x} \\ &\quad + \iint f_{X,Y}(\mathbf{x}, y) \log \frac{f_{X,Y}(\mathbf{x}, y)}{f_Y(y)} d\mathbf{x} dy \\ &= - \iint f_{X,Y}(\mathbf{x}, y) \log f_X(\mathbf{x}) d\mathbf{x} dy \\ &\quad + \iint f_{X,Y}(\mathbf{x}, y) \log \frac{f_{X,Y}(\mathbf{x}, y)}{f_Y(y)} d\mathbf{x} dy \\ &= \iint f_X(\mathbf{x}) f_Y(y) \frac{f_{X,Y}(\mathbf{x}, y)}{f_X(\mathbf{x}) f_Y(y)} \\ &\quad \cdot \log \frac{f_{X,Y}(\mathbf{x}, y)}{f_X(\mathbf{x}) f_Y(y)} d\mathbf{x} dy \\ &\geq \left(\iint f_X(\mathbf{x}) f_Y(y) \frac{f_{X,Y}(\mathbf{x}, y)}{f_X(\mathbf{x}) f_Y(y)} d\mathbf{x} dy \right) \\ &\quad \cdot \log \left(\iint f_X(\mathbf{x}) f_Y(y) \frac{f_{X,Y}(\mathbf{x}, y)}{f_X(\mathbf{x}) f_Y(y)} d\mathbf{x} dy \right) \\ &= 0, \end{aligned}$$

by Jensen's inequality, under the hypothesis that the differential entropies appearing are finite. Equality holds if and only if $\frac{f_{X,Y}(\mathbf{x}, y)}{f_X(\mathbf{x}) f_Y(y)}$ is constant almost everywhere *under the product density* $f_X(\mathbf{x}) f_Y(y)$, which is equivalent to $f_{X,Y}(\mathbf{x}, y) = f_X(\mathbf{x}) f_Y(y)$ almost everywhere, i.e. X and Y being independent. Constancy under the joint law alone does not characterise independence. $\blacksquare$

Proposition 1.23. (Chain rule) If X, Y have a joint density and the entropies $h(X, Y)$, $h(X)$, $h(Y|X)$ are all finite, then $h(X, Y) = h(X) + h(Y|X)$. $\circ$

Proof. We compute

$$\begin{aligned} h(X, Y) &= - \iint f_{X,Y}(\mathbf{x}, y) \log f_{X,Y}(\mathbf{x}, y) d\mathbf{x} dy \\ &= - \iint f_{X,Y}(\mathbf{x}, y) \log f_X(\mathbf{x}) d\mathbf{x} dy \\ &\quad - \iint f_{X,Y}(\mathbf{x}, y) \log \frac{f_{X,Y}(\mathbf{x}, y)}{f_X(\mathbf{x})} d\mathbf{x} dy \end{aligned}$$

$$\begin{aligned}
 &= - \int f_X(\mathbf{x}) \log f_X(\mathbf{x}) d\mathbf{x} \\
 &\quad - \iint f_{X,Y}(\mathbf{x}, y) \log \frac{f_{X,Y}(\mathbf{x}, y)}{f_X(\mathbf{x})} d\mathbf{x} dy \\
 &= h(X) + h(Y|X).
 \end{aligned}$$

■

Proposition 1.24. (Gaussian maximises differential entropy) Let $\mathbf{X}$ be a random vector in $\mathbb{R}^d$ with a density, zero mean, finite nonsingular covariance matrix $K_X = E[\mathbf{X}\mathbf{X}^T] \succ 0$, and finite differential entropy. Then

$$h(\mathbf{X}) \leq \frac{d}{2} \log(2\pi e) + \frac{1}{2} \log |K_X|,$$

with equality if and only if $\mathbf{X} \sim \mathcal{N}(0, K_X)$.

○

Remark 1.6. The zero-mean assumption is immaterial for entropy-maximisation problems: h is translation invariant, and under a second-moment upper bound $E[\mathbf{X}\mathbf{X}^T] \preceq K$ with $K \succ 0$, centering is admissible — if $\mathbf{X}$ has mean μ , then $\mathbf{X} - \mu$ has covariance $K_X = E[\mathbf{X}\mathbf{X}^T] - \mu\mu^T \preceq K$, the same entropy, and a second moment no larger. Thus the entropy-maximising law for the constraint $E[\mathbf{X}\mathbf{X}^T] \preceq K$ is the centered Gaussian with covariance K . For a *fixed* covariance constraint without a mean constraint, every mean gives the same maximum entropy (translation invariance).

○

Proof. Whitening reduces to the covariance- I case: applying the invertible affine map $\mathbf{x} \mapsto K_X^{-1/2} \mathbf{x}$ turns the constraint into covariance I , preserves the sign of the difference between the two sides by the scaling identity $h(A\mathbf{X}) = h(\mathbf{X}) + \log |\det A|$, and is a bijection between densities with covariance K_X and densities with covariance I (singular covariances are excluded precisely so that this whitening is valid). So assume $K_X = I$ and let

$$g(\mathbf{x}) \triangleq (2\pi)^{-\frac{d}{2}} e^{-\frac{1}{2} \|\mathbf{x}\|^2}$$

be the density of the d -dimensional standard Gaussian distribution. If $\mathbf{X}$ has density f , then, computing the binary logarithm of g term by term,

$$\log g(\mathbf{x}) = -\frac{d}{2} \log(2\pi) - \frac{\|\mathbf{x}\|^2}{2 \ln 2}.$$

Consequently,

$$\begin{aligned}
 D(f||g) &= \int f(\mathbf{x}) \log \frac{f(\mathbf{x})}{g(\mathbf{x})} d\mathbf{x} \\
 &= - \int f(\mathbf{x}) \log g(\mathbf{x}) d\mathbf{x} - h(\mathbf{X}) \\
 &= \frac{d}{2} \log(2\pi) + \frac{E\|\mathbf{X}\|^2}{2 \ln 2} - h(\mathbf{X})
 \end{aligned}$$

$$\begin{aligned}
&= \frac{d}{2} \log(2\pi) + \frac{d}{2 \ln 2} - h(\mathbf{X}) \\
&= \frac{d}{2} \log(2\pi e) - h(\mathbf{X}).
\end{aligned}$$

Here $E \|\mathbf{X}\|^2 = \text{tr}(I) = d$, and $\frac{d}{2 \ln 2} = \frac{d}{2} \log e$. Since $D(f\|g) \geq 0$, we obtain $h(\mathbf{X}) \leq \frac{d}{2} \log(2\pi e)$, with equality if and only if $f = g$, i.e. $\mathbf{X}$ is standard Gaussian. The change-of-variables identity gives the general covariance case. ■

1.6 Maximum entropy distributions

The entropy-maximisation problem of the heuristics (the Boltzmann distribution) admits a far-reaching generalisation: among all distributions on a set S with a prescribed average of some function f , the entropy maximiser is a one-parameter exponential family. This single principle underlies the geometric, Gaussian, exponential, and Boltzmann distributions, and its proof is a one-line application of the non-negativity of KL divergence.

Theorem 1.25. (Maximum entropy under a moment constraint) Let S be a measurable subset of $\mathbb{R}^d$, or a finite or countable set; in the discrete case the reference measure is the counting measure, and in the continuous case it is Lebesgue measure restricted to S . Let $f : S \rightarrow \mathbb{R}$ be measurable and $c \in \mathbb{R}$, and suppose there is a distribution q on S of the form

$$q(x) = a e^{-bf(x)}, \quad b \geq 0,$$

(with $a > 0$ the normalising constant) such that $E_q[f(X)] = c$ and q is supported in S . Then q maximises the entropy H (in the discrete case) or the differential entropy h (in the continuous case) among all admissible distributions p with support in S and $E_p[f(X)] \leq c$. It is the *unique* maximiser, provided all the expectations appearing in the proof are finite. For an *equality* moment constraint $E_p[f(X)] = c$ (rather than the inequality constraint), the multiplier need not satisfy $b \geq 0$. ○

Proof. We give the proof in the discrete case; the continuous case is identical with sums replaced by integrals and H by h . The conclusion is invariant under rescaling the entropy by a positive constant, so we work with natural logarithms. Let p be any admissible distribution with support in S and $E_p[f(X)] \leq c = E_q[f(X)]$ (admissibility includes finiteness of the expectations used below). Using $\ln q(x) = \ln a - bf(x)$,

$$\begin{aligned}
D(p\|q) &= \sum_{x \in S} p(x) \ln \frac{p(x)}{q(x)} \\
&= H(q) - H(p) + \sum_{x \in S} (q(x) - p(x)) \ln q(x) \\
&= H(q) - H(p) - b (E_q[f(X)] - E_p[f(X)]) \leq H(q) - H(p),
\end{aligned}$$

because $b \geq 0$ and $E_p[f(X)] \leq E_q[f(X)]$. Since $D(p\|q) \geq 0$, we obtain $H(p) \leq H(q)$, i.e. q is a maximiser. Moreover, if $H(p) = H(q)$ then the chain forces $D(p\|q) = 0$, hence $p = q$: the maximiser is unique. (For the equality constraint the same computation with b of either sign gives $H(q) - H(p) = D(p\|q)$, which is ≥ 0 ; no sign condition on b is needed.) ■

Example 1.26. (Geometric distribution) Let X take values in the non-negative integers $\{0, 1, 2, \dots\}$ with $E[X] \leq \lambda$ ($\lambda > 0$). The maximiser of $H(X)$ is the geometric distribution $q(x) = (1 - c)c^x$ with $c = \frac{\lambda}{1+\lambda}$ (so that $E_q[X] = \frac{c}{1-c} = \lambda$). Indeed $q(x) = (1 - c)e^{x \ln c}$ (note the natural logarithm: $e^{x \log_2 c}$ would be the wrong family), which is of the form $ae^{-bf(x)}$ with $f(x) = x$ and $b = -\ln c \geq 0$, so Theorem 1.25 applies. Its entropy is

$$H(\text{Geom}(\lambda)) = (1 + \lambda) \log(1 + \lambda) - \lambda \log \lambda.$$

Example 1.27. (Gaussian under a second-moment constraint) For a continuous random variable X with $E[X^2] \leq \sigma^2$, Theorem 1.25 with $f(x) = x^2$ reproduces the Gaussian maximiser $q(x) = ae^{-bx^2}$, in agreement with the Gaussian-maximality proposition of the differential-entropy section. The same optimiser can be guessed directly by calculus of variations, using the functional of the whole density

$$\mathcal{L}(p) = - \int p \ln p - \mu \left(\int p - 1 \right) - \lambda \left(\int x^2 p - \sigma^2 \right),$$

in which the normalisation and moment terms are integrals of p against the fixed functions 1 and x^2 , with the global constants $-\mu$ and $-\lambda\sigma^2$ merely shifting the functional. The interior Euler equation, obtained by varying p at each point, is

$$-\ln p - 1 - \mu - \lambda x^2 = 0,$$

whose solution is $p(x) = ae^{-bx^2}$, a Gaussian density.

Remark 1.7 (Infinite entropy). The uniform bound $H(X) \leq \log |X|$ only yields a finite maximum on a finite alphabet. On a countably infinite alphabet, entropy can be infinite: for the heavy-tailed distribution

$$p(k) = \frac{c}{(k+1) \ln^2(k+1)}, \quad k \geq 1$$

(where c is the normalising constant and $\ln$ is the natural logarithm), one has $H(X) = \infty$. Finiteness of the alphabet, or of a suitable moment constraint, is therefore often needed to guarantee a finite, attained maximum. ○

1.7 Properties of mutual information

Definition 1.28. (Mutual information) The mutual information of two (possibly infinitely supported) random variables X, Y is

$$I(X; Y) \triangleq D(P_{XY} \| P_X P_Y).$$

○

Proposition 1.29. $I(X; Y) \geq 0$, with equality if and only if X and Y are independent. ○

Proof. This follows from the non-negativity of KL divergence. Equality holds if and only if $p(x, y) = p(x)p(y)$, i.e. X and Y are independent. ■

Proposition 1.30. $I(X; Y) = I(Y; X)$. ○

Proof. This is immediate from the definition. ■

Proposition 1.31. If X, Y are finitely supported random variables, then

$$I(X; Y) = H(X) - H(X|Y).$$

○

Proof. We have

$$\begin{aligned} I(X; Y) &= \sum_{x,y} p(x, y) \log \frac{p(x, y)}{p(x)p(y)} \\ &= - \sum_{x,y} p(x, y) \log p(x) + \sum_{x,y} p(x, y) \log \frac{p(x, y)}{p(y)} \\ &= H(X) - H(X|Y). \end{aligned}$$

■

Proposition 1.32. If $\mathbf{X}, \mathbf{Y}$ are real random vectors with *finite* differential and conditional differential entropies, then

$$I(\mathbf{X}; \mathbf{Y}) = h(\mathbf{X}) - h(\mathbf{X}|\mathbf{Y}).$$

○

Proof. We have

$$\begin{aligned} I(\mathbf{X}; \mathbf{Y}) &= \iint f_{\mathbf{X}, \mathbf{Y}}(\mathbf{x}, \mathbf{y}) \log \frac{f_{\mathbf{X}, \mathbf{Y}}(\mathbf{x}, \mathbf{y})}{f_{\mathbf{X}}(\mathbf{x})f_{\mathbf{Y}}(\mathbf{y})} d\mathbf{x} d\mathbf{y} \\ &= - \iint f_{\mathbf{X}, \mathbf{Y}}(\mathbf{x}, \mathbf{y}) \log f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x} d\mathbf{y} \\ &\quad + \iint f_{\mathbf{X}, \mathbf{Y}}(\mathbf{x}, \mathbf{y}) \log \frac{f_{\mathbf{X}, \mathbf{Y}}(\mathbf{x}, \mathbf{y})}{f_{\mathbf{Y}}(\mathbf{y})} d\mathbf{x} d\mathbf{y} \\ &= h(\mathbf{X}) - h(\mathbf{X}|\mathbf{Y}). \end{aligned}$$

■

Definition 1.33. (Conditional mutual information) Let X, Y, Z be random variables taking values in standard Borel spaces. The conditional mutual information of X and Y given Z is defined as

$$\begin{aligned} I(X; Y|Z) &\triangleq \mathbb{E}_Z [I(X; Y|Z = z)] \\ &= \mathbb{E}_Z \left[D \left(P_{XY|Z=z} \| P_{X|Z=z} P_{Y|Z=z} \right) \right], \end{aligned}$$

where $I(X; Y|Z = z)$ is the mutual information between X and Y conditioned on $Z = z$, defined through a regular conditional law $P_{XY|Z=z}$ (regular conditional laws exist because all three variables take values in standard Borel spaces, in particular for the finite alphabets of these notes). $\circ$

Proposition 1.34. (Chain rule) $I(X, Y; Z) = I(Y; Z|X) + I(X; Z)$. $\circ$

Proof. For finitely supported (X, Y, Z) the computation is

$$\begin{aligned} I(X, Y; Z) &= \sum_{x,y,z} p(x, y, z) \log \frac{p(x, y, z)}{p(x, y)p(z)} \\ &= \sum_{x,y,z} p(x, y, z) \log \frac{p(y, z|x)}{p(y|x)p(z)} \\ &= \sum_{x,y,z} p(x, y, z) \log \frac{p(y, z|x)}{p(y|x)p(z|x)} + \sum_{x,y,z} p(x, y, z) \log \frac{p(z|x)}{p(z)} \\ &= \sum_x p(x) I(Y; Z|X = x) + I(X; Z) \\ &= I(Y; Z|X) + I(X; Z). \end{aligned}$$

For general random variables taking values in standard Borel spaces, the same identity is the multiplicative decomposition of Radon–Nikodym derivatives with respect to product measures, stated here in the absolute-continuity setting: if $P_{XYZ} \ll P_{XY} \otimes P_Z$, $P_{XZ} \ll P_X \otimes P_Z$, and $P_{YZ|X} \ll P_{Y|X} \otimes P_{Z|X}$ for P_X -almost every x (all of which hold under the finiteness hypotheses used here), then with $r = \frac{dP_{XYZ}}{d(P_{XY} \otimes P_Z)}$ and $r' = \frac{dP_{XZ}}{d(P_X \otimes P_Z)}$,

$$r(x, y, z) = \frac{dP_{YZ|X}}{d(P_{Y|X} \otimes P_{Z|X})}(y, z) r'(x, z)$$

for P_{XYZ} -almost every (x, y, z) (using regular conditional laws, which exist on standard Borel spaces), and taking the P_{XYZ} -expectation of log on both sides gives exactly $I(Y; Z|X) + I(X; Z)$. If some of the three quantities are infinite, the chain rule is a separately stated extended relative-entropy chain rule; the identity above is valid whenever the quantities involved are finite, and in the extended sense it holds provided no $+\infty$ is subtracted from $+\infty$. $\blacksquare$

Proposition 1.35. If $I(X; Z|Y)$ is finite, then $I(W, X; Z|Y) = I(X; Z|Y)$ if and only if $W \rightarrow (X, Y) \rightarrow Z$ is a Markov chain. $\circ$

Proof. By the chain rule,

$$I(W, X; Z|Y) = I(W; Z|X, Y) + I(X; Z|Y).$$

Since $I(X; Z|Y)$ is finite, the subtraction is legitimate, and $I(W, X; Z|Y) = I(X; Z|Y)$ holds if and only if $I(W; Z|X, Y) = 0$, which is equivalent to W and Z being conditionally independent given (X, Y) , i.e. $W \rightarrow (X, Y) \rightarrow Z$ is Markov. ■

Remark 1.8. The finiteness hypothesis excludes undefined extended-value arithmetic such as $+\infty - \infty$: if $I(X; Z|Y) = +\infty$ the claim can fail or become vacuous (both sides are then infinite for some non-Markov chains as well). For finitely supported random variables the hypothesis holds automatically, and this is the setting in which the proposition is used in these notes. ○

Proposition 1.36. (Data processing inequality) If $X \rightarrow Y \rightarrow Z$ forms a Markov chain, then $I(Y; Z) \geq I(X; Z)$. ○

Proof. Since $X \rightarrow Y \rightarrow Z$ is Markov,

$$I(Y; Z) = I(X, Y; Z).$$

By the chain rule,

$$I(X, Y; Z) = I(X; Z) + I(Y; Z|X) \geq I(X; Z).$$

■

Remark 1.9 (d-separation: reading Markov relations off a graph). Most converse steps in these notes need conditional independences of the form $A \perp\!\!\!\perp B \mid C$. They are read off a *directed acyclic graph* (DAG) by the rule of *d-separation*, provided the joint law factorises over the graph:

$$p(z_V) = \prod_{v \in V} p(z_v \mid z_{\text{pa}(v)}),$$

where $\text{pa}(v)$ denotes the parents of v . Both stochastic kernels and deterministic functions are represented by *directed parent edges* (deterministic nodes have degenerate conditionals), so a message $M = f(Y^n)$ is drawn with arrows $Y_j \rightarrow M$. Then $A \perp\!\!\!\perp B \mid C$ holds whenever every path between A and B is *blocked* by C : a path is blocked if an intermediate node on it belongs to C without being a collider, or if some collider on it has no conditioned descendant (neither itself nor any descendant lies in C); a collider is a node at which two path edges meet head-to-head. A precise statement and proof — importing the cited direction of the ancestral-moralisation equivalence — is given in §9.4 of the Appendix, together with the caution that the factorisation hypothesis must be verified for the actual joint law.

The clean application in these notes is the *coordinate star* of the one-helper converse (Chapter 7): orient the graph as $Y_j \rightarrow X_j$ and $Y_j \rightarrow M_2$ with the factorisation $\prod_j p(y_j)p(x_j \mid y_j)\mathbf{1}\{M_2 = f_2(y^n)\}$, under which every X_j is a leaf, so conditioning on Y^{i-1} separates X^{i-1} from everything else and conditioning on Y_i separates X_i from everything else:

$$X^{i-1} \perp\!\!\!\perp (M_2, Y_i^n, X_i) \mid Y^{i-1}, \quad X_i \perp\!\!\!\perp (M_2, Y^{i-1}, X^{i-1}) \mid Y_i$$

(Corollary 9.19 of the Appendix). Other converses have their own generative graphs: memoryless channels are drawn with directed edges from the current inputs into the current outputs, and causal systems such as the relay of Chapter 8 use their time-expanded generative graph, in which each node's parents are exactly the variables that determine its conditional law. Reading conditional independences off such graphs is the board-level way to verify these steps without chasing symbol-level computations. $\circ$

1.8 Properties of Fisher information

Definition 1.37. (Fisher information) Let $\mathbf{X}$ be a random vector in $\mathbb{R}^d$ with density $f_{\mathbf{X}}$. The Fisher information of $\mathbf{X}$ is

$$J(\mathbf{X}) \triangleq \mathbb{E} [\|\rho_{\mathbf{X}}(\mathbf{X})\|^2],$$

where $\rho_{\mathbf{X}} \triangleq \frac{\nabla f_{\mathbf{X}}}{f_{\mathbf{X}}} = \nabla \ln f_{\mathbf{X}}$ is the *score function* of $\mathbf{X}$. $\circ$

Remark 1.10 (Regularity conventions). For score-based computations we work initially with positive C^2 densities, finite Fisher information, and enough decay (or explicit cutoff arguments) to justify the integrations by parts used below; regularised extensions can then be cited. A precise general definition, when needed, is

$$J(f) = 4 \int_{\mathbb{R}^d} \|\nabla \sqrt{f}\|^2 dx$$

whenever $\sqrt{f} \in H^1(\mathbb{R}^d)$, and $J(f) = +\infty$ otherwise. A pointwise derivative of $\ln f$ in the interior of the support of f is not a substitute for this boundary-sensitive convention. $\circ$

Remark 1.11 (Convention for this section). Throughout the sections on Fisher information and the entropy power inequality we work with *natural* logarithms, so that entropy is measured in nats and the classical identities (de Bruijn, Stam) take their standard form with no conversion constants. At the end of the chapter the results are re-expressed in base 2, which is the default elsewhere in these notes; the entropy-power form of the (conditional) EPI is in fact independent of the base of the logarithm. The Fisher information $J(\mathbf{X})$ is *base-independent* as defined: the score function is fixed to be $\rho_{\mathbf{X}} = \nabla \ln f_{\mathbf{X}}$ with the natural logarithm, and changing the unit in which *entropies* are measured does not change J . (In bits, de Bruijn's identity carries a conversion constant, $\frac{d}{dt} h_2(\mathbf{X} + \sqrt{t}\mathbf{Z}) = \frac{1}{2 \ln 2} J(\mathbf{X} + \sqrt{t}\mathbf{Z})$.) $\circ$

Proposition 1.38. (de Bruijn's identity) Let $\mathbf{X}$ be a random vector in $\mathbb{R}^d$ with finite second moment, and let $\mathbf{Z}$ be a d -dimensional standard Gaussian random vector independent of $\mathbf{X}$. Then $\mathbf{X} + \sqrt{t}\mathbf{Z}$ has finite differential entropy for every $t > 0$, and for any $t > 0$,

$$\frac{1}{2} J(\mathbf{X} + \sqrt{t}\mathbf{Z}) = \frac{\partial}{\partial t} h(\mathbf{X} + \sqrt{t}\mathbf{Z}).$$

$\circ$

Proof. Throughout this computation ∇^2 denotes the *Laplacian* (the trace of the Hessian), which is what appears in the heat equation. Let f_t and g_t denote the densities of $\mathbf{X} + \sqrt{t}\mathbf{Z}$ and $\sqrt{t}\mathbf{Z}$, respectively. Note that

$$\begin{aligned}\frac{1}{2}\nabla^2 g_t(\mathbf{x}) &= \left(\frac{\|\mathbf{x}\|^2}{2t^2} - \frac{d}{2t} \right) g_t(\mathbf{x}) \\ &= \frac{\partial}{\partial t} g_t(\mathbf{x}),\end{aligned}$$

and hence

$$\begin{aligned}\frac{1}{2}\nabla^2 f_t(\mathbf{y}) &= \frac{1}{2}\nabla^2 \mathbb{E} [g_t(\mathbf{y} - \mathbf{X})] \\ &= \mathbb{E} \left[\frac{1}{2}\nabla^2 g_t(\mathbf{y} - \mathbf{X}) \right] \\ &= \mathbb{E} \left[\frac{\partial}{\partial t} g_t(\mathbf{y} - \mathbf{X}) \right] \\ &= \frac{\partial}{\partial t} \mathbb{E} [g_t(\mathbf{y} - \mathbf{X})] \\ &= \frac{\partial}{\partial t} f_t(\mathbf{y}).\end{aligned}$$

Then

$$\begin{aligned}\frac{\partial}{\partial t} h(\mathbf{X} + \sqrt{t}\mathbf{Z}) &= \frac{\partial}{\partial t} \left(- \int f_t \log f_t \right) \\ &= - \int (\log f_t + 1) \frac{\partial}{\partial t} f_t \\ &= - \int (\log f_t) \frac{\partial}{\partial t} f_t \\ &= - \frac{1}{2} \int (\log f_t) \nabla^2 f_t \\ &= - \frac{1}{2} \int (\nabla \cdot ((\log f_t) \nabla f_t) - \langle \nabla \log f_t, \nabla f_t \rangle) \\ &= \frac{1}{2} \int \langle \nabla \log f_t, \nabla f_t \rangle \\ &= \frac{1}{2} \int \frac{\|\nabla f_t\|^2}{f_t} \\ &= \frac{1}{2} J(\mathbf{X} + \sqrt{t}\mathbf{Z}).\end{aligned}$$

The formal steps (differentiating under the integral, dropping the divergence term $\nabla \cdot ((\log f_t) \nabla f_t)$, and the integration by parts) are justified by a standard cutoff/limiting argument: for $t > 0$ the density f_t is the convolution of $P_{\mathbf{X}}$ with the positive Gaussian kernel, hence is positive and smooth. Its tails, however, need not decay rapidly: a density proportional to $(1 + x^2)^{-2}$ has finite second moment and retains polynomial tails after smoothing. The cutoffs converge by

moment/score estimates under the regularity conventions stated at the beginning of this section, not by a blanket fast-decay claim. ■

Remark 1.12. The proof uses the fact that the density of $\mathbf{X} + \sqrt{t}\mathbf{Z}$ solves the heat equation $\frac{\partial}{\partial t}f_t = \frac{1}{2}\nabla^2 f_t$ with initial condition given by the distribution of $\mathbf{X}$. The family of distributions of $\mathbf{X} + \sqrt{t}\mathbf{Z}$ is therefore sometimes referred to as the heat flow generated by $\mathbf{X}$. ○

Proposition 1.39. Let $\mathbf{X}$ be a random vector in $\mathbb{R}^d$ with finite second moment, and let $\mathbf{Z}$ be a d -dimensional standard Gaussian random vector independent of $\mathbf{X}$. Then, for any $t > 0$,

$$\frac{\partial}{\partial t}h(\mathbf{X} + \sqrt{t}\mathbf{Z})$$

is non-increasing in t . ○

Proof. Fix $t_1 > t_2 > 0$ and let $\delta \triangleq t_1 - t_2 > 0$. Let $\mathbf{Z}_1, \mathbf{Z}_2, \mathbf{Z}_3$ be independent standard Gaussian random vectors, independent of $\mathbf{X}$. It suffices to show that, for every $\varepsilon > 0$,

$$\begin{aligned} h(\mathbf{X} + \sqrt{t_2}\mathbf{Z}_1 + \sqrt{\varepsilon}\mathbf{Z}_3) - h(\mathbf{X} + \sqrt{t_2}\mathbf{Z}_1) \\ \geq h(\mathbf{X} + \sqrt{t_2}\mathbf{Z}_1 + \sqrt{\delta}\mathbf{Z}_2 + \sqrt{\varepsilon}\mathbf{Z}_3) - h(\mathbf{X} + \sqrt{t_2}\mathbf{Z}_1 + \sqrt{\delta}\mathbf{Z}_2). \end{aligned}$$

Using $h(U + V) - h(U) = I(V; U + V)$ when V is independent of U , this can be rewritten as

$$I(\mathbf{Z}_3; \mathbf{X} + \sqrt{t_2}\mathbf{Z}_1 + \sqrt{\varepsilon}\mathbf{Z}_3) \geq I(\mathbf{Z}_3; \mathbf{X} + \sqrt{t_2}\mathbf{Z}_1 + \sqrt{\delta}\mathbf{Z}_2 + \sqrt{\varepsilon}\mathbf{Z}_3).$$

Note that

$$\mathbf{Z}_3 \longrightarrow \mathbf{X} + \sqrt{t_2}\mathbf{Z}_1 + \sqrt{\varepsilon}\mathbf{Z}_3 \longrightarrow \mathbf{X} + \sqrt{t_2}\mathbf{Z}_1 + \sqrt{\delta}\mathbf{Z}_2 + \sqrt{\varepsilon}\mathbf{Z}_3.$$

By the data-processing inequality,

$$I(\mathbf{Z}_3; \mathbf{X} + \sqrt{t_2}\mathbf{Z}_1 + \sqrt{\varepsilon}\mathbf{Z}_3) \geq I(\mathbf{Z}_3; \mathbf{X} + \sqrt{t_2}\mathbf{Z}_1 + \sqrt{\delta}\mathbf{Z}_2 + \sqrt{\varepsilon}\mathbf{Z}_3),$$

which is exactly the desired inequality. Dividing by ε and letting $\varepsilon \downarrow 0$ then gives

$$\left. \frac{\partial}{\partial t}h(\mathbf{X} + \sqrt{t}\mathbf{Z}) \right|_{t=t_2} \geq \left. \frac{\partial}{\partial t}h(\mathbf{X} + \sqrt{t}\mathbf{Z}) \right|_{t=t_1},$$

since $\mathbf{X} + \sqrt{t_2}\mathbf{Z}_1 + \sqrt{\varepsilon}\mathbf{Z}_3 \stackrel{d}{=} \mathbf{X} + \sqrt{t_2 + \varepsilon}\mathbf{Z}$ and $\mathbf{X} + \sqrt{t_2}\mathbf{Z}_1 + \sqrt{\delta}\mathbf{Z}_2 + \sqrt{\varepsilon}\mathbf{Z}_3 \stackrel{d}{=} \mathbf{X} + \sqrt{t_1 + \varepsilon}\mathbf{Z}$, which proves that $\frac{\partial}{\partial t}h(\mathbf{X} + \sqrt{t}\mathbf{Z})$ is non-increasing in t . ■

Remark 1.13. The monotonicity of $t \mapsto \frac{\partial}{\partial t}h(\mathbf{X} + \sqrt{t}\mathbf{Z})$ shows that the limit $\lim_{t \downarrow 0^+} \frac{\partial}{\partial t}h(\mathbf{X} + \sqrt{t}\mathbf{Z})$ always exists and, because the derivative is nonnegative by de Bruijn's identity, its value lies in $[0, +\infty]$, not in an interval allowing $-\infty$. Identifying this limit with $\frac{1}{2}J(\mathbf{X})$ requires the appropriate endpoint regularity (or an extended Fisher-information theorem): under mild regularity conditions ensuring differentiability at $t = 0$,

$$\frac{1}{2}J(\mathbf{X}) = \lim_{t \downarrow 0^+} \frac{\partial}{\partial t}h(\mathbf{X} + \sqrt{t}\mathbf{Z}),$$

where $J(\mathbf{X})$ denotes the Fisher information of $\mathbf{X}$. ○

Proposition 1.40. $J(a\mathbf{X}) = a^{-2}J(\mathbf{X})$ for all $a > 0$. ○

Proof. This is a direct computation from the score function and does not require the de Bruijn endpoint $t = 0$: the density of $a\mathbf{X}$ is $f_{a\mathbf{X}}(\mathbf{y}) = a^{-d}f_{\mathbf{X}}(\mathbf{y}/a)$, so its score satisfies $\rho_{a\mathbf{X}}(a\mathbf{x}) = a^{-1}\rho_{\mathbf{X}}(\mathbf{x})$. Taking squared norms and expectations,

$$J(a\mathbf{X}) = \mathbb{E} \|\rho_{a\mathbf{X}}(a\mathbf{X})\|^2 = a^{-2} \mathbb{E} \|\rho_{\mathbf{X}}(\mathbf{X})\|^2 = a^{-2}J(\mathbf{X}).$$

The same scaling follows from the Sobolev definition $J(f) = 4 \int \|\nabla \sqrt{f}\|^2$ by the chain rule for gradients. ■

Proposition 1.41. If $\mathbf{X}$ has a density, finite second moment, and finite differential entropy $h(\mathbf{X})$, then

$$h(\mathbf{X}) = \frac{d}{2} \log(2\pi e) - \frac{1}{2} \int_0^\infty \left(J(\mathbf{X} + \sqrt{t} \mathbf{Z}) - \frac{d}{1+t} \right) dt,$$

where $\mathbf{Z}$ is a standard Gaussian random vector independent of $\mathbf{X}$. (Finite covariance alone is not enough: the left-hand side requires $\mathbf{X}$ to have a density, and the integral is understood through the limiting procedure in the proof.) ○

Proof. Recall that for every $t > 0$,

$$\frac{1}{2}J(\mathbf{X} + \sqrt{t} \mathbf{Z}) = \frac{\partial}{\partial t} h(\mathbf{X} + \sqrt{t} \mathbf{Z}).$$

Integrating from $\varepsilon > 0$ to t gives

$$h(\mathbf{X} + \sqrt{t} \mathbf{Z}) - h(\mathbf{X} + \sqrt{\varepsilon} \mathbf{Z}) = \frac{1}{2} \int_\varepsilon^t J(\mathbf{X} + \sqrt{s} \mathbf{Z}) ds.$$

Subtracting $\frac{d}{2} \log(1+t)$ from both sides and adding $\frac{d}{2} \log(1+\varepsilon)$,

$$\begin{aligned} & h(\mathbf{X} + \sqrt{\varepsilon} \mathbf{Z}) - \frac{d}{2} \log(1+\varepsilon) \\ &= h(\mathbf{X} + \sqrt{t} \mathbf{Z}) - \frac{d}{2} \log(1+t) - \frac{1}{2} \int_\varepsilon^t \left(J(\mathbf{X} + \sqrt{s} \mathbf{Z}) - \frac{d}{1+s} \right) ds. \end{aligned}$$

The first limit taken is $t \rightarrow \infty$. Let K be the covariance matrix of $\mathbf{X}$, and let λ be the largest eigenvalue of K . Since the Gaussian distribution maximises differential entropy among distributions with a given covariance (proved in the previous section),

$$\begin{aligned} h(\mathbf{X} + \sqrt{t} \mathbf{Z}) &\leq \frac{d}{2} \log(2\pi e) + \frac{1}{2} \log |K + tI| \\ &\leq \frac{d}{2} \log(2\pi e) + \frac{d}{2} \log(\lambda + t). \end{aligned}$$

On the other hand,

$$\begin{aligned} h(\mathbf{X} + \sqrt{t} \mathbf{Z}) &\geq h(\mathbf{X} + \sqrt{t} \mathbf{Z} | \mathbf{X}) \\ &= h(\sqrt{t} \mathbf{Z}) \\ &= \frac{d}{2} \log(2\pi e) + \frac{d}{2} \log t. \end{aligned}$$

Subtracting $\frac{d}{2} \log(1+t)$ from each side and letting $t \rightarrow \infty$ then yields, via the squeeze theorem,

$$h(\mathbf{X} + \sqrt{t} \mathbf{Z}) - \frac{d}{2} \log(1+t) \rightarrow \frac{d}{2} \log(2\pi e).$$

Inserting this into the displayed identity gives

$$h(\mathbf{X} + \sqrt{\varepsilon} \mathbf{Z}) - \frac{d}{2} \log(1+\varepsilon) = \frac{d}{2} \log(2\pi e) - \frac{1}{2} \int_{\varepsilon}^{\infty} \left(J(\mathbf{X} + \sqrt{s} \mathbf{Z}) - \frac{d}{1+s} \right) ds.$$

Finally let $\varepsilon \downarrow 0$, using the standard entropy-continuity-at-zero result for Gaussian smoothing, $h(\mathbf{X} + \sqrt{\varepsilon} \mathbf{Z}) \rightarrow h(\mathbf{X})$; since $h(\mathbf{X})$ is finite by hypothesis, this limit is legitimate, and the claim follows. $\blacksquare$

Proposition 1.42. (Gaussian minimises Fisher information) Let $\mathbf{X}$ be a random vector in $\mathbb{R}^d$ with a positive C^2 density, finite Fisher information, and invertible covariance matrix $K_{\mathbf{X}}$, satisfying the integration-by-parts convention below. Then

$$J(\mathbf{X}) \geq \text{tr}(K_{\mathbf{X}}^{-1}),$$

with equality if and only if $\mathbf{X}$ is Gaussian. $\circ$

Proof. Without loss of generality, we may assume that $\mathbf{X}$ has zero mean. Let f be the density of $\mathbf{X}$, and let g be the Gaussian density with zero mean and covariance $K_{\mathbf{X}}$. Consider

$$\mathbb{E} \left[\left\| (\nabla \log f)(\mathbf{X}) - (\nabla \log g)(\mathbf{X}) \right\|^2 \right] \geq 0.$$

Expanding gives

$$\begin{aligned} 0 &\leq \mathbb{E} \left[\left\| (\nabla \log f)(\mathbf{X}) \right\|^2 \right] + \mathbb{E} \left[\left\| (\nabla \log g)(\mathbf{X}) \right\|^2 \right] \\ &\quad - 2 \mathbb{E} \left[\langle (\nabla \log f)(\mathbf{X}), (\nabla \log g)(\mathbf{X}) \rangle \right]. \end{aligned}$$

Now,

$$\begin{aligned} \mathbb{E} \left[\left\| (\nabla \log f)(\mathbf{X}) \right\|^2 \right] &= J(\mathbf{X}), \\ \mathbb{E} \left[\left\| (\nabla \log g)(\mathbf{X}) \right\|^2 \right] &= \mathbb{E} \left[\left\| -K_{\mathbf{X}}^{-1} \mathbf{X} \right\|^2 \right] \\ &= \mathbb{E} \left[\mathbf{X}^T K_{\mathbf{X}}^{-2} \mathbf{X} \right] \end{aligned}$$

$$\begin{aligned}
&= \mathbb{E} [\text{tr} (\mathbf{X}\mathbf{X}^T K_X^{-2})] \\
&= \text{tr} (\mathbb{E}[\mathbf{X}\mathbf{X}^T] K_X^{-2}) \\
&= \text{tr}(K_X^{-1}),
\end{aligned}$$

and

$$\begin{aligned}
\mathbb{E} [\langle (\nabla \log f)(\mathbf{X}), (\nabla \log g)(\mathbf{X}) \rangle] &= \int \langle \nabla f, \nabla \log g \rangle d\mathbf{x} \\
&= \int (\nabla \cdot (f \nabla \log g) - f \nabla^2 \log g) d\mathbf{x} \\
&= - \int f \nabla^2 \log g d\mathbf{x} \\
&= \mathbb{E} [- \nabla^2 \log g(\mathbf{X})] \\
&= \text{tr}(K_X^{-1}).
\end{aligned}$$

Substituting, we find

$$J(\mathbf{X}) + \text{tr}(K_X^{-1}) - 2 \text{tr}(K_X^{-1}) \geq 0,$$

so $J(\mathbf{X}) \geq \text{tr}(K_X^{-1})$. (Throughout this computation we work with positive C^2 densities and finite Fisher information, so that the integration by parts is justified; the general case follows by the standard smooth regularisation. Under the same convention the identity $\mathbb{E}[\rho_{\mathbf{X}}(\mathbf{X})\mathbf{X}^T] = -I$ for zero-mean $\mathbf{X}$ is the special case $g \propto \text{Gaussian}$.)

For the equality characterisation, the squared norm expanded above equals $\mathbb{E} \|\rho_{\mathbf{X}}(\mathbf{X}) + K_X^{-1}\mathbf{X}\|^2$. Equality therefore gives $\nabla \ln f(\mathbf{x}) = -K_X^{-1}\mathbf{x}$ almost everywhere. Under the positive C^2 regularity hypothesis, continuity extends this identity to all of $\mathbb{R}^d$, and integrating along rays identifies f with the Gaussian density of covariance K_X (normalisability forces the proportionality constant). ■

Remark 1.14. The following heat-flow derivation of Gaussian entropy maximality is an *alternative* derivation, not an independent proof: the integral representation used above already relied on Gaussian maximality (in the large- t squeeze). With that caveat, if $\lambda_1, \dots, \lambda_d$ are the eigenvalues of K_X , then

$$\begin{aligned}
h(\mathbf{X}) &= \frac{d}{2} \log(2\pi e) - \frac{1}{2} \int_0^\infty \left(J(\mathbf{X} + \sqrt{t}\mathbf{Z}) - \frac{d}{1+t} \right) dt \\
&\leq \frac{d}{2} \log(2\pi e) - \frac{1}{2} \sum_{i=1}^d \int_0^\infty \left(\frac{1}{\lambda_i + t} - \frac{1}{1+t} \right) dt \\
&= \frac{d}{2} \log(2\pi e) + \frac{1}{2} \sum_{i=1}^d \log \lambda_i \\
&= \frac{d}{2} \log(2\pi e) + \frac{1}{2} \log |K_X|.
\end{aligned}$$

○

1.9 Entropy power inequality

The entropy power inequality provides a bound on the entropy of the sum of two independent random variables in terms of their individual entropies. It was proposed by Shannon in (1948); a complete proof was first given by Stam (1959), with Blachman later providing an independent proof, and several alternative proofs appear later in this chapter.

Theorem 1.43. (Entropy power inequality) Let $\mathbf{X}, \mathbf{Y}$ be independent random vectors in $\mathbb{R}^d$ such that $h(\mathbf{X})$, $h(\mathbf{Y})$ and $h(\mathbf{X} + \mathbf{Y})$ are finite. Recall that in this section h is measured in nats, so the entropy power of $\mathbf{X}$ is $\mathcal{N}(\mathbf{X}) \triangleq e^{\frac{2}{d}h(\mathbf{X})}$. Then

$$\mathcal{N}(\mathbf{X} + \mathbf{Y}) \geq \mathcal{N}(\mathbf{X}) + \mathcal{N}(\mathbf{Y}).$$

Since $\mathcal{N}(\mathbf{X}) = e^{2h_{\text{nats}}(\mathbf{X})/d} = 2^{2h_2(\mathbf{X})/d}$, the identical inequality holds with 2 in place of e when entropies are measured in bits: the entropy-power form is independent of the base of the logarithm. $\circ$

Lemma 1.44. (Lieb's equivalent form of the EPI) For any $0 \leq \lambda \leq 1$ and independent random vectors $\mathbf{X}, \mathbf{Y}$ with finite entropies,

$$h\left(\sqrt{1-\lambda}\mathbf{X} + \sqrt{\lambda}\mathbf{Y}\right) \geq (1-\lambda)h(\mathbf{X}) + \lambda h(\mathbf{Y}).$$

$\circ$

Proof of equivalence of the EPI and Lieb's EPI. Write $A = \mathcal{N}(\mathbf{X})$, $B = \mathcal{N}(\mathbf{Y})$. Assume first Lieb's form, and choose the optimising weight

$$\lambda = \frac{B}{A+B}, \quad 1-\lambda = \frac{A}{A+B}.$$

Since $h(a\mathbf{X}) = h(\mathbf{X}) + d \ln a$ in nats,

$$\begin{aligned} h(\mathbf{X} + \mathbf{Y}) &\geq (1-\lambda)h\left(\frac{\mathbf{X}}{\sqrt{1-\lambda}}\right) + \lambda h\left(\frac{\mathbf{Y}}{\sqrt{\lambda}}\right) \\ &= (1-\lambda)h(\mathbf{X}) + \lambda h(\mathbf{Y}) - \frac{d}{2}[(1-\lambda)\ln(1-\lambda) + \lambda\ln\lambda] \\ &= \frac{d}{2}\left[(1-\lambda)\ln A + \lambda\ln B\right] - \frac{d}{2}[(1-\lambda)\ln(1-\lambda) + \lambda\ln\lambda] \\ &= \frac{d}{2}\left[\frac{A}{A+B}\ln A + \frac{B}{A+B}\ln B - \frac{A}{A+B}\ln\frac{A}{A+B} - \frac{B}{A+B}\ln\frac{B}{A+B}\right] \\ &= \frac{d}{2}\left[\frac{A}{A+B}\ln(A+B) + \frac{B}{A+B}\ln(A+B)\right] \\ &= \frac{d}{2}\ln(A+B), \end{aligned}$$

which is exactly $\mathcal{N}(\mathbf{X} + \mathbf{Y}) \geq A + B$, the EPI.

Conversely, assume the EPI. Then, for $0 < \lambda < 1$,

$$\begin{aligned} h\left(\sqrt{1-\lambda}\mathbf{X} + \sqrt{\lambda}\mathbf{Y}\right) &\geq \frac{d}{2} \ln\left(\mathcal{N}\left(\sqrt{1-\lambda}\mathbf{X}\right) + \mathcal{N}\left(\sqrt{\lambda}\mathbf{Y}\right)\right) \\ &= \frac{d}{2} \ln\left((1-\lambda)\mathcal{N}(\mathbf{X}) + \lambda\mathcal{N}(\mathbf{Y})\right) \\ &\geq \frac{d}{2} \left((1-\lambda)\ln\mathcal{N}(\mathbf{X}) + \lambda\ln\mathcal{N}(\mathbf{Y})\right) \quad (\text{concavity of } \ln) \\ &= (1-\lambda)h(\mathbf{X}) + \lambda h(\mathbf{Y}), \end{aligned}$$

using $\mathcal{N}(\sqrt{1-\lambda}\mathbf{X}) = (1-\lambda)\mathcal{N}(\mathbf{X})$ and the same for $\mathbf{Y}$ (in nats). The endpoints $\lambda \in \{0, 1\}$ are immediate by direct substitution. This proves the equivalence. ■

Lemma 1.45. (Fisher information form of the EPI) For any $0 \leq \lambda \leq 1$ and independent random vectors $\mathbf{X}, \mathbf{Y}$ in $\mathbb{R}^d$ with positive C^2 densities and finite Fisher information,

$$J\left(\sqrt{1-\lambda}\mathbf{X} + \sqrt{\lambda}\mathbf{Y}\right) \leq (1-\lambda)J(\mathbf{X}) + \lambda J(\mathbf{Y}).$$

○

Proof. Write $a = \sqrt{1-\lambda}$, $b = \sqrt{\lambda}$ and $\mathbf{S} = a\mathbf{X} + b\mathbf{Y}$, with $0 < \lambda < 1$ (the endpoints are immediate). The convolution score identity for $\mathbf{S}$ is

$$\rho_{\mathbf{S}}(\mathbf{S}) = \mathbb{E}\left[a\rho_{\mathbf{X}}(\mathbf{X}) + b\rho_{\mathbf{Y}}(\mathbf{Y}) \mid \mathbf{S}\right],$$

which follows by differentiating the convolution formula for the density of $\mathbf{S}$ (justified for positive C^2 densities). Both scores have mean zero under the integration-by-parts convention, since $\mathbb{E}[\rho_{\mathbf{X}}(\mathbf{X})] = \int \nabla f_{\mathbf{X}} = 0$. Applying the conditional Jensen inequality to the squared norm of a conditional expectation,

$$\begin{aligned} J(\mathbf{S}) &= \mathbb{E}\left\|\mathbb{E}\left[a\rho_{\mathbf{X}}(\mathbf{X}) + b\rho_{\mathbf{Y}}(\mathbf{Y}) \mid \mathbf{S}\right]\right\|^2 \\ &\leq \mathbb{E}\left\|a\rho_{\mathbf{X}}(\mathbf{X}) + b\rho_{\mathbf{Y}}(\mathbf{Y})\right\|^2 \\ &= a^2 J(\mathbf{X}) + b^2 J(\mathbf{Y}) + 2ab \mathbb{E}\langle\rho_{\mathbf{X}}(\mathbf{X}), \rho_{\mathbf{Y}}(\mathbf{Y})\rangle \\ &= (1-\lambda)J(\mathbf{X}) + \lambda J(\mathbf{Y}), \end{aligned}$$

where the cross term vanishes by independence: $\mathbb{E}\langle\rho_{\mathbf{X}}(\mathbf{X}), \rho_{\mathbf{Y}}(\mathbf{Y})\rangle = \langle\mathbb{E}\rho_{\mathbf{X}}(\mathbf{X}), \mathbb{E}\rho_{\mathbf{Y}}(\mathbf{Y})\rangle = 0$. This direct proof avoids the Taylor expansion at $t = 0$ used in the mutual-information argument below.

Equality analysis (needed for the vector equality cases later). Equality in the conditional Jensen step means

$$a\rho_{\mathbf{X}}(\mathbf{x}) + b\rho_{\mathbf{Y}}(\mathbf{y}) = \rho_{\mathbf{S}}(a\mathbf{x} + b\mathbf{y})$$

for almost every $(\mathbf{x}, \mathbf{y})$; under the smooth positive hypotheses the identity extends by continuity to all $(\mathbf{x}, \mathbf{y})$. Differentiating in $\mathbf{x}$ and in $\mathbf{y}$ gives the two separate identities

$$a \nabla \rho_X(\mathbf{x}) = a \nabla \rho_S(a\mathbf{x} + b\mathbf{y}), \quad b \nabla \rho_Y(\mathbf{y}) = b \nabla \rho_S(a\mathbf{x} + b\mathbf{y}).$$

Since $a, b > 0$, divide the two identities by their coefficients:

$$\nabla \rho_X(\mathbf{x}) = \nabla \rho_S(a\mathbf{x} + b\mathbf{y}) = \nabla \rho_Y(\mathbf{y}).$$

Because $\mathbf{x}$ and $\mathbf{y}$ vary independently, all these matrices equal a single constant matrix, so $\nabla \rho_X = \nabla \rho_Y = \nabla \rho_S$ is constant. Normalisability then forces $\mathbf{X}$ and $\mathbf{Y}$ to be Gaussian densities with the same positive definite covariance (a density with constant Hessian of $\ln f$ is a Gaussian density). Conversely, Gaussian $\mathbf{X}, \mathbf{Y}$ with equal covariances attain equality. ■

Alternative proof via mutual information, with endpoint assumption. Let $\mathbf{Z}$ be a standard Gaussian random vector independent of $\mathbf{X}$ and $\mathbf{Y}$. Consider

$$\begin{aligned} I(\mathbf{Z}; \sqrt{1-\lambda}\mathbf{X} + \sqrt{\lambda}\mathbf{Y} + \sqrt{t}\mathbf{Z}) \\ \leq I(\mathbf{Z}; \mathbf{X} + \sqrt{(1-\lambda)t}\mathbf{Z}, \mathbf{Y} + \sqrt{\lambda t}\mathbf{Z}) \end{aligned}$$

since

$$\mathbf{Z} \rightarrow (\mathbf{X} + \sqrt{(1-\lambda)t}\mathbf{Z}, \mathbf{Y} + \sqrt{\lambda t}\mathbf{Z}) \rightarrow \sqrt{1-\lambda}\mathbf{X} + \sqrt{\lambda}\mathbf{Y} + \sqrt{t}\mathbf{Z}$$

is a Markov chain. By the chain rule,

$$\begin{aligned} I(\mathbf{Z}; \mathbf{X} + \sqrt{(1-\lambda)t}\mathbf{Z}, \mathbf{Y} + \sqrt{\lambda t}\mathbf{Z}) \\ = I(\mathbf{Z}; \mathbf{X} + \sqrt{(1-\lambda)t}\mathbf{Z} \mid \mathbf{Y} + \sqrt{\lambda t}\mathbf{Z}) + I(\mathbf{Z}; \mathbf{Y} + \sqrt{\lambda t}\mathbf{Z}). \end{aligned}$$

Adding the non-negative term $I(\mathbf{Y} + \sqrt{\lambda t}\mathbf{Z}; \mathbf{X} + \sqrt{(1-\lambda)t}\mathbf{Z})$ to the right-hand side and applying the chain rule again yields

$$\begin{aligned} I(\mathbf{Z}; \sqrt{1-\lambda}\mathbf{X} + \sqrt{\lambda}\mathbf{Y} + \sqrt{t}\mathbf{Z}) \\ \leq I(\mathbf{Z}, \mathbf{Y} + \sqrt{\lambda t}\mathbf{Z}; \mathbf{X} + \sqrt{(1-\lambda)t}\mathbf{Z}) + I(\mathbf{Z}; \mathbf{Y} + \sqrt{\lambda t}\mathbf{Z}) \\ = I(\mathbf{Z}; \mathbf{X} + \sqrt{(1-\lambda)t}\mathbf{Z}) + I(\mathbf{Z}; \mathbf{Y} + \sqrt{\lambda t}\mathbf{Z}), \end{aligned}$$

where we used that

$$\mathbf{Y} + \sqrt{\lambda t}\mathbf{Z} \rightarrow \mathbf{Z} \rightarrow \mathbf{X} + \sqrt{(1-\lambda)t}\mathbf{Z}$$

is Markov.

It remains to compare the derivatives at $t = 0$. For this alternative proof, assume additionally finite differential entropies of $\mathbf{X}, \mathbf{Y}, \mathbf{S}$ and validity of the right-endpoint de Bruijn identities for those three variables. (The direct score proof does not require these additional entropy

hypotheses; smooth positive densities with finite Fisher information do not by themselves ensure finite entropy.) We have

$$\begin{aligned} \left. \frac{\partial}{\partial t} I(\mathbf{Z}; \mathbf{X} + \sqrt{(1-\lambda)t} \mathbf{Z}) \right|_{t=0} &= \left. \frac{\partial}{\partial t} (h(\mathbf{X} + \sqrt{(1-\lambda)t} \mathbf{Z}) - h(\mathbf{X})) \right|_{t=0} \\ &= \frac{1}{2} (1-\lambda) J(\mathbf{X}), \end{aligned}$$

and similarly,

$$\begin{aligned} \left. \frac{\partial}{\partial t} I(\mathbf{Z}; \mathbf{Y} + \sqrt{\lambda t} \mathbf{Z}) \right|_{t=0} &= \frac{1}{2} \lambda J(\mathbf{Y}), \\ \left. \frac{\partial}{\partial t} I(\mathbf{Z}; \sqrt{1-\lambda} \mathbf{X} + \sqrt{\lambda} \mathbf{Y} + \sqrt{t} \mathbf{Z}) \right|_{t=0} &= \frac{1}{2} J(\sqrt{1-\lambda} \mathbf{X} + \sqrt{\lambda} \mathbf{Y}). \end{aligned}$$

Therefore, as $t \rightarrow 0$,

$$\begin{aligned} &I(\mathbf{Z}; \sqrt{1-\lambda} \mathbf{X} + \sqrt{\lambda} \mathbf{Y} + \sqrt{t} \mathbf{Z}) \\ &\quad - I(\mathbf{Z}; \mathbf{X} + \sqrt{(1-\lambda)t} \mathbf{Z}) - I(\mathbf{Z}; \mathbf{Y} + \sqrt{\lambda t} \mathbf{Z}) \\ &= t \cdot \frac{1}{2} \left(J(\sqrt{1-\lambda} \mathbf{X} + \sqrt{\lambda} \mathbf{Y}) - (1-\lambda) J(\mathbf{X}) - \lambda J(\mathbf{Y}) \right) + o(t) \end{aligned}$$

is non-positive. Dividing by t and letting $t \downarrow 0$ yields

$$J(\sqrt{1-\lambda} \mathbf{X} + \sqrt{\lambda} \mathbf{Y}) \leq (1-\lambda) J(\mathbf{X}) + \lambda J(\mathbf{Y}),$$

as claimed. ■

From Fisher Information Inequality Back to Lieb's EPI

Assume the Fisher information inequality

$$J(\sqrt{1-\lambda} \mathbf{X} + \sqrt{\lambda} \mathbf{Y}) \leq (1-\lambda) J(\mathbf{X}) + \lambda J(\mathbf{Y}), \quad 0 < \lambda < 1, \quad (1.4)$$

for all independent $\mathbf{X}, \mathbf{Y}$ with positive C^2 densities and finite Fisher information (the hypotheses of Lemma 1.45). We show that this implies Lieb's form of the EPI for all independent $\mathbf{X}, \mathbf{Y}$ with densities, finite second moments, and finite differential entropies (so that the integral representation of h below applies to each of the variables that appear); the endpoint value $g(0)$ in the argument below is understood as the limit as $t \downarrow 0$.

$$h(\sqrt{1-\lambda} \mathbf{X} + \sqrt{\lambda} \mathbf{Y}) \geq (1-\lambda) h(\mathbf{X}) + \lambda h(\mathbf{Y}). \quad (1.5)$$

Let $\mathbf{Z}$ be standard Gaussian in $\mathbb{R}^d$ independent of $(\mathbf{X}, \mathbf{Y})$, and set

$$\mathbf{X}_t := \mathbf{X} + \sqrt{t} \mathbf{Z}_1, \quad \mathbf{Y}_t := \mathbf{Y} + \sqrt{t} \mathbf{Z}_2, \quad \mathbf{S}_t := \sqrt{1-\lambda} \mathbf{X}_t + \sqrt{\lambda} \mathbf{Y}_t,$$

where Z_1, Z_2 are i.i.d. copies of Z , independent of (X, Y) . Then X_t, Y_t remain independent for all $t \geq 0$, and

$$\sqrt{1-\lambda}X + \sqrt{\lambda}Y + \sqrt{t}Z \stackrel{d}{=} S_t, \quad X + \sqrt{t}Z \stackrel{d}{=} X_t, \quad Y + \sqrt{t}Z \stackrel{d}{=} Y_t.$$

Applying (1.4) to the smoothed pair (X_t, Y_t) (which has positive smooth densities) gives, for all $t > 0$,

$$J(S_t) \leq (1-\lambda)J(X_t) + \lambda J(Y_t). \quad (1.6)$$

The integral argument below only uses positive t , so the value at $t = 0$ (where the smoothness hypotheses need not hold) is not needed.

Using the proposition, for any U with a density, finite second moment, and finite differential entropy,

$$h(U) = \frac{d}{2} \log(2\pi e) - \frac{1}{2} \int_0^\infty \left(J(U + \sqrt{t}Z) - \frac{d}{1+t} \right) dt.$$

Applying this with $U = \sqrt{1-\lambda}X + \sqrt{\lambda}Y$, and with $U = X, Y$, and using the distributional identities above, we obtain

$$\begin{aligned} h(\sqrt{1-\lambda}X + \sqrt{\lambda}Y) &= \frac{d}{2} \log(2\pi e) - \frac{1}{2} \int_0^\infty \left(J(S_t) - \frac{d}{1+t} \right) dt, \\ h(X) &= \frac{d}{2} \log(2\pi e) - \frac{1}{2} \int_0^\infty \left(J(X_t) - \frac{d}{1+t} \right) dt, \\ h(Y) &= \frac{d}{2} \log(2\pi e) - \frac{1}{2} \int_0^\infty \left(J(Y_t) - \frac{d}{1+t} \right) dt. \end{aligned}$$

Therefore

$$\begin{aligned} &h(\sqrt{1-\lambda}X + \sqrt{\lambda}Y) - (1-\lambda)h(X) - \lambda h(Y) \\ &= -\frac{1}{2} \int_0^\infty \left(J(S_t) - (1-\lambda)J(X_t) - \lambda J(Y_t) \right) dt. \end{aligned}$$

By (1.6), the integrand is ≤ 0 for all $t > 0$, so the integral is ≤ 0 , and hence

$$h(\sqrt{1-\lambda}X + \sqrt{\lambda}Y) \geq (1-\lambda)h(X) + \lambda h(Y),$$

which is Lieb's form (1.5) of the EPI.

Theorem 1.46. (Stam's inequality) Let X, Y be independent random vectors in $\mathbb{R}^d$ with finite Fisher information. Then

$$\frac{1}{J(X+Y)} \geq \frac{1}{J(X)} + \frac{1}{J(Y)},$$

with equality if and only if X and Y are Gaussian with *proportional* covariance matrices, $\text{Cov}(Y) = \rho \text{Cov}(X)$ for some $\rho > 0$. (For scalars, $d = 1$, proportionality is automatic and the condition reduces to both variables being Gaussian.) $\circ$

Proof. Apply the Fisher information form of the EPI to the rescaled variables $\mathbf{X}/\sqrt{1-\lambda}$ and $\mathbf{Y}/\sqrt{\lambda}$. Writing $\mathbf{X} + \mathbf{Y} = \sqrt{1-\lambda} \frac{\mathbf{X}}{\sqrt{1-\lambda}} + \sqrt{\lambda} \frac{\mathbf{Y}}{\sqrt{\lambda}}$ gives, for every $\lambda \in (0, 1)$,

$$\begin{aligned} J(\mathbf{X} + \mathbf{Y}) &\leq (1-\lambda) J\left(\frac{\mathbf{X}}{\sqrt{1-\lambda}}\right) + \lambda J\left(\frac{\mathbf{Y}}{\sqrt{\lambda}}\right) \\ &= (1-\lambda)^2 J(\mathbf{X}) + \lambda^2 J(\mathbf{Y}). \end{aligned}$$

The right-hand side, as a function of λ , is convex (a sum of two convex quadratics); its minimum over $\lambda \in [0, 1]$ is attained at

$$\lambda = \frac{J(\mathbf{X})}{J(\mathbf{X}) + J(\mathbf{Y})},$$

giving

$$J(\mathbf{X} + \mathbf{Y}) \leq \frac{J(\mathbf{X}) J(\mathbf{Y})}{J(\mathbf{X}) + J(\mathbf{Y})}.$$

Taking reciprocals yields Stam's inequality. Equality in the Fisher-information inequality applied to the rescaled inputs requires Gaussian inputs with *identical* covariance matrices; hence $\mathbf{X}$ and $\mathbf{Y}$ themselves must be Gaussian with *proportional* covariance matrices. Conversely, Gaussian inputs with proportional covariances attain equality at the optimising weight. ■

Remark 1.15 (Entropy-power form of Stam's inequality). Stam's inequality is the Fisher-information analogue of the entropy power inequality: with the entropy power $N(\cdot) \triangleq e^{\frac{2}{d}h(\cdot)}$ (a base-independent quantity), both $N(\cdot)$ and the reciprocal Fisher information $1/J(\cdot)$ are *superadditive*: for independent $\mathbf{X}, \mathbf{Y}$,

$$N(\mathbf{X} + \mathbf{Y}) \geq N(\mathbf{X}) + N(\mathbf{Y}), \quad \frac{1}{J(\mathbf{X} + \mathbf{Y})} \geq \frac{1}{J(\mathbf{X})} + \frac{1}{J(\mathbf{Y})},$$

each with equality precisely for Gaussian inputs with *proportional* covariance matrices (under the stated regularity and nondegeneracy hypotheses). The first display is exactly the EPI; the two inequalities are two faces of the same fact, connected through the de Bruijn identity and the heat-flow representation of h used earlier in this section. ○

Remark 1.16 (Further reading). The entropy-power and Fisher-information inequalities fit into a wider circle of ideas that generalises them in two directions. First, the *Brascamp–Lieb* programme: the EPI and the Brascamp–Lieb inequality admit a unified treatment through information measures [Anantharam et al. \(2022\)](#), with extensions to finite Abelian groups. Second, the *heat-flow* direction: for *scalar* random variables ($d = 1$) the Fisher information is log-convex along the heat flow [Ledoux et al. \(2021\)](#), and the signs of the derivatives of entropy along Markov semi-groups organise several of the identities of this chapter (de Bruijn's identity being the first derivative). The log-convexity statement of [Ledoux et al. \(2021\)](#) is proved there in one dimension; no arbitrary-dimensional statement is attributed to that paper. Both circles of ideas interact with the sub-additivity theme of §1.13: the Brascamp–Lieb inequalities are sub-additivity statements for a family of linear maps, and the heat-flow monotonicities are the infinitesimal form of the data-processing inequality. The sharp form of the data-processing inequality, given by the contraction coefficient, is developed in the distributed source coding chapter. ○

Alternative proof of Lieb's EPI. We give a second, self-contained derivation of Lieb's inequality (1.5) from the Fisher information inequality (1.4), via a decreasing-function monotonicity argument. Assume $\mathbf{X}, \mathbf{Y}$ have densities, finite second moments, and finite differential entropies (the hypotheses under which the integral representation of h and the heat-flow identities above hold). Let $\mathbf{Z}_1, \mathbf{Z}_2, \mathbf{Z}$ be i.i.d. standard Gaussians, independent of $(\mathbf{X}, \mathbf{Y})$, and for $t \in [0, 1]$ set

$$\mathbf{X}_t := \sqrt{1-t} \mathbf{X} + \sqrt{t} \mathbf{Z}_1, \quad \mathbf{Y}_t := \sqrt{1-t} \mathbf{Y} + \sqrt{t} \mathbf{Z}_2.$$

The pair $\mathbf{X}_t, \mathbf{Y}_t$ remains independent for all $t \in [0, 1]$, and

$$\sqrt{1-\lambda} \mathbf{X}_t + \sqrt{\lambda} \mathbf{Y}_t = \sqrt{1-t} \left(\sqrt{1-\lambda} \mathbf{X} + \sqrt{\lambda} \mathbf{Y} \right) + \sqrt{t} \left(\sqrt{1-\lambda} \mathbf{Z}_1 + \sqrt{\lambda} \mathbf{Z}_2 \right)$$

has the same distribution as $\sqrt{1-t} \left(\sqrt{1-\lambda} \mathbf{X} + \sqrt{\lambda} \mathbf{Y} \right) + \sqrt{t} \mathbf{Z}$, since $\sqrt{1-\lambda} \mathbf{Z}_1 + \sqrt{\lambda} \mathbf{Z}_2$ is standard Gaussian. Define

$$g(t) := h\left(\sqrt{1-\lambda} \mathbf{X}_t + \sqrt{\lambda} \mathbf{Y}_t\right) - (1-\lambda) h(\mathbf{X}_t) - \lambda h(\mathbf{Y}_t).$$

At $t = 1$ we have $\mathbf{X}_1 = \mathbf{Z}_1, \mathbf{Y}_1 = \mathbf{Z}_2$, and $\sqrt{1-\lambda} \mathbf{Z}_1 + \sqrt{\lambda} \mathbf{Z}_2$ is standard Gaussian, so $g(1) = 0$. We now show $g'(t) \leq 0$ on $(0, 1)$. Using the scaling identity $h(\sqrt{1-t} \mathbf{W} + \sqrt{t} \mathbf{Z}) = h(\mathbf{W} + \sqrt{\frac{t}{1-t}} \mathbf{Z}) + \frac{d}{2} \log(1-t)$ and the change of variables $s = \frac{t}{1-t}$ (so $ds/dt = \frac{1}{(1-t)^2}$), the de Bruijn identity gives

$$\begin{aligned} g'(t) &= \frac{1}{2(1-t)^2} \left[J\left(\sqrt{1-\lambda} \mathbf{X} + \sqrt{\lambda} \mathbf{Y} + \sqrt{s} \mathbf{Z}\right) - (1-\lambda) J\left(\mathbf{X} + \sqrt{s} \mathbf{Z}\right) - \lambda J\left(\mathbf{Y} + \sqrt{s} \mathbf{Z}\right) \right] \\ &= \frac{1}{2(1-t)} \left[J\left(\sqrt{1-\lambda} \mathbf{X}_t + \sqrt{\lambda} \mathbf{Y}_t\right) - (1-\lambda) J(\mathbf{X}_t) - \lambda J(\mathbf{Y}_t) \right] \leq 0, \end{aligned}$$

where the second line uses $J(\mathbf{W} + \sqrt{s} \mathbf{Z}) = (1-t) J(\sqrt{1-t} \mathbf{W} + \sqrt{t} \mathbf{Z})$, and the final inequality is the Fisher information inequality (1.4) applied to the independent pair $(\mathbf{X}_t, \mathbf{Y}_t)$. Hence g is non-increasing on $[0, 1]$, so $g(0) \geq g(1) = 0$, which is exactly Lieb's inequality (1.5). (The value $g(0)$ is understood as the limit as $t \downarrow 0$.) ■

Remark 1.17 (Scope of the proofs above). The heat-flow proofs given here assume densities, finite second moments, and finite differential entropies (the hypotheses of the integral representation of h and of the Fisher inequality). The general finite-entropy EPI without a second-moment assumption is covered by the cited general results [Anantharam et al. \(2022\)](#). The Gaussian-optimality sketch below is presented in its covariance-constrained setting; the extension beyond that setting is not reproduced here. ○

1.9.1 Conditional EPI

There is a conditional version of the EPI.

Theorem 1.47. (Conditional EPI, finite auxiliary) Let X, Y be random vectors in $\mathbb{R}^d$ and let U be a random variable taking finitely many values such that X and Y are *conditionally independent given U* , i.e. $X \rightarrow U \rightarrow Y$ forms a Markov chain. For every u of positive probability, assume that X, Y are independent under the conditional law, have conditional densities, and that the entropies of $X, Y, X + Y$ under that conditional law are finite. Then, with h measured in nats (the convention of this section),

$$e^{\frac{2}{d}h(X+Y|U)} \geq e^{\frac{2}{d}h(X|U)} + e^{\frac{2}{d}h(Y|U)}.$$

○

Remark 1.18. The conditional independence hypothesis $X \rightarrow U \rightarrow Y$ cannot be dropped: with only marginal independence of X and Y the conditional EPI is false in general. The restriction to finitely-valued U keeps the proof elementary. For a general U on a standard Borel space, also assume the three conditional entropy functions are finite almost surely and integrable, and apply the ordinary EPI under the regular conditional law of (X, Y) given $U = u$ for almost every u (under which X and Y are conditionally independent), then take expectations: with $A(u) = \frac{2}{d}h(X|U = u)$ and $B(u) = \frac{2}{d}h(Y|U = u)$,

$$\frac{2}{d}h(X + Y|U = u) \geq \ln(e^{A(u)} + e^{B(u)}),$$

and since $(a, b) \mapsto \ln(e^a + e^b)$ is convex, Jensen's inequality over u gives the claim. Note that conditioning on a coarsening of U does *not* generally preserve conditional independence of X and Y , so the extension cannot be justified by naive quantisation. (If bits are preferred, label all finite- U displays h_2 and use base 2; the two choices are not mixed.) ○

Proof. For each u of positive probability, the ordinary EPI applies to the independent pair $(X|U = u, Y|U = u)$:

$$\frac{2}{d}h(X + Y|U = u) \geq \ln\left(e^{\frac{2}{d}h(X|U=u)} + e^{\frac{2}{d}h(Y|U=u)}\right).$$

Averaging over u and using convexity of log-sum-exp,

$$\begin{aligned} \frac{2}{d}h(X + Y|U) &= \sum_u \Pr(U = u) \frac{2}{d}h(X + Y|U = u) \\ &\geq \sum_u \Pr(U = u) \ln\left(e^{\frac{2}{d}h(X|U=u)} + e^{\frac{2}{d}h(Y|U=u)}\right) \\ &\geq \ln\left(e^{\sum_u \Pr(U=u) \frac{2}{d}h(X|U=u)} + e^{\sum_u \Pr(U=u) \frac{2}{d}h(Y|U=u)}\right) \\ &= \ln\left(e^{\frac{2}{d}h(X|U)} + e^{\frac{2}{d}h(Y|U)}\right), \end{aligned}$$

where we used Jensen's inequality in the second step. Exponentiating both sides gives exactly the displayed conclusion. ■

Lemma 1.48. (Elementary compression inequality) Let (X, Y, Z) be random variables. Then

$$H(X, Y, Z) + H(Z) \leq H(X, Z) + H(Y, Z).$$

○

Proof. We have

$$\begin{aligned} H(X, Z) + H(Y, Z) - H(X, Y, Z) - H(Z) &= H(X|Z) - H(X|Y, Z) \\ &\geq 0, \end{aligned}$$

since conditioning reduces entropy. ■

Definition 1.49. (Compression) Let n be a positive integer and let $\{\alpha_T\}_T, \{\beta_T\}_T$ be two finite collections of non-negative real numbers indexed by $T \subseteq [1 : n]$. We call $\{\beta_T\}_T$ an *elementary compression* of $\{\alpha_T\}_T$ if there exist $A, B \subseteq [1 : n]$ with $A \not\subseteq B$ and $B \not\subseteq A$, and $0 \leq \delta \leq \min\{\alpha_A, \alpha_B\}$ such that for all $T \subseteq [1 : n]$,

$$\beta_T = \begin{cases} \alpha_T - \delta & \text{if } T = A \text{ or } T = B, \\ \alpha_T + \delta & \text{if } T = A \cup B \text{ or } T = A \cap B, \\ \alpha_T & \text{otherwise.} \end{cases}$$

The result of a finite sequence of elementary compressions of $\{\alpha_T\}_T$ is called a *compression* of $\{\alpha_T\}_T$. ○

An *n-fractional multiset* is a finite collection $\{\beta_T\}_{T \subseteq [1:n]}$ of non-negative real numbers indexed by subsets $T \subseteq [1 : n]$; it is *integer* when all β_T are non-negative integers. A *fractional cover* is an *n-fractional multiset* satisfying the coverage condition $\sum_{T \ni i} \beta_T \geq 1$ for all i (the term is reserved for this coverage condition). For a collection of random variables $X_1, \dots, X_n$ we write $X_T \triangleq (X_i)_{i \in T}$, with the convention $X_\emptyset$ constant and $H(X_\emptyset) = 0$. As shown in [Balister and Bollobás \(2012\)](#), the relation “is a compression of” defines a partial order on the collection of *n-fractional multisets*. An *n-fractional multiset* $\{\beta_T\}_T$ is minimal under this partial order (i.e. cannot be further compressed) if and only if the set $\{T \subseteq [1 : n] : \beta_T \neq 0\}$ is totally ordered under set inclusion.

The following lemma, a submodularity inequality, follows quickly from the definitions.

Lemma 1.50. Let $X_1, \dots, X_n$ be any collection of random variables. Then:

- (i) $H(X_A) + H(X_B) \geq H(X_{A \cup B}) + H(X_{A \cap B})$ for all $A, B \subseteq [1 : n]$.
- (ii) $\sum_{T \subseteq [1:n]} \alpha_T H(X_T) \geq \sum_{T \subseteq [1:n]} \beta_T H(X_T)$ for any *n-fractional multisets* $\{\alpha_T\}, \{\beta_T\}$ such that $\{\beta_T\}$ is a compression of $\{\alpha_T\}$.
- (iii) $\sum_{T \subseteq [1:n]} \beta_T H(X_T) \geq H(X_{[1:n]})$ whenever $\{\beta_T\}$ is an *n-fractional multiset* satisfying $\sum_{T \subseteq [1:n]: T \ni i} \beta_T \geq 1$ for all i .

○

Proof. For (i), let $A, B \subseteq [1 : n]$. Then

$$\begin{aligned} H(X_B) - H(X_{A \cap B}) &= H(X_{B \setminus A} | X_{A \cap B}) \\ &\geq H(X_{B \setminus A} | X_A) \\ &= H(X_{B \setminus A}, X_A) - H(X_A) \\ &= H(X_{A \cup B}) - H(X_A), \end{aligned}$$

where the inequality uses the fact that conditioning reduces entropy. Rearranging gives the desired inequality.

If $\{\beta_T\}$ is an elementary compression of $\{\alpha_T\}$, the inequality in (ii) follows from (i) by expanding and cancelling equal terms on both sides. Since a general compression is obtained as a sequence of elementary compressions, (ii) follows by induction on the number of steps.

For (iii), we argue by induction on n . The base case $n = 1$ is trivial. Assume the claim holds for $n - 1$. By (i),

$$H(X_{[1:n-1]}) + H(X_{T \cup \{n\}}) \geq H(X_{[1:n]}) + H(X_T)$$

for all $T \subseteq [1 : n - 1]$. Let β_T ($T \subseteq [1 : n]$) satisfy $\sum_{T \ni i} \beta_T \geq 1$ for all $i = 1, \dots, n$. Then

$$\begin{aligned} \sum_{T \subseteq [1:n]} \beta_T H(X_T) &= \sum_{T \subseteq [1:n-1]} (\beta_T H(X_T) + \beta_{T \cup \{n\}} H(X_{T \cup \{n\}})) \\ &= \sum_{T \subseteq [1:n-1]} (\beta_T + \beta_{T \cup \{n\}}) H(X_T) \\ &\quad + \sum_{T \subseteq [1:n-1]} \beta_{T \cup \{n\}} (H(X_{T \cup \{n\}}) - H(X_T)) \\ &\geq \sum_{T \subseteq [1:n-1]} (\beta_T + \beta_{T \cup \{n\}}) H(X_T) + (H(X_{[1:n]}) - H(X_{[1:n-1]})), \end{aligned}$$

where the last inequality uses (i) with $A = [1 : n - 1]$ and $B = T \cup \{n\}$, namely $H(X_{T \cup \{n\}}) - H(X_T) \geq H(X_{[1:n]}) - H(X_{[1:n-1]})$, together with $\sum_{T \subseteq [1:n-1]} \beta_{T \cup \{n\}} \geq 1$ and $H(X_n | X_{[1:n-1]}) \geq 0$. The family $\{\beta_T + \beta_{T \cup \{n\}}\}_{T \subseteq [1:n-1]}$ is an $(n - 1)$ -fractional multiset that still satisfies the coverage condition, so by the induction hypothesis,

$$\sum_{T \subseteq [1:n-1]} (\beta_T + \beta_{T \cup \{n\}}) H(X_T) \geq H(X_{[1:n-1]}).$$

Combining yields $\sum_{T \subseteq [1:n]} \beta_T H(X_T) \geq H(X_{[1:n]})$, which completes the proof. ■

1.10 Further inequalities

Here, we provide several inequalities in linear algebra using the properties of differential entropy. Throughout this section differential entropies are measured in *natural* logarithms

(nats), matching the Fisher-information and EPI sections, so the Gaussian entropy formula is $\frac{d}{2} \ln(2\pi e) + \frac{1}{2} \ln |K|$ and no base-conversion constants appear.

Lemma 1.51. (Minkowski determinant inequality) For $K_1, K_2 \in \mathbb{S}_+^n$ (positive semidefinite $n \times n$ matrices),

$$|K_1 + K_2|^{1/n} \geq |K_1|^{1/n} + |K_2|^{1/n}.$$

○

Proof. Prove the claim first for $K_1, K_2 \in \mathbb{S}_{++}^n$ positive definite. Let $\mathbf{X}_1 \sim \mathcal{N}(0, K_1)$ and $\mathbf{X}_2 \sim \mathcal{N}(0, K_2)$ be independent. In nats the entropy power of a Gaussian is $\mathcal{N}(\mathbf{X}_i) = e^{2h(\mathbf{X}_i)/n} = (2\pi e) |K_i|^{1/n}$, i.e. the Gaussian entropy powers are proportional to $(\det K_i)^{1/n}$; the EPI then gives

$$|K_1 + K_2|^{1/n} \geq |K_1|^{1/n} + |K_2|^{1/n}$$

exactly. For positive semidefinite inputs, apply the proved inequality to $K_1 + \varepsilon I$ and $K_2 + \varepsilon I$ and let $\varepsilon \downarrow 0$; continuity of the determinant gives the claim in general. ■

The Minkowski sum of two sets $A, B \subseteq \mathbb{R}^d$ is

$$A + B \triangleq \{x + y : x \in A, y \in B\}.$$

Lemma 1.52. For compact sets $A, B \subseteq \mathbb{R}^d$ of positive volume,

$$\text{vol}(A)^{2/d} + \text{vol}(B)^{2/d} \leq \text{vol}(A + B)^{2/d}.$$

○

Proof. Let $\mathbf{X}_1 \sim \text{Unif}(A)$ and $\mathbf{X}_2 \sim \text{Unif}(B)$ be independent. Then $h(\mathbf{X}_1) = \ln \text{vol}(A)$ and $h(\mathbf{X}_2) = \ln \text{vol}(B)$. The sum $\mathbf{X}_1 + \mathbf{X}_2$ is supported in the compact set $A + B$ (the Minkowski sum of compact sets is compact), so $h(\mathbf{X}_1 + \mathbf{X}_2) \leq \ln \text{vol}(A + B)$. Applying the EPI to $\mathbf{X}_1, \mathbf{X}_2$ yields the claimed inequality. Compactness is assumed to avoid qualifications about the measurability of the Minkowski sum; more general measurable-set statements require the appropriate measurability or outer-measure conventions. ■

Lemma 1.53. The map $K \mapsto \log |K|$ is concave on the cone of positive *definite* matrices; on singular positive semidefinite matrices the natural extended-real extension is defined by $\log \det K = -\infty$ (determinant continuity alone does not make the logarithm a real-valued function on the singular cone). ○

Proof. Prove the claim first for positive definite K_1, K_2 . Let $\mathbf{X}_1 \sim \mathcal{N}(0, K_1)$ and $\mathbf{X}_2 \sim \mathcal{N}(0, K_2)$, and let $Q \in \{1, 2\}$ be independent of $(\mathbf{X}_1, \mathbf{X}_2)$ with $\Pr(Q = 1) = \alpha$. Then $\mathbf{X}_Q$ has covariance $\alpha K_1 + (1 - \alpha) K_2$. By conditioning reduces entropy,

$$h(\mathbf{X}_Q) \geq h(\mathbf{X}_Q | Q) = \alpha h(\mathbf{X}_1) + (1 - \alpha) h(\mathbf{X}_2).$$

The two conditional entropies are Gaussian entropies, $h(\mathbf{X}_i) = \frac{d}{2} \ln(2\pi e) + \frac{1}{2} \ln |K_i|$, and for the *mixture* $\mathbf{X}_Q$ (which is not Gaussian in general) we use the Gaussian maximality proposition with its covariance $\alpha K_1 + (1 - \alpha)K_2$:

$$h(\mathbf{X}_Q) \leq \frac{d}{2} \ln(2\pi e) + \frac{1}{2} \ln |\alpha K_1 + (1 - \alpha)K_2|.$$

Combining the two displays cancels the $(2\pi e)$ -terms and gives

$$\alpha \ln |K_1| + (1 - \alpha) \ln |K_2| \leq \ln |\alpha K_1 + (1 - \alpha)K_2|,$$

which proves concavity for $K_1, K_2 \succ 0$. For positive semidefinite inputs, apply the proved inequality to $K_1 + \varepsilon I$ and $K_2 + \varepsilon I$ (which are positive definite) and let $\varepsilon \downarrow 0$; continuity of the determinant gives the claim in general. ■

Lemma 1.54. (Hadamard's inequality) If K is positive semidefinite, then

$$|K| \leq \prod_{i=1}^d K_{ii}.$$

○

Proof. Prove the claim first for $K \succ 0$. Let $\mathbf{X} \sim \mathcal{N}(0, K)$. Entropy subadditivity gives

$$h(\mathbf{X}) \leq \sum_{i=1}^d h(X_i),$$

with equality if and only if the components are independent. Substituting $h(\mathbf{X}) = \frac{d}{2} \ln(2\pi e) + \frac{1}{2} \ln |K|$ and $h(X_i) = \frac{1}{2} \ln(2\pi e) + \frac{1}{2} \ln K_{ii}$ yields

$$\frac{d}{2} \ln(2\pi e) + \frac{1}{2} \ln |K| \leq \sum_{i=1}^d \left(\frac{1}{2} \ln(2\pi e) + \frac{1}{2} \ln K_{ii} \right),$$

which simplifies to $|K| \leq \prod_{i=1}^d K_{ii}$, with equality if and only if K is diagonal. For a singular $K \succeq 0$, apply the proved inequality to $K + \varepsilon I \succ 0$ and let $\varepsilon \downarrow 0$. For nonsingular K , equality means diagonality; for singular K , equality occurs exactly in the trivial case $|K| = 0 = \prod_i K_{ii}$ when K has a zero row/column. ■

Lemma 1.55. (Schur complement) Suppose $K \succeq 0$ and is partitioned as

$$K = \begin{bmatrix} K_{11} & K_{12} \\ K_{12}^T & K_{22} \end{bmatrix}.$$

If $K_{11} \succ 0$, then

$$|K| = |K_{11}| |K_{22} - K_{12}^T K_{11}^{-1} K_{12}|.$$

○

Proof. Block elimination gives

$$\begin{pmatrix} I & 0 \\ -K_{12}^T K_{11}^{-1} & I \end{pmatrix} K \begin{pmatrix} I & -K_{11}^{-1} K_{12} \\ 0 & I \end{pmatrix} = \begin{pmatrix} K_{11} & 0 \\ 0 & K_{22} - K_{12}^T K_{11}^{-1} K_{12} \end{pmatrix}.$$

Both multiplying matrices are triangular with unit diagonal, hence have determinant one, so

$$|K| = |K_{11}| |K_{22} - K_{12}^T K_{11}^{-1} K_{12}|.$$

This holds even when K or its Schur complement is singular. ■

Lemma 1.56. (Double Markovity (Ahlsvede–Körner ’74)) Let A, B, C be random variables taking values in finite sets. Then

$$A \rightarrow B \rightarrow C \quad \text{and} \quad A \rightarrow C \rightarrow B$$

are both Markov chains if and only if there exist functions f and g such that

$$f(B) = g(C) \quad \text{a.s.}$$

and A is independent of (B, C) given $f(B)$ (equivalently, given $(f(B), g(C))$). ○

Proof. Assume first that there is a random variable U and functions f, g such that $U = f(B) = g(C)$ a.s. and $A \perp\!\!\!\perp (B, C) | U$. Then

$$I(A; C|B) = I(A; C|B, U) \leq I(A; B, C|U) = 0,$$

and similarly $I(A; B|C) = 0$, so $A \rightarrow B \rightarrow C$ and $A \rightarrow C \rightarrow B$ hold.

Conversely, assume $A \rightarrow B \rightarrow C$ and $A \rightarrow C \rightarrow B$. We first treat the case where

$$p(b, c) > 0 \quad \text{for all } b, c.$$

From $A \rightarrow B \rightarrow C$ and $A \rightarrow C \rightarrow B$ we have, for all a, b, c ,

$$p(a|b, c) = p(a|b) = p(a|c).$$

Summing over c gives $p(a|b) = p(a)$, so $A \perp\!\!\!\perp (B, C)$. In this case we may take U to be a constant; the conclusion holds trivially with any constant functions f, g .

Now consider the general finite case. Assume w.l.o.g. that $p(b) > 0$ for all b and $p(c) > 0$ for all c . For each c , define the conditional distribution

$$p_{A|c}(\cdot) \triangleq p(A \in \cdot | C = c),$$

and similarly for each b define $p_{A|b}$. Let

$$\mathcal{P}_C \triangleq \{p_{A|c} : c \in C\}, \quad \mathcal{P}_B \triangleq \{p_{A|b} : b \in \mathcal{B}\}$$

be subsets of the probability simplex on $\mathcal{A}$.

Claim: $\mathcal{P}_C = \mathcal{P}_B$.

We prove both inclusions. First, suppose $p_1 \in \mathcal{P}_C \setminus \mathcal{P}_B$, so $p_1 = p_{A|c_1}$ for some c_1 . Choose b_1 with $p(b_1, c_1) > 0$. Using $A \rightarrow C \rightarrow B$ and $A \rightarrow B \rightarrow C$ we obtain

$$p_1 = p_{A|c_1} = p_{A|c_1, b_1} = p_{A|b_1},$$

so $p_1 \in \mathcal{P}_B$, a contradiction. Hence $\mathcal{P}_C \subseteq \mathcal{P}_B$. The reverse inclusion is proved by interchanging the roles of B and C : if $p_2 \in \mathcal{P}_B \setminus \mathcal{P}_C$, say $p_2 = p_{A|b_2}$, choose c_2 with $p(b_2, c_2) > 0$ and use

$$p_2 = p_{A|b_2} = p_{A|b_2, c_2} = p_{A|c_2},$$

so $p_2 \in \mathcal{P}_C$, a contradiction. Thus $\mathcal{P}_C = \mathcal{P}_B$, and we denote this common set by $\mathcal{P}$.

For each $p \in \mathcal{P}$ define the “level sets”

$$C_p \triangleq \{c : p_{A|c} = p\}, \quad B_p \triangleq \{b : p_{A|b} = p\}.$$

Let $p_1, \dots, p_k$ enumerate the elements of $\mathcal{P}$. Define

$$g(c) = i \quad \text{if } c \in C_{p_i}, \quad f(b) = i \quad \text{if } b \in B_{p_i}.$$

Set $U \triangleq f(B)$ and $V \triangleq g(C)$.

Step 1: $U = V$ a.s.

If not, there exists (b_1, c_1) with $p(b_1, c_1) > 0$ and $f(b_1) \neq g(c_1)$, say $f(b_1) = j$, $g(c_1) = i$. Then

$$p_i = p_{A|c_1} = p_{A|b_1, c_1} = p_{A|b_1} = p_j,$$

contradicting $i \neq j$. Hence $f(B) = g(C)$ almost surely; write this common random variable as U .

Step 2: $A \perp\!\!\!\perp (B, C) \mid U$.

Fix any i in the range of U . The two events $\{f(B) = i\} = \{B \in B_{p_i}\}$ and $\{g(C) = i\} = \{C \in C_{p_i}\}$ coincide almost surely (this is exactly Step 1 with both inclusions proved there: $f(B) = g(C)$ a.s. means that $\{B \in B_{p_i}\}$ equals $\{C \in C_{p_i}\}$ up to null sets, and also that there is no mass on pairs (b, c) with $f(b) \neq g(c)$); they are *not* asserted to be a disjoint union. For any (b, c) with $f(b) = g(c) = i$ and $p(b, c) > 0$, we have

$$p_{A|b, c} = p_{A|b} = p_{A|c} = p_i.$$

In particular, given $U = i$,

$$p_{A|B=b, C=c, U=i} = p_i$$

for all (b, c) with $p(b, c, U = i) > 0$. Therefore

$$p_{A|B, C, U}(a \mid b, c, i) = p_i(a)$$

depends only on i , not on (b, c) . This is exactly the statement that

$$A \perp\!\!\!\perp (B, C) \mid U.$$

Thus there exists a random variable $U = f(B) = g(C)$ such that A is independent of (B, C) given U , as claimed. ■

Remark 1.19 (Extension to Polish spaces). The lemma extends to the case where A, B, C take values in Polish spaces (complete separable metric spaces) equipped with their Borel σ -algebras.

Assume $A \rightarrow B \rightarrow C$ and $A \rightarrow C \rightarrow B$, and that all conditional distributions $P_{A|B}$ and $P_{A|C}$ exist as regular conditional probabilities (which they do on Polish spaces). For each b , let $P_{A|B=b}$ be a probability measure on the state space of A , and similarly for c .

Define the measurable maps

$$\Phi_B : b \mapsto P_{A|B=b}, \quad \Phi_C : c \mapsto P_{A|C=c},$$

taking values in the (Polish) space of probability measures on the state space of A , endowed with the topology of weak convergence.

The Markov relations $A \rightarrow B \rightarrow C$ and $A \rightarrow C \rightarrow B$ imply that for $P_{B,C}$ -almost every (b, c) ,

$$P_{A|B=b, C=c} = P_{A|B=b} = P_{A|C=c}.$$

Thus $\Phi_B(b) = \Phi_C(c)$ for $P_{B,C}$ -a.e. (b, c) . In particular, the random variables $\Phi_B(B)$ and $\Phi_C(C)$ are almost surely equal. Define

$$U \triangleq \Phi_B(B) = \Phi_C(C) \quad \text{a.s.}$$

Then U is measurable with respect to both $\sigma(B)$ and $\sigma(C)$, and plays the role of $f(B) = g(C)$ in the finite case.

Finally, viewing U as a probability-measure-valued random variable, the tower property gives, for every bounded measurable ϕ ,

$$\mathbb{E}[\phi(A) \mid U] = \int \phi(a) U(da),$$

which justifies $P_{A|U} = U$ (that is, the conditional law of A given U is U itself, a.s.). Combined with the identity $P_{A|B,C} = P_{A|U}$ (holding P -a.s.) this implies

$$P_{A|B,C,U} = P_{A|U},$$

which is equivalent to $A \perp\!\!\!\perp (B, C) \mid U$. This yields the same conclusion as in the finite case, with f and g now understood as measurable maps into the space of conditional laws of A . ○

Proposition 1.57. (Fano's inequality) Let $X, \hat{X}$ be random variables taking values in a finite alphabet $\mathcal{X}$ with $|\mathcal{X}| \geq 2$. Then

$$H(X|\hat{X}) \leq H_b(\Pr(X \neq \hat{X})) + \Pr(X \neq \hat{X}) \log(|\mathcal{X}| - 1),$$

where $H_b(\cdot)$ is the binary entropy function. ○

Proof. Define the error indicator

$$E \triangleq \begin{cases} 1, & \text{if } X \neq \hat{X}, \\ 0, & \text{if } X = \hat{X}. \end{cases}$$

Then

$$\begin{aligned} H(X|\hat{X}) &= H(X, E|\hat{X}) && (E \text{ is a function of } (X, \hat{X})) \\ &= H(E|\hat{X}) + H(X|\hat{X}, E) && (\text{chain rule}) \\ &= H(E|\hat{X}) + \Pr(E = 1)H(X|\hat{X}, E = 1) \\ &\quad + \Pr(E = 0)H(X|\hat{X}, E = 0) \\ &= H(E|\hat{X}) + \Pr(E = 1)H(X|\hat{X}, E = 1) && (\text{since } X = \hat{X} \text{ when } E = 0) \\ &\leq H(E|\hat{X}) + \Pr(E = 1) \log(|\mathcal{X}| - 1) && (X \neq \hat{X} \text{ when } E = 1) \\ &\leq H(E) + \Pr(E = 1) \log(|\mathcal{X}| - 1) \\ &= H_b(\Pr(X \neq \hat{X})) \\ &\quad + \Pr(X \neq \hat{X}) \log(|\mathcal{X}| - 1). \end{aligned}$$

■

Lemma 1.58. (Jensen's inequality for conditional expectation) Let X be a real random variable with $E|X| < \infty$. Let ϕ be convex and finite-valued on an open interval I with $X \in I$ almost surely, and assume $\phi(X)$ is integrable from below (i.e. $E[\phi(X)^-] < \infty$, with $E[\phi(X)] = +\infty$ permitted). Then

$$E[\phi(X) | \mathcal{F}] \geq \phi(E[X | \mathcal{F}])$$

almost surely, for every sub- σ -field $\mathcal{F}$. ○

Proof. A finite convex function on the open interval I is continuous there. Choose a subgradient at each rational point of I ; the resulting *countable* family Ξ of affine minorants $ax + b$ has supremum ϕ on I , by continuity and the local boundedness of subgradients. Thus $\phi(x) \geq ax + b$ for all $(a, b) \in \Xi$ and all $x \in I$. Taking conditional expectations,

$$E[\phi(X) | \mathcal{F}] \geq E[aX + b | \mathcal{F}] = a E[X | \mathcal{F}] + b$$

almost surely, for all $(a, b) \in \Xi$. Because Ξ is countable, the almost-sure inequalities hold on a common probability-one set, and taking the supremum over $(a, b) \in \Xi$ yields

$$E[\phi(X) | \mathcal{F}] \geq \phi(E[X | \mathcal{F}]),$$

as claimed; the conditional expectation of the integrable I -valued random variable X belongs to I almost surely, so the right-hand side is well defined. For vector-valued X , the same proof runs on an open convex subset of $\mathbb{R}^d$ with inner products and a countable dense subset. ■

1.11 Gaussian optimality via rotation: general idea

We now describe the general idea behind rotation-based proofs of Gaussian optimality for information-theoretic functionals.

Let f_1 be a functional of the marginal law μ_X , and let f_2 be a functional of the joint law $\mu_{X_1 X_2}$ on $\mathcal{X} \times \mathcal{X}$. Assume:

- *Subadditivity*: For all pairs (X_1, X_2) ,

$$f_2(\mu_{X_1 X_2}) \leq f_1(\mu_{X_1}) + f_1(\mu_{X_2}),$$

with equality exactly for independent X_1, X_2 .

- *Rotational invariance*: f_2 is invariant under the normalised sum/difference transform: for $X_{\pm} = \frac{X_1 \pm X_2}{\sqrt{2}}$, $f_2(\mu_{X_1 X_2}) = f_2(\mu_{X_+ X_-})$.

Let $\mathcal{S}$ be a specified admissible class of laws on a real vector space, closed under the normalised sum/difference transform (the phrase *closed under rotations* alone would be ambiguous), and suppose f_1 has a finite attained maximum V on $\mathcal{S}$. If the sum and difference marginals of independent copies of any admissible law are admissible, then any maximiser must be Gaussian.

Indeed, let $\mu_{X^*} \in \mathcal{S}$ be a maximiser and set $V \triangleq f_1(\mu_{X^*})$. Let X_1, X_2 be independent copies of X^* , and define

$$X_+ = \frac{X_1 + X_2}{\sqrt{2}}, \quad X_- = \frac{X_1 - X_2}{\sqrt{2}}.$$

Then

$$\begin{aligned} 2V &= f_1(\mu_{X^*}) + f_1(\mu_{X^*}) \\ &= f_2(\mu_{X_1 X_2}) && \text{(equality case of subadditivity: independence)} \\ &= f_2(\mu_{X_+ X_-}) && \text{(rotational invariance)} \\ &\leq f_1(\mu_{X_+}) + f_1(\mu_{X_-}) && \text{(subadditivity)} \\ &\leq V + V, \end{aligned}$$

because X_+ and X_- are admissible candidates in $\mathcal{S}$ (their laws are sum/difference marginals of independent copies of the admissible X^*). Hence all inequalities must be equalities. The equality case of subadditivity then forces X_+ and X_- to be independent, and Bernstein's theorem (with equal covariance, as stated below) implies that X^* (and thus $X_1, X_2, X_{\pm}$) is Gaussian. In any two-letter equality statement the admissibility constraints must be stated for each supremum; here they are exactly the closure conditions above.

Remark 1.20. Many single-letterisation arguments in multi-user information theory amount to proving subadditivity of an associated functional. In Gaussian noise settings, combining such subadditivity with rotational invariance often yields Gaussian optimality with minimal additional work. Note that subadditivity and convexity are distinct properties: subadditivity of a functional does not by itself supply the convexity or compactness hypotheses that a min–max theorem requires, and vice versa. $\circ$

1.11.1 Gaussian maximises differential entropy under a covariance constraint: alternate proof

Lemma 1.59. (Gaussian maximises entropy under a covariance constraint) Under the constraint $E[XX^T] \preceq K$ with $K \succ 0$, among laws with a density and finite differential entropy the Gaussian vector $X \sim \mathcal{N}(0, K)$ maximises the differential entropy $h(X)$. $\circ$

Remark 1.21. Since h is translation invariant, the maximum is taken over zero-mean random vectors without loss of generality: for a zero-mean X , $E[XX^T]$ is its covariance matrix, and a non-zero mean only tightens the covariance constraint (see the remark after the Gaussian maximality proposition in the differential entropy section). $\circ$

Another point of view is to phrase the problem as the optimisation

$$\max_{E[XX^T] \preceq K} h(X).$$

Proving Lemma 1.59 is equivalent to establishing Gaussian optimality for this optimisation problem. One powerful technique for such Gaussian optimality is the *rotation trick*. Before presenting the second proof, we recall a lemma of Bernstein.

Lemma 1.60. (Bernstein (1941)) If $X \perp\!\!\!\perp Y$ and $(X - Y) \perp\!\!\!\perp (X + Y)$, then X and Y must be Gaussian with the same covariance matrix. $\circ$

Proof sketch. Consider the characteristic function of $(X - Y, X + Y)$:

$$\begin{aligned} \phi_{(X-Y, X+Y)}(t_1, t_2) &= E \left[e^{it_1^T(X-Y) + it_2^T(X+Y)} \right] \\ &= E \left[e^{i(t_1+t_2)^T X} \right] E \left[e^{i(t_2-t_1)^T Y} \right] \\ &= \phi_X(t_1 + t_2) \phi_Y(t_2 - t_1). \end{aligned} \tag{1.7}$$

Independence of $(X - Y)$ and $(X + Y)$ also gives

$$\begin{aligned} \phi_{(X-Y, X+Y)}(t_1, t_2) &= E \left[e^{it_1^T(X-Y)} \right] E \left[e^{it_2^T(X+Y)} \right] \\ &= \phi_X(t_1) \phi_Y(-t_1) \phi_X(t_2) \phi_Y(t_2). \end{aligned} \tag{1.8}$$

Equating (1.7) and (1.8) yields

$$\phi_X(t_1 + t_2) \phi_Y(t_2 - t_1) = \phi_X(t_1) \phi_X(t_2) \phi_Y(-t_1) \phi_Y(t_2), \tag{1.9}$$

for all (t_1, t_2) . The only solutions to this functional equation (subject to the usual properties of characteristic functions: continuity, boundedness, and $\phi(0) = 1$) are Gaussian characteristic functions of the form

$$\phi_X(t) = \exp(-t^T C t + i a^T t), \quad \phi_Y(t) = \exp(-t^T C t + i b^T t),$$

for some real vectors a, b and positive semidefinite matrix C . This proves the lemma. $\blacksquare$

Exercise 1: Solve the functional equation (1.9) in full detail and complete the proof of Lemma 1.60. You may use that $\phi_X(t)$ is uniformly continuous in t , satisfies $|\phi_X(t)| \leq 1$, and obeys $\phi_X(-t) = \phi_X^*(t)$ (complex conjugate).

Alternative proof. Let $\mathbf{X}^*$ achieve the maximum of the optimisation problem above (see Remark 1.22 for existence). Take two independent copies $\mathbf{X}_1$ and $\mathbf{X}_2$ of $\mathbf{X}^*$. Then

$$h(\mathbf{X}_1, \mathbf{X}_2) = h(\mathbf{X}_1) + h(\mathbf{X}_2) = 2h(\mathbf{X}^*).$$

Consider the orthogonal linear transformation

$$\mathbf{X}_+ = \frac{\mathbf{X}_1 + \mathbf{X}_2}{\sqrt{2}}, \quad \mathbf{X}_- = \frac{\mathbf{X}_1 - \mathbf{X}_2}{\sqrt{2}},$$

or in matrix form

$$\begin{pmatrix} \mathbf{X}_- \\ \mathbf{X}_+ \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} \mathbf{X}_1 \\ \mathbf{X}_2 \end{pmatrix}.$$

Since the transformation matrix is orthogonal (hence determinant ± 1), we have

$$h(\mathbf{X}_+, \mathbf{X}_-) = h(\mathbf{X}_1, \mathbf{X}_2) = 2h(\mathbf{X}^*).$$

On the other hand, entropy satisfies (strict) sub-additivity and,

$$h(\mathbf{X}_+, \mathbf{X}_-) = h(\mathbf{X}_+) + h(\mathbf{X}_-) - I(\mathbf{X}_+; \mathbf{X}_-),$$

and $I(\mathbf{X}_+; \mathbf{X}_-) \geq 0$, so

$$h(\mathbf{X}_+) + h(\mathbf{X}_-) \geq 2h(\mathbf{X}^*). \quad (1.10)$$

Compute the covariance of $\mathbf{X}_-$ (recall that $\mathbf{X}^*$ is zero-mean, so the second moment equals the covariance):

$$\mathbb{E}[\mathbf{X}_- \mathbf{X}_-^T] = \frac{1}{2} \left(\mathbb{E}[\mathbf{X}_1 \mathbf{X}_1^T] + \mathbb{E}[\mathbf{X}_2 \mathbf{X}_2^T] \right) = \mathbb{E}[\mathbf{X}^* (\mathbf{X}^*)^T] \preceq K.$$

By optimality of $\mathbf{X}^*$, we must have $h(\mathbf{X}_-) \leq h(\mathbf{X}^*)$. Similarly,

$$\mathbb{E}[\mathbf{X}_+ \mathbf{X}_+^T] \preceq K$$

and hence $h(\mathbf{X}_+) \leq h(\mathbf{X}^*)$. Thus

$$h(\mathbf{X}_+) + h(\mathbf{X}_-) \leq 2h(\mathbf{X}^*). \quad (1.11)$$

Combining (1.10) and (1.11) yields

$$h(\mathbf{X}_+) = h(\mathbf{X}_-) = h(\mathbf{X}^*),$$

and forces $I(\mathbf{X}_+; \mathbf{X}_-) = 0$, i.e. $\mathbf{X}_+ \perp\!\!\!\perp \mathbf{X}_-$. By Bernstein's lemma (Lemma 1.60), such a pair can only occur when $\mathbf{X}^*$ is Gaussian, say with covariance $K_* = \mathbb{E}[\mathbf{X}^*(\mathbf{X}^*)^T] \preceq K$. It remains to identify the optimal covariance: among the candidate Gaussian laws (centered, covariance $K_* \preceq K$) the entropy $\frac{d}{2} \ln(2\pi e) + \frac{1}{2} \ln |K_*|$ is maximised by the largest K_* in the Loewner order, namely $K_* = K$, because the determinant is monotone on the cone: if $K_1 \succeq K_2 \succeq 0$ then $|K_1| \geq |K_2|$ (apply the Minkowski determinant inequality just proved to K_2 and $K_1 - K_2$, or compare the ordered eigenvalues). Hence the optimiser is $\mathcal{N}(0, K)$. ■

Remark 1.22. The existence of maximisers requires an argument. Since this is the first time we use the rotation trick, and it always assumes the existence of a maximiser, we outline why such a maximiser exists.

Under the bounded second-moment constraint $\mathbb{E}[\mathbf{X}\mathbf{X}^T] \preceq K$ with $K \succ 0$, the centered admissible laws form a weakly compact class: the moment constraint passes to a weak limit by lower semicontinuity, and the zero mean passes by uniform integrability. Differential entropy, extended by $-\infty$ on singular laws, is upper semicontinuous on this uniformly moment-bounded class. Since $K \succ 0$ permits a finite-entropy Gaussian candidate and gives a finite entropy upper bound, an admissible weak limit of a maximising sequence attains the finite maximum. Indeed, for such a weak limit $\mathbf{X}_*$ of a maximising sequence,

$$h(\mathbf{X}_*) \geq \limsup_n h(\mathbf{X}_n) = V, \quad h(\mathbf{X}_*) \leq V,$$

so $\mathbf{X}_*$ attains the maximum; no convergence of the sequence of entropies is needed. The rotation argument then identifies that maximiser as Gaussian.

In the smoothing variant of this argument, the source and the smoothed output must be distinguished: $\mathbf{X} \sim \mathcal{N}(0, K)$ gives $\mathbf{X} + \delta \mathbf{Z} \sim \mathcal{N}(0, K + \delta^2 I)$ (the amplitude is δ , so the covariance grows by $\delta^2 I$), and entropy continuity after smoothing holds under the uniform moment bound. The full justification of the regularised limiting passage, including exact attainment of the unregularised supremum, is the existence argument of Anantharam–Jog–Nair 2022, to which we refer for the technical details. ○

1.11.2 Alternate proof of the EPI via Gaussian extremality (advanced, optional)

This subsection presents the general Gaussian-extremality route to the (conditional) EPI, following Anantharam–Jog–Nair 2022. It is advanced and optional on a first reading: the heat-flow proof above already provides a complete, elementary route through the EPI, and only the statement of the Gaussian-extremality theorem at the end of this subsection is used elsewhere in these notes.

More generally, suppose $a, b > 0$ and Q_X, Q_Y are positive semidefinite matrices. Consider

$$\sup_{\substack{U, X, Y: X \rightarrow U \rightarrow Y \\ 0 \preceq K_X \preceq Q_X, 0 \preceq K_Y \preceq Q_Y}} \left(a h(X|U) + b h(Y|U) - h\left(M \begin{pmatrix} X \\ Y \end{pmatrix} \middle| U\right) \right),$$

for a fixed matrix M . The Gaussian-extremality theorem of Anantharam–Jog–Nair 2022 identifies this unrestricted supremum with the corresponding supremum over independent Gaussian inputs with constant auxiliary. Attainment is asserted only under hypotheses that establish finiteness and extremizer existence; for the EPI specialisation with positive definite covariance bounds, a common positive definite Gaussian covariance supplies an attained value of zero (see the remark below). The Gaussian value then coincides with

$$\sup_{\substack{K_X \preceq Q_X \\ K_Y \preceq Q_Y}} \frac{a}{2} \log((2\pi e)^d |K_X|) + \frac{b}{2} \log((2\pi e)^d |K_Y|) - \frac{1}{2} \log\left((2\pi e)^k \left| M \begin{pmatrix} K_X & 0 \\ 0 & K_Y \end{pmatrix} M^T \right| \right),$$

where k is the number of rows of M (i.e. the dimension of the output of the linear map M ; we require M to have full row rank so that the output differential entropy is taken with respect to Lebesgue measure in the declared output dimension k).

Remark 1.23 (Equality of suprema is not unregularized attainment). The Gaussian-extremality statement is a statement about the *equality of suprema*: the unrestricted supremum coincides with the supremum over Gaussian inputs. It does not, by itself, assert that the unregularized general problem is *attained* for every $a, b > 0$ and every M , and arbitrary semidefinite upper bounds can force singular inputs. A scalar counterexample is

$$M = \begin{bmatrix} 1 & 1 \end{bmatrix}, \quad a = b = \frac{1}{4}, \quad X, Y \text{ independent } \mathcal{N}(0, t), \quad 0 < t \leq 1,$$

for which

$$a h(X) + b h(Y) - h(X + Y) = -\frac{1}{4} \ln t + \text{constant} \longrightarrow +\infty.$$

Thus finiteness and extremizer existence must be stated separately whenever they are needed. For the actual EPI specialisation ($Q_X, Q_Y \succ 0$) the situation is well behaved: choose a common $K \succ 0$ with $K \preceq Q_X, Q_Y$; independent $\mathcal{N}(0, K)$ inputs give a Lieb objective of exactly zero, and once the Gaussian reduction and the log-determinant bound below are established this exhibits an attained optimum with constant U . $\circ$

Remark 1.24. To avoid technical issues, we may introduce small Gaussian perturbations. Specifically, we replace X by $X + \sqrt{\delta} Z_a$ and Y by $Y + \sqrt{\delta} Z_b$, where Z_a, Z_b are independent standard Gaussians independent of (X, Y) and $\delta > 0$ is small. The regularisation must *match the optimisation class*: the same δ and ε are used for both coordinates and for both copies in the two-letter problem, the admissible class is the set of joint laws with the fixed marginals and the Markov structure, and the extra strictness entropy perturbation (the $-\varepsilon h(\cdot|U)$ term) is chosen so that the regularised problem is attained and the unregularised value is approached. The adaptation between the two regularisations (δ fixed, $\varepsilon \downarrow 0$, then $\delta \downarrow 0$) and the final limiting passage are technical steps delegated to the cited proof Anantharam et al. (2022), not an inference from pointwise convergence alone. $\circ$

Sketch of Gaussian optimality. Given fixed marginals p_X, p_Y , define the functional as a supremum over *admissible joint laws* (U, X, Y) with these fixed marginals and the Markov structure $X \rightarrow U \rightarrow Y$ (i.e. over admissible couplings, not over arbitrary laws):

$$F(p_X, p_Y) \triangleq \sup_{U: X \rightarrow U \rightarrow Y} \left(a h(X + \sqrt{\delta} Z_a \mid U) + b h(Y + \sqrt{\delta} Z_b \mid U) \right. \\ \left. - h \left(M \begin{pmatrix} X + \sqrt{\delta} Z_a \\ Y + \sqrt{\delta} Z_b \end{pmatrix} \middle| U \right) - \varepsilon h \left(\begin{pmatrix} X + \sqrt{\delta} Z_a + \sqrt{\delta} Z_c \\ Y + \sqrt{\delta} Z_b + \sqrt{\delta} Z_d \end{pmatrix} \middle| U \right) \right),$$

where Z_a, Z_b, Z_c, Z_d are *mutually independent standard Gaussian random vectors* of the required dimensions, independent of all source variables (U, X, Y) .

We now show that $F(p_X, p_Y)$ is sub-additive, in the sense that

$$F^{(2)}(p_{X_1 X_2}, p_{Y_1 Y_2}) \leq F(p_{X_1}, p_{Y_1}) + F(p_{X_2}, p_{Y_2}).$$

Here

$$F^{(2)}(p_{X_1 X_2}, p_{Y_1 Y_2}) \triangleq \sup_{U: (X_1, X_2) \rightarrow U \rightarrow (Y_1, Y_2)} \left(a h(X_1 + \sqrt{\delta} Z_{a1}, X_2 + \sqrt{\delta} Z_{a2} \mid U) \right. \\ \left. + b h(Y_1 + \sqrt{\delta} Z_{b1}, Y_2 + \sqrt{\delta} Z_{b2} \mid U) \right. \\ \left. - h \left(M \begin{pmatrix} X_1 + \sqrt{\delta} Z_{a1} \\ Y_1 + \sqrt{\delta} Z_{b1} \end{pmatrix}, M \begin{pmatrix} X_2 + \sqrt{\delta} Z_{a2} \\ Y_2 + \sqrt{\delta} Z_{b2} \end{pmatrix} \middle| U \right) \right. \\ \left. - \varepsilon h \left(\begin{pmatrix} X_1 + \sqrt{\delta} Z_{a1} + \sqrt{\delta} Z_{c1} \\ Y_1 + \sqrt{\delta} Z_{b1} + \sqrt{\delta} Z_{d1} \end{pmatrix}, \right. \right. \\ \left. \left. \begin{pmatrix} X_2 + \sqrt{\delta} Z_{a2} + \sqrt{\delta} Z_{c2} \\ Y_2 + \sqrt{\delta} Z_{b2} + \sqrt{\delta} Z_{d2} \end{pmatrix} \middle| U \right) \right),$$

where $Z_{a1}, Z_{a2}, Z_{b1}, Z_{b2}, Z_{c1}, Z_{c2}, Z_{d1}, Z_{d2}$ are independent standard Gaussian random variables, independent of everything else.

Observe that the inner functional can be decomposed as

$$a h(X_1 + \sqrt{\delta} Z_{a1}, X_2 + \sqrt{\delta} Z_{a2} \mid U) + b h(Y_1 + \sqrt{\delta} Z_{b1}, Y_2 + \sqrt{\delta} Z_{b2} \mid U) \\ - h \left(M \begin{pmatrix} X_1 + \sqrt{\delta} Z_{a1} \\ Y_1 + \sqrt{\delta} Z_{b1} \end{pmatrix}, M \begin{pmatrix} X_2 + \sqrt{\delta} Z_{a2} \\ Y_2 + \sqrt{\delta} Z_{b2} \end{pmatrix} \middle| U \right) \\ - \varepsilon h \left(\begin{pmatrix} X_1 + \sqrt{\delta} Z_{a1} + \sqrt{\delta} Z_{c1} \\ Y_1 + \sqrt{\delta} Z_{b1} + \sqrt{\delta} Z_{d1} \end{pmatrix}, \begin{pmatrix} X_2 + \sqrt{\delta} Z_{a2} + \sqrt{\delta} Z_{c2} \\ Y_2 + \sqrt{\delta} Z_{b2} + \sqrt{\delta} Z_{d2} \end{pmatrix} \middle| U \right) \\ = \left(a h(X_1 + \sqrt{\delta} Z_{a1} \mid U) + b h(Y_1 + \sqrt{\delta} Z_{b1} \mid U) \right.$$

$$\begin{aligned}
& -h\left(M\left(\begin{matrix} X_1 + \sqrt{\delta} Z_{a1} \\ Y_1 + \sqrt{\delta} Z_{b1} \end{matrix}\right) \middle| U\right) - \varepsilon h\left(\begin{matrix} X_1 + \sqrt{\delta} Z_{a1} + \sqrt{\delta} Z_{c1} \\ Y_1 + \sqrt{\delta} Z_{b1} + \sqrt{\delta} Z_{d1} \end{matrix} \middle| U\right) \\
& + \left(a h(X_2 + \sqrt{\delta} Z_{a2} \mid U, X_1 + \sqrt{\delta} Z_{a1}) + b h(Y_2 + \sqrt{\delta} Z_{b2} \mid U, Y_1 + \sqrt{\delta} Z_{b1})\right. \\
& \quad \left.- h\left(M\left(\begin{matrix} X_2 + \sqrt{\delta} Z_{a2} \\ Y_2 + \sqrt{\delta} Z_{b2} \end{matrix}\right) \middle| U, M\left(\begin{matrix} X_1 + \sqrt{\delta} Z_{a1} \\ Y_1 + \sqrt{\delta} Z_{b1} \end{matrix}\right)\right)\right. \\
& \quad \left.- \varepsilon h\left(\begin{matrix} X_2 + \sqrt{\delta} Z_{a2} + \sqrt{\delta} Z_{c2} \\ Y_2 + \sqrt{\delta} Z_{b2} + \sqrt{\delta} Z_{d2} \end{matrix} \middle| U, \begin{matrix} X_1 + \sqrt{\delta} Z_{a1} + \sqrt{\delta} Z_{c1} \\ Y_1 + \sqrt{\delta} Z_{b1} + \sqrt{\delta} Z_{d1} \end{matrix}\right)\right).
\end{aligned}$$

By the distributional assumptions,

$$\begin{aligned}
h(X_2 + \sqrt{\delta} Z_{a2} \mid U, X_1 + \sqrt{\delta} Z_{a1}) &= h(X_2 + \sqrt{\delta} Z_{a2} \mid U, X_1 + \sqrt{\delta} Z_{a1}, Y_1 + \sqrt{\delta} Z_{b1}), \\
h(Y_2 + \sqrt{\delta} Z_{b2} \mid U, Y_1 + \sqrt{\delta} Z_{b1}) &= h(Y_2 + \sqrt{\delta} Z_{b2} \mid U, X_1 + \sqrt{\delta} Z_{a1}, Y_1 + \sqrt{\delta} Z_{b1}),
\end{aligned}$$

and, by the chain rule (using that $M\left(\begin{smallmatrix} X_1 + \sqrt{\delta} Z_{a1} \\ Y_1 + \sqrt{\delta} Z_{b1} \end{smallmatrix}\right)$ is a function of $(X_1 + \sqrt{\delta} Z_{a1}, Y_1 + \sqrt{\delta} Z_{b1})$),

$$\begin{aligned}
& h\left(M\left(\begin{matrix} X_2 + \sqrt{\delta} Z_{a2} \\ Y_2 + \sqrt{\delta} Z_{b2} \end{matrix}\right) \middle| U, M\left(\begin{matrix} X_1 + \sqrt{\delta} Z_{a1} \\ Y_1 + \sqrt{\delta} Z_{b1} \end{matrix}\right)\right) \\
& = h\left(M\left(\begin{matrix} X_2 + \sqrt{\delta} Z_{a2} \\ Y_2 + \sqrt{\delta} Z_{b2} \end{matrix}\right) \middle| U, X_1 + \sqrt{\delta} Z_{a1}, Y_1 + \sqrt{\delta} Z_{b1}\right) \\
& \quad + I\left(M\left(\begin{matrix} X_2 + \sqrt{\delta} Z_{a2} \\ Y_2 + \sqrt{\delta} Z_{b2} \end{matrix}\right); X_1 + \sqrt{\delta} Z_{a1}, Y_1 + \sqrt{\delta} Z_{b1} \middle| U, M\left(\begin{matrix} X_1 + \sqrt{\delta} Z_{a1} \\ Y_1 + \sqrt{\delta} Z_{b1} \end{matrix}\right)\right),
\end{aligned}$$

and similarly

$$\begin{aligned}
& h\left(\begin{matrix} X_2 + \sqrt{\delta} Z_{a2} + \sqrt{\delta} Z_{c2} \\ Y_2 + \sqrt{\delta} Z_{b2} + \sqrt{\delta} Z_{d2} \end{matrix} \middle| U, \begin{matrix} X_1 + \sqrt{\delta} Z_{a1} + \sqrt{\delta} Z_{c1} \\ Y_1 + \sqrt{\delta} Z_{b1} + \sqrt{\delta} Z_{d1} \end{matrix}\right) \\
& = h\left(\begin{matrix} X_2 + \sqrt{\delta} Z_{a2} + \sqrt{\delta} Z_{c2} \\ Y_2 + \sqrt{\delta} Z_{b2} + \sqrt{\delta} Z_{d2} \end{matrix} \middle| U, X_1 + \sqrt{\delta} Z_{a1}, Y_1 + \sqrt{\delta} Z_{b1}\right) \\
& \quad + I\left(\begin{matrix} X_2 + \sqrt{\delta} Z_{a2} + \sqrt{\delta} Z_{c2} \\ Y_2 + \sqrt{\delta} Z_{b2} + \sqrt{\delta} Z_{d2} \end{matrix}; X_1 + \sqrt{\delta} Z_{a1}, Y_1 + \sqrt{\delta} Z_{b1} \middle| U, \begin{matrix} X_1 + \sqrt{\delta} Z_{a1} + \sqrt{\delta} Z_{c1} \\ Y_1 + \sqrt{\delta} Z_{b1} + \sqrt{\delta} Z_{d1} \end{matrix}\right).
\end{aligned}$$

Substituting these identities into the decomposition above, the second big bracket can be written as

$$\underbrace{\left(\text{same form as } F(p_{X_2}, p_{Y_2}) \text{ with auxiliary } U_2\right)}_{\leq F(p_{X_2}, p_{Y_2})} - \Delta,$$

where the *gap term* Δ is

$$\begin{aligned} \Delta \triangleq & I \left(M \left(\begin{matrix} X_2 + \sqrt{\delta} Z_{a2} \\ Y_2 + \sqrt{\delta} Z_{b2} \end{matrix} ; X_1 + \sqrt{\delta} Z_{a1}, Y_1 + \sqrt{\delta} Z_{b1} \right) \middle| U, M \left(\begin{matrix} X_1 + \sqrt{\delta} Z_{a1} \\ Y_1 + \sqrt{\delta} Z_{b1} \end{matrix} \right) \right) \\ & + \varepsilon I \left(\left(\begin{matrix} X_2 + \sqrt{\delta} Z_{a2} + \sqrt{\delta} Z_{c2} \\ Y_2 + \sqrt{\delta} Z_{b2} + \sqrt{\delta} Z_{d2} \end{matrix} \right) ; X_1 + \sqrt{\delta} Z_{a1}, Y_1 + \sqrt{\delta} Z_{b1} \right) \middle| U, \left(\begin{matrix} X_1 + \sqrt{\delta} Z_{a1} + \sqrt{\delta} Z_{c1} \\ Y_1 + \sqrt{\delta} Z_{b1} + \sqrt{\delta} Z_{d1} \end{matrix} \right) \right). \end{aligned}$$

Both mutual information terms are non-negative, hence $\Delta = \Delta(U) \geq 0$; the gap depends on the auxiliary U .

Now identify

$$U_1 \triangleq U, \quad U_2 \triangleq (U, X_1 + \sqrt{\delta} Z_{a1}, Y_1 + \sqrt{\delta} Z_{b1}).$$

The first big bracket is at most $F(p_{X_1}, p_{Y_1})$, and the second big bracket is at most $F(p_{X_2}, p_{Y_2}) - \Delta(U)$. Therefore, for every admissible U ,

$$\text{inner functional} \leq F(p_{X_1}, p_{Y_1}) + F(p_{X_2}, p_{Y_2}) - \Delta(U).$$

Taking the supremum over U yields either the plain subadditivity

$$F^{(2)}(p_{X_1 X_2}, p_{Y_1 Y_2}) \leq F(p_{X_1}, p_{Y_1}) + F(p_{X_2}, p_{Y_2}),$$

or the stronger form with the loss

$$F^{(2)}(p_{X_1 X_2}, p_{Y_1 Y_2}) \leq F(p_{X_1}, p_{Y_1}) + F(p_{X_2}, p_{Y_2}) - \inf_U \Delta(U),$$

but one must *not* retain an unspecified free Δ after optimising over U . The gap term $\Delta(U) \geq 0$ is used in the rotation argument below at a particular constructed auxiliary V .

The envelope $F(p_X, p_Y)$ is defined with fixed marginals, and the rotation step below changes those marginals. Define the *global* covariance-constrained value

$$v = v_{\delta, \varepsilon}(Q_X, Q_Y) \triangleq \sup_{P_X, P_Y: K_X \preceq Q_X, K_Y \preceq Q_Y} F_{\delta, \varepsilon}(P_X, P_Y),$$

over the appropriate admissible, centered laws under the covariance conventions specified above, and let (U, X, Y) be a maximising joint law for this global problem. (Existence of such a maximiser requires justification; the covariance constraint together with the added Gaussian perturbations ensure tightness and continuity properties that guarantee existence. For details, we refer the reader to [Anantharam et al. \(2022\)](#) or [Geng and Nair \(2014\)](#).)

Take two independent copies

$$(U_1, X_1, Y_1), \quad (U_2, X_2, Y_2),$$

each distributed as (U, X, Y) and independent of each other. Define the rotated variables

$$X_{\pm} \triangleq \frac{X_1 \pm X_2}{\sqrt{2}}, \quad Y_{\pm} \triangleq \frac{Y_1 \pm Y_2}{\sqrt{2}}.$$

Since this is an orthogonal linear transformation, the two-letter functional evaluated at

$$((X_1, Y_1), (X_2, Y_2)) \quad \text{and at} \quad ((X_+, Y_+), (X_-, Y_-))$$

has the same value.

By construction, $F(p_X, p_Y) = v$ (the chosen law attains the global supremum v), and the independent copies give $F(p_{X_1}, p_{Y_1}) = F(p_{X_2}, p_{Y_2}) = v$. Let $V \triangleq (U_1, U_2)$ be the auxiliary obtained from the two independent copies of the optimiser, and rotate the sources together with their independent Gaussian noises (so that the fresh noises of the rotated blocks remain independent standard Gaussians). Write $\Phi^{(2)}(V)$ for the *particular inner objective* of the two-letter functional, evaluated at this constructed auxiliary V (not its supremum). Apply the fixed-auxiliary decomposition to the *rotated* pair $((X_+, Y_+), (X_-, Y_-))$ with this particular auxiliary V :

$$\Phi_{\text{rot}}^{(2)}(V) \leq F(p_{X_+}, p_{Y_+}) + F(p_{X_-}, p_{Y_-}) - \Delta(V),$$

where $\Delta(V) \geq 0$ is the gap term evaluated at this auxiliary (a sum of non-negative conditional mutual informations).

On the other hand, since (X_+, Y_+) and (X_-, Y_-) are obtained from two independent copies of the maximiser by an orthogonal transform (applied to sources and noises together), the particular inner objective at V attains the value $2v$:

$$\Phi_{\text{rot}}^{(2)}(V) = 2v,$$

the equality case of subadditivity (independence) giving exactly $2v$ at V . The rotated marginals remain covariance-admissible, because $\text{Cov}(X_{\pm}) = \text{Cov}(X)$ and $\text{Cov}(Y_{\pm}) = \text{Cov}(Y)$ for the normalised sum/difference transform of independent copies; hence by the global optimality of v ,

$$F(p_{X_+}, p_{Y_+}) \leq v, \quad F(p_{X_-}, p_{Y_-}) \leq v.$$

Therefore

$$2v = \Phi_{\text{rot}}^{(2)}(V) \leq F(p_{X_+}, p_{Y_+}) + F(p_{X_-}, p_{Y_-}) - \Delta(V) \leq v + v - \Delta(V) = 2v - \Delta(V),$$

which forces $\Delta(V) = 0$ at this constructed auxiliary. No independence of the rotated blocks is assumed at this stage; their independence is a later equality consequence. (Equality of the *optimised* $F^{(2)}$ with $2v$ follows afterwards from the lower bound furnished by this witness together with the plain subadditivity upper bound; it is not assumed first to justify the fixed-witness loss estimate.)

In particular we have

$$I\left(\begin{pmatrix} X_- + \sqrt{\delta} Z_{a-} + \sqrt{\delta} Z_{c-} \\ Y_- + \sqrt{\delta} Z_{b-} + \sqrt{\delta} Z_{d-} \end{pmatrix}; X_+ + \sqrt{\delta} Z_{a+}, Y_+ + \sqrt{\delta} Z_{b+} \middle| (U_1, U_2), \begin{pmatrix} X_+ + \sqrt{\delta} Z_{a+} + \sqrt{\delta} Z_{c+} \\ Y_+ + \sqrt{\delta} Z_{b+} + \sqrt{\delta} Z_{d+} \end{pmatrix}\right) = 0.$$

Therefore, the vanishing of this conditional mutual information is equivalent to the Markov chain

$$\begin{aligned} \begin{pmatrix} X_- + \sqrt{\delta} Z_{a-} + \sqrt{\delta} Z_{c-} \\ Y_- + \sqrt{\delta} Z_{b-} + \sqrt{\delta} Z_{d-} \end{pmatrix} &\longrightarrow \left((U_1, U_2), \begin{pmatrix} X_+ + \sqrt{\delta} Z_{a+} + \sqrt{\delta} Z_{c+} \\ Y_+ + \sqrt{\delta} Z_{b+} + \sqrt{\delta} Z_{d+} \end{pmatrix} \right) \\ &\longrightarrow (X_+ + \sqrt{\delta} Z_{a+}, Y_+ + \sqrt{\delta} Z_{b+}). \end{aligned}$$

On the other hand, by construction (Z_{c+}, Z_{d+}) is independent of everything else, so we also have the Markov chain

$$\begin{aligned} \begin{pmatrix} X_- + \sqrt{\delta} Z_{a-} + \sqrt{\delta} Z_{c-} \\ Y_- + \sqrt{\delta} Z_{b-} + \sqrt{\delta} Z_{d-} \end{pmatrix} &\longrightarrow \left((U_1, U_2), X_+ + \sqrt{\delta} Z_{a+}, Y_+ + \sqrt{\delta} Z_{b+} \right) \\ &\longrightarrow \begin{pmatrix} X_+ + \sqrt{\delta} Z_{a+} + \sqrt{\delta} Z_{c+} \\ Y_+ + \sqrt{\delta} Z_{b+} + \sqrt{\delta} Z_{d+} \end{pmatrix}. \end{aligned}$$

Thus, conditional on (U_1, U_2) , we have a double Markovity situation:

$$A \rightarrow B \rightarrow C \quad \text{and} \quad A \rightarrow C \rightarrow B,$$

where

$$A \equiv \begin{pmatrix} X_- + \sqrt{\delta} Z_{a-} + \sqrt{\delta} Z_{c-} \\ Y_- + \sqrt{\delta} Z_{b-} + \sqrt{\delta} Z_{d-} \end{pmatrix}, \quad B \equiv \begin{pmatrix} X_+ + \sqrt{\delta} Z_{a+} \\ Y_+ + \sqrt{\delta} Z_{b+} \end{pmatrix}, \quad C \equiv \begin{pmatrix} X_+ + \sqrt{\delta} Z_{a+} + \sqrt{\delta} Z_{c+} \\ Y_+ + \sqrt{\delta} Z_{b+} + \sqrt{\delta} Z_{d+} \end{pmatrix}.$$

We apply the double-Markovity lemma (Lemma 1.56) conditionally on $(U_1, U_2) = v$. For almost every fixed v , the conditional joint density of (B, C) given v is strictly positive on all of $\mathbb{R}^{2d} \times \mathbb{R}^{2d}$:

$$p_{B,C|v}(b, c) = p_{B|v}(b) \varphi_\delta(c - b) > 0,$$

where φ_δ is the Gaussian density of the added noise (this uses that $C = B + \text{independent Gaussian noise}$, so $C = B + Z$ has positive joint density). Marginal full support alone would *not* suffice: for example $B = C$ Gaussian has full marginal support and nontrivial common information. Positive joint density and Fubini's theorem force a common measurable function of B and C to be constant almost surely, for almost every fixed v ; hence the latent variable of the double-Markovity lemma is trivial, and

$$A \perp\!\!\!\perp (B, C) \mid (U_1, U_2),$$

and in particular,

$$\begin{pmatrix} X_- + \sqrt{\delta} Z_{a-} + \sqrt{\delta} Z_{c-} \\ Y_- + \sqrt{\delta} Z_{b-} + \sqrt{\delta} Z_{d-} \end{pmatrix} \perp\!\!\!\perp (X_+ + \sqrt{\delta} Z_{a+}, Y_+ + \sqrt{\delta} Z_{b+}) \mid (U_1, U_2).$$

Lemma 1.61 (proved below) states: if Z is Gaussian and independent of (X, Y) , and $X \perp\!\!\!\perp (Y + Z)$, then $X \perp\!\!\!\perp Y$; the same argument works for random vectors, applied to whole blocks. We apply the vector version to the rotated variables: the Gaussian noise block $(Z_{a-}, Z_{b-}, Z_{c-}, Z_{d-})$ is

independent of $(X_+, Y_+, X_-, Y_-, U_1, U_2)$, and since independent Gaussian noise was added to both blocks, the Gaussian factor of each block is cancelled separately (equivalently, remove the two noises successively), giving

$$(X_-, Y_-) \perp\!\!\!\perp (X_+, Y_+) \mid (U_1, U_2).$$

Thus, at the maximiser, the rotated “sum” and “difference” variables are conditionally independent given (U_1, U_2) .

By construction of the functional, the Markov condition $X \rightarrow U \rightarrow Y$ holds at the maximiser, so in particular

$$X \perp\!\!\!\perp Y \mid U_1,$$

and similarly for U_2 on the second copy. Together with the conditional independence of (X_+, Y_+) and (X_-, Y_-) given (U_1, U_2) , this implies that for almost every (u_1, u_2) the pair of *independent* copies $(X_1, Y_1), (X_2, Y_2) \mid \{U_1 = u_1, U_2 = u_2\}$ (which need not be identically distributed: u_1 and u_2 are independent values of the auxiliary) satisfies the assumptions of Bernstein’s theorem (Lemma 1.60) for suitable linear combinations (sum and difference). Hence, for almost every (u_1, u_2) , the conditional laws of $(X, Y) \mid \{U_1 = u_1\}$ and $(X, Y) \mid \{U_2 = u_2\}$ are jointly Gaussian with the *same* covariance matrix, and X and Y are independent given $U_1 = u_1$, so this covariance is block-diagonal.

Finally, suppose $K(u_1) \neq K(u_2)$ on a set of pairs (u_1, u_2) of positive $P_{U_1 U_2}$ -probability. For almost every such pair, the conditional independence established above makes the conditional cross-covariance vanish:

$$\text{Cov}((X_+, Y_+), (X_-, Y_-) \mid U_1 = u_1, U_2 = u_2) = 0.$$

On the other hand, from the covariance calculation,

$$\text{Cov}((X_+, Y_+), (X_-, Y_-) \mid U_1 = u_1, U_2 = u_2) = \frac{1}{2} (K_{u_1} - K_{u_2}),$$

where

$$K_{u_1} \triangleq \text{Cov}((X_1, Y_1) \mid U_1 = u_1), \quad K_{u_2} \triangleq \text{Cov}((X_2, Y_2) \mid U_2 = u_2).$$

Equating the two expressions yields

$$\frac{1}{2} (K_{u_1} - K_{u_2}) = 0 \quad \text{almost surely,}$$

hence

$$K_{u_1} = K_{u_2} \quad \text{for } P_{U_1 U_2}\text{-almost every } (u_1, u_2).$$

This holds for $P_{U_1 U_2}$ -almost every (u_1, u_2) , so the conditional covariance matrix

$$K_u \triangleq \text{Cov}((X, Y) \mid U = u)$$

is essentially independent of u .

We have shown that K_u is essentially constant in u ; that is, there exists a fixed block-diagonal matrix

$$K = \begin{pmatrix} K_X & 0 \\ 0 & K_Y \end{pmatrix}$$

such that

$$\text{Cov}((X, Y) \mid U = u) = K \quad \text{for almost every } u.$$

The conditional *means* may still vary with u ; replace the conditional laws by centered Gaussians with the common block-diagonal covariance K . Translation does not change the conditional objective (entropies and the linear-combination term are translation invariant), and the total (unconditional) covariance of the centered Gaussian family equals K , so

$$K_X \preceq Q_X, \quad K_Y \preceq Q_Y,$$

i.e. the common conditional covariance satisfies the same covariance constraints as in the optimisation problem. This is exactly why the auxiliary U may be taken to be constant.

Letting first $\varepsilon \rightarrow 0$ and then $\delta \rightarrow 0$, the regularised functional converges to the original one, and Gaussian optimality passes to the limit; this limiting passage is a technical step delegated to the cited proof [Anantharam et al. \(2022\)](#), not an inference from pointwise convergence alone. In particular, we obtain:

Consider

$$\sup_{\substack{U, X, Y: X \rightarrow U \rightarrow Y \\ 0 \preceq K_X \preceq Q_X, 0 \preceq K_Y \preceq Q_Y}} \left(a h(X \mid U) + b h(Y \mid U) - h\left(M \begin{pmatrix} X \\ Y \end{pmatrix} \mid U\right) \right),$$

for a full-row-rank matrix M . Under the admissibility hypotheses of the Gaussian-extremality theorem, the unrestricted supremum equals the corresponding supremum over independent Gaussian inputs with a constant auxiliary. This is an equality of suprema, not an assertion of attainment for every $a, b > 0$; attainment is asserted separately when finiteness and extremizer existence have been established, as in the EPI specialisation with positive definite covariance bounds. The Gaussian candidates are independent pairs (X, Y) with block-diagonal covariance $\text{diag}(K_X, K_Y)$ and a constant auxiliary; their supremum is the determinant optimisation displayed below. No attainment of this general covariance optimisation is claimed here; attainment in the EPI specialisation is established separately:

$$\sup_{\substack{0 \prec K_X \preceq Q_X \\ 0 \prec K_Y \preceq Q_Y}} \frac{a}{2} \log((2\pi e)^d |K_X|) + \frac{b}{2} \log((2\pi e)^d |K_Y|) - \frac{1}{2} \log\left((2\pi e)^k \left| M \begin{pmatrix} K_X & 0 \\ 0 & K_Y \end{pmatrix} M^T \right| \right),$$

where k is the number of rows of M (i.e. the dimension of the output MX), the covariance optimisation is over positive definite K_X, K_Y satisfying the upper bounds with nonsingular output covariance, and singular limits, if needed, are taken separately (differences of divergent log-determinants are not evaluated).

In the particular case $a = \lambda$, $b = 1 - \lambda$ (with $0 < \lambda < 1$) and

$$M = \begin{bmatrix} \sqrt{\lambda} I_d & \sqrt{1 - \lambda} I_d \end{bmatrix},$$

the optimisation problem becomes

$$\sup_{\substack{U, X, Y: X \rightarrow U \rightarrow Y \\ 0 \preceq K_X \preceq Q_X, 0 \preceq K_Y \preceq Q_Y}} \left(\lambda h(X | U) + (1 - \lambda) h(Y | U) - h(\sqrt{\lambda} X + \sqrt{1 - \lambda} Y | U) \right).$$

By the Gaussian optimality argument above, the unrestricted supremum coincides with the Gaussian supremum, and for this EPI specialisation (with $Q_X, Q_Y \succ 0$) it is attained by jointly Gaussian (X, Y) with $X \perp\!\!\!\perp Y$, covariance

$$\text{Cov}((X, Y)) = \begin{pmatrix} K_X & 0 \\ 0 & K_Y \end{pmatrix}, \quad 0 \prec K_X \preceq Q_X, 0 \prec K_Y \preceq Q_Y,$$

and U constant. Using the entropy representation, the optimal Gaussian value can be written as

$$\sup_{\substack{0 \prec K_X \preceq Q_X \\ 0 \prec K_Y \preceq Q_Y}} \frac{\lambda}{2} \log((2\pi e)^d |K_X|) + \frac{1 - \lambda}{2} \log((2\pi e)^d |K_Y|) - \frac{1}{2} \log((2\pi e)^d |\lambda K_X + (1 - \lambda) K_Y|).$$

Cancelling the common constants, this reduces to

$$\sup_{\substack{0 \prec K_X \preceq Q_X \\ 0 \prec K_Y \preceq Q_Y}} \frac{\lambda}{2} \log |K_X| + \frac{1 - \lambda}{2} \log |K_Y| - \frac{1}{2} \log |\lambda K_X + (1 - \lambda) K_Y| \leq 0.$$

Here, the last inequality follows from the concavity of the log-determinant: for any $0 < \lambda < 1$ and positive definite K_X, K_Y ,

$$\lambda \log |K_X| + (1 - \lambda) \log |K_Y| \leq \log |\lambda K_X + (1 - \lambda) K_Y|.$$

Therefore, for every U such that $X \rightarrow U \rightarrow Y$, under the hypotheses of the preceding conditional EPI theorem (conditional densities and finite conditional entropies of X, Y and $\sqrt{\lambda} X + \sqrt{1 - \lambda} Y$ given U), we have

$$\lambda h(X | U) + (1 - \lambda) h(Y | U) - h(\sqrt{\lambda} X + \sqrt{1 - \lambda} Y | U) \leq 0.$$

This is Lieb's equivalent form of the conditional entropy power inequality (recall that h is measured in nats in this section).

The following *concavity form* of the entropy power is equivalent to Lieb's inequality above (both are equivalent to the EPI); in terms of entropy power it reads

$$e^{\frac{2}{d} h(\sqrt{\lambda} X + \sqrt{1 - \lambda} Y | U)} \geq \lambda e^{\frac{2}{d} h(X | U)} + (1 - \lambda) e^{\frac{2}{d} h(Y | U)},$$

which is the conditional entropy power inequality in its entropy-power form. Since $e^{\frac{2}{d}h_{\text{nat}}(\cdot)} = 2^{\frac{2}{d}h_2(\cdot)}$, the entropy-power form is independent of the base of the logarithm, and in base 2 it reads

$$2^{\frac{2}{d}h_2(\sqrt{\lambda}X + \sqrt{1-\lambda}Y|U)} \geq \lambda 2^{\frac{2}{d}h_2(X|U)} + (1-\lambda) 2^{\frac{2}{d}h_2(Y|U)}.$$

On the equivalence. The equivalence of Lieb's family and the entropy-power family uses *all* weights and rescalings, as in the equivalence proof above: exponentiating a single fixed-weight Lieb inequality only produces a geometric-mean lower bound, $e^{\frac{2}{d}h(\sqrt{1-\lambda}X + \sqrt{\lambda}Y|U)} \geq e^{\frac{2}{d}h(X|U) \cdot (1-\lambda)} e^{\frac{2}{d}h(Y|U) \cdot \lambda}$, and does not by itself produce the arithmetic sum in the (conditional) EPI. ■

Lemma 1.61. Let Z be Gaussian and independent of (X, Y) . If $X \perp\!\!\!\perp (Y + Z)$, then $X \perp\!\!\!\perp Y$. The same statement holds for random vectors $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$ (with $\mathbf{Z}$ Gaussian) with the identical proof, using vector arguments t_1, t_2 ; if independent Gaussian noise is added to *both* blocks, cancel the Gaussian factor on each block, or remove the noises successively. ○

Proof. Consider the characteristic function of $(X, Y + Z)$. Since $X \perp\!\!\!\perp (Y + Z)$ and $Y \perp\!\!\!\perp Z$,

$$\mathbb{E} \left[e^{i(t_1X + t_2(Y+Z))} \right] = \mathbb{E} \left[e^{it_1X} \right] \mathbb{E} \left[e^{it_2(Y+Z)} \right] = \mathbb{E} \left[e^{it_1X} \right] \mathbb{E} \left[e^{it_2Y} \right] \mathbb{E} \left[e^{it_2Z} \right].$$

On the other hand, since $(X, Y) \perp\!\!\!\perp Z$,

$$\mathbb{E} \left[e^{i(t_1X + t_2(Y+Z))} \right] = \mathbb{E} \left[e^{i(t_1X + t_2Y)} \right] \mathbb{E} \left[e^{it_2Z} \right].$$

Equating the two expressions and using that $\mathbb{E}[e^{it_2Z}] \neq 0$ for Gaussian Z , we obtain

$$\mathbb{E} \left[e^{i(t_1X + t_2Y)} \right] = \mathbb{E} \left[e^{it_1X} \right] \mathbb{E} \left[e^{it_2Y} \right]$$

for all (t_1, t_2) , which proves independence of X and Y . ■

1.11.3 Geometric interpretation of the EPI

The EPI admits a geometric analogue reminiscent of the Brunn–Minkowski theorem. The Brunn–Minkowski theorem states that for non-empty compact sets $A, B \subseteq \mathbb{R}^n$,

$$(\text{vol}(A + B))^{1/n} \geq (\text{vol}(A))^{1/n} + (\text{vol}(B))^{1/n}.$$

If X, Y are uniform on compact sets A, B , then $h(X) = \ln \text{vol}(A)$ and $h(Y) = \ln \text{vol}(B)$. Applying the EPI to (X, Y) gives

$$\text{vol}(A + B)^{2/n} \geq \text{vol}(A)^{2/n} + \text{vol}(B)^{2/n},$$

which is similar in form to Brunn–Minkowski, but *weaker*: writing $a = \text{vol}(A)^{1/n}$ and $b = \text{vol}(B)^{1/n}$, the EPI inequality gives $\text{vol}(A + B)^{1/n} \geq \sqrt{a^2 + b^2}$, while Brunn–Minkowski gives the

stronger $\text{vol}(A + B)^{1/n} \geq a + b$ (the exponent $2/n$ of the EPI form is *larger* than the exponent $1/n$, which is what makes that form the weaker statement). These are useful analogies, not an identification of the two operations.

For Gaussian laws, the EPI becomes Minkowski's determinant inequality: taking $X \sim \mathcal{N}(0, K_1)$ and $Y \sim \mathcal{N}(0, K_2)$ independent, the Gaussian entropy powers are proportional to $\det(K_1)^{1/n}$ and $\det(K_2)^{1/n}$ respectively (and that of $X + Y$ to $\det(K_1 + K_2)^{1/n}$), so

$$\det(K_1 + K_2)^{1/n} \geq \det(K_1)^{1/n} + \det(K_2)^{1/n},$$

which says that Gaussian addition adds *covariance matrices*. This is different from the Minkowski addition of the corresponding ellipsoids $A_i = \{x : x^T K_i^{-1} x \leq 1\}$: the Minkowski sum $A_1 + A_2$ of ellipsoids is generally *not* an ellipsoid, and it does not generally add the shape matrices K_i . The two operations (covariance addition and Minkowski addition of ellipsoids) do not coincide, so no coincident ellipsoid statement is claimed.

1.12 Perturbation method

We now review some useful properties of mutual information using a perturbation method. Recall that $I(X; Y)$ is a functional of the input distribution P_X and the channel $W_{Y|X}$, which we will denote by $I(P_X; W)$ in this subsection (following Csiszár–Körner).

Lemma 1.62. $I(X; Y)$ is concave in P_X . ○

Proof. Note that

$$H(Y|X) = \sum_x P_X(x) H(Y|X = x)$$

is linear in P_X , while $H(Y) = H(P_X W_{Y|X})$ is concave in P_X because entropy is concave in the underlying distribution. Thus

$$I(X; Y) = H(Y) - H(Y|X)$$

is the difference of a concave function and a linear function, hence concave in P_X . ■

Lemma 1.63. $I(X; Y)$ is convex in the channel $W_{Y|X}$. ○

Proof. We use a perturbation argument. First, consider entropy. Let

$$H(P_X) = - \sum_x P_X(x) \log P_X(x),$$

and perturb P_X as

$$P_X^\epsilon(x) = P_X(x)(1 + \epsilon L(x)),$$

with $\sum_x P_X(x)L(x) = 0$ and $1 + \varepsilon L(x) > 0$ for all x . Throughout this computation P_X is an *interior* law, $P_X(x) > 0$ for all x , and L is bounded, so that P_X^ε is a valid distribution for all $|\varepsilon|$ small enough; laws with zero-mass coordinates are handled by continuity from the interior. A Taylor expansion, using $\log(1 + u) = \frac{u}{\ln 2} - \frac{u^2}{2 \ln 2} + O(u^3)$, gives

$$\begin{aligned} H(P_X^\varepsilon) &= - \sum_x P_X(x)(1 + \varepsilon L(x)) \log(P_X(x)(1 + \varepsilon L(x))) \\ &= H(P_X) - \varepsilon \sum_x P_X(x)L(x) \log P_X(x) - \frac{\varepsilon^2}{2 \ln 2} \sum_x P_X(x)L(x)^2 + O(\varepsilon^3). \end{aligned}$$

Thus

$$\left. \frac{d^2}{d\varepsilon^2} H(P_X^\varepsilon) \right|_{\varepsilon=0} = -\frac{1}{\ln 2} \mathbb{E}[L(X)^2] \leq 0,$$

showing that $H(P_X)$ is concave (the factor $1/\ln 2$ is the conversion constant between base-2 logarithms, in which H is measured, and natural logarithms, in which the Taylor expansion of $\log(1 + u)$ is taken).

Now consider a perturbation of the joint distribution:

$$P_{XY}^\varepsilon(x, y) = P_{XY}(x, y)(1 + \varepsilon L(x, y)),$$

subject to

$$\sum_y P_{XY}(x, y)L(x, y) = 0 \quad \forall x,$$

so that $P_X^\varepsilon(x) = P_X(x)$ for all x . Then

$$P_Y^\varepsilon(y) = \sum_x P_{XY}^\varepsilon(x, y) = P_Y(y)(1 + \varepsilon \mathbb{E}[L | Y = y]).$$

Write $I(P_{XY})$ for $I(X; Y)$ under the joint law P_{XY} . Then

$$I(P_{XY}^\varepsilon) = H(P_X) + H(P_Y^\varepsilon) - H(P_{XY}^\varepsilon),$$

and differentiating twice (again with the $1/\ln 2$ conversion factor in the second-order terms) yields

$$\left. \frac{d^2}{d\varepsilon^2} I(P_{XY}^\varepsilon) \right|_{\varepsilon=0} = \frac{1}{\ln 2} \left(\mathbb{E}[L^2] - \mathbb{E}[\mathbb{E}[L | Y]^2] \right).$$

By Jensen's inequality or Cauchy-Schwarz for conditional expectation,

$$\mathbb{E}[L^2] \geq \mathbb{E}[\mathbb{E}[L | Y]^2],$$

so the second derivative is non-negative. Hence $I(X; Y)$ is convex in the channel (for fixed P_X). ■

Remark 1.25. Perturbation (variational) arguments for information measures go back to the early information theory literature; the specific use of second-order perturbation expansions in this form appears in the work of A. Gohari (around 2009, see [Geng and Nair \(2014\)](#)), and similar expansions can be found in earlier convexity arguments for $I(X; Y)$ and related functionals. In particular,

$$\left. \frac{d^2}{d\varepsilon^2} H(P_X^\varepsilon) \right|_{\varepsilon=0} = -\frac{1}{\ln 2} \mathbb{E}[L(X)^2]$$

is a compact way to capture concavity of entropy. $\circ$

The concavity-in- P_X / convexity-in-the-channel dichotomy is one of the ways in which the hypothesis of the min-max corollary (Corollary 9.16 of the Appendix) is established: convexity of the relevant lower-set rate region in the specific application. For that particular corollary no compactness or continuity of the distribution class is required — the compact connected space in Terkelsen’s theorem is the weight simplex, and affine dependence on the weights supplies the topological hypotheses automatically. Sub-additivity is a distinct property used to single-letterise the resulting bound; it does not by itself establish lower-set convexity. See Section 1.13 for how these pieces assemble into converses and Gaussian optimality.

1.13 Sub-additivity: the key principle behind outer bounds and converses

Theorem 1.64: Sub-additivity $\Rightarrow$ single-letter converse / Gaussian optimality

Let f_1 be a functional of a single-letter law μ_X and f_2 a functional of the two-letter law $\mu_{X_1 X_2}$. Suppose that:

- (strict) sub-additivity: for all pairs (X_1, X_2) ,

$$f_2(\mu_{X_1 X_2}) \leq f_1(\mu_{X_1}) + f_1(\mu_{X_2}),$$

with equality if and only if $X_1 \perp\!\!\!\perp X_2$;

- rotational invariance: f_2 is invariant under the normalised sum/difference transform, $f_2(\mu_{X_1 X_2}) = f_2(\mu_{X_+ X_-})$ for $X_\pm = \frac{X_1 \pm X_2}{\sqrt{2}}$;
- f_1 has a finite attained maximum V on a specified admissible class $\mathcal{S}$, closed under the normalised sum/difference transform, and the sum and difference marginals of independent copies of any admissible law are admissible.

Then (a) the multi-letter value collapses to the single-letter value,

$$\sup_{(X_1, X_2) \text{ admissible}} f_2(\mu_{X_1 X_2}) = 2 \sup_{X \in \mathcal{S}} f_1(\mu_X),$$

where the two-letter supremum is over the specified admissible class (including independent copies of admissible laws, as in the closure conditions above), and (b) every maximiser of f_1 in $\mathcal{S}$ is Gaussian. Throughout, the laws are on a real vector space.

Vocabulary. A *converse* for a coding problem is an upper bound on what is achievable; when the bound is stated as a set of admissible rate tuples it is called an *outer bound*. The passage from an n -letter bound $nR \leq F^{(n)} + o(n)$ to a single-letter bound $R \leq F^*$ (with $F^* = \sup F$) is called *single-letterisation*; it is valid whenever $F^{(n)}$ is sub-additive. When F is not sub-additive but its supremum is obtained by averaging over the coordinates, the argument introduces the *time-sharing variable* Q , uniform on $[1 : n]$ (selecting a random coordinate), and the single-letter region is expressed through this auxiliary Q . The Markov structure used to decompose $F^{(n)}$ into single-letter terms is the *Csiszár–Körner trick*.

We now state the principle in a form general enough for the later chapters, and then illustrate it with the examples of this chapter.

1.13.1 The abstract template

Let $F^{(n)}$ be a functional defined on n -letter laws $p_{X^n Y^n}$ (for two-letter laws we write $F^{(2)}$), and let $F \triangleq F^{(1)}$ be the corresponding single-letter functional.

Definition 1.65. (Sub-additivity) We say that $F^{(n)}$ is *sub-additive* (with respect to the class of joint laws under consideration) if for every admissible $p_{X^n Y^n}$,

$$F^{(n)}(p_{X^n Y^n}) \leq \sum_{i=1}^n F(p_{X_i Y_i}),$$

where $p_{X_i Y_i}$ denotes the i -th marginal pair. If in addition each single-letter value is bounded by a constant $F^* \triangleq \sup_{p_{XY}} F(p_{XY})$, then

$$\frac{1}{n} F^{(n)}(p_{X^n Y^n}) \leq F^*.$$

○

The role of sub-additivity is the following. In a converse we typically arrive at an inequality of the form

$$nR \leq F^{(n)}(p_{X^n Y^n}) + o(n),$$

where $p_{X^n Y^n}$ is the joint law induced by the code. Sub-additivity then collapses the n -letter functional to the single-letter supremum:

$$R \leq F^*,$$

and, when F depends on the auxiliary law through some channel, the outer bound is the resulting single-letter region. Three remarks are worth making immediately:

- *Where the summands come from.* The inequality $F^{(n)} \leq \sum_i F(p_{X_i Y_i})$ is never free: it requires the Markov/independence structure of the setting. For memoryless channels it is the chain $(X^{n \setminus i}, Y^{i-1}) \rightarrow X_i \rightarrow Y_i$ that makes the “Csiszár–Körner trick” work; in Gaussian settings it is the chain combined with Gaussian optimality; in the perturbation arguments of this chapter it is the double Markovity lemma.
- *The Cesàro variable Q .* When the functional is not sub-additive but its *supremum* is *convexified* by averaging, the same argument produces the time-sharing random variable Q (uniform on $[1 : n]$), and the single-letter region is expressed with the auxiliary Q ; the multiple access channel in Chapter 3 is the prototype.
- *The min–max bridge.* Many outer bounds have the shape

$$\sup_p \min_{\lambda \in \Lambda} \sum_i \lambda_i T_i(p) \leq C.$$

The corollary of Terkelsen’s theorem in the appendix (Theorem 9.15, in the form of Corollary 9.16) applies once convexity of the componentwise lower-set region $\mathcal{A} = \{a : a_i \leq T_i(p) \text{ for some } p\}$ has been established, for example through simultaneous concavity on a common convex domain or an admissible time-sharing construction. Compactness of the weight simplex and affine dependence on the weights supply the topological hypotheses automatically; no compactness of the distribution class is assumed in that corollary. Subadditivity is a separate property used to single-letterise the resulting bound; it does not by itself establish lower-set convexity (two admissible parameter values with $T(p_0) = (1, 0)$ and $T(p_1) = (0, 1)$ have a nonconvex lower-set region even when each coordinate extends to an exactly additive multi-letter functional). The UV and UVW outer bounds for the broadcast channel (Chapter 4) and the genie-based bounds for the interference channel (Chapter 5) are organised this way.

Remark 1.26 (Submodularity of entropy vs. sub-additivity of functionals). The compression material earlier in this chapter (the elementary compression inequality and Lemma 1.50) is a *set-function* technique: it asserts inequalities such as $H(X_A) + H(X_B) \geq H(X_{A \cup B}) + H(X_{A \cap B})$, used for *discrete* outer bounds. The *functional* sub-additivity $F^{(2)} \leq F + F$ of the present section is a different — though related — object: a product-space inequality used for *single-letterisation*. The two share the word “sub-” but operate at different levels; this section concerns the latter. $\circ$

Remark 1.27 (Sub-additivity, the auxiliary U , and the min–max theorem). In the alternate proof of the EPI, the functional

$$F(p_X, p_Y) = \sup_{U: X \rightarrow U \rightarrow Y} \left(a h(X|U) + b h(Y|U) - h(M[X; Y] | U) \right)$$

is a *supremum over an auxiliary U* . For a converse one must be able to pull this supremum out of the objective, or exchange it with a min over dual weights. This is exactly what the min–max corollary of the appendix provides: once the componentwise lower-set region of the objectives T_i is convex (which must be verified in the specific application; the objectives are linear in the auxiliary law $p(u)$ only when the conditional laws of (X, Y) given U are held fixed, and they are *not* linear in the joint law $p(u, x, y)$ when those conditional laws vary), Corollary 9.16

(Terkelsen's min-max, Theorem 9.15) permits the exchange $\sup_{p(u)} \min_{\lambda} = \min_{\lambda} \sup_{p(u)}$, with the compact weight simplex supplying the topological hypotheses automatically. The convexity of $\mathcal{A}$ is one ingredient that the perturbation method (concavity of I in P_X , convexity in the channel) is used to establish in the applications; combined with the sub-additivity of F , this collapse is what turns the n -letter functional into the single-letter (Gaussian) value. $\bigcirc$

1.13.2 Example: the single-letterisation of $I(X^n; Y^n)$

For a *no-feedback product channel* (the law of Y_i given the inputs is memoryless and Y_i depends on the inputs only through X_i , with no feedback), the Markov chain $(X^{n \setminus i}, Y^{i-1}) \rightarrow X_i \rightarrow Y_i$ holds, and the mutual information is sub-additive:

$$\begin{aligned} I(X^n; Y^n) &= \sum_{i=1}^n \left(I(X^n; Y_i | Y^{i-1}) + I(Y^{i-1}; Y_i) - I(Y^{i-1}; Y_i) \right) \\ &= \sum_{i=1}^n \left(I(X_i; Y_i) - I(Y^{i-1}; Y_i) \right) \leq \sum_{i=1}^n I(X_i; Y_i). \end{aligned}$$

This computation is a central *converse step* of the point-to-point channel coding theorem for a memoryless channel (it turns the Fano-based n -letter bound into the single-letter bound $I(X; Y)$; the remaining converse ingredients — Fano's inequality, the code structure, and the channel's memorylessness — are carried out in Chapter 2), and a conditional version of it is a central step of the converse for the multiple access channel (Chapter 3). We will refer to this computation as *the sub-additivity step* of a converse.

1.13.3 Example: the functional of the perturbation proof of the EPI

In the alternate proof of the conditional EPI, the functional

$$F(p_X, p_Y) \triangleq \sup_{U: X \rightarrow U \rightarrow Y} \left(a h(X|U) + b h(Y|U) - h(M[X; Y] | U) \right)$$

was shown to be sub-additive,

$$F^{(2)}(p_{X_1 X_2}, p_{Y_1 Y_2}) \leq F(p_{X_1}, p_{Y_1}) + F(p_{X_2}, p_{Y_2}),$$

with an explicit non-negative *gap* $\Delta(U)$ depending on the auxiliary and quantifying the loss in the fixed-auxiliary decomposition. The rotation trick then forces the gap to vanish at the constructed auxiliary $V = (U_1, U_2)$ of two independent copies of a maximiser, double Markovity converts the vanishing gap into conditional independence of the rotated pair, and Bernstein's theorem upgrades independence to Gaussianity. This is the template we will reuse, in simplified form, whenever a Gaussian optimality statement is needed: *sub-additivity + rotation + double Markovity $\Rightarrow$ Gaussian optimality*. The supremum over the auxiliary U is exchanged with the dual minimum via the min-max theorem, as explained in Remark 1.27, and the hypotheses of that theorem are checked separately in each application.

1.13.4 Roadmap

The following points describe the *intended role* of the principle in later chapters; the detailed hypotheses are checked there, chapter by chapter.

- Chapter 2 (point-to-point): the sub-additivity step $I(X^n; Y^n) \leq \sum_i I(X_i; Y_i)$ is a central converse step; the multi-letter remark explains why the same trick fails for the MAC (the marginal functional $p(x_1) \mapsto I(X_1; Y)$ is *not* sub-additive), which is repaired by time-sharing.
- Chapter 4 (broadcast): the UV outer bound is a sub-additivity statement for a family of mutual-information functionals, and the UVW bound adds the double Markovity lemma; the finite-vector lower-set min-max corollary of the appendix converts them to the classical form of the bounds, once convexity of the relevant lower-set region is verified in that chapter.
- Chapter 5 (interference): the genie-based outer bounds (giving away Y_2 to receiver 1, etc.) are sub-additivity arguments followed by Gaussian optimality; the strong-interference and weak-interference converses are the cleanest instances.
- Chapter 8 (relay): the cut-set bound is the purest form of the principle — each cut defines a functional, sub-additivity gives the bound, and the min-max corollary handles the union over cuts, again once convexity of the relevant lower-set region is verified in that chapter.
- The appendix min-max theorem and the maximal-coupling/Sanov machinery provide the analytic backbone for the steps above. Neither sub-additivity alone nor a min-max theorem alone proves every later multi-user converse; each chapter supplies and verifies the remaining hypotheses.

2 Point-to-point information theory

In this chapter, we first review the notion of typicality, a fundamental property of random variables that plays a central role throughout this field. We then study two of the most important results in information theory: the source coding theorem and the channel coding theorem. These theorems were first introduced and proved by Shannon in his seminal paper [Shannon \(1948\)](#).

2.1 Typicality

Most of network information theory is based on the concentration of the multinomial distribution: The empirical distribution of n random variables $X_1, \dots, X_n$ i.i.d. according to some distribution $p(x)$ should be close to $p(x)$ when n is large. One may think of the sequence X^n as the outcomes of repeating an experiment, and hence their empirical distribution, which is obtained by counting the occurrences of each possible outcome, should be close to the “real” underlying distribution when the number of experiments is large. We call sequences that have an empirical distribution close to $p(x)$ *typical* with respect to $p(x)$. In fact, we have the following properties of typical sequences:

- There are approximately $2^{nH(X)}$ typical sequences.
- The probability that X^n is equal to a given typical sequence is approximately $2^{-nH(X)}$.
- The law of large numbers: The probability that X^n is typical is close to 1.
- The probability that X^n is typical with respect to some other distribution $\hat{p}(x)$ is approximately $2^{-nD(\hat{p}||p)}$, where for a fixed tolerance the exponential rate is an infimum over the typicality neighbourhood and the exact exponent $D(\hat{p}||p)$ is recovered only by shrinking the tolerance after the limit (Proposition 2.7).

Hence, we see that the entropy arises naturally in counting typical sequences and that the KL divergence is a natural measure of atypicality.

Definition 2.1. (Type) The type, also called the empirical distribution, of a sequence $x^n \in \mathcal{X}^n$ is the distribution $\pi(\cdot|x^n)$ on $\mathcal{X}$ defined by

$$\pi(x|x^n) \triangleq \frac{1}{n} |\{i \in [1 : n] : x_i = x\}|.$$

○

Sanov's theorem describes the large-deviation behavior of empirical distributions. In particular, it quantifies the exponential decay of the probability that the empirical distribution of i.i.d. samples belongs to a prescribed set of distributions.

Theorem 2.2. (Sanov's theorem) Let $X_1, \dots, X_n$ be random variables i.i.d. according to a distribution $p(x)$ on a finite alphabet $\mathcal{X}$, and write $\widehat{P}_n \triangleq \pi(\cdot|X^n)$ for the empirical distribution. For every set $\mathcal{P}$ of distributions on $\mathcal{X}$,

$$\Pr(\widehat{P}_n \in \mathcal{P}) \leq (n+1)^{|\mathcal{X}|} \cdot 2^{-n \inf_{q \in \mathcal{P}} D(q\|p)},$$

with the conventions $\inf \emptyset = +\infty$ and $\log 0 = -\infty$.

Put $S \triangleq \text{supp } p$ and let $\Delta(S)$ denote the probability simplex on S , and set $A \triangleq \mathcal{P} \cap \Delta(S)$. If A is nonempty and equals the closure, within $\Delta(S)$, of its interior relative to $\Delta(S)$, then

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log_2 \Pr(\widehat{P}_n \in \mathcal{P}) = -\min_{q \in A} D(q\|p).$$

If $A = \emptyset$, then $\Pr(\widehat{P}_n \in \mathcal{P}) = 0$ for every n , because empirical laws sampled from p always lie in $\Delta(S)$. This support qualification is essential: with the ambient-simplex regular-closedness hypothesis alone the limit statement is false when p has zero coordinates. $\circ$

Proof. For any type q of sequences in $\mathcal{X}^n$, the number of sequences of type q is at most $2^{nH(q)}$ (Lemma 2.3), and each such sequence has p -probability $\prod_{x \in \mathcal{X}} p(x)^{nq(x)} = 2^{-n(H(q)+D(q\|p))}$; hence

$$\Pr(\widehat{P}_n = q) \leq 2^{-nD(q\|p)}.$$

Summing over the types $q \in \mathcal{P} \cap \mathcal{T}_n$, of which there are at most $(n+1)^{|\mathcal{X}|}$, gives the upper bound.

For the lower bound, all empirical laws under p lie in $\Delta(S)$, on which $D(\cdot\|p)$ is finite and continuous. The set A is closed in the compact simplex $\Delta(S)$, so it has a minimizer q^* of $D(\cdot\|p)$, an *information projection* of p onto A . Choose laws u_j in the interior of A relative to $\Delta(S)$ with $u_j \rightarrow q^*$, and for each j choose $r_j > 0$ so that the ball of radius r_j about u_j within $\Delta(S)$ is contained in A . For fixed j and n , put $a_x = \lfloor nu_j(x) \rfloor$ for $x \in S$ and $r = n - \sum_{x \in S} a_x$. Add one count to each of r distinct coordinates whose fractional part $nu_j(x) - a_x$ is positive, and leave the other coordinates unchanged. There are enough such coordinates because their fractional parts sum to the integer r and each is less than one. With $q_{n,j}(x) = a_x/n$ after these increments, $\sum_x q_{n,j}(x) = 1$, $\text{supp } q_{n,j} \subseteq \text{supp } u_j \subseteq S$, and $|q_{n,j}(x) - u_j(x)| \leq 1/n$. Thus $\|q_{n,j} - u_j\|_{\text{TV}} \leq |S|/(2n)$. Choose increasing N_j so that this bound is below the interior radius r_j for all $n \geq N_j$, take $j = j(n) \rightarrow \infty$ slowly with $n \geq N_{j(n)}$, and set $q_n = q_{n,j(n)}$. Then $q_n \in A \cap \mathcal{T}_n$ eventually and $q_n \rightarrow q^*$. (This support-preserving diagonal rounding is the step that the density of $\bigcup_n \mathcal{T}_n$ in the full simplex alone does not supply.)

By the multinomial lower bound of Lemma 2.3,

$$\Pr(\widehat{P}_n = q_n) \geq \left(\sqrt{2\pi n} e^{\frac{1}{12}}\right)^{-|\mathcal{X}|} 2^{-nD(q_n \| p)}.$$

Hence

$$\frac{1}{n} \log_2 \Pr(\widehat{P}_n \in \mathcal{P}) \geq -D(q_n \| p) - \frac{|\mathcal{X}|}{n} \log_2 \left(\sqrt{2\pi n} e^{\frac{1}{12}}\right) \rightarrow -D(q^* \| p)$$

as $n \rightarrow \infty$, by continuity of $D(\cdot \| p)$ on $\Delta(S)$. In the upper bound, $\inf_{q \in \mathcal{P}} D(q \| p) = \min_{q \in A} D(q \| p) = D(q^* \| p)$, because every law outside $\Delta(S)$ assigns positive mass to a p -null symbol and has divergence $+\infty$ (with the convention $\log 0 = -\infty$). Squeezing gives the result. ■

Lemma 2.3. (Multinomial estimate) Let $\hat{p}$ be a type of sequences in $\mathcal{X}^n$, i.e. a distribution on $\mathcal{X}$ such that $n\hat{p}(x)$ is an integer for every $x \in \mathcal{X}$. The number of sequences of type $\hat{p}$ satisfies

$$\begin{aligned} \left(\sqrt{2\pi n} e^{\frac{1}{12}}\right)^{-|\mathcal{X}|} 2^{nH(\hat{p})} &\leq |\{x^n \in \mathcal{X}^n : \pi(\cdot | x^n) = \hat{p}\}| \\ &\leq 2^{nH(\hat{p})}. \end{aligned}$$

○

Proof. The number of sequences of type $\hat{p}$ is the multinomial coefficient $\frac{n!}{\prod_{x \in T} (n\hat{p}(x))!}$ with $T = \text{supp } \hat{p}$. For the upper bound, every sequence x^n of type $\hat{p}$ has $\hat{p}$ -probability $\prod_{x \in T} \hat{p}(x)^{n\hat{p}(x)} = 2^{-nH(\hat{p})}$, so

$$\begin{aligned} 1 &\geq \sum_{x^n : \pi(\cdot | x^n) = \hat{p}} \prod_{x \in T} \hat{p}(x)^{n\hat{p}(x)} \\ &= |\{x^n : \pi(\cdot | x^n) = \hat{p}\}| \cdot 2^{-nH(\hat{p})}, \end{aligned}$$

which gives the upper bound. For the lower bound, let $T = \text{supp } \hat{p}$, $s = |T|$, $M = |\mathcal{X}|$. For every $x \in T$, $n\hat{p}(x)$ is a positive integer and hence is at least one. The following Stirling calculation therefore applies to every n -type, without an additional large- n condition. Applying the Stirling bounds $\sqrt{2\pi m} \left(\frac{m}{e}\right)^m \leq m! \leq \sqrt{2\pi m} \left(\frac{m}{e}\right)^m e^{\frac{1}{12m}}$ to $n!$ and to each $(n\hat{p}(x))!$ with $x \in T$ gives

$$\begin{aligned} |\{x^n : \pi(\cdot | x^n) = \hat{p}\}| &= \frac{n!}{\prod_{x \in T} (n\hat{p}(x))!} \\ &\geq \frac{\sqrt{2\pi n}}{\prod_{x \in T} \sqrt{2\pi n\hat{p}(x)}} \exp\left(-\sum_{x \in T} \frac{1}{12n\hat{p}(x)}\right) 2^{nH(\hat{p})} \\ &\geq (\sqrt{2\pi n})^{1-s} e^{-s/12} 2^{nH(\hat{p})} \\ &\geq \left(\sqrt{2\pi n} e^{\frac{1}{12}}\right)^{-M} 2^{nH(\hat{p})}. \end{aligned}$$

The upper bound obtained above already rests on a direct probability argument, so the lemma is self-contained. ■

Definition 2.4. (Typical set) Let $p(x)$ be a distribution on $\mathcal{X}$. Let $\epsilon > 0$. The set of ϵ -typical sequences in $\mathcal{X}^n$ is defined by

$$\mathcal{T}_\epsilon^{(n)}(X) \triangleq \{x^n \in \mathcal{X}^n : |\pi(x|x^n) - p(x)| \leq \epsilon p(x) \text{ for all } x \in \mathcal{X}\}.$$

○

Definition 2.5. (Joint typical set) Let $p(x, y)$ be a distribution on $\mathcal{X} \times \mathcal{Y}$, and let $\epsilon > 0$. The set of jointly ϵ -typical sequences in $\mathcal{X}^n \times \mathcal{Y}^n$ is defined by

$$\mathcal{T}_\epsilon^{(n)}(X, Y) \triangleq \{(x^n, y^n) \in \mathcal{X}^n \times \mathcal{Y}^n : |\pi(x, y|x^n, y^n) - p(x, y)| \leq \epsilon p(x, y) \text{ for all } (x, y) \in \mathcal{X} \times \mathcal{Y}\},$$

where $\pi(x, y|x^n, y^n) \triangleq \frac{1}{n} |\{i \in [1 : n] : x_i = x, y_i = y\}|$ is the joint type of (x^n, y^n) . Equivalently, $\mathcal{T}_\epsilon^{(n)}(X, Y) = \mathcal{T}_\epsilon^{(n)}(Z)$ where $Z \triangleq (X, Y)$ is treated as a single random variable on the product alphabet $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ with pmf $p(z) = p(x, y)$. The same convention defines $\mathcal{T}_\epsilon^{(n)}(X_1, \dots, X_k)$ for a tuple of random variables via the joint pmf $p(x_1, \dots, x_k)$ on $\mathcal{X}_1 \times \dots \times \mathcal{X}_k$. ○

Proposition 2.6. (Law of large numbers) Let $X_1, \dots, X_n$ be random variables i.i.d. according to a distribution $p(x)$ on $\mathcal{X}$. Let $\epsilon > 0$. Then

$$\Pr(X^n \notin \mathcal{T}_\epsilon^{(n)}(X)) \rightarrow 0$$

as $n \rightarrow \infty$. ○

Proof. For any $x \in \mathcal{X}$, note that $n\pi(x|X^n)$ follows the binomial distribution with parameters n , $p(x)$, and therefore $\pi(x|X^n)$ has mean $p(x)$ and variance $\frac{1}{n}p(x)(1-p(x))$. For x with $p(x) > 0$, Chebyshev's inequality gives

$$\Pr(|\pi(x|X^n) - p(x)| > \epsilon p(x)) \leq \frac{1 - p(x)}{n\epsilon^2 p(x)}.$$

When $p(x) = 0$ the empirical count of x is identically zero, so the corresponding event has probability 0 (and no Chebyshev bound is needed). Thus,

$$\begin{aligned} \Pr(X^n \notin \mathcal{T}_\epsilon^{(n)}(X)) &= \Pr(|\pi(x|X^n) - p(x)| > \epsilon p(x) \text{ for some } x \in \mathcal{X}) \\ &\leq \sum_{x \in \mathcal{X}} \Pr(|\pi(x|X^n) - p(x)| > \epsilon p(x)) \\ &\leq \frac{1}{n\epsilon^2} \sum_{\substack{x \in \mathcal{X} \\ p(x) > 0}} \frac{1 - p(x)}{p(x)} \\ &\rightarrow 0 \end{aligned}$$

as $n \rightarrow \infty$. ■

Proposition 2.7. (Atypicality under a mismatched law) Let $p(z)$ and $\hat{p}(z)$ be distributions on an alphabet $\mathcal{Z}$, and let $Z_1, \dots, Z_n$ be random variables i.i.d. according to $p(z)$. For $0 < \epsilon < 1$,

$$\Pr \left(Z^n \in \mathcal{T}_\epsilon^{(n)}(\hat{Z}) \right) \leq (n+1)^{|\mathcal{Z}|-1} \cdot 2^{-n(D(\hat{p}||p) - \delta(\epsilon))},$$

where $\delta(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$ is defined as follows. If $\hat{p} \not\ll p$, the probability is identically zero: the empirical law would have to assign positive mass to a p -null symbol. In that case the bound holds with any $\delta(\epsilon) \rightarrow 0$, and we set $\delta(\epsilon) = 0$. Otherwise, with $S \triangleq \text{supp } \hat{p}$, the finite constant

$$L \triangleq \sum_{z \in S} \hat{p}(z) \left| \log_2 \frac{\hat{p}(z)}{p(z)} \right|, \quad \delta(\epsilon) \triangleq \epsilon L - \log_2(1 - \epsilon)$$

is well defined and tends to 0 as $\epsilon \rightarrow 0$. For pairs, one takes $Z = (X, Y)$ to be a single random variable on the alphabet $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$: with $p = p_X p_Y$ and $\hat{p} = p_{XY}$ (which always satisfies $p_{XY} \ll p_X p_Y$), the bound reads

$$\Pr \left((X^n, Y^n) \in \mathcal{T}_\epsilon^{(n)}(X, Y) \right) \leq (n+1)^{|\mathcal{X}||\mathcal{Y}|-1} \cdot 2^{-n(I(X;Y) - \delta(\epsilon))},$$

since $D(p_{XY}||p_X p_Y) = I(X; Y)$. ○

Proof. Let $\mathcal{P}$ be the set of all distributions $q(z)$ on $\mathcal{Z}$ such that $|q(z) - \hat{p}(z)| \leq \epsilon \hat{p}(z)$ for all $z \in \mathcal{Z}$. For $\hat{p} \not\ll p$ the event has probability zero, so assume $\hat{p} \ll p$. By Sanov's theorem applied to the alphabet $\mathcal{Z}$,

$$\begin{aligned} \Pr \left(Z^n \in \mathcal{T}_\epsilon^{(n)}(\hat{Z}) \right) &= \Pr (\pi(\cdot|Z^n) \in \mathcal{P}) \\ &\leq |\mathcal{P} \cap \mathcal{T}_n| \cdot 2^{-n \inf_{q(z) \in \mathcal{P}} D(q||p)} \\ &\leq \binom{n+|\mathcal{Z}|-1}{|\mathcal{Z}|-1} \cdot 2^{-n \inf_{q(z) \in \mathcal{P}} D(q||p)} \\ &\leq (n+1)^{|\mathcal{Z}|-1} \cdot 2^{-n \inf_{q(z) \in \mathcal{P}} D(q||p)}, \end{aligned}$$

where the second inequality counts the types of sequences in $\mathcal{Z}^n$ (a type is determined by the integers $nq(z)$, so there are at most $\binom{n+|\mathcal{Z}|-1}{|\mathcal{Z}|-1}$ types), and the third uses $\binom{n+|\mathcal{Z}|-1}{|\mathcal{Z}|-1} \leq (n+1)^{|\mathcal{Z}|-1}$. Every admissible $q \in \mathcal{P}$ satisfies $q(z) \leq \hat{p}(z)(1 + \epsilon)$ for all z (so $q(z) = 0$ outside $S = \text{supp } \hat{p}$) and $q(z) \geq \hat{p}(z)(1 - \epsilon)$ for all z . Hence

$$\begin{aligned} D(q||p) &= \sum_{z \in S} q(z) \log_2 \frac{q(z)}{p(z)} \\ &\geq \sum_{z \in S} q(z) \log_2 \frac{\hat{p}(z)(1 - \epsilon)}{p(z)} \\ &= \sum_{z \in S} q(z) \log_2 \frac{\hat{p}(z)}{p(z)} + \log_2(1 - \epsilon) \\ &\geq \sum_{z \in S} \hat{p}(z) \log_2 \frac{\hat{p}(z)}{p(z)} - \epsilon \sum_{z \in S} \hat{p}(z) \left| \log_2 \frac{\hat{p}(z)}{p(z)} \right| + \log_2(1 - \epsilon) \end{aligned}$$

$$= D(\hat{p} \| p) - \delta(\epsilon),$$

where the last inequality uses $|q(z) - \hat{p}(z)| \leq \epsilon \hat{p}(z)$: on coordinates with $\log_2(\hat{p}/p) \geq 0$ it uses $q \geq (1 - \epsilon)\hat{p}$, and on coordinates with $\log_2(\hat{p}/p) < 0$ it uses $q \leq (1 + \epsilon)\hat{p}$. Therefore

$$\Pr \left(Z^n \in \mathcal{T}_\epsilon^{(n)}(\hat{Z}) \right) \leq (n+1)^{|\mathcal{Z}|-1} \cdot 2^{-n(D(\hat{p} \| p) - \delta(\epsilon))}.$$

The polynomial prefactor $(n+1)^{|\mathcal{Z}|-1}$ is the packing penalty of the support-qualified type count; it does not affect the exponent. Since $\delta(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$, this makes the support convention and the packing penalty explicit; for fixed tolerance the actual exponent is an infimum over the typicality neighbourhood, and shrinking the tolerance recovers $D(\hat{p} \| p)$ in the limit. ■

Proposition 2.8. (Probability for typical sequences) Let $X_1, \dots, X_n$ be random variables i.i.d. according to a distribution $p(x)$ on $\mathcal{X}$. Let $\epsilon > 0$. If $x^n \in \mathcal{T}_\epsilon^{(n)}(X)$, then

$$2^{-nH(X)(1+\epsilon)} \leq \Pr(X^n = x^n) \leq 2^{-nH(X)(1-\epsilon)}.$$

○

Proof. Note that

$$\Pr(X^n = x^n) = \prod_{x \in \text{supp } p} p(x)^{n\pi(x|x^n)}.$$

(For the fixed typical sequence x^n , typicality forces $\pi(x|x^n) = 0$ at every coordinate x with $p(x) = 0$, so those coordinates contribute the factor 1 and may be omitted; this avoids implicit 0^0 powers.) Then, by $x^n \in \mathcal{T}_\epsilon^{(n)}(X)$,

$$2^{-nH(X)(1+\epsilon)} = \prod_{x \in \text{supp } p} p(x)^{np(x)(1+\epsilon)} \leq \Pr(X^n = x^n) \leq \prod_{x \in \text{supp } p} p(x)^{np(x)(1-\epsilon)} = 2^{-nH(X)(1-\epsilon)}.$$

■

Proposition 2.9. (Cardinality of typical set) Let $p(x)$ be a distribution on $\mathcal{X}$. Let $\epsilon > 0$. Then

$$(1 - \epsilon)2^{nH(X)(1-\epsilon)} \leq \left| \mathcal{T}_\epsilon^{(n)}(X) \right| \leq 2^{nH(X)(1+\epsilon)}$$

for n large enough.

○

Proof. Consider

$$\begin{aligned} 1 &\geq \Pr \left(X^n \in \mathcal{T}_\epsilon^{(n)}(X) \right) \\ &= \sum_{x^n \in \mathcal{T}_\epsilon^{(n)}(X)} \Pr(X^n = x^n) \end{aligned}$$

$$\begin{aligned}
 &\geq \sum_{x^n \in \mathcal{T}_\epsilon^{(n)}(X)} 2^{-nH(X)(1+\epsilon)} \\
 &= |\mathcal{T}_\epsilon^{(n)}(X)| \cdot 2^{-nH(X)(1+\epsilon)}.
 \end{aligned}$$

This shows the upper bound.

Now we show the lower bound. By the law of large numbers,

$$\begin{aligned}
 \Pr(X^n \in \mathcal{T}_\epsilon^{(n)}(X)) &= 1 - \Pr(X^n \notin \mathcal{T}_\epsilon^{(n)}(X)) \\
 &\geq 1 - \epsilon
 \end{aligned}$$

for n large enough. Also,

$$\begin{aligned}
 \Pr(X^n \in \mathcal{T}_\epsilon^{(n)}(X)) &= \sum_{x^n \in \mathcal{T}_\epsilon^{(n)}(X)} \Pr(X^n = x^n) \\
 &\leq \sum_{x^n \in \mathcal{T}_\epsilon^{(n)}(X)} 2^{-nH(X)(1-\epsilon)} \\
 &= |\mathcal{T}_\epsilon^{(n)}(X)| \cdot 2^{-nH(X)(1-\epsilon)}.
 \end{aligned}$$

Thus,

$$|\mathcal{T}_\epsilon^{(n)}(X)| \geq (1 - \epsilon) 2^{nH(X)(1-\epsilon)}$$

for n large enough. ■

2.2 Channel coding

Consider the point-to-point communication system depicted in [Figure 2.1](#).




Figure 2.1: Point-to-point communication system

There is a transmitter who wants to communicate a message $M \in \mathcal{M}$ to a receiver. The transmitter encodes the message into a codeword $X^n \in \mathcal{X}^n$ which is then passed through a channel described by a stochastic matrix W . The receiver receives the channel output $Y^n \in \mathcal{Y}^n$ and decodes it into an estimate $\hat{M} \in \mathcal{M}$ of the transmitted message.

We shall focus only on memoryless channels. A channel with input X^n and output Y^n is called memoryless if $(X^{n \setminus i}, Y^{i-1}) \rightarrow X_i \rightarrow Y_i$ forms a Markov chain for $i \in [1 : n]$, or in other words,

the current output Y_i depends only on the current input X_i . For a memoryless channel, the output satisfies

$$\begin{aligned} p(y^n|x^n) &= \prod_{i=1}^n p(y_i|x^n, y^{i-1}) \\ &= \prod_{i=1}^n p(y_i|x_i). \end{aligned} \quad ((X^n \setminus i, Y^{i-1}) \rightarrow X_i \rightarrow Y_i \text{ is Markov})$$

We further assume the channel is *stationary* and used *without feedback*, so that the same transition rule applies at every time. Then a point-to-point memoryless channel is represented by a single stochastic matrix $W : \mathcal{X} \rightarrow \mathcal{Y}$, where $\mathcal{X}$ and $\mathcal{Y}$ are called the input alphabet and output alphabet, respectively, and the n -letter channel is the stationary product kernel

$$W^{\otimes n}(y^n|x^n) \triangleq \prod_{i=1}^n W(y_i|x_i).$$

(The Markov condition by itself would also permit time-dependent kernels W_i ; stationarity is what reduces the n channel uses to powers of one matrix.)

Our goal is to communicate reliably at the largest possible rate.

An (n, R) code for a memoryless channel $W : \mathcal{X} \rightarrow \mathcal{Y}$ consists of

- a message set $\mathcal{M} \triangleq [1 : 2^{nR}]$,
- a mapping $x^n : \mathcal{M} \rightarrow \mathcal{X}^n$ called the encoder,
- a mapping $\hat{m} : \mathcal{Y}^n \rightarrow \mathcal{M}$ called the decoder.

We assume that the message sent M is a random variable uniformly distributed over the message set $\mathcal{M}$. The codeword sent X^n is given by $X^n \triangleq x^n(M)$. The codeword received Y^n is a random variable such that

$$p(y^n|x^n, m) = \prod_{i=1}^n W(y_i|x_i).$$

The decoded message $\hat{M}$ is a random variable on $\mathcal{M}$ given by $\hat{M} \triangleq \hat{m}(Y^n)$.

$R \geq 0$ is said to be an achievable rate for a memoryless channel if there exists a sequence of (n, R) codes whose (average) probability of error

$$\Pr(\hat{M} \neq M) = \frac{1}{|\mathcal{M}|} \sum_{m \in \mathcal{M}} \Pr(\hat{M} \neq m | M = m)$$

goes to 0 as $n \rightarrow \infty$. We say that such a sequence of codes achieves rate R . The supremum of all achievable rates is called the capacity of the channel.

Figure 2.2 illustrates the spaces and the mappings used in the point-to-point communication.




Figure 2.2: Spaces and mappings in a point-to-point communication

Remark 2.1. One should use $|\mathcal{M}| = 2^{\lceil nR \rceil}$ so that the message set size is an integer. For simplicity, we write $|\mathcal{M}| = 2^{nR}$ throughout. This only affects additive $o(n)$ terms (e.g., $\log |\mathcal{M}| = nR + o(n)$) and has no impact on the proofs of this section. $\circ$

Remark 2.2. Other than the average probability of error, there are other metrics to measure the performance of a code. One such metric is the maximum probability of error

$$\max_{m \in \mathcal{M}} \Pr(\hat{M} \neq m | M = m).$$

The supremum of all rates R for which there exists a sequence of (n, R) codes whose maximum probability of error goes to zero as $n \rightarrow \infty$ is called the worst-case capacity of the channel. $\circ$

Remark 2.3. The supremum of all rates R for which there exists a sequence of (n, R) codes such that the probability of error $\Pr(\hat{M} \neq M) = 0$ is called the zero-error capacity. Finding the zero-error capacity of a channel is a combinatorial problem: It depends only on the bipartite graph $\{(x, y) \in \mathcal{X} \times \mathcal{Y} : W(y|x) > 0\}$. In general, finding the zero-error capacity of a channel is a hard problem, and we only know the zero-error capacity of a few channels. $\circ$

Remark 2.4. We have the following inclusions: average probability of error capacity region $\supseteq$ maximum probability of error capacity region $\supseteq$ zero-error capacity region. For the point-to-point channel, the capacity regions with the average and maximum probability of error are equal. Indeed, if a code has average error ϵ_n , then at least half of its messages have conditional error at most $2\epsilon_n$; retaining only those messages, and restricting the decoder to output only retained messages, yields a code with maximal error at most $2\epsilon_n$ and rate at least $R - \frac{1}{n}$ (the retained message set has size at least $\frac{1}{2}2^{nR}$). Letting $\epsilon_n \rightarrow 0$ along a code sequence, together with a small rate slack, proves the equality of the two capacity suprema. $\circ$

Lemma 2.10. (Sub-additivity (Csiszár–Körner single-letterisation)) For a memoryless channel, i.e. one satisfying $(X^{n \setminus i}, Y^{i-1}) \rightarrow X_i \rightarrow Y_i$ for every $i \in [1 : n]$,

$$\begin{aligned}
 I(X^n; Y^n) &= \sum_{i=1}^n I(X^n; Y_i | Y^{i-1}) && \text{(chain rule)} \\
 &\leq \sum_{i=1}^n (I(X^n; Y_i | Y^{i-1}) + I(Y^{i-1}; Y_i)) && (I(Y^{i-1}; Y_i) \geq 0) \\
 &= \sum_{i=1}^n I(X^n, Y^{i-1}; Y_i) && \text{(chain rule)} \\
 &= \sum_{i=1}^n I(X_i; Y_i) && ((X^{n \setminus i}, Y^{i-1}) \rightarrow X_i \rightarrow Y_i \text{ is Markov}).
 \end{aligned}$$

This is the *sub-additivity step* of §1.13: it bounds the n -letter mutual information by the sum (hence, after dividing by n , the Cesàro average) of single-letter terms, whose supremum is the capacity C . $\circ$

Theorem 2.11. (Channel coding theorem) The capacity of a memoryless channel W is given by $C \triangleq \max_{p(x)} I(X; Y)$, where $(X, Y) \sim p(x)W(y|x)$. $\circ$

There are two parts to the proof of the channel coding theorem:

- **Achievability:** If $R < C$, then R is an achievable rate.
- **Converse:** If R is an achievable rate, then $R \leq C$.

We are going to start with the converse since it is much easier in this case.

Proof of converse. Consider a sequence of codes achieving rate R . We want to show $R \leq C$. Note that each choice of n induces a distribution on $(M, X^n, Y^n, \hat{M})$:

$$p(m, x^n, y^n, \hat{m}) = p(m)p(x^n|m) \left(\prod_{i=1}^n W(y_i|x_i) \right) p(\hat{m}|y^n)$$

with M uniformly distributed on $\mathcal{M} = [1 : 2^{nR}]$. Then

$$\begin{aligned}
 \log |\mathcal{M}| &= H(M) \\
 &= I(M; \hat{M}) + H(M|\hat{M}) && \text{(chain rule)} \\
 &\leq I(M; \hat{M}) + \Pr(\hat{M} \neq M) \log |\mathcal{M}| + 1. && \text{(Fano's inequality)}
 \end{aligned}$$

Since R is an achievable rate, we have $\Pr(\hat{M} \neq M) \rightarrow 0$ as $n \rightarrow \infty$. Recall that $\log |\mathcal{M}| = nR$. Thus

$$nR \leq I(M; \hat{M}) + o(n).$$

Now consider

$$\begin{aligned} I(M; \hat{M}) &\leq I(X^n; \hat{M}) && (M \rightarrow X^n \rightarrow \hat{M} \text{ is Markov}) \\ &\leq I(X^n; Y^n). && (X^n \rightarrow Y^n \rightarrow \hat{M} \text{ is Markov}) \end{aligned}$$

Also, by the sub-additivity of $I(X^n; Y^n)$ (Lemma 2.10),

$$\begin{aligned} I(X^n; Y^n) &\leq \sum_{i=1}^n I(X_i; Y_i) \\ &= nI(X_Q; Y_Q|Q), \end{aligned}$$

where in the last step we have introduced a random variable Q uniformly distributed on $[1 : n]$ and independent of the full block variables $(M, X^n, Y^n, \hat{M})$ (the Cesàro variable of §1.13). Now,

$$\begin{aligned} I(X_Q; Y_Q|Q) &\leq I(X_Q; Y_Q|Q) + I(Q; Y_Q) \\ &= I(Q, X_Q; Y_Q) && (\text{chain rule}) \\ &= I(X_Q; Y_Q) && (Q \rightarrow X_Q \rightarrow Y_Q \text{ is Markov}) \\ &\leq C \end{aligned}$$

since $C = \max_{p(x)} I(X; Y)$ by definition and $p_{Y_Q|X_Q}(y|x) = W(y|x)$. Here $\frac{1}{n} \sum_i I(X_i; Y_i) = I(X_Q; Y_Q|Q)$ averages the mutual informations of the *column* input laws p_{X_i} , while $I(X_Q; Y_Q)$ evaluates the mutual information at the *pooled* law $p_{X_Q} = \frac{1}{n} \sum_i p_{X_i}$; the two quantities need not be equal, and the inequality $I(X_Q; Y_Q|Q) \leq I(X_Q; Y_Q)$ (equivalently $I(Q; Y_Q) \geq 0$) is precisely the gap between them. To conclude, we have

$$\begin{aligned} nR &\leq I(M; \hat{M}) + o(n) \\ &\leq I(X^n; Y^n) + o(n) \\ &\leq nI(X_Q; Y_Q) + o(n) \\ &\leq nC + o(n). \end{aligned}$$

Dividing both sides by n and letting $n \rightarrow \infty$ completes the proof. ■

Remark 2.5. In the proof of the converse, we have used that

$$R \leq I(X_Q; Y_Q) + o(1)$$

where Q is a uniformly distributed random variable on $[1 : n]$ independent of the full block variables $(M, X^n, Y^n, \hat{M})$. Note that the distribution of X_Q is given by

$$\begin{aligned} p_{X_Q}(x) &= \sum_{i=1}^n p_Q(i) p_{X_Q|Q}(x|i) \\ &= \frac{1}{n} \sum_{i=1}^n p_{X_i}(x) \end{aligned}$$

$$\begin{aligned}
 &= \frac{1}{n} \frac{1}{|\mathcal{M}|} \sum_{i=1}^n \sum_{m \in \mathcal{M}} p_{X_i|M}(x|m) \\
 &= \frac{1}{n} \frac{1}{|\mathcal{M}|} |\{(i, m) \in [1 : n] \times \mathcal{M} : X_i|_{M=m} = x\}|.
 \end{aligned}$$

So X_Q is the empirical distribution of the codebook. This gives the important insight that any rate achieved by a sequence of codes is asymptotically bounded above by the quantity $I(X; Y)$ evaluated at the empirical distribution of the codebook. This is an upper bound only, not an identification of the two mutual-information values: a sharper converse upper bound is the column average $I(X_Q; Y_Q|Q)$; the pooled value $I(X_Q; Y_Q)$ can exceed it. Neither upper bound alone asserts that an arbitrary codebook achieves the indicated mutual-information value. Near-capacity code sequences have pooled mutual information approaching capacity. Therefore, to achieve a rate close to the capacity, the codebook should have an empirical distribution that maximizes $I(X; Y)$.

This is the sub-additivity principle of §1.13 in action: Lemma 2.10 bounds the n -letter mutual information $I(X^n; Y^n)$ by the sum $\sum_{i=1}^n I(X_i; Y_i)$ of single-letter terms, and the uniform time index Q identifies this sum (divided by n) with $I(X_Q; Y_Q|Q)$ — an average of $I(X; Y)$ over the columns of the codebook, i.e. the empirical distribution of the codebook. For the point-to-point channel the averaging over Q is not essential, since each summand $I(X_i; Y_i)$ is already at most C ; it is retained here only to mirror the template of the multiple access channel converse of Chapter 3, where the uniform Q is the time-sharing variable that genuinely appears in the single-letter region. $\circ$

We will prove achievability by showing, using probabilistic method, that the required codebooks exist, without explicitly constructing the codebooks. This is a very important distinction between information theory and coding theory. Coding theory studies the design and structure of codebooks, together with efficient encoding and decoding algorithms. Meanwhile, information theory cares about theoretical limits, i.e. the conditions under which codebooks exist.

We will prove achievability using random coding and joint typicality decoding. The idea is as follows: Suppose that the codebook consists of codewords $x^n \in \mathcal{X}^n$ that have an empirical distribution close to $p(x)$, or in other words, typical with respect to $p(x)$. Suppose that the decoder receives $y^n \in \mathcal{Y}^n$ at the channel output. If x^n is the codeword sent, then (x^n, y^n) will be typical with respect to $p(x, y)$ with high probability. If x^n is not the codeword sent, then (x^n, y^n) will be typical with respect to $p(x)p(y)$ with high probability and will be typical with respect to $p(x, y)$ with probability approximately $2^{-nD(p(x, y) \| p(x)p(y))} = 2^{-nI(X; Y)}$. Therefore, by testing whether (x^n, y^n) is typical with respect to $p(x, y)$ for each codeword x^n , the decoder rules out each wrong codeword with high probability, and a union bound shows it can handle codebooks with up to about $2^{nI(X; Y)}$ codewords. Now, instead of constructing concrete codebooks, we consider the ensemble of all codebooks, distributed as if they were generated according to i.i.d. copies of $p(x)$. As long as the ensemble average probability of error goes to 0, there will exist concrete codebooks required.

Proof of achievability. We want to show that any rate $R < C$ is achievable.

Random codebook generation: Since $R < C$, we can take a distribution $p(x)$ on $\mathcal{X}$ such that $R < I(X; Y)$, where the distribution on (X, Y) induced is given by $p(x, y) = p(x)W(y|x)$. For each $m \in \mathcal{M}$ and $i \in [1 : n]$, define i.i.d. random variables $X_i(m)$ distributed according to $p(x)$:

$$p(x^n(1), \dots, x^n(|\mathcal{M}|)) = \prod_{m \in \mathcal{M}} \prod_{i=1}^n p_X(x_i(m)).$$

Encoding: The encoder maps any message $m \in \mathcal{M}$ to the codeword $X^n(m)$. The generated codebook is assumed to be known to both the transmitter and the receiver.

Decoding: The decoder maps any received codeword y^n to the $\hat{m}$ if there is a unique $\hat{m} \in \mathcal{M}$ such that $(X^n(\hat{m}), y^n) \in \mathcal{T}_\epsilon^{(n)}(X, Y)$, or some predefined default message $\hat{m} = 1$ otherwise.

Probability of error analysis: We consider the overall joint distribution across all the random codebooks of the random variables induced by the channel, i.e. the distribution on $(M, X^n(1), \dots, X^n(|\mathcal{M}|), Y^n, \hat{M})$, which is given by

$$\begin{aligned} p(m, x^n(1), \dots, x^n(|\mathcal{M}|), y^n, \hat{m}) \\ = p(m) p(x^n(1), \dots, x^n(|\mathcal{M}|)) W^{\otimes n}(y^n | x^n(m)) \mathbf{1}\left\{\hat{m} = g_{x^n(1), \dots, x^n(|\mathcal{M}|)}(y^n)\right\}, \end{aligned}$$

where $g_{x^n(1), \dots, x^n(|\mathcal{M}|)}$ denotes the typicality decoder determined by the generated codebook, and

$$p(x^n(1), \dots, x^n(|\mathcal{M}|)) = \prod_{m \in \mathcal{M}} \prod_{i=1}^n p_X(x_i(m)).$$

Note that the decoder kernel is $P_{\hat{M}|Y^n, \mathcal{C}}$, conditioning on the *entire* generated codebook $\mathcal{C}$, not merely on the received word Y^n : the decoding regions depend on all the codewords. In the fixed deterministic-code converse below, conditioning on Y^n alone is legitimate because the codebook is then fixed.

Recall that, by construction, if M is the unique $\hat{m}$ such that $(X^n(\hat{m}), y^n) \in \mathcal{T}_\epsilon^{(n)}(X, Y)$ then $\hat{M} = M$. Thus, if the error event $\hat{M} \neq M$ occurred, then at least one of the following two events must also have occurred:

- $\mathcal{E}_0$: $(X^n(M), Y^n) \notin \mathcal{T}_\epsilon^{(n)}(X, Y)$.
- $\mathcal{E}_1$: $(X^n(\hat{m}), Y^n) \in \mathcal{T}_\epsilon^{(n)}(X, Y)$ for some $\hat{m} \neq M$.

Note that, by symmetry, both $\Pr(\mathcal{E}_0 | M = m)$ and $\Pr(\mathcal{E}_1 | M = m)$ do not depend on the choice of m . Now we consider the distribution on $(X^n(\hat{m}), Y^n)$ conditioned on $M = m$,

$$p(x^n(\hat{m}), y^n | m) = \begin{cases} \prod_{i=1}^n p_X(x_i(\hat{m})) W(y_i | x_i(\hat{m})), & \text{if } \hat{m} = m, \\ \prod_{i=1}^n p_X(x_i(\hat{m})) p_Y(y_i), & \text{if } \hat{m} \neq m. \end{cases}$$

In other words, conditioned on $M = m$, $(X^n(\hat{m}), Y^n)$ is distributed according to n i.i.d. copies $p(x, y)$ (respectively, $p(x)p(y)$) if $\hat{m} = m$ (respectively, $\hat{m} \neq m$).

Now consider

$$\begin{aligned}\Pr(\mathcal{E}_0) &= \Pr(\mathcal{E}_0 | M = m) \\ &= \Pr\left((X^n(M), Y^n) \notin \mathcal{T}_\epsilon^{(n)}(X, Y) \mid M = m\right) \\ &\rightarrow 0\end{aligned}$$

as $n \rightarrow \infty$, by the law of large numbers (Proposition 2.6). Also,

$$\begin{aligned}\Pr(\mathcal{E}_1) &= \Pr(\mathcal{E}_1 | M = m) \\ &= \Pr\left(\bigcup_{\substack{\hat{m} \in \mathcal{M} \\ \hat{m} \neq m}} \left\{(X^n(\hat{m}), Y^n) \in \mathcal{T}_\epsilon^{(n)}(X, Y)\right\} \mid M = m\right) \\ &\leq \sum_{\substack{\hat{m} \in \mathcal{M} \\ \hat{m} \neq m}} \Pr\left((X^n(\hat{m}), Y^n) \in \mathcal{T}_\epsilon^{(n)}(X, Y) \mid M = m\right) \quad (\text{union bound}) \\ &\stackrel{(a)}{\leq} |\mathcal{M}| \cdot (n+1)^{|\mathcal{X}||\mathcal{Y}|-1} \cdot 2^{-n(D(p_{XY} \| p_X p_Y) - \delta(\epsilon))} \quad (\text{Proposition 2.7}) \\ &= 2^{nR - n(D(p_{XY} \| p_X p_Y) - \delta(\epsilon)) + o(n)} \\ &= 2^{-n(I(X; Y) - R - \delta(\epsilon)) + o(n)} \\ &\rightarrow 0\end{aligned}$$

as $n \rightarrow \infty$, since $R < I(X; Y) - \delta(\epsilon)$ for small enough ϵ (recall $\delta(\epsilon) \rightarrow 0$). Note that for every $\hat{m} \neq m$, the codeword $X^n(\hat{m})$ is generated independently of $X^n(m)$. The claimed independence is read off the fixed-message ensemble: condition on $M = m$; for a fixed wrong index $\hat{m} \neq m$,

$$p(x^n(\hat{m}), x^n(m), y^n \mid m) = p_X^{\otimes n}(x^n(\hat{m})) p_X^{\otimes n}(x^n(m)) W^{\otimes n}(y^n \mid x^n(m)),$$

because the two codewords are independent i.i.d. draws and the output is generated through the memoryless channel from the transmitted codeword only. Marginalising $x^n(m)$ gives $p(x^n(\hat{m}), y^n \mid m) = \prod_i p(x_i)p(y_i)$, the product law needed by the packing estimate. Hence, (a) follows from Proposition 2.7. Then,

$$\begin{aligned}\Pr(\hat{M} \neq M) &\leq \Pr(\mathcal{E}_0 \cup \mathcal{E}_1) \\ &\leq \Pr(\mathcal{E}_0) + \Pr(\mathcal{E}_1) \\ &\rightarrow 0\end{aligned}$$

as $n \rightarrow \infty$.

We have shown that the average probability of error, taken over the random codebook ensemble, satisfies

$$\Pr(\hat{M} \neq M) \rightarrow 0.$$

Since this average is taken over all codebooks, there must exist at least one codebook whose probability of error is no larger than the ensemble average. Hence, for every n , we may choose a deterministic codebook $(x^n(1), \dots, x^n(|\mathcal{M}|))$ satisfying

$$\Pr(\hat{M} \neq M | X^n(1) = x^n(1), \dots, X^n(|\mathcal{M}|) = x^n(|\mathcal{M}|)) \leq \Pr(\hat{M} \neq M).$$

Since the right-hand side tends to 0 as $n \rightarrow \infty$, the resulting sequence of deterministic (n, R) codes also has probability of error tending to 0. Therefore, R is achievable. ■

2.3 Examples

2.3.1 Binary symmetric channel




Figure 2.3: The binary symmetric channel $\text{BSC}(p)$

Definition 2.12. The binary symmetric channel $\text{BSC}(p)$ with crossover probability $0 \leq p \leq 1$ is the channel $W : \{0, 1\} \rightarrow \{0, 1\}$ given by

$$W(y|x) = \begin{cases} 1-p, & \text{if } y = x, \\ p, & \text{otherwise.} \end{cases}$$

○

Proposition 2.13. The capacity of $\text{BSC}(p)$ is $1 - H_2(p)$, where $H_2(p) \triangleq -p \log_2 p - (1-p) \log_2 (1-p)$ denotes the binary entropy. ○

Proof. We can write $Y = X \oplus Z$ where $\oplus$ denotes the xor operation, Z is independent of X and distributed as $\Pr(Z = 0) = 1 - p$ and $\Pr(Z = 1) = p$. Then the capacity is

$$\begin{aligned} \max_{p(x)} I(X; Y) &= \max_{p(x)} (H(Y) - H(Y|X)) \\ &= \max_{p(x)} (H(Y) - H(X \oplus Z|X)) \\ &= \max_{p(x)} (H(Y) - H(Z|X)) \\ &= \max_{p(x)} H(Y) - H(Z) \quad (X, Z \text{ are independent}) \end{aligned}$$

$$\begin{aligned}
 &= 1 - H(Z) \\
 &= 1 - H_2(p),
 \end{aligned}$$

where the maximization is attained by $\Pr(X = 0) = \Pr(X = 1) = \frac{1}{2}$. ■

2.3.2 Binary erasure channel




Figure 2.4: The binary erasure channel $\text{BEC}(p)$

Definition 2.14. The binary erasure channel $\text{BEC}(p)$ with erasure probability $0 \leq p \leq 1$ is the channel $W : \{0, 1\} \rightarrow \{0, 1, ?\}$, where $? \neq 0, 1$ is called the erasure symbol, given by

$$W(y|x) = \begin{cases} 1-p, & \text{if } y = x, \\ p, & \text{if } y = ?, \\ 0, & \text{otherwise.} \end{cases}$$

○

Proposition 2.15. The capacity of $\text{BEC}(p)$ is $1 - p$. ○

Proof. The capacity is

$$\begin{aligned}
 \max_{p(x)} I(X; Y) &= \max_{p(x)} (H(X) - H(X|Y)) \\
 &= \max_{p(x)} \left(H(X) - p \underbrace{H(X|Y=?)}_{=H(X)} - (1-p) \underbrace{H(X|Y \neq ?)}_{=0} \right) \\
 &= \max_{p(x)} (H(X) - pH(X)) \\
 &= 1 - p,
 \end{aligned}$$

where the maximization is attained by $\Pr(X = 0) = \Pr(X = 1) = \frac{1}{2}$. ■

2.3.3 Product channel

Suppose that there are n channels $W_i : \mathcal{X}_i \rightarrow \mathcal{Y}_i$ in parallel. Suppose that the transmitter and the receiver can use these channels simultaneously. This can be regarded as that they are using one big resultant channel $W : \mathcal{X}_1 \times \cdots \times \mathcal{X}_n \rightarrow \mathcal{Y}_1 \times \cdots \times \mathcal{Y}_n$, called the product channel. What is its capacity?




Figure 2.5: Parallel channels

Proposition 2.16. The capacity C of the product channel $W \triangleq W_1 \otimes \cdots \otimes W_n$, i.e.

$$W(y^n|x^n) = \prod_{i=1}^n W_i(y_i|x_i),$$

is given by

$$C = \sum_{i=1}^n C_i,$$

where C_i is the capacity of the channel W_i . ○

Proof. For any (X^n, Y^n) jointly distributed according to

$$p(x^n, y^n) = p(x^n)W(y^n|x^n)$$

we have

$$\begin{aligned} I(X^n; Y^n) &= \sum_{i=1}^n I(X^n; Y_i | Y^{i-1}) && \text{(chain rule)} \\ &= \sum_{i=1}^n (I(X^n, Y^{i-1}; Y_i) - I(Y^{i-1}; Y_i)) && \text{(chain rule)} \\ &= \sum_{i=1}^n (I(X_i; Y_i) - I(Y^{i-1}; Y_i)) && \text{(memoryless Markov chain)} \end{aligned}$$

$$\begin{aligned}
 &\leq \sum_{i=1}^n I(X_i; Y_i) && \text{(sub-additivity, Lemma 2.10)} \\
 &\leq \sum_{i=1}^n C_i.
 \end{aligned}$$

Maximizing over all $p(x^n)$ gives $C \leq \sum_{i=1}^n C_i$. Now, note that if $X_1, \dots, X_n$ are independent, then $Y_1, \dots, Y_n$ are also independent, which gives

$$\begin{aligned}
 C &\geq I(X^n; Y^n) \\
 &= \sum_{i=1}^n (I(X_i; Y_i) - I(Y^{i-1}; Y_i)) \\
 &= \sum_{i=1}^n I(X_i; Y_i).
 \end{aligned}$$

Maximizing over all $p(x^n) = p(x_1) \dots p(x_n)$ gives $C \geq \sum_{i=1}^n C_i$. ■

Remark 2.6. In other words, correlation across parallel channels does not improve on the capacity. The optimal strategy is to use the channels independently. ○

2.3.4 Direct sum channel

Suppose that there are n channels $W_i : \mathcal{X}_i \rightarrow \mathcal{Y}_i$ in parallel. Suppose that the transmitter and the receiver are allowed to use one channel at a time. This can be regarded as that they are using an overall channel $W : \mathcal{X}_1 \cup \dots \cup \mathcal{X}_n \rightarrow \mathcal{Y}_1 \cup \dots \cup \mathcal{Y}_n$, called the direct sum channel, where $\cup$ denotes disjoint union. What is its capacity?




Figure 2.6: Parallel channels with disjoint alphabet

Proposition 2.17. The capacity C of the direct sum channel $W \triangleq W_1 \oplus \dots \oplus W_n$, i.e.

$$W(y|x) = \begin{cases} W_i(y|x), & \text{if } x \in \mathcal{X}_i \text{ and } y \in \mathcal{Y}_i, \\ 0, & \text{otherwise.} \end{cases}$$

is given by

$$2^C = \sum_{i=1}^n 2^{C_i},$$

where C_i is the capacity of the channel W_i . ○

Proof. Suppose (X, Y) are jointly distributed according to

$$p(x, y) = p(x)W(y|x).$$

Consider the random variable Q defined as the unique element in $[1 : n]$ such that $X \in \mathcal{X}_Q$. Then,

$$\begin{aligned} I(X; Y) &= I(Q, X; Y) && (Q \rightarrow X \rightarrow Y \text{ is Markov}) \\ &= I(Q; Y) + I(X; Y|Q) && (\text{chain rule}) \\ &= H(Q) - H(Q|Y) + \sum_q p(q)I(X; Y|Q = q) \\ &= H(Q) + \sum_q p(q)I(X; Y|Q = q). && (Q \text{ is a function of } Y) \end{aligned}$$

In the last equality we used $H(Q|Y) = 0$, since the output alphabets $\mathcal{Y}_i$ are disjoint, and observing Y uniquely determines which channel was used. Therefore, the capacity is given by

$$\begin{aligned} C &= \max_{p(x)} I(X; Y) \\ &= \max_{p(q), p(x|q)} \left(H(Q) + \sum_q p(q)I(X; Y|Q = q) \right) \\ &= \max_{p(q)} \left(H(Q) + \sum_q p(q) \max_{p(x|q)} I(X; Y|Q = q) \right) \\ &= \max_{p(q)} \left(H(Q) + \sum_q p(q)C_q \right) \\ &= \max_{p(q)} \sum_q p(q) \log \frac{2^{C_q}}{p(q)} \\ &= \max_{p(q)} \left(\log \sum_q 2^{C_q} - D(p(q) \parallel \hat{p}(q)) \right) && (\text{with } \hat{p}(q) \triangleq \frac{2^{C_q}}{\sum_q 2^{C_q}}) \\ &= \log \sum_q 2^{C_q}. \end{aligned}$$

■

Remark 2.7. A product channel permits simultaneous use of every component and has capacity $\sum_i C_i$. A direct-sum channel uses one branch at a time, but its disjoint output alphabets reveal which branch was chosen. Consequently it has capacity

$$\max_{\alpha} \left[H(\alpha) + \sum_i \alpha_i C_i \right] = \log_2 \sum_i 2^{C_i}.$$

The additional term $H(\alpha)$ is information conveyed by the branch selection. The direct-sum channel is still a DMC, so its block mutual information remains temporally sub-additive; this formula is not a failure of the channel-coding sub-additivity lemma. $\circ$

2.4 Source coding

Consider the source coding system depicted in [Figure 2.7](#).




Figure 2.7: Source coding system

The user observes a source sequence $X^n \in \mathcal{X}^n$ and encodes it into a message $M \in \mathcal{M}$ that can be decoded to an estimate $\hat{X}^n \in \mathcal{X}^n$ of the source sequence. We assume that the source is memoryless, i.e. the source can be represented by a distribution $p(x)$ on $\mathcal{X}$ and the source sequence is i.i.d. according to $p(x)$.

One can think of the message as a compression of the source sequence. Our goal is to compress at the best rate.

An (n, R) code for a memoryless source $p(x)$ consists of

- a message set $\mathcal{M} \triangleq [1 : 2^{nR}]$,
- a mapping $m : \mathcal{X}^n \rightarrow \mathcal{M}$ called the encoder,
- a mapping $\hat{x}^n : \mathcal{M} \rightarrow \mathcal{X}^n$ called the decoder.

We assume that the source sequence $X^n \in \mathcal{X}^n$ is i.i.d. according to $p(x)$. The encoded source M is given by $M \triangleq m(X^n)$. The source estimate $\hat{X}^n$ is given by $\hat{X}^n \triangleq \hat{x}^n(M)$.

$R \geq 0$ is said to be an achievable rate for a memoryless source if there exists a sequence of (n, R) codes whose probability of error $\Pr(\hat{X}^n \neq X^n)$ goes to 0 as $n \rightarrow \infty$. We say that such a sequence of codes achieves rate R . The infimum of all achievable rates is called the optimal rate of the source.

The source coding theorem states that a memoryless source can be compressed to approximately $H(X)$ bits per source symbol.

Theorem 2.18. (Source coding theorem) The optimal rate of a memoryless source $p(x)$ is $H(X)$. $\circ$

There are two parts to the proof of the source coding theorem:

- **Achievability:** If $R > H(X)$, then R is an achievable rate.
- **Converse:** If R is an achievable rate, then $R \geq H(X)$.

Proof of achievability. The achievability follows from the fact that almost all source sequences are typical and there are about $2^{nH(X)}$ such sequences. Formally, since $H(X) < R$, we can choose an $\epsilon > 0$ small enough such that $R > H(X)(1 + \epsilon)$. Then,

$$\begin{aligned} |\mathcal{T}_\epsilon^{(n)}(X)| &\leq 2^{nH(X)(1+\epsilon)} \\ &\leq 2^{nR} \end{aligned}$$

for n large enough, which implies that there is an injection $\mathcal{T}_\epsilon^{(n)}(X) \rightarrow [1 : 2^{nR}]$. The encoder is obtained by extending the domain of this map to $\mathcal{X}^n$ by mapping all $x^n \notin \mathcal{T}_\epsilon^{(n)}(X)$ to some predefined default message $m = 1$. Now, all source sequences in $\mathcal{T}_\epsilon^{(n)}(X)$ can be decoded without error. Hence

$$\begin{aligned} \Pr(\hat{X}^n \neq X^n) &\leq \Pr(X^n \notin \mathcal{T}_\epsilon^{(n)}(X)) \\ &\rightarrow 0 \end{aligned}$$

as $n \rightarrow \infty$, by the law of large numbers (Proposition 2.6). This shows the achievability. $\blacksquare$

Proof of converse. Consider a sequence of codes achieving rate R . We want to show $R \geq H(X)$. Then,

$$\begin{aligned} \log |\mathcal{M}| &\geq H(M) \\ &\geq H(M) - H(M|X^n) \\ &= I(M; X^n) \\ &\geq I(\hat{X}^n; X^n) && (X^n \rightarrow M \rightarrow \hat{X}^n \text{ is Markov}) \\ &= H(X^n) - H(X^n|\hat{X}^n) \\ &\geq H(X^n) - \Pr(\hat{X}^n \neq X^n) n \log |\mathcal{X}| - 1. && (\text{Fano's inequality}) \end{aligned}$$

Since R is an achievable rate, we have $\Pr(\hat{X}^n \neq X^n) \rightarrow 0$ as $n \rightarrow \infty$. Recall that $\log |\mathcal{M}| = nR$. Also note that $H(X^n) = nH(X)$ since $X_1, \dots, X_n$ are i.i.d. according to $p(x)$. Therefore,

$$nR \geq nH(X) - o(n).$$

Dividing both sides by n and letting $n \rightarrow \infty$ yields $R \geq H(X)$. $\blacksquare$

Remark 2.8 (Fano plus sub-additivity). The two converses of this chapter rest on the same two steps, which together are the organising principle of §1.13. First, Fano's inequality converts a vanishing probability of error into an entropic bound ($H(M|\hat{M}) = o(n)$ for channel coding, $H(X^n|\hat{X}^n) = o(n)$ for source coding). Second, the resulting n -letter information quantity collapses to a sum of single-letter terms: for channel coding this is exactly the sub-additivity of $I(X^n; Y^n)$ (Lemma 2.10), while for source coding the collapse $H(X^n) = \sum_{i=1}^n H(X_i) = nH(X)$ is immediate from the i.i.d. structure and needs no Cesàro variable. The multiple access channel converse of Chapter 3 applies the same sub-additivity step in its conditional form $I(X_1^n; Y^n|X_2^n) \leq \sum_i I(X_{1i}; Y_i|X_{2i})$ and joint form $I(X_1^n, X_2^n; Y^n) \leq \sum_i I(X_{1i}, X_{2i}; Y_i)$. $\bigcirc$

3 Multiple access channel

The multiple access channel (MAC) models the scenario in which multiple transmitters want to send a message to a single receiver. In particular, we focus on the case where there are two transmitters and the channel is memoryless.




Figure 3.1: A multiple access channel

A two-user memoryless multiple access channel is represented by a stochastic matrix $W : \mathcal{X}_1 \times \mathcal{X}_2 \rightarrow \mathcal{Y}$. Transmitter ℓ ($\ell = 1, 2$) encodes a message $M_\ell \in \mathcal{M}_\ell$ into a codeword $X_\ell^n \in \mathcal{X}_\ell^n$ which is then passed through the channel. The receiver receives the channel output $Y^n \in \mathcal{Y}^n$ and decodes it into an estimate $(\hat{M}_1, \hat{M}_2) \in \mathcal{M}_1 \times \mathcal{M}_2$ of the transmitted messages.

An (n, R_1, R_2) code consists of

- message sets $\mathcal{M}_1 \triangleq [1 : 2^{nR_1}]$ and $\mathcal{M}_2 \triangleq [1 : 2^{nR_2}]$,
- encoder mappings $x_1^n : \mathcal{M}_1 \rightarrow \mathcal{X}_1^n$ and $x_2^n : \mathcal{M}_2 \rightarrow \mathcal{X}_2^n$,
- a decoder mapping $\hat{m} : \mathcal{Y}^n \rightarrow \mathcal{M}_1 \times \mathcal{M}_2$.

As before, one may replace 2^{nR_ℓ} by $2^{\lceil nR_\ell \rceil}$ ($\ell = 1, 2$) to obtain integer message set sizes. This only affects $o(n)$ terms. We assume that the message sent by transmitter ℓ is a random variable M_ℓ independent across ℓ and uniformly distributed over the message set $\mathcal{M}_\ell$. The codeword sent by transmitter ℓ ($\ell = 1, 2$) is given by $X_\ell^n \triangleq x_\ell^n(M_\ell)$. The codeword received Y^n is a random variable such that

$$p(y^n | x_1^n, x_2^n, m_1, m_2) = \prod_{i=1}^n W(y_i | x_{1i}, x_{2i}).$$

The decoded message pair $(\hat{M}_1, \hat{M}_2)$ is a random variable on $\mathcal{M}_1 \times \mathcal{M}_2$ given by $(\hat{M}_1, \hat{M}_2) \triangleq \hat{m}(Y^n)$.

(R_1, R_2) is said to be an achievable rate if there exists a sequence of (n, R_1, R_2) codes whose probability of error

$$\begin{aligned} \Pr((\hat{M}_1, \hat{M}_2) \neq (M_1, M_2)) &= \Pr(\hat{M}_1 \neq M_1 \vee \hat{M}_2 \neq M_2) \\ &= \frac{1}{|\mathcal{M}_1 \times \mathcal{M}_2|} \sum_{m_1 \in \mathcal{M}_1, m_2 \in \mathcal{M}_2} \Pr(\hat{M}_1 \neq m_1 \vee \hat{M}_2 \neq m_2 | M_1 = m_1, M_2 = m_2) \end{aligned}$$

goes to 0 as $n \rightarrow \infty$. We say that such a sequence of codes achieves rate (R_1, R_2) . The closure of the set of all achievable rate pairs is called the capacity region of the channel.

3.1 Capacity region

Theorem 3.1. The capacity region of a multiple access channel $W_{Y|X_1, X_2}$ is the set of all *nonnegative* rate pairs $(R_1, R_2) \in \mathbb{R}_+^2$ satisfying

$$\begin{aligned} R_1 &\leq I(X_1; Y|X_2, Q), \\ R_2 &\leq I(X_2; Y|X_1, Q), \\ R_1 + R_2 &\leq I(X_1, X_2; Y|Q) \end{aligned} \tag{3.1}$$

for some distribution $p(q)p(x_1|q)p(x_2|q)$. We will later show that restricting the auxiliary random variable Q to satisfy $|\mathcal{Q}| \leq 2$ does not change the capacity region. $\circ$

There are two parts to the proof of the capacity region of the multiple access channel:

- **Converse:** If (R_1, R_2) is achievable, then

$$\begin{aligned} R_1 &\leq I(X_1; Y|X_2, Q), \\ R_2 &\leq I(X_2; Y|X_1, Q), \\ R_1 + R_2 &\leq I(X_1, X_2; Y|Q) \end{aligned}$$

for some distribution $p(q)p(x_1|q)p(x_2|q)$. Then, by the Fenchel–Bunt extension of Carathéodory’s theorem, the region remains unchanged after introducing the cardinality bound on Q .

- **Achievability:** We first show that if

$$\begin{aligned} R_1 &< I(X_1; Y|X_2), \\ R_2 &< I(X_2; Y|X_1), \\ R_1 + R_2 &< I(X_1, X_2; Y) \end{aligned}$$

for some distribution $p(x_1)p(x_2)$, then (R_1, R_2) is achievable. Once this achievable region is shown, we will see, using a time-sharing random variable Q , we can obtain the larger achievable region

$$R_1 < I(X_1; Y|X_2, Q),$$

$$R_2 < I(X_2; Y|X_1, Q),$$

$$R_1 + R_2 < I(X_1, X_2; Y|Q)$$

for some distribution $p(q)p(x_1|q)p(x_2|q)$.

3.2 Cardinality bound

The basic technique for bounding the cardinality of auxiliary random variables is to use Carathéodory's theorem which states that any point contained in the convex hull of a subset S of $\mathbb{R}^d$ can be written as a convex combination of at most $d + 1$ points in S . The key to the proof of Carathéodory's theorem is that any $d + 2$ or more points in $\mathbb{R}^d$ are affinely dependent and hence their convex combination can always be reduced to that of a proper subset of them.

It was first noticed by Fenchel–Bunt that Carathéodory's theorem can further be refined to that the convex combination can further be reduced to that of d points if S has at most d connected components. The idea is that if a point can be expressed as the convex combination of a minimum of $d + 1$ points in S , then these $d + 1$ points define a slicing of S into $d + 1$ disconnected parts.

Definition 3.2. (Affine dependence) Vectors $x_1, \dots, x_n$ are said to be affinely dependent if there exist scalars $\lambda_1, \dots, \lambda_n$ not all zero such that $\sum_{i=1}^n \lambda_i x_i = 0$ and $\sum_{i=1}^n \lambda_i = 0$, and are said to be affinely independent otherwise. $\circ$

Proposition 3.3. Vectors $x_1, \dots, x_n$ are affinely dependent if and only if $x_1 - x_n, \dots, x_{n-1} - x_n$ are linearly dependent. $\circ$

Proof. If $x_1, \dots, x_n$ are affinely dependent, then there exist scalars $\lambda_1, \dots, \lambda_n$ not all zero such that $\sum_{i=1}^n \lambda_i x_i = 0$ and $\sum_{i=1}^n \lambda_i = 0$, which implies that $\lambda_1, \dots, \lambda_{n-1}$ are not all zero (otherwise λ_n would also be zero) and $\sum_{i=1}^{n-1} \lambda_i (x_i - x_n) = 0$.

If $x_1 - x_n, \dots, x_{n-1} - x_n$ are linearly dependent, then there exist scalars $\lambda_1, \dots, \lambda_{n-1}$ not all zero such that $\sum_{i=1}^{n-1} \lambda_i (x_i - x_n) = 0$, which implies that $\sum_{i=1}^n \lambda_i x_i = 0$ and $\sum_{i=1}^n \lambda_i = 0$, if we set $\lambda_n \triangleq -\sum_{i=1}^{n-1} \lambda_i$. $\blacksquare$

Proposition 3.4. Any set of $d + 2$ or more vectors in $\mathbb{R}^d$ is affinely dependent. $\circ$

Proof. This follows from the fact that any set of $d + 1$ or more vectors in $\mathbb{R}^d$ is linearly dependent. $\blacksquare$

Theorem 3.5. (Carathéodory's theorem) Any point contained in the convex hull of a subset S of $\mathbb{R}^d$ can be written as a convex combination of at most $d + 1$ points in S . $\circ$

Proof. Suppose otherwise that x is a point contained in the convex hull of a subset S that can only be written as a convex combination of a minimum of $n \geq d + 2$ points $x_1, \dots, x_n$ in S , i.e. $x = \sum_{i=1}^n \lambda_i x_i$ for some $\lambda_1, \dots, \lambda_n > 0$ with $\sum_{i=1}^n \lambda_i = 1$. Then $x_1, \dots, x_n$ are affinely dependent, i.e. there exist $\mu_1, \dots, \mu_n$ not all zero such that $\sum_{i=1}^n \mu_i x_i = 0$ and $\sum_{i=1}^n \mu_i = 0$. Observe that $k \triangleq \min \left\{ \frac{\lambda_i}{\mu_i} : \mu_i > 0 \right\}$ is the maximum k such that $\lambda_i - k\mu_i \geq 0$ for $i = 1, \dots, n$. Then $\lambda_i - k\mu_i \geq 0$ with equality holding for some j . Now $x = \sum_{i=1}^n (\lambda_i - k\mu_i) x_i$ is a convex combination of at most $n - 1$ points in S , contradicting the minimality of n . ■

Theorem 3.6. (Fenchel–Bunt extension to Carathéodory’s theorem) Suppose $S \subseteq \mathbb{R}^d$ has at most d connected components. Then any point contained in the convex hull of S can be written as a convex combination of at most d points in S . ○

Proof. Without loss of generality, we may assume 0 is contained in the convex hull of S and it suffices to show that 0 can be written as a convex combination of at most d points in S . Suppose otherwise. By Carathéodory’s theorem, 0 can be written as a convex combination of $d + 1$ points $x_1, \dots, x_{d+1}$ in S , i.e. $\sum_{i=1}^{d+1} \lambda_i x_i = 0$ for some $\lambda_1, \dots, \lambda_{d+1} > 0$ with $\sum_{i=1}^{d+1} \lambda_i = 1$. Then $x_1, \dots, x_{d+1}$ are affinely independent, since otherwise we would be able to reduce the convex combination to that of at most d points. Furthermore, any d of $x_1, \dots, x_{d+1}$ are linearly independent; otherwise the relation $\sum_{i=1}^{d+1} \lambda_i x_i = 0$ would allow one of the points to be removed.

Let $C_i \triangleq \left\{ -\sum_{\ell \neq i} \alpha_\ell x_\ell : \alpha_\ell > 0 \right\}$, or in other words, C_i is the interior of the cone generated by $-x_1, \dots, -x_{i-1}, -x_{i+1}, \dots, -x_{d+1}$. We verify the three cone facts.

- $x_i \in C_i$: the $d \times (d + 1)$ column matrix of the x_ℓ has rank d and its kernel is spanned by the strictly positive coefficient vector λ ; from $\sum_\ell \lambda_\ell x_\ell = 0$ one solves $x_i = -\sum_{\ell \neq i} (\lambda_\ell / \lambda_i) x_\ell$ with positive coefficients, so x_i is an interior point of C_i .
- $\partial C_i \cap S = \emptyset$: any d columns of the matrix are linearly independent (a nontrivial relation among d columns would combine with the positive kernel vector to drop one point, contradicting minimality). If some point of S lay on ∂C_i , it would lie in the cone generated by at most $d - 1$ of the $-x_\ell$, i.e. zero would be a convex combination of that point and at most $d - 1$ of the x_ℓ ’s, again contradicting minimality.
- $C_i \cap C_j = \emptyset$ for any $i \neq j$: if C_i and C_j intersected, equating the two cone representations would give a kernel vector of the column matrix with opposite signs at positions i and j ; such a vector cannot be a multiple of the strictly positive kernel vector.

Now, the connected component of S that contains x_i must be a subset of C_i , since otherwise it would intersect both C_i and $S \setminus C_i$, which in turn implies that $\partial C_i \cap S \neq \emptyset$ by connectedness. But then S has at least $d + 1$ connected components, since C_i ’s are pairwise disjoint. This contradicts with the assumption that S has at most d connected components. ■

Reduction of the cardinality of $\mathcal{Q}$.

Let $\mathcal{C}$ denote the capacity region, i.e., the set of rate pairs (R_1, R_2) such that

$$R_1 \leq I(X_1; Y | X_2, Q),$$

$$\begin{aligned} R_2 &\leq I(X_2; Y|X_1, Q), \\ R_1 + R_2 &\leq I(X_1, X_2; Y|Q) \end{aligned}$$

for some distribution $p(q)p(x_1|q)p(x_2|q)$. The converse says that every achievable rate pair (R_1, R_2) is contained in C . We would like to obtain a bound of the cardinality of Q .

A simple application of Fenchel–Bunt extension to Carathéodory’s theorem gives $|Q| \leq 3$. This is done by replacing the distribution $p(q)p(x_1|q)p(x_2|q)$ without changing the right-hand sides of the inequalities. The set

$$\{(I(X_1; Y|X_2), I(X_2; Y|X_1), I(X_1, X_2; Y)) \in \mathbb{R}^3 : p(x_1)p(x_2)\}$$

is a connected subset of $\mathbb{R}^3$ since it is the image of a connected set under the continuous function

$$p(x_1)p(x_2) \mapsto (I(X_1; Y|X_2), I(X_2; Y|X_1), I(X_1, X_2; Y)).$$

Note that its convex hull is

$$\{(I(X_1; Y|X_2, Q), I(X_2; Y|X_1, Q), I(X_1, X_2; Y|Q)) \in \mathbb{R}^3 : p(q)p(x_1|q)p(x_2|q)\},$$

which, by Fenchel–Bunt extension to Carathéodory’s theorem, is equal to

$$\{(I(X_1; Y|X_2, Q), I(X_2; Y|X_1, Q), I(X_1, X_2; Y|Q)) \in \mathbb{R}^3 : p(q)p(x_1|q)p(x_2|q), |Q| \leq 3\}.$$

Therefore, C remains the same after imposing the constraint $|Q| \leq 3$.

We can further improve the cardinality bound to $|Q| \leq 2$. Let $\mathcal{R}$ denote the set of nonnegative rate pairs $(R_1, R_2) \in \mathbb{R}_+^2$ satisfying

$$\begin{aligned} R_1 &\leq I(X_1; Y|X_2), \\ R_2 &\leq I(X_2; Y|X_1), \\ R_1 + R_2 &\leq I(X_1, X_2; Y). \end{aligned} \tag{3.2}$$

for some distribution $p(x_1)p(x_2)$. We claim that C is the convex hull of $\mathcal{R}$.

Lemma 3.7. Let C be the capacity region (3.1) and let $\mathcal{R}$ be the region defined in (3.2). Then

$$C = \text{ConvHull}(\mathcal{R}).$$

○

Proof. It is immediate that $\mathcal{R} \subseteq C$ (by letting Q be constant). The region C is convex: mixing two auxiliary distributions $p(q)p(x_1|q)p(x_2|q)$ and $p'(q)p'(x_1|q)p'(x_2|q)$ yields a third law of the same form under which all three conditional mutual informations average (algebraic convexity; no operational achievability is needed for this). Hence $\text{ConvHull}(\mathcal{R}) \subseteq C$, and it is sufficient to prove $C \subseteq \text{ConvHull}(\mathcal{R})$.

To prove the reverse inclusion, it suffices to show that every point in C is a convex combination of points in $\mathcal{R}$. For a fixed distribution $p(q)p(x_1|q)p(x_2|q)$, every point in the corresponding MAC region is a convex combination of the points

- $(0, 0)$,
- $(I(X_1; Y|X_2, Q), 0)$,
- $(0, I(X_2; Y|X_1, Q))$,
- $(I(X_1; Y|X_2, Q), I(X_2; Y|Q))$,
- $(I(X_1; Y|Q), I(X_2; Y|X_1, Q))$.

The first three points are clearly in $\text{ConvHull}(\mathcal{R})$. Consider $(R_1, R_2) = (I(X_1; Y|X_2, Q), I(X_2; Y|Q))$. For each $q \in \mathcal{Q}$, define $(R_{1q}, R_{2q}) \triangleq (I(X_1; Y|X_2, Q = q), I(X_2; Y|Q = q))$. Then

$$\begin{aligned}
 R_{1q} &= I(X_1; Y|X_2, Q = q), \\
 R_{2q} &= I(X_2; Y|Q = q) \\
 &\leq I(X_2; Y, X_1|Q = q) \\
 &= I(X_2; Y|X_1, Q = q) + I(X_1; X_2|Q = q) && \text{(chain rule)} \\
 &= I(X_2; Y|X_1, Q = q), && (X_1 \rightarrow Q \rightarrow X_2 \text{ is Markov}),
 \end{aligned}$$

and

$$\begin{aligned}
 R_{1q} + R_{2q} &= I(X_1; Y|X_2, Q = q) + I(X_2; Y|Q = q) \\
 &= I(X_1, X_2; Y|Q = q), && \text{(chain rule)}.
 \end{aligned}$$

Hence $(R_{1q}, R_{2q}) \in \mathcal{R}$ for every q , and $(R_1, R_2) = \sum_{q \in \mathcal{Q}} p(q)(R_{1q}, R_{2q})$, so $(R_1, R_2) \in \text{ConvHull}(\mathcal{R})$.

Similarly, consider

$$(R_1, R_2) = (I(X_1; Y|Q), I(X_2; Y|X_1, Q)), \quad (R_{1q}, R_{2q}) \triangleq (I(X_1; Y|Q = q), I(X_2; Y|X_1, Q = q)).$$

Then

$$\begin{aligned}
 R_{1q} &= I(X_1; Y|Q = q) \\
 &\leq I(X_1; Y, X_2|Q = q) \\
 &= I(X_1; Y|X_2, Q = q) + I(X_1; X_2|Q = q) && \text{(chain rule)} \\
 &= I(X_1; Y|X_2, Q = q), && (X_1 \rightarrow Q \rightarrow X_2 \text{ is Markov}),
 \end{aligned}$$

while

$$\begin{aligned}
 R_{2q} &= I(X_2; Y|X_1, Q = q), \\
 R_{1q} + R_{2q} &= I(X_1; Y|Q = q) + I(X_2; Y|X_1, Q = q) \\
 &= I(X_1, X_2; Y|Q = q), && \text{(chain rule)}.
 \end{aligned}$$

Therefore $(R_{1q}, R_{2q}) \in \mathcal{R}$ for every q , and since $(R_1, R_2) = \sum_{q \in \mathcal{Q}} p(q)(R_{1q}, R_{2q})$, we have $(R_1, R_2) \in \text{ConvHull}(\mathcal{R})$.

We conclude that $C \subseteq \text{ConvHull}(\mathcal{R})$. ■

Now, note that $\mathcal{R}$ is a connected subset of $\mathbb{R}^2$: it is *star-shaped* about the origin, since every fixed-input pentagon in the union contains the origin together with the segment joining any of its points to the origin, and it is compact because the product-input parameter set is compact and the information functions are continuous (with the nonnegativity restriction making the feasible-rate graph closed and bounded). By Fenchel–Bunt extension to Carathéodory’s theorem, every point in C can be written as a convex combination of at most 2 points in $\mathcal{R}$. This implies that for any $(R_1, R_2) \in C$, there exists a distribution $p(q)p(x_1|q)p(x_2|q)$ with $|\mathcal{Q}| \leq 2$ such that

$$\begin{aligned} R_1 &\leq I(X_1; Y|X_2, Q), \\ R_2 &\leq I(X_2; Y|X_1, Q), \\ R_1 + R_2 &\leq I(X_1, X_2; Y|Q). \end{aligned}$$

In other words, C remains the same after imposing the constraint $|\mathcal{Q}| \leq 2$. This convex-hull structure ($C = \text{ConvHull}(\mathcal{R})$) is also why time-sharing realizes the full region: §3.5 below shows that every finite convex combination of points of $\mathcal{R}$ belongs to the capacity region, and $\text{ConvHull}(\mathcal{R})$ is exactly the time-sharing region in the closed-region sense.

3.3 Converse

Proof of converse. Suppose a sequence of (n, R_1, R_2) codes is given. The joint distribution of the random variables $(M_1, M_2, X_1^n, X_2^n, Y^n, \hat{M}_1, \hat{M}_2)$ is given by

$$\begin{aligned} &p(m_1, m_2, x_1^n, x_2^n, y^n, \hat{m}_1, \hat{m}_2) \\ &= p(m_1)p(m_2)p(x_1^n|m_1)p(x_2^n|m_2) \left(\prod_{i=1}^n W(y_i|x_{1i}, x_{2i}) \right) p(\hat{m}_1, \hat{m}_2|y^n), \end{aligned}$$

with (M_1, M_2) uniformly distributed on $\mathcal{M}_1 \times \mathcal{M}_2 \triangleq [1 : 2^{nR_1}] \times [1 : 2^{nR_2}]$.

Consider

$$\begin{aligned} \log |\mathcal{M}_1| &= H(M_1) \\ &= H(M_1|M_2) && (M_1, M_2 \text{ are independent}) \\ &= I(M_1; \hat{M}_1|M_2) + H(M_1|\hat{M}_1, M_2) && (\text{chain rule}) \\ &\leq I(M_1; \hat{M}_1|M_2) + H(M_1|\hat{M}_1) \\ &\leq I(M_1; \hat{M}_1|M_2) + \Pr(\hat{M}_1 \neq M_1) \log |\mathcal{M}_1| + 1. && (\text{Fano's inequality}) \end{aligned}$$

By the assumption, we have $\Pr(\hat{M}_1 \neq M_1) \rightarrow 0$ as $n \rightarrow \infty$. Recall that $\log |\mathcal{M}_1| = nR_1$. Thus

$$nR_1 \leq I(M_1; \hat{M}_1|M_2) + o(n).$$

Now consider

$$\begin{aligned}
 I(M_1; \hat{M}_1 | M_2) &= I(M_1; \hat{M}_1, M_2) && (I(M_1; M_2) = 0 \text{ since } M_1, M_2 \text{ are independent}) \\
 &\leq I(X_1^n; \hat{M}_1, M_2) && (M_1 \rightarrow X_1^n \rightarrow (\hat{M}_1, M_2) \text{ is Markov}) \\
 &\leq I(X_1^n; Y^n, X_2^n) && (X_1^n \rightarrow (Y^n, X_2^n) \rightarrow (\hat{M}_1, M_2) \text{ is Markov}) \\
 &= I(X_1^n; Y^n, X_2^n) - I(X_1^n; X_2^n) && (X_1^n, X_2^n \text{ are independent}) \\
 &= I(X_1^n; Y^n | X_2^n). && (\text{chain rule})
 \end{aligned}$$

(The chain $X_1^n \rightarrow (Y^n, X_2^n) \rightarrow (\hat{M}_1, M_2)$ is valid under the separate-encoder factorization, even if encoder 2 is not injective: given X_2^n , the posterior law of M_2 does not depend on (X_1^n, Y^n) , so no injectivity assumption is needed.) Also,

$$\begin{aligned}
 &I(X_1^n; Y^n | X_2^n) \\
 &= \sum_{i=1}^n I(X_1^n; Y_i | X_2^n, Y^{i-1}) && (\text{chain rule}) \\
 &\leq \sum_{i=1}^n \left(I(X_1^n; Y_i | X_2^n, Y^{i-1}) + I(X_2^{n \setminus i}, Y^{i-1}; Y_i | X_{2i}) \right) && (\text{sub-additivity step}) \\
 &= \sum_{i=1}^n I(X_1^n, X_2^{n \setminus i}, Y^{i-1}; Y_i | X_{2i}) && (\text{chain rule}) \\
 &= \sum_{i=1}^n I(X_{1i}; Y_i | X_{2i}) && (\text{sub-additivity, §1.13}) \\
 &= nI(X_{1Q}; Y_Q | X_{2Q}, Q),
 \end{aligned}$$

where the first inequality is the sub-additivity step of §1.13 applied to the conditional functional $I(X_1^n; Y^n | X_2^n)$ (adding the non-negative term $I(X_2^{n \setminus i}, Y^{i-1}; Y_i | X_{2i})$), the second-to-last equality uses the memoryless Markov chain $(X_1^{n \setminus i}, X_2^{n \setminus i}, Y^{i-1}) \rightarrow (X_{1i}, X_{2i}) \rightarrow Y_i$, and in the last step we have introduced a random variable Q uniformly distributed on $[1 : n]$, which can be interpreted as a time-sharing variable. Thus,

$$R_1 \leq I(X_{1Q}; Y_Q | X_{2Q}, Q) + o(1).$$

Similarly we have

$$R_2 \leq I(X_{2Q}; Y_Q | X_{1Q}, Q) + o(1).$$

Now we show the bound for $R_1 + R_2$. Consider

$$\begin{aligned}
 \log |\mathcal{M}_1 \times \mathcal{M}_2| &= H(M_1, M_2) \\
 &= I(M_1, M_2; \hat{M}_1, \hat{M}_2) + H(M_1, M_2 | \hat{M}_1, \hat{M}_2) && (\text{chain rule}) \\
 &\leq I(M_1, M_2; \hat{M}_1, \hat{M}_2) \\
 &\quad + \Pr(\hat{M}_1 \neq M_1 \vee \hat{M}_2 \neq M_2) \log |\mathcal{M}_1 \times \mathcal{M}_2| + 1. && (\text{Fano's inequality})
 \end{aligned}$$

By the assumption, we have

$$\begin{aligned} \Pr(\hat{M}_1 \neq M_1 \vee \hat{M}_2 \neq M_2) &\leq \Pr(\hat{M}_1 \neq M_1) + \Pr(\hat{M}_2 \neq M_2) \\ &\rightarrow 0 \end{aligned}$$

as $n \rightarrow \infty$. Recall that $\log |\mathcal{M}_1 \times \mathcal{M}_2| = n(R_1 + R_2)$. Thus

$$n(R_1 + R_2) \leq I(M_1, M_2; \hat{M}_1, \hat{M}_2) + o(n).$$

Now consider

$$\begin{aligned} I(M_1, M_2; \hat{M}_1, \hat{M}_2) &\leq I(X_1^n, X_2^n; \hat{M}_1, \hat{M}_2) \quad ((M_1, M_2) \rightarrow (X_1^n, X_2^n) \rightarrow (\hat{M}_1, \hat{M}_2) \text{ is Markov}) \\ &\leq I(X_1^n, X_2^n; Y^n). \quad ((X_1^n, X_2^n) \rightarrow Y^n \rightarrow (\hat{M}_1, \hat{M}_2) \text{ is Markov}) \end{aligned}$$

Also,

$$\begin{aligned} &I(X_1^n, X_2^n; Y^n) \\ &= \sum_{i=1}^n I(X_1^n, X_2^n; Y_i | Y^{i-1}) \quad (\text{chain rule}) \\ &\leq \sum_{i=1}^n (I(X_1^n, X_2^n; Y_i | Y^{i-1}) + I(Y^{i-1}; Y_i)) \quad (\text{sub-additivity step}) \\ &= \sum_{i=1}^n I(X_1^n, X_2^n, Y^{i-1}; Y_i) \quad (\text{chain rule}) \\ &= \sum_{i=1}^n I(X_{1i}, X_{2i}; Y_i) \quad (\text{sub-additivity, §1.13}) \\ &= nI(X_{1Q}, X_{2Q}; Y_Q | Q). \end{aligned}$$

(The first inequality adds the non-negative term $I(Y^{i-1}; Y_i)$; the last equality uses the memoryless Markov chain $(X_1^{n \setminus i}, X_2^{n \setminus i}, Y^{i-1}) \rightarrow (X_{1i}, X_{2i}) \rightarrow Y_i$.) Thus,

$$R_1 + R_2 \leq I(X_{1Q}, X_{2Q}; Y_Q | Q) + o(1).$$

To conclude, we have

$$\begin{aligned} R_1 &\leq I(X_{1Q}; Y_Q | X_{2Q}, Q) + o(1), \\ R_2 &\leq I(X_{2Q}; Y_Q | X_{1Q}, Q) + o(1), \\ R_1 + R_2 &\leq I(X_{1Q}, X_{2Q}; Y_Q | Q) + o(1), \end{aligned} \tag{3.3}$$

where Q is uniformly distributed on $[1 : n]$ and independent of the messages and the full block sequences $(M_1, M_2, X_1^n, X_2^n, Y^n, \hat{M}_1, \hat{M}_2)$. (The indexed variables X_{1Q}, X_{2Q}, Y_Q themselves depend on Q ; only the block sequences and messages are independent of it.) Now, observe that

$$I(X_{1Q}, X_{2Q} | Q) = \sum_q p(q) I(X_{1Q}, X_{2Q} | Q = q)$$

$$\begin{aligned}
&= \sum_q p(q) I(X_{1q}; X_{2q}) \\
&= 0,
\end{aligned}$$

which implies that X_{1Q} and X_{2Q} are conditionally independent given Q . Moreover,

$$\begin{aligned}
p(Y_Q = y \mid X_{1Q} = x_1, X_{2Q} = x_2, Q = q) &= p(Y_q = y \mid X_{1q} = x_1, X_{2q} = x_2, Q = q) \\
&= p(Y_q = y \mid X_{1q} = x_1, X_{2q} = x_2) \\
&= W(y \mid x_1, x_2),
\end{aligned}$$

where the last equality follows from the memorylessness of the channel. Consequently, the joint distribution of (Q, X_{1Q}, X_{2Q}, Y_Q) factors as

$$p(q, x_1, x_2, y) = p(q)p(x_1 \mid q)p(x_2 \mid q)W(y \mid x_1, x_2).$$

Therefore, for each n , there exists a distribution $p(q)p(x_1 \mid q)p(x_2 \mid q)$ such that (R_1, R_2) satisfies (3.3). Equivalently, taking a common $\epsilon_n \downarrow 0$ dominating the finitely many slack terms, $((R_1 - \epsilon_n)^+, (R_2 - \epsilon_n)^+) \in C$ (the positive part keeps the approximating points nonnegative, and C contains the origin). Since C is compact (and hence closed) and $((R_1 - \epsilon_n)^+, (R_2 - \epsilon_n)^+) \rightarrow (R_1, R_2)$ as $n \rightarrow \infty$, it follows that $(R_1, R_2) \in C$. ■

3.4 Achievability

Proof of achievability. It is enough to show the case where Q is constant. For the achievability when Q is non-trivial follows from a time-sharing argument. We will show the required achievability by random coding and joint typicality decoding.

We will show that every rate pair (R_1, R_2) satisfying the *strict* inequalities

$$R_1 < I(X_1; Y \mid X_2), \quad R_2 < I(X_2; Y \mid X_1), \quad R_1 + R_2 < I(X_1, X_2; Y)$$

for some product distribution $p(x_1)p(x_2)$ is achievable. Boundary points of the closed region (3.2) then belong to the capacity region by closure (the capacity region is closed by definition; the random-coding estimates below establish only strict rate margins). Zero-rate users are handled by singleton message sets, and all error classes requiring an incorrect message for such a user are simply absent.

Random codebook generation: Fix distributions $p(x_1), p(x_2)$ on $\mathcal{X}_1, \mathcal{X}_2$, respectively. This induces a distribution on (X_1, X_2, Y) by $p(x_1, x_2, y) = p(x_1)p(x_2)W(y \mid x_1, x_2)$. For each $m_1 \in \mathcal{M}_1$, $m_2 \in \mathcal{M}_2$ and $i \in [1 : n]$, define i.i.d. random variables $X_{1i}(m_1)$ and $X_{2i}(m_2)$ distributed according to $p(x_1)$ and $p(x_2)$ accordingly.

$$p(x_1^n(1), \dots, x_1^n(|\mathcal{M}_1|),$$

$$\begin{aligned}
 & x_2^n(1), \dots, x_2^n(|\mathcal{M}_2|) \\
 &= \prod_{m_1 \in \mathcal{M}_1} \prod_{i=1}^n p_{X_1}(x_{1i}(m_1)) \prod_{m_2 \in \mathcal{M}_2} \prod_{i=1}^n p_{X_2}(x_{2i}(m_2)).
 \end{aligned}$$

Encoding: Encoder ℓ ($\ell = 1, 2$) maps any message $m_\ell \in \mathcal{M}_\ell$ to the codeword $X_\ell^n(m_\ell)$. The generated codebook is assumed to be known to both the transmitter and the receiver.

Decoding: The decoder maps any received sequence y^n to the $(\hat{m}_1, \hat{m}_2)$ if there is a unique $(\hat{m}_1, \hat{m}_2) \in \mathcal{M}_1 \times \mathcal{M}_2$ such that $(X_1^n(\hat{m}_1), X_2^n(\hat{m}_2), y^n) \in \mathcal{T}_\epsilon^{(n)}(X_1, X_2, Y)$, or some predefined default message $(\hat{m}_1, \hat{m}_2) = (1, 1)$ otherwise.

Probability of error analysis: We consider the overall joint distribution across all the random codebooks of the random variables induced by the channel,

$$(M_1, M_2, X_1^n(1), \dots, X_1^n(|\mathcal{M}_1|), X_2^n(1), \dots, X_2^n(|\mathcal{M}_2|), Y^n, \hat{M}_1, \hat{M}_2),$$

which is jointly distributed according to

$$\begin{aligned}
 & p(m_1, m_2, x_1^n(1), \dots, x_1^n(|\mathcal{M}_1|), x_2^n(1), \dots, x_2^n(|\mathcal{M}_2|), y^n, \hat{m}_1, \hat{m}_2) \\
 &= p(m_1) p(m_2) p(x_1^n(1), \dots, x_1^n(|\mathcal{M}_1|), x_2^n(1), \dots, x_2^n(|\mathcal{M}_2|)) W^{\otimes n}(y^n | x_1^n(m_1), x_2^n(m_2)) \\
 & \quad \times \mathbf{1}\{(\hat{m}_1, \hat{m}_2) = g_{x_1^n(1), \dots, x_1^n(|\mathcal{M}_1|), x_2^n(1), \dots, x_2^n(|\mathcal{M}_2|)}(y^n)\},
 \end{aligned}$$

where $g_{x_1^n(1), \dots, x_1^n(|\mathcal{M}_1|), x_2^n(1), \dots, x_2^n(|\mathcal{M}_2|)}$ denotes the typicality decoder determined by the generated pair of codebooks, and

$$\begin{aligned}
 & p(x_1^n(1), \dots, x_1^n(|\mathcal{M}_1|), x_2^n(1), \dots, x_2^n(|\mathcal{M}_2|)) \\
 &= \prod_{m_1 \in \mathcal{M}_1} \prod_{i=1}^n p_{X_1}(x_{1i}(m_1)) \prod_{m_2 \in \mathcal{M}_2} \prod_{i=1}^n p_{X_2}(x_{2i}(m_2)). \\
 & W^{\otimes n}(y^n | x_1^n(m_1), x_2^n(m_2)) = \prod_{i=1}^n W(y_i | x_{1i}(m_1), x_{2i}(m_2)).
 \end{aligned}$$

Note that the decoder kernel is $P_{\hat{M}_1, \hat{M}_2 | Y^n, C_1, C_2}$, conditioning on the *entire* generated pair of codebooks (C_1, C_2) , not merely on the received word Y^n : the decoding regions depend on all the codewords. In the fixed deterministic-code converse, conditioning on Y^n alone is legitimate because the codebooks are then fixed.

Recall that, by construction, if (M_1, M_2) is the unique $(\hat{m}_1, \hat{m}_2)$ such that $(X_1^n(\hat{m}_1), X_2^n(\hat{m}_2), y^n) \in \mathcal{T}_\epsilon^{(n)}(X_1, X_2, Y)$ then $(\hat{M}_1, \hat{M}_2) = (M_1, M_2)$. Thus, if the error event $(\hat{M}_1, \hat{M}_2) \neq (M_1, M_2)$ occurred, then at least one of the following four events must also have occurred:

- $\mathcal{E}_0$: $(X_1^n(M_1), X_2^n(M_2), Y^n) \notin \mathcal{T}_\epsilon^{(n)}(X_1, X_2, Y)$.
- $\mathcal{E}_1$: $(X_1^n(\hat{m}_1), X_2^n(M_2), Y^n) \in \mathcal{T}_\epsilon^{(n)}(X_1, X_2, Y)$ for some $\hat{m}_1 \neq M_1$.

- $\mathcal{E}_2$: $(X_1^n(M_1), X_2^n(\hat{m}_2), Y^n) \in \mathcal{T}_\epsilon^{(n)}(X_1, X_2, Y)$ for some $\hat{m}_2 \neq M_2$.
- $\mathcal{E}_3$: $(X_1^n(\hat{m}_1), X_2^n(\hat{m}_2), Y^n) \in \mathcal{T}_\epsilon^{(n)}(X_1, X_2, Y)$ for some $\hat{m}_1 \neq M_1$ and $\hat{m}_2 \neq M_2$.

Note that, by symmetry, each of $\Pr(\mathcal{E}_j | M_1 = m_1, M_2 = m_2)$ does not depend on the choice of (m_1, m_2) . Also note that $(X_{1i}(\hat{m}_1), X_{2i}(\hat{m}_2), Y_i)$ conditioned on $(M_1, M_2) = (m_1, m_2)$ is independent across $i \in [1 : n]$ and has joint distribution summarized in the following table:

Event	$\hat{m}_1$	$\hat{m}_2$	Joint distribution
$\mathcal{E}_0$	$= m_1$	$= m_2$	$p(x_1)p(x_2)p(y x_1, x_2) = p(x_1, x_2, y)$
$\mathcal{E}_1$	$\neq m_1$	$= m_2$	$p(x_1)p(x_2)p(y x_2) = p(x_1)p(x_2, y)$
$\mathcal{E}_2$	$= m_1$	$\neq m_2$	$p(x_1)p(x_2)p(y x_1) = p(x_2)p(x_1, y)$
$\mathcal{E}_3$	$\neq m_1$	$\neq m_2$	$p(x_1)p(x_2)p(y)$

An application of the law of large numbers (Proposition 2.6) gives

$$\Pr(\mathcal{E}_0) = \Pr(\mathcal{E}_0 | M_1 = m_1, M_2 = m_2) \rightarrow 0$$

as $n \rightarrow \infty$. Also,

$$\begin{aligned} \Pr(\mathcal{E}_1) &= \Pr(\mathcal{E}_1 | M_1 = m_1, M_2 = m_2) \\ &\leq |\mathcal{M}_1| \cdot 2^{-n(D(p(x_1, x_2, y) \| p(x_1)p(x_2, y)) - O(\epsilon)) + o(n)} && \text{(Proposition 2.7)} \\ &= 2^{-n(I(X_1; X_2, Y) - R_1 - O(\epsilon)) + o(n)} \\ &= 2^{-n(I(X_1; Y | X_2) + I(X_1; X_2) - R_1 - O(\epsilon)) + o(n)} && \text{(chain rule)} \\ &= 2^{-n(I(X_1; Y | X_2) - R_1 - O(\epsilon)) + o(n)} && (X_1, X_2 \text{ are independent}) \\ &\rightarrow 0 \end{aligned}$$

as $n \rightarrow \infty$, since $R_1 < I(X_1; Y | X_2) - O(\epsilon)$ for small enough ϵ . (The $O(\epsilon)$ terms here and below hide the subset-dependent packing penalties $\delta_S(\epsilon) \rightarrow 0$ of Proposition 2.7; only the finitely many events $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3$ occur, so taking the maximum of their finitely many penalties gives a single small ϵ that serves all three bounds at once.) Similarly,

$$\begin{aligned} \Pr(\mathcal{E}_2) &= \Pr(\mathcal{E}_2 | M_1 = m_1, M_2 = m_2) \\ &\leq |\mathcal{M}_2| \cdot 2^{-n(D(p(x_1, x_2, y) \| p(x_2)p(x_1, y)) - O(\epsilon)) + o(n)} && \text{(Proposition 2.7)} \\ &= 2^{-n(I(X_2; X_1, Y) - R_2 - O(\epsilon)) + o(n)} \\ &= 2^{-n(I(X_2; Y | X_1) + I(X_2; X_1) - R_2 - O(\epsilon)) + o(n)} && \text{(chain rule)} \\ &= 2^{-n(I(X_2; Y | X_1) - R_2 - O(\epsilon)) + o(n)} && (X_1, X_2 \text{ are independent}) \\ &\rightarrow 0 \end{aligned}$$

as $n \rightarrow \infty$, since $R_2 < I(X_2; Y | X_1) - O(\epsilon)$ for small enough ϵ . Also,

$$\Pr(\mathcal{E}_3) = \Pr(\mathcal{E}_3 | M_1 = m_1, M_2 = m_2)$$

$$\begin{aligned}
&\leq |\mathcal{M}_1 \times \mathcal{M}_2| \cdot 2^{-n(D(p(x_1, x_2, y) \| p(x_1)p(x_2)p(y)) - O(\epsilon)) + o(n)} \quad (\text{Proposition 2.7}) \\
&= 2^{-n(I(X_1, X_2; Y) - R_1 - R_2 - O(\epsilon)) + o(n)} \\
&\rightarrow 0
\end{aligned}$$

as $n \rightarrow \infty$, since $R_1 + R_2 < I(X_1, X_2; Y) - O(\epsilon)$ for small enough ϵ . Therefore,

$$\begin{aligned}
\Pr(\hat{M} \neq M) &\leq \Pr(\mathcal{E}_0 \cup \mathcal{E}_1 \cup \mathcal{E}_2 \cup \mathcal{E}_3) \\
&\leq \Pr(\mathcal{E}_0) + \Pr(\mathcal{E}_1) + \Pr(\mathcal{E}_2) + \Pr(\mathcal{E}_3) \\
&\rightarrow 0
\end{aligned}$$

as $n \rightarrow \infty$, where $M = (M_1, M_2)$ and $\hat{M} = (\hat{M}_1, \hat{M}_2)$.

We have shown that the average probability of error, taken over the random codebook ensemble, satisfies

$$\Pr(\hat{M} \neq M) \rightarrow 0.$$

Since this average is taken over all codebooks, there must exist at least one codebook whose probability of error is no larger than the ensemble average. Hence, for every n , we choose a deterministic codebook $(x_1^n(1), \dots, x_1^n(|\mathcal{M}_1|), x_2^n(1), \dots, x_2^n(|\mathcal{M}_2|))$ satisfying

$$\Pr(\hat{M} \neq M | X_1^n(m_1) = x_1^n(m_1), X_2^n(m_2) = x_2^n(m_2), \forall m_1 \in \mathcal{M}_1, \forall m_2 \in \mathcal{M}_2) \leq \Pr(\hat{M} \neq M).$$

■

3.5 Time-sharing

This section makes the convex-hull structure of the capacity region operational. Lemma 3.7 in §3.2 already established that $C = \text{ConvHull}(\mathcal{R})$; here we verify that every finite convex combination of rate pairs from $\mathcal{R}$ belongs to the capacity region, so that time sharing realizes the convex hull in the closed-region sense. This convexification is the achievable-side counterpart of the single-letterisation step of §1.13, which bounds the block region by single-letter conditional functionals in the MAC converse of §3.3.

Assume $|\mathcal{Q}| < \infty$ and that $p(q) > 0$ for every $q \in \mathcal{Q}$ (discard zero-mass atoms). Suppose that for each $q \in \mathcal{Q}$, the rate pair $(R_{1q}, R_{2q}) \in \mathcal{R}$ is achievable. Let $p(q)$ be a probability distribution on $\mathcal{Q}$, and define

$$(R_1, R_2) \triangleq \left(\sum_{q \in \mathcal{Q}} p(q) R_{1q}, \sum_{q \in \mathcal{Q}} p(q) R_{2q} \right).$$

We show that (R_1, R_2) belongs to the capacity region. First assume the constituent rate pairs are achievable, and construct concatenated codes whose rates converge to the indicated convex average. For arbitrary constituent points of $\mathcal{R}$, approximate them by achievable rate pairs as

established in the preceding achievability argument (using singleton message sets for zero-rate users), apply the finite construction, and then pass to the limit using closure of the capacity region. For each $q \in \mathcal{Q}$, let $n_q \triangleq \lfloor np(q) \rfloor$. Since (R_{1q}, R_{2q}) is achievable, for every q there exists a sequence of (n_q, R_{1q}, R_{2q}) codes whose probability of error tends to zero as $n_q \rightarrow \infty$:

- Message sets $\mathcal{M}_{1q} \times \mathcal{M}_{2q} \triangleq [1 : 2^{n_q R_{1q}}] \times [1 : 2^{n_q R_{2q}}]$,
- Encoders $x_{1q}^{n_q} : \mathcal{M}_{1q} \rightarrow \mathcal{X}_1^{n_q}, x_{2q}^{n_q} : \mathcal{M}_{2q} \rightarrow \mathcal{X}_2^{n_q}$,
- Decoders $\hat{m} : \mathcal{Y}^{n_q} \rightarrow \mathcal{M}_{1q} \times \mathcal{M}_{2q}$.

Construct a new code by concatenating these codes. Specifically, divide the transmission block into $|\mathcal{Q}|$ subblocks, where the q -th subblock has length n_q , and use the (n_q, R_{1q}, R_{2q}) code on that subblock:

- Message sets $\mathcal{M}_1 \times \mathcal{M}_2 \triangleq \left(\prod_{q \in \mathcal{Q}} [1 : 2^{n_q R_{1q}}] \right) \times \left(\prod_{q \in \mathcal{Q}} [1 : 2^{n_q R_{2q}}] \right)$,
- Encoders $x_1^{N_n} : \mathcal{M}_1 \rightarrow \mathcal{X}_1^{N_n}, x_2^{N_n} : \mathcal{M}_2 \rightarrow \mathcal{X}_2^{N_n}$,
- Decoders $\hat{m} : \mathcal{Y}^{N_n} \rightarrow \mathcal{M}_1 \times \mathcal{M}_2$,

where $N_n \triangleq \sum_{q \in \mathcal{Q}} n_q$ is blocklength. (If every physical blocklength is required, pad the leftover $n - N_n$ channel uses with fixed symbols; padding changes the rates by $O(1/n)$, hence not their limits, and it does not increase the error probability.) Then,

$$N_n = \sum_{q \in \mathcal{Q}} \lfloor np(q) \rfloor = n - O(1),$$

since $\mathcal{Q}$ is finite and $\sum_q p(q) = 1$. Hence,

$$\begin{aligned} \frac{1}{N_n} \log |\mathcal{M}_1| &\rightarrow R_1, \\ \frac{1}{N_n} \log |\mathcal{M}_2| &\rightarrow R_2. \end{aligned}$$

Let $P_{e,q}^{(n_q)}$ denote the probability of error of the q -th constituent code. By the union bound,

$$P_e^{(N_n)} \leq \sum_{q \in \mathcal{Q}} P_{e,q}^{(n_q)}.$$

Since $|\mathcal{Q}| < \infty$ and $P_{e,q}^{(n_q)} \rightarrow 0$ for every q , it follows that $P_e^{(N_n)} \rightarrow 0$. The concatenated codes have rates converging to (R_1, R_2) and probability of error tending to zero. Consequently (R_1, R_2) belongs to the capacity region. Every target rate pair with a fixed positive slack below this limit can be obtained by restricting the message sets for sufficiently large blocklengths. Closure of the capacity region then yields the full time-sharing region.

Therefore every finite convex combination of points of $\mathcal{R}$ belongs to the capacity region. Since, for the MAC, the cardinality bound $|\mathcal{Q}| \leq 2$ is sufficient, this completes the time-sharing argument needed for the characterization of the capacity region.

Remark 3.1. The time-sharing argument can be extended to the case where $\mathcal{Q}$ is infinite. Although this is not needed for the MAC because the cardinality bound $|\mathcal{Q}| \leq 2$ is sufficient, the extension is conceptually useful.

For *countably supported* $\mathcal{Q}$, choose a finite subset $F \subseteq \mathcal{Q}$ with $\sum_{q \in F} p(q) > 1 - \epsilon$ and proceed as follows. Since achievable rates are bounded (they lie in compact region $\mathcal{R}$), we can assume $\sup_q R_{1q}, \sup_q R_{2q} \leq C$ for some $C < \infty$. Define $p_F(q) = p(q)/p(F)$ for $q \in F$. Then $(R_1^F, R_2^F) = \left(\sum_{q \in F} p_F(q) R_{1q}, \sum_{q \in F} p_F(q) R_{2q} \right)$ is a finite convex combination of points of $\mathcal{R}$, hence belongs to the capacity region by the finite time-sharing construction, and

$$\begin{aligned} |R_1 - R_1^F| &\leq \sum_{q \notin F} p(q) R_{1q} + \left| 1 - \frac{1}{p(F)} \right| \sum_{q \in F} p(q) R_{1q} \\ &\leq C(1 - p(F)) + \frac{1 - p(F)}{p(F)} C p(F) = 2C(1 - p(F)) \leq 2C\epsilon. \end{aligned}$$

The same bound holds for $|R_2 - R_2^F|$. Letting $\epsilon \rightarrow 0$ yields $(R_1^F, R_2^F) \rightarrow (R_1, R_2)$. Since the *capacity region* is closed, it follows that (R_1, R_2) belongs to the capacity region.

For *arbitrary measurable* $\mathcal{Q}$, the finite-subset argument does not apply: a nonatomic law gives zero mass to every finite set. Instead, regard $r(Q) \triangleq (R_1(Q), R_2(Q))$ as a bounded random vector taking values in the compact rate region. Partition its bounded range into finitely many small cells and choose a representative from each cell of positive mass; the weighted representatives approximate $E r(Q)$ uniformly, and their limits lie in the closed capacity region. Equivalently, barycenters of measures supported on the compact set $\mathcal{R}$ lie in $\text{ConvHull}(\mathcal{R})$, to which the finite-dimensional cardinality argument applies directly. $\circ$

3.6 Extension to multiuser multiple access channel




Figure 3.2: A k -user multiple access channel

Theorem 3.8. The capacity region of a multiple access channel $W_{Y|X_1, \dots, X_k}$ is the set of *nonnegative* rate tuples $(R_1, \dots, R_k) \in \mathbb{R}_+^k$ such that

$$\sum_{\ell \in S} R_\ell \leq I(X_S; Y | X_{S^c}, Q)$$

for all $S \subseteq [1 : k]$, for some distribution $p(q)p(x_1|q) \dots p(x_k|q)$ with $|Q| \leq k$. (The nonnegativity restriction is operationally necessary and makes the region compact; with upper constraints alone the written sets would contain arbitrary negative rates.) $\circ$

The proof is similar to the two-user case.

Converse

Proof of converse. Suppose a sequence of $(n, R_1, \dots, R_k)$ codes is given. Each choice of n induces the random variables

$$(M_1, \dots, M_k, X_1^n, \dots, X_k^n, Y^n, \hat{M}_1, \dots, \hat{M}_k)$$

distributed according to

$$p(m_1) \dots p(m_k) p(x_1^n | m_1) \dots p(x_k^n | m_k) \left(\prod_{i=1}^n W(y_i | x_{1i}, \dots, x_{ki}) \right) p(\hat{m}_1, \dots, \hat{m}_k | y^n).$$

with $(M_1, \dots, M_k)$ uniformly distributed on $\mathcal{M}_1 \times \dots \times \mathcal{M}_k \triangleq [1 : 2^{nR_1}] \times \dots \times [1 : 2^{nR_k}]$. Let $S \subseteq [1 : k]$. Consider

$$\begin{aligned} n \sum_{\ell \in S} R_\ell &= \log \left| \prod_{\ell \in S} \mathcal{M}_\ell \right| \\ &= H(M_S) && (M_i \text{'s are independent}) \\ &= H(M_S | M_{S^c}) && (M_S, M_{S^c} \text{ are independent}) \\ &= I(M_S; \hat{M}_S | M_{S^c}) + H(M_S | \hat{M}_S, M_{S^c}) && (\text{chain rule}) \\ &\leq I(M_S; \hat{M}_S | M_{S^c}) + H(M_S | \hat{M}_S). \end{aligned}$$

An application of Fano's inequality gives $H(M_S | \hat{M}_S) = o(n)$ and hence

$$n \sum_{\ell \in S} R_\ell \leq I(M_S; \hat{M}_S | M_{S^c}) + o(n).$$

Now,

$$\begin{aligned} I(M_S; \hat{M}_S | M_{S^c}) &= I(M_S; \hat{M}_S, M_{S^c}) && (I(M_S; M_{S^c}) = 0 \text{ since } M_S, M_{S^c} \text{ are independent}) \\ &\leq I(X_S^n; \hat{M}_S, M_{S^c}) && (M_S \rightarrow X_S^n \rightarrow (\hat{M}_S, M_{S^c}) \text{ is Markov}) \\ &\leq I(X_S^n; Y^n, X_{S^c}^n) && (X_S^n \rightarrow (Y^n, X_{S^c}^n) \rightarrow (\hat{M}_S, M_{S^c}) \text{ is Markov}) \end{aligned}$$

$$\begin{aligned}
&= I(X_S^n; Y^n, X_{S^c}^n) - I(X_S^n; X_{S^c}^n) \quad (X_S^n, X_{S^c}^n \text{ are independent}) \\
&= I(X_S^n; Y^n | X_{S^c}^n). \quad (\text{chain rule})
\end{aligned}$$

Also,

$$\begin{aligned}
&I(X_S^n; Y^n | X_{S^c}^n) \\
&= \sum_{i=1}^n I(X_S^n; Y_i | X_{S^c}^n, Y^{i-1}) \quad (\text{chain rule}) \\
&\leq \sum_{i=1}^n \left(I(X_S^n; Y_i | X_{S^c}^n, Y^{i-1}) + I(X_{S^c}^{n \setminus i}, Y^{i-1}; Y_i | X_{S^c}^i) \right) \quad (\text{sub-additivity step}) \\
&= \sum_{i=1}^n I(X_S^n, X_{S^c}^{n \setminus i}, Y^{i-1}; Y_i | X_{S^c}^i) \quad (\text{chain rule}) \\
&= \sum_{i=1}^n I(X_{S_i}; Y_i | X_{S^c}^i) \quad (\text{sub-additivity, §1.13}) \\
&= nI(X_{SQ}; Y_Q | X_{S^cQ}, Q),
\end{aligned}$$

where $X_{S^c}^i$ denotes the i -th symbol of the vector $X_{S^c}^n$; the first inequality adds the non-negative term $I(X_{S^c}^{n \setminus i}, Y^{i-1}; Y_i | X_{S^c}^i)$, and the second-to-last equality uses the memoryless Markov chain $(X_1^{n \setminus i}, \dots, X_k^{n \setminus i}, Y^{i-1}) \rightarrow (X_{1i}, \dots, X_{ki}) \rightarrow Y_i$. Finally, Q is a random variable uniformly distributed on $[1 : n]$, independent of all messages, of the full block sequences $X_1^n, \dots, X_k^n, Y^n$, and of the estimates $\hat{M}_1, \dots, \hat{M}_k$ (the indexed variables X_{SQ}, Y_Q themselves depend on Q). Therefore, we have

$$\sum_{\ell \in S} R_\ell \leq I(X_{SQ}; Y_Q | X_{S^cQ}, Q) + o(1)$$

for all $S \subseteq [1 : k]$. The required distribution is then obtained by the similar argument as in the two-user case. $\blacksquare$

Remark 3.2. Converse of this form is a *product-input* cut-set-form bound: the separate no-feedback encoders force the inputs to be independent. The usual network cut-set outer bound permits correlated channel inputs and need not be tight for a MAC. For example, in $Y = X_1 + X_2 + Z$ with unit powers, let the jointly Gaussian inputs have correlation $\rho = 1/4$; the correlated-input cut-set quantities are

$$I(X_1; Y | X_2) = I(X_2; Y | X_1) = \frac{1}{2} \log_2(31/16), \quad I(X_1, X_2; Y) = \frac{1}{2} \log_2(7/2).$$

The symmetric point $R_1 = R_2 = \frac{1}{4} \log_2(7/2)$ satisfies the individual bounds because $7/2 \leq (31/16)^2$, but its sum exceeds the actual independent-encoder sum capacity $\frac{1}{2} \log_2 3$. Thus unrestricted cut-set tightness should not be asserted; the product-input restriction forced by the actual code is essential. $\circ$

Achievability

Proof of achievability. It is enough to show the case where Q is constant. For the achievability when Q is non-trivial follows from a time-sharing argument. We will show the required achievability by random coding and joint typicality decoding.

Fix distributions $p(x_1), \dots, p(x_k)$ on $\mathcal{X}_1, \dots, \mathcal{X}_k$, respectively. This induces a distribution on $(X_1, \dots, X_k, Y)$ by

$$p(x_1, \dots, x_k, y) = p(x_1) \dots p(x_k) W(y|x_1, \dots, x_k).$$

Generate the random codebook $X_{\ell i}(m_\ell)$ ($\ell \in [1 : k]$, $m_\ell \in \mathcal{M}_\ell$, $i \in [1 : n]$) according to

$$\prod_{\ell=1}^k \prod_{m_\ell \in \mathcal{M}_\ell} \prod_{i=1}^n p_{X_\ell}(x_{\ell i}(m_\ell)).$$

Encoder ℓ maps any message $m_\ell \in \mathcal{M}_\ell$ to the codeword $X_\ell^n(m_\ell)$. The decoder maps any received codeword y^n to the unique tuple $(\hat{m}_1, \dots, \hat{m}_k)$ such that $(X_1^n(\hat{m}_1), \dots, X_k^n(\hat{m}_k), y^n) \in \mathcal{T}_\epsilon^{(n)}(X_1, \dots, X_k, Y)$ if such tuple exists, or some predefined default message tuple $(\hat{m}_1, \dots, \hat{m}_k) = (1, \dots, 1)$ otherwise.

The error event $(\hat{M}_1, \dots, \hat{M}_k) \neq (M_1, \dots, M_k)$ occurs only if at least one of the following 2^k events have occurred (for zero-rate users use singleton message sets, and omit all error classes that require an incorrect message for such a user, since no wrong message exists):

- $\mathcal{E}_\emptyset: (X_1^n(M_1), \dots, X_k^n(M_k), Y^n) \notin \mathcal{T}_\epsilon^{(n)}(X_1, \dots, X_k, Y)$.
- $\mathcal{E}_S$ ($\emptyset \neq S \subseteq [1 : k]$): $(X_1^n(\hat{m}_1), \dots, X_k^n(\hat{m}_k), Y^n) \in \mathcal{T}_\epsilon^{(n)}(X_1, \dots, X_k, Y)$ for some $\hat{m}_\ell \neq M_\ell$ ($\ell \in S$) and $\hat{m}_\ell = M_\ell$ ($\ell \notin S$).

Note that, by symmetry, each of $\Pr(\mathcal{E}_S | M_1 = m_1, \dots, M_k = m_k)$ ($S \subseteq [1 : k]$) does not depend on the choice of $(m_1, \dots, m_k)$. Also note that $(X_{1i}(\hat{m}_1), \dots, X_{ki}(\hat{m}_k), Y_i)$ conditioned on $(M_1, \dots, M_k) = (m_1, \dots, m_k)$ is independent across $i \in [1 : n]$ and has joint distribution $p(x_1) \dots p(x_k) p(y|x_{S^c}) = p(x_S) p(x_{S^c}, y)$ where $S = \{\ell \in [1 : k] : \hat{m}_\ell \neq m_\ell\}$.

An application of the law of large numbers (Proposition 2.6) gives

$$\begin{aligned} \Pr(\mathcal{E}_\emptyset) &= \Pr(\mathcal{E}_\emptyset | M_1 = m_1, \dots, M_k = m_k) \\ &\rightarrow 0 \end{aligned}$$

as $n \rightarrow \infty$. Also, for $\emptyset \neq S \subseteq [1 : k]$,

$$\begin{aligned} \Pr(\mathcal{E}_S) &= \Pr(\mathcal{E}_S | M_1 = m_1, \dots, M_k = m_k) \\ &\leq \left| \prod_{\ell \in S} \mathcal{M}_\ell \right| \cdot 2^{-n(D(p(x_1, \dots, x_k, y) \| p(x_S) p(x_{S^c}, y)) - O(\epsilon)) + o(n)} \quad (\text{Proposition 2.7}) \\ &= 2^{-n(I(X_S; X_{S^c}, Y) - \sum_{\ell \in S} R_\ell - O(\epsilon)) + o(n)} \\ &= 2^{-n(I(X_S; Y | X_{S^c}) + I(X_S; X_{S^c}) - \sum_{\ell \in S} R_\ell - O(\epsilon)) + o(n)} \quad (\text{chain rule}) \end{aligned}$$

$$\begin{aligned}
&= 2^{-n(I(X_S; Y | X_{S^c}) - \sum_{\ell \in S} R_\ell - O(\epsilon)) + o(n)} \\
&\rightarrow 0
\end{aligned}
\quad (X_S, X_{S^c} \text{ are independent})$$

as $n \rightarrow \infty$, since $\sum_{\ell \in S} R_\ell < I(X_S; Y | X_{S^c}) - O(\epsilon)$ for small enough ϵ ; a single small ϵ serves every subset S at once, obtained as in the two-user case by taking the maximum of the finitely many subset-dependent packing penalties. Therefore,

$$\begin{aligned}
\Pr(\hat{M} \neq M) &\leq \Pr\left(\bigcup_{S \subseteq [1:k]} \mathcal{E}_S\right) \\
&\leq \sum_{S \subseteq [1:k]} \Pr(\mathcal{E}_S) \\
&\rightarrow 0
\end{aligned}$$

as $n \rightarrow \infty$. Finally, the sequence of codes achieving rate $(R_1, \dots, R_k)$ is picked as follows: For each n , choose $x_\ell^n(m_\ell)$ ($\ell \in [1:k]$, $m_\ell \in \mathcal{M}_\ell$) such that

$$\Pr(\hat{M} \neq M | X_\ell^n(m_\ell) = x_\ell^n(m_\ell), \forall \ell \in [1:k], \forall m_\ell \in \mathcal{M}_\ell) \leq \Pr(\hat{M} \neq M).$$

■

Cardinality bound:

The same argument as in the two-sender case applies here, with one correction to the definition. Let $\mathcal{C}$ denote the capacity region, and let $\mathcal{R}$ be the set of all nonnegative rate tuples $(R_1, \dots, R_k) \in \mathbb{R}_+^k$ satisfying

$$\sum_{\ell \in S} R_\ell \leq I(X_S; Y | X_{S^c})$$

for every nonempty $S \subseteq [1:k]$, under a *product* input distribution $p(x_1) \cdots p(x_k)$. (Do not use a mixture law $p(q) \prod_j p(x_j | q)$ and then omit Q from the information bounds: such a mixture can make the users marginally correlated, contrary to the no- Q independent-input region being convexified.)

Convex-hull identity via polymatroids. For a product input law put $f(S) = I(X_S; Y | X_{S^c})$ and $f(\emptyset) = 0$. This is a nonnegative, monotone, submodular rank function: independence of the inputs gives, for example,

$$f(S) + f(T) - f(S \cup T) - f(S \cap T) = I(X_{S \setminus T}; X_{T \setminus S} | Y, X_{(S \cup T)^c}) \geq 0,$$

and the remaining submodularity identities follow by the same computation. Its MAC polytope is the polymatroid $P(f) = \{r \geq 0 : r(S) \leq f(S)\}$. For weights $w_{\pi(1)} \geq \dots \geq w_{\pi(k)} \geq 0$, let $S_j = \{\pi(1), \dots, \pi(j)\}$ and $w_{\pi(k+1)} = 0$; the greedy formula is

$$\max_{r \in P(f)} w \cdot r = \sum_{j=1}^k (w_{\pi(j)} - w_{\pi(j+1)}) f(S_j).$$

To see the greedy formula without invoking a citation, take $S_0 = \emptyset$ and let $v_{\pi(j)} = f(S_j) - f(S_{j-1})$. Monotonicity gives $v \geq 0$, and submodularity gives, for any $T \subseteq [1:k]$,

$$\sum_{j:\pi(j) \in T} [f(S_j) - f(S_{j-1})] \leq \sum_{j:\pi(j) \in T} [f(T \cap S_j) - f(T \cap S_{j-1})] = f(T),$$

so $v \in P(f)$. For every feasible $r \in P(f)$ the telescoping identity $w \cdot r = \sum_{j=1}^k (w_{\pi(j)} - w_{\pi(j+1)})r(S_j)$ is bounded above by $\sum_{j=1}^k (w_{\pi(j)} - w_{\pi(j+1)})f(S_j)$ because the weights are non-negative and $r(S_j) \leq f(S_j)$, and v attains this bound since $v(S_j) = f(S_j)$. For negative weights set the associated coordinates to zero and apply the same formula to the positively weighted subset. Thus every support function is linear in the rank function for a fixed weight order. For an auxiliary law, $\tilde{f} = \sum_q p(q)f_q$, and the support functions imply $P(\tilde{f}) = \sum_q p(q)P(f_q)$, where the right side is the weighted Minkowski sum of polytopes. This proves that every point of the conditional MAC region is an average of no- Q region points; conversely such averages are realized by introducing the auxiliary. Therefore $C = \text{ConvHull}(\mathcal{R})$.

Since the corrected $\mathcal{R}$ is star-shaped (every polytope in the union contains the origin and the segment from any of its points to the origin) and connected in $\mathbb{R}^k$, the Fenchel–Bunt extension of Carathéodory’s theorem implies that every point of C can be expressed as a convex combination of at most k points in $\mathcal{R}$. Therefore, restricting the time-sharing random variable to satisfy $|Q| \leq k$ does not change the capacity region.

3.7 Capacity of the Gaussian multiple access channel

In a wireless setting, the channel is often modeled as being corrupted by additive white Gaussian noise. Consider the Gaussian multiple access channel

$$Y = AX_1 + BX_2 + Z,$$

where $X_1, X_2, Y, Z \in \mathbb{R}^k$, A and B are fixed channel gain matrices, and $Z \sim \mathcal{N}(0, I_k)$ is a standard Gaussian random vector independent of (X_1, X_2) .

Unlike the discrete setting, the input alphabets are continuous. Without a power constraint, the capacity region would be unbounded, *provided at least one channel gain is nonzero*: by allowing codewords of arbitrarily large energy one may increase the distance between codewords, making the additive Gaussian noise negligible and permitting arbitrarily large rates. (If $A = B = 0$ the output contains no message information, and the capacity region remains the origin even without a power constraint.)

We therefore impose the average power constraints

$$\begin{aligned} \mathbb{E}[\|X_1\|^2] &\leq kP_1, \\ \mathbb{E}[\|X_2\|^2] &\leq kP_2. \end{aligned}$$

Similar to the discrete MAC, the capacity region of the Gaussian MAC is the set of rate pairs (R_1, R_2) satisfying

$$\begin{aligned} R_1 &\leq I(\mathbf{X}_1; \mathbf{Y} | \mathbf{X}_2, Q), \\ R_2 &\leq I(\mathbf{X}_2; \mathbf{Y} | \mathbf{X}_1, Q), \\ R_1 + R_2 &\leq I(\mathbf{X}_1, \mathbf{X}_2; \mathbf{Y} | Q), \end{aligned}$$

for some distribution of the form $p(q)p(\mathbf{x}_1|q)p(\mathbf{x}_2|q)$ satisfying the power constraints.

For the converse, we impose the *per-codeword block-average* power constraint

$$\frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_{j,i}(m_j)\|^2 \leq kP_j \quad \text{for every } m_j \in [1 : 2^{nR_j}], \quad j = 1, 2,$$

and show that any sequence of (n, R_1, R_2) codes whose probability of error tends to zero must have its rate pair (R_1, R_2) in the above region. (An average-over-messages version is a different operational formulation; if it is intended, the same entropy bounds hold under the resulting average constraint.) The proof is essentially identical to the discrete MAC converse: Fano's inequality and the sub-additivity step (§1.13) give the same three inequalities, which are then bounded by the entropy computations below. Note that the block-average bound does *not* imply any per-coordinate bound (a codeword may put all its energy in a single symbol); only the time-average conclusion follows: with $Q \sim \text{unif}[1 : n]$ independent of the block variables,

$$\mathbb{E}[\|\mathbf{X}_{j,Q}\|^2] = \frac{1}{n} \sum_{q=1}^n \mathbb{E}[\|\mathbf{X}_{j,q}\|^2] \leq kP_j.$$

For the achievability, one uses Gaussian codebooks: by the Gaussian-optimality lemma of Chapter 1 (Lemma 1.59), Gaussian inputs maximise every mutual information term that appears, and the SIC vertices, together with time sharing between decoding orders and downward rate reduction, realize the resulting region in the capacity sense (see §3.8 below); this replaces the discrete joint-typicality decoding, which does not apply verbatim to continuous alphabets.

We now simplify this characterization by bounding the three information quantities with *averaged conditional covariances*. The time-shared input law is of the form $p(q)p(\mathbf{x}_1|q)p(\mathbf{x}_2|q)$, so $\mathbf{X}_1 \perp\!\!\!\perp \mathbf{X}_2 \mid Q$ (conditional independence; note this need not give unconditional independence — for example, Q Gaussian with $\mathbf{X}_1 = \mathbf{X}_2 = Q$ is conditionally independent given Q but not marginally independent — which is why the sum bound must be derived from the conditional computation rather than copied from the individual unconditional one).

Proposition 3.9. (Gaussian domination of a time-shared input law) For $\mathbf{Y} = A\mathbf{X}_1 + B\mathbf{X}_2 + \mathbf{Z}$, let $\mathbf{Z} \sim N(0, I_k)$ be independent of $(Q, \mathbf{X}_1, \mathbf{X}_2)$, let $\mathbf{X}_1 \perp\!\!\!\perp \mathbf{X}_2 \mid Q$, and suppose $\mathbb{E} \|\mathbf{X}_j\|^2 \leq kP_j$. Define $K_j(q) = \text{Cov}(\mathbf{X}_j \mid Q = q)$ and $\bar{K}_j = \mathbb{E} K_j(Q)$. Then $\bar{K}_j \succeq 0$, $\text{tr } \bar{K}_j \leq kP_j$, and

$$\begin{aligned} I(\mathbf{X}_1; \mathbf{Y} \mid \mathbf{X}_2, Q) &\leq \frac{1}{2} \log_2 \det(I + A\bar{K}_1A^T), \\ I(\mathbf{X}_2; \mathbf{Y} \mid \mathbf{X}_1, Q) &\leq \frac{1}{2} \log_2 \det(I + B\bar{K}_2B^T), \end{aligned}$$

$$I(\mathbf{X}_1, \mathbf{X}_2; \mathbf{Y} \mid Q) \leq \frac{1}{2} \log_2 \det(I + A\bar{K}_1 A^T + B\bar{K}_2 B^T).$$

All three bounds are simultaneously attained by independent centered Gaussians with covariances $\bar{K}_1, \bar{K}_2$ and constant Q . $\circ$

Proof. Conditional covariance has trace at most the conditional second moment, so averaging gives $\text{tr } \bar{K}_j \leq kP_j$. For each q , conditional independence makes the output covariance $I + AK_1(q)A^T + BK_2(q)B^T$, and Gaussian entropy maximality plus subtraction of $h(\mathbf{Z})$ give

$$I(\mathbf{X}_1, \mathbf{X}_2; \mathbf{Y} \mid Q) \leq \frac{1}{2} \mathbb{E} \log_2 \det[I + AK_1(Q)A^T + BK_2(Q)B^T].$$

Log-determinant concavity (Jensen) gives the sum bound with $\bar{K}_1, \bar{K}_2$. Conditioning also on the other user's input removes that input by translation and gives the two individual bounds by the same argument. The Gaussian noise ensures every output covariance is positive definite even when input covariances are singular. Independent centered Gaussian candidates give equality in all bounds. $\blacksquare$

We conclude that the capacity region of the Gaussian MAC is the set of nonnegative rate pairs (R_1, R_2) satisfying

$$\begin{aligned} R_1 &\leq \frac{1}{2} \log |I_k + AK_1 A^T|, \\ R_2 &\leq \frac{1}{2} \log |I_k + BK_2 B^T|, \\ R_1 + R_2 &\leq \frac{1}{2} \log |I_k + AK_1 A^T + BK_2 B^T|, \end{aligned}$$

for some $K_1 \succeq 0, K_2 \succeq 0$ with $\text{tr}(K_1) \leq kP_1$ and $\text{tr}(K_2) \leq kP_2$.

Achievability and power compliance. Gaussian entropy maximality is an upper-bound/extremality argument, not a coding achievability proof by itself; the achievability is the continuous cost-constrained MAC coding theorem (equivalently an information-density random-coding argument) with a power slack. For a direct sketch, use independent Gaussian codebooks with covariances $(1 - \tau)K_j$, so each active power constraint has slack. A randomly drawn codeword violates its block constraint with probability tending to zero by the LLN, so the expected fraction of violating words in each codebook tends to zero. Jointly with the ensemble error bound, choose codebooks for which the error and both violating fractions tend to zero. Discard the violating words separately from each codebook: the retained message sets remain Cartesian, rates change by $o(1)$, and the retained average error is at most the original average error divided by the product of the retained fractions, hence still tends to zero, while every retained word satisfies the power bound. Finally send $\tau \downarrow 0$ and use closure. (For a desired exact interior rate, generate at a slightly larger interior rate before this restriction.)

Finally, the time-sharing random variable Q is unnecessary for the Gaussian MAC. Indeed, the maps $K_j \mapsto \log |I + AK_j A^T|$ and $(K_1, K_2) \mapsto \log |I + AK_1 A^T + BK_2 B^T|$ are concave (log-determinant concavity, Lemma 1.53 of Chapter 1), and the projection onto rate coordinates of the concavity-defined joint (r, K_1, K_2) feasible set is convex, so Gaussian inputs already achieve the whole (convex) region without time-sharing.

3.8 Simple implementation on the AWGN MAC

Consider the scalar Gaussian MAC, $Y = aX_1 + bX_2 + Z$, where $Z \sim \mathcal{N}(0, 1)$, and suppose the transmitters satisfy the power constraints $\mathbb{E}[X_j^2] \leq P_j$ for $j = 1, 2$ (the scalar case of the model above, with $k = 1$, $A = a$, $B = b$). The capacity region becomes

$$\begin{aligned} R_1 &\leq \frac{1}{2} \log(1 + a^2 P_1), \\ R_2 &\leq \frac{1}{2} \log(1 + b^2 P_2), \\ R_1 + R_2 &\leq \frac{1}{2} \log(1 + a^2 P_1 + b^2 P_2). \end{aligned} \tag{3.4}$$

The corner points of this region can be achieved by *successive interference cancellation* (SIC). For example, the receiver may first decode X_1 while treating $bX_2 + Z$ as effective noise, at rate

$$R_1 \leq I(X_1; Y) = \frac{1}{2} \log \left(1 + \frac{a^2 P_1}{1 + b^2 P_2} \right).$$

After recovering X_1 , the receiver subtracts aX_1 from the received signal and then decodes X_2 , at rate

$$R_2 \leq I(X_2; bX_2 + Z) = \frac{1}{2} \log(1 + b^2 P_2).$$

Therefore we achieve the corner point

$$(R_1, R_2) = \left(\frac{1}{2} \log \left(1 + \frac{a^2 P_1}{1 + b^2 P_2} \right), \frac{1}{2} \log(1 + b^2 P_2) \right).$$

Observe that

$$\begin{aligned} R_1 + R_2 &= \frac{1}{2} \log \left(1 + \frac{a^2 P_1}{1 + b^2 P_2} \right) + \frac{1}{2} \log(1 + b^2 P_2) \\ &= \frac{1}{2} \log(1 + a^2 P_1 + b^2 P_2), \end{aligned}$$

which is precisely the maximum achievable sum rate. The other corner point is obtained by reversing the decoding order.

One may also consider a time-division scheme. Suppose transmitter 1 uses the channel for an α -fraction of the time with power P_1/α , while transmitter 2 uses the remaining $1 - \alpha$ fraction with power $P_2/(1 - \alpha)$. The resulting rates are

$$\begin{aligned} R_1 &= \frac{\alpha}{2} \log \left(1 + \frac{a^2 P_1}{\alpha} \right), \\ R_2 &= \frac{1 - \alpha}{2} \log \left(1 + \frac{b^2 P_2}{1 - \alpha} \right). \end{aligned}$$

As α varies, these rate pairs trace a curve connecting the two single-user points

$$\left(\frac{1}{2} \log(1 + a^2 P_1), 0\right) \quad \text{and} \quad \left(0, \frac{1}{2} \log(1 + b^2 P_2)\right).$$

For $u \triangleq a^2 P_1 > 0$ and $v \triangleq b^2 P_2 > 0$, the TDMA sum rate is $\frac{1}{2} [\alpha \log(1 + u/\alpha) + (1 - \alpha) \log(1 + v/(1 - \alpha))] \leq \frac{1}{2} \log(1 + u + v)$, with equality precisely at $\alpha = u/(u + v)$ by strict concavity of log; the endpoints $\alpha = 0, 1$ are treated by limits (single-user rates) and the case $u + v = 0$ is trivial. Note that this time-division scheme generally achieves only a subset of the Gaussian MAC capacity region: it meets the sum-rate boundary at a single point (the optimal α). By contrast, the two SIC corner points are the two *endpoints* of the sum-rate face; time-sharing between them, together with downward rate reduction, recovers the entire pentagon (the two corner points alone do not attain the whole face).

3.9 Weighted sum-rate of the Gaussian MAC

For a weight $\lambda \geq 1$, the λ -weighted sum-rate is the support function of the capacity region in the direction $(\lambda, 1)$:

$$\max_{(R_1, R_2) \in C} \lambda R_1 + R_2.$$

For $\lambda > 1$, an optimizer is the corner where the individual constraint on R_1 and the sum-rate constraint are both tight (the corner obtained by decoding user 1 last); at $\lambda = 1$ the entire maximizing dominant face is optimal, and the indicated corner may be *chosen* as an optimizer. In both cases

$$\max_{\substack{K_1 \geq 0, K_2 \geq 0 \\ \text{tr}(K_1) \leq k P_1 \\ \text{tr}(K_2) \leq k P_2}} \frac{\lambda - 1}{2} \log |I + A K_1 A^T| + \frac{1}{2} \log |I + A K_1 A^T + B K_2 B^T|.$$

In the scalar case ($k = 1, A = a, B = b$) of (3.4), both users use full power and the optimum is explicit:

$$\max_{(R_1, R_2) \in C} \lambda R_1 + R_2 = \frac{\lambda - 1}{2} \log(1 + a^2 P_1) + \frac{1}{2} \log(1 + a^2 P_1 + b^2 P_2),$$

while for $0 < \lambda < 1$ the symmetric corner (user 1 decoded first) is optimal, giving

$$\max_{(R_1, R_2) \in C} \lambda R_1 + R_2 = \frac{\lambda}{2} \log \left(1 + \frac{a^2 P_1}{1 + b^2 P_2} \right) + \frac{1}{2} \log(1 + b^2 P_2).$$

For general matrix channels this chapter does not give a general closed-form solution: the objective is concave in (K_1, K_2) (log-determinant concavity, Lemma 1.53 of Chapter 1), so the weighted sum-rate is obtained by convex optimisation over a convex compact covariance set,

but the two covariance matrices are coupled through the sum $AK_1A^T + BK_2B^T$, and the trace constraints act on K_1, K_2 rather than on the transformed matrices $\tilde{A} = AK_1A^T$, $\tilde{B} = BK_2B^T$ (they pass through only when A, B are orthogonal/unitary). The eigenvalue bound of §3.10 below controls the determinant of that sum; opposite-order simultaneous diagonalisation makes the fixed-spectrum determinant bound sharp, but it alone does not remove the anisotropic gain and covariance-cost constraints from the weighted optimisation.

3.10 Bounds for the determinant of the sum of Hermitian matrices

Fiedler (1971) gave a bound for the determinant of the sum of Hermitian matrices in terms of their eigenvalues. We state the upper bound for two positive semi-definite matrices only.

Theorem 3.10. (Fiedler (1971)) Let $A, B \succeq 0$ be Hermitian, with eigenvalues in nondecreasing order $\lambda_1(A) \leq \dots \leq \lambda_n(A)$ and $\lambda_1(B) \leq \dots \leq \lambda_n(B)$. Then

$$\det(A + B) \leq \prod_{i=1}^n (\lambda_i(A) + \lambda_{n+1-i}(B)).$$

○

The bound is tight in terms of the eigenvalues of A and B (it is attained when A and B are simultaneously diagonalisable with their eigenvalues in opposite order).

Remark 3.3. The companion *lower* bound is the Minkowski determinant inequality of Chapter 1, $\det(A + B)^{1/n} \geq \det(A)^{1/n} + \det(B)^{1/n}$, which is proved there via the entropy power inequality applied to Gaussians. ○

Proof. Without loss of generality, assume A has distinct eigenvalues and that $A, B \succ 0$; the general positive semi-definite case follows by the limiting argument $A + \delta I, B + \delta I$ with $\delta \downarrow 0$ (and the distinctness of the eigenvalues of A is obtained by a generic small perturbation). Write $A = Q_A \Lambda_A Q_A^H$ and $B = Q_B \Lambda_B Q_B^H$, where Q_A, Q_B are unitary and Λ_A, Λ_B are diagonal. Then

$$\det(A + B) = \det(\Lambda_A + Q_A^H Q_B \Lambda_B Q_B^H Q_A).$$

Since the unitary group U_n is compact and the determinant is continuous, the maximum of $Q \mapsto \det(\Lambda_A + Q \Lambda_B Q^H)$ is attained at some $Q_0 \in U_n$:

$$\max_{Q \in U_n} \det(\Lambda_A + Q \Lambda_B Q^H) = \det(\Lambda_A + Q_0 \Lambda_B Q_0^H),$$

and therefore

$$\det(A + B) \leq \det(\Lambda_A + Q_0 \Lambda_B Q_0^H).$$

Set $D \triangleq Q_0 \Lambda_B Q_0^H$ and $C \triangleq \Lambda_A + D$; note $C \succ 0$ since $\Lambda_A \succ 0$ and $D \succeq 0$, so C^{-1} is defined. We claim that C and D commute. Suppose not, and consider the smooth family $Q_\epsilon \triangleq Q_0 \exp(\epsilon S)$ of unitaries, where S is an arbitrary skew-Hermitian matrix. Since Q_0 is a maximiser, the derivative of $\epsilon \mapsto \det(\Lambda_A + Q_\epsilon \Lambda_B Q_\epsilon^H)$ at $\epsilon = 0$ vanishes for every skew-Hermitian S . Expanding

$$Q_\epsilon \Lambda_B Q_\epsilon^H = Q_0 e^{\epsilon S} \Lambda_B e^{-\epsilon S} Q_0^H = D + \epsilon [T, D] + O(\epsilon^2),$$

where $T \triangleq Q_0 S Q_0^H$ is skew-Hermitian and $[T, D] \triangleq TD - DT$, and using the lemma $\det(I + \epsilon M) = 1 + \epsilon \operatorname{tr}(M) + O(\epsilon^2)$ below, we get

$$\begin{aligned} \det(\Lambda_A + Q_\epsilon \Lambda_B Q_\epsilon^H) &= \det(C + \epsilon [T, D] + O(\epsilon^2)) \\ &= \det(C) (1 + \epsilon \operatorname{tr}(C^{-1} [T, D]) + O(\epsilon^2)). \end{aligned}$$

By the cyclicity of the trace,

$$\begin{aligned} \operatorname{tr}(C^{-1} [T, D]) &= \operatorname{tr}(C^{-1} TD) - \operatorname{tr}(C^{-1} DT) \\ &= \operatorname{tr}(DC^{-1} T) - \operatorname{tr}(C^{-1} DT) = \operatorname{tr}((DC^{-1} - C^{-1} D) T). \end{aligned}$$

This must vanish for every skew-Hermitian T (conjugation by Q_0 is a bijection of the skew-Hermitian matrices). Put $L \triangleq DC^{-1} - C^{-1} D$; since C^{-1} and D are Hermitian, $L^H = C^{-1} D - DC^{-1} = -L$, so L is itself skew-Hermitian. The vanishing condition $\operatorname{tr}(LT) = 0$ for every skew-Hermitian T then applies to $T = -L$: $0 = -\operatorname{tr}(L^2) = \operatorname{tr}(L^H L) = \|L\|_F^2$, hence $L = 0$, i.e. $DC^{-1} = C^{-1} D$, so $CD = DC$.

Thus C and D commute and are simultaneously unitarily diagonalisable: there is a unitary V such that $V^H C V$ and $V^H D V$ are diagonal. Because Λ_A has distinct eigenvalues and $C = \Lambda_A + D$, the matrix $V^H \Lambda_A V$ is also diagonal, so V permutes the eigenvalues of Λ_A (up to phases). Hence

$$\det(A + B) \leq \det(C) = \prod_{i=1}^n (\lambda_{\sigma(i)}(A) + \lambda_{\tau(i)}(B))$$

for some permutations σ, τ of $[1 : n]$. Finally, for nondecreasing sequences the product $\prod_i (a_i + b_{\pi(i)})$ is maximised by the *opposite* pairing: if $i < j$ and $\pi(i) < \pi(j)$ then swapping $\pi(i), \pi(j)$ increases the product, since

$$(a_i + b_{\pi(j)})(a_j + b_{\pi(i)}) - (a_i + b_{\pi(i)})(a_j + b_{\pi(j)}) = (a_j - a_i)(b_{\pi(j)} - b_{\pi(i)}) \geq 0.$$

Iterating gives $\prod_i (a_i + b_{\pi(i)}) \leq \prod_i (a_i + b_{n+1-i})$. Applying this with $a_i = \lambda_i(A)$, $b_j = \lambda_j(B)$ yields

$$\det(A + B) \leq \prod_{i=1}^n (\lambda_i(A) + \lambda_{n+1-i}(B)),$$

as claimed. ■

Lemma 3.11. For any square matrix M , $\det(I + \epsilon M) = 1 + \epsilon \operatorname{tr}(M) + O(\epsilon^2)$. ○

Lemma 3.12. (Unitary matrices and skew-Hermitian matrices) A matrix U is unitary if and only if there exists a skew-Hermitian matrix S (i.e. $S^H = -S$) such that $U = \exp(S)$. $\circ$

Proof. Suppose $U = \exp(S)$ with S skew-Hermitian. Then $U^H = \exp(S^H) = \exp(-S)$, so $UU^H = \exp(S) \exp(-S) = I$; hence U is unitary. Conversely, if U is unitary, then $U = Q\Lambda Q^H$ with Q unitary and $\Lambda = \text{diag}(e^{i\theta_1}, \dots, e^{i\theta_n})$. Taking $S = Q \text{diag}(i\theta_1, \dots, i\theta_n) Q^H$ gives a skew-Hermitian matrix with $U = \exp(S)$. $\blacksquare$

3.11 Multi-letter characterization of the capacity region

Let $\mathcal{A}_n$ be the set of all *nonnegative* rate pairs $(R_1, R_2) \in \mathbb{R}_+^2$ satisfying

$$\begin{aligned} nR_1 &\leq I(X_1^n; Y^n), \\ nR_2 &\leq I(X_2^n; Y^n), \end{aligned}$$

for some distribution $p(x_1^n)p(x_2^n)$, where $p(y^n|x_1^n, x_2^n) = \prod_{i=1}^n W(y_i|x_{1i}, x_{2i})$. Note that no sum-rate constraint is listed: as the achievability below shows, the two marginal constraints together imply $nR_1 + nR_2 \leq I(X_1^n, X_2^n; Y^n)$, so the sum-rate constraint is redundant in $\mathcal{A}_n$. (These *marginal-rectangle* regions $\mathcal{A}_n$ are in general different from the one-letter *conditional-pentagon* union $\mathcal{R}$ of (3.2); the two should not be identified.) The following theorem states that the capacity region of a multiple-access channel can be characterized in terms of the regions $\mathcal{A}_n$.

Theorem 3.13. The capacity region of a multiple-access channel $W(y | x_1, x_2)$ is given by the closure of $\bigcup_{n=1}^{\infty} \mathcal{A}_n$. $\circ$

Proof. Converse: Given a sequence of (n, R_1, R_2) codes, consider

$$\begin{aligned} nR_1 &= \log |\mathcal{M}_1| \\ &= H(M_1) \\ &= I(M_1; \hat{M}_1) + H(M_1 | \hat{M}_1) && \text{(chain rule)} \\ &= I(M_1; \hat{M}_1) + o(n) && \text{(Fano's inequality)} \\ &\leq I(X_1^n; \hat{M}_1) + o(n) && (M_1 \rightarrow X_1^n \rightarrow \hat{M}_1 \text{ is Markov}) \\ &\leq I(X_1^n; Y^n) + o(n), && (X_1^n \rightarrow Y^n \rightarrow \hat{M}_1 \text{ is Markov}) \end{aligned}$$

and similarly

$$nR_2 \leq I(X_2^n; Y^n) + o(n).$$

Thus, there exists a sequence of non-negative real numbers $\epsilon_n \rightarrow 0$ such that $((R_1 - \epsilon_n)^+, (R_2 - \epsilon_n)^+) \in \mathcal{A}_n \subseteq \bigcup_{m=1}^{\infty} \mathcal{A}_m$ (the positive part keeps the points nonnegative; $\mathcal{A}_n$ contains the origin). Taking the limit $n \rightarrow \infty$ shows that (R_1, R_2) belongs to the closure of $\bigcup_{n=1}^{\infty} \mathcal{A}_n$.

Achievability: We show that $\mathcal{A}_n \subseteq C$ via the length- n superchannel, applying the finite-alphabet MAC coding theorem already proved in §3.3–§3.5. For independent block inputs $A = X_1^n$ and $B = X_2^n$,

$$I(A; Y^n) + I(B; Y^n) = I(A, B; Y^n) - I(A; B | Y^n) \leq I(A, B; Y^n),$$

and moreover $I(A; Y^n) \leq I(A; Y^n | B)$ (conditioning on B can only help user 1), and symmetrically. Thus each marginal rectangle of $\mathcal{A}_n$ lies inside the MAC information region of the length- n superchannel $(A, B) \mapsto Y^n$. Repeating that superchannel across independent blocks and applying the MAC coding theorem shows $\mathcal{A}_n \subseteq C$ for every n , with closure handling boundary rates and the zero-rate coordinates handled by singleton message sets. Hence $\bigcup_n \mathcal{A}_n \subseteq C$, and since the capacity region is closed, $\bigcup_n \mathcal{A}_n \subseteq C$; together with the converse this gives $C = \overline{\bigcup_n \mathcal{A}_n}$. ■

Remark 3.4. There are two distinct single-letter objects: $\mathcal{R}$, the union of conditional MAC pentagons over product input laws, and $\mathcal{A}_1$, the union of marginal-information rectangles. They need not coincide. The example below proves that the marginal functional $I(X_1^n; Y^n)$ is not temporally sub-additive in general. Therefore it cannot be replaced by a sum of the corresponding one-coordinate marginal informations. The conditional and joint functionals in the MAC converse remain sub-additive. Separately, time sharing gives $C = \text{ConvHull}(\mathcal{R})$. This convexification statement is distinct from the multi-letter marginal characterization $C = \overline{\bigcup_n \mathcal{A}_n}$; neither implies a universal recipe that arbitrary failures of sub-additivity are cured by time sharing. To see where sub-additivity fails, consider $(X_{11}, X_{12}, X_{21}, X_{22}, Y_1, Y_2)$ distributed according to

$$p(x_{11}, x_{12}, x_{21}, x_{22}, y_1, y_2) = p(x_{11}, x_{12})p(x_{21}, x_{22})W(y_1|x_{11}, x_{21})W(y_2|x_{12}, x_{22}).$$

The following explicit example shows the failure is real, not merely an artifact of a failed Markov manipulation. Let X_{11}, X_{12}, U be independent fair bits, take $X_{21} = X_{22} = U$, and use the XOR MAC at both times: $Y_i = X_{1i} \oplus U$. Then

$$I(X_{11}; Y_1) = I(X_{12}; Y_2) = 0,$$

because each output is an independent fair bit independent of the corresponding input, while Y^2 is uniform on four pairs and its entropy given X_1^2 is one bit, so

$$I(X_1^2; Y^2) = 2 - 1 = 1 > 0 = I(X_{11}; Y_1) + I(X_{12}; Y_2).$$

Thus the marginal functional is genuinely not temporally sub-additive. (The point-to-point contrast remains instructive: for a single user, $I(X^n; Y^n)$ is sub-additive, Lemma 2.10 of Chapter 2, because the two-letter form admits the Markov chain $(X_1, Y_1) \rightarrow X_2 \rightarrow Y_2$, which gives

$$\begin{aligned} I(X_1, X_2; Y_1, Y_2) &= I(X_1, X_2; Y_1) + I(X_1, X_2; Y_2|Y_1) && \text{(chain rule)} \\ &\leq I(X_1, X_2; Y_1) + I(X_1, X_2; Y_2|Y_1) + I(Y_1; Y_2) \\ &= I(X_1, X_2; Y_1) + I(X_1, X_2, Y_1; Y_2) && \text{(chain rule)} \end{aligned}$$

$$\begin{aligned}
 &= I(X_1; Y_1) + I(X_1, X_2, Y_1; Y_2) && (X_2 \rightarrow X_1 \rightarrow Y_1 \text{ is Markov}) \\
 &= I(X_1; Y_1) + I(X_2; Y_2). && ((X_1, Y_1) \rightarrow X_2 \rightarrow Y_2 \text{ is Markov})
 \end{aligned}$$

By contrast, the *conditional* functional $I(X_1^n; Y^n | X_2^n)$ and the *joint* functional $I(X_1^n, X_2^n; Y^n)$ are sub-additive — this is exactly the single-letterisation used in the converse of §3.3 — which is why the sum-rate and single-user-with-genie constraints single-letterise while the marginal constraints do not. Note also that even if both marginal block functionals were temporally sub-additive, the time average would generally yield the convex hull of the one-letter marginal region, not necessarily the unconvexified region; the relevant operational convexification statement is $C = \text{ConvHull}(\mathcal{R})$ with the conditional pentagon union, and is distinct from the regularized marginal characterization. $\bigcirc$

4 Broadcast channel

The broadcast channel (BC) was formulated by Cover in his 1972 paper (Cover 1972). At the time wireless communication was not yet a thing, but he modeled it as an interesting question. The broadcast channel models the scenario in which a single transmitter wants to send a message to multiple receivers so that each receiver gets a certain part of the message. In particular, we focus on the case where there are two receivers and the channel is memoryless.




Figure 4.1: A broadcast channel

A two-user memoryless broadcast channel is represented by a stochastic matrix $p_{Y_1, Y_2 | X} : \mathcal{X} \rightarrow \mathcal{Y}_1 \times \mathcal{Y}_2$. The transmitter encodes a message that consists of a common message $M_0 \in \mathcal{M}_0$ and private messages $M_1 \in \mathcal{M}_1, M_2 \in \mathcal{M}_2$ into a codeword $X^n \in \mathcal{X}^n$ which is then passed through the channel. Receiver ℓ ($\ell = 1, 2$) receives the channel output $Y_\ell^n \in \mathcal{Y}_\ell^n$ and decodes it into an estimate $(\hat{M}_{0\ell}, \hat{M}_\ell) \in \mathcal{M}_0 \times \mathcal{M}_\ell$ of the transmitted messages.

An (n, R_0, R_1, R_2) code consists of

- message sets $\mathcal{M}_0 \triangleq [1 : 2^{nR_0}]$, $\mathcal{M}_1 \triangleq [1 : 2^{nR_1}]$, $\mathcal{M}_2 \triangleq [1 : 2^{nR_2}]$,
- an encoder mapping $x^n : (\mathcal{M}_0, \mathcal{M}_1, \mathcal{M}_2) \rightarrow \mathcal{X}^n$,
- decoder mappings $(\hat{m}_{01}, \hat{m}_1) : \mathcal{Y}_1^n \rightarrow \mathcal{M}_0 \times \mathcal{M}_1$ and $(\hat{m}_{02}, \hat{m}_2) : \mathcal{Y}_2^n \rightarrow \mathcal{M}_0 \times \mathcal{M}_2$.

We assume that the message M_ℓ is a random variable independent across ℓ and uniformly distributed over the message set $\mathcal{M}_\ell$. The codeword sent is given by $X^n \triangleq x^n(M_0, M_1, M_2)$. The codewords received Y_1^n, Y_2^n are random variables such that

$$p(y_1^n, y_2^n | x^n, m_0, m_1, m_2) = \prod_{i=1}^n p(y_{1i}, y_{2i} | x_i).$$

The decoded message $(\hat{M}_{0\ell}, \hat{M}_\ell)$ by receiver ℓ ($\ell = 1, 2$) is a random variable on $\mathcal{M}_0 \times \mathcal{M}_\ell$ given by $(\hat{M}_{0\ell}, \hat{M}_\ell) \triangleq (\hat{m}_{0\ell}(Y_\ell^n), \hat{m}_\ell(Y_\ell^n))$.

(R_0, R_1, R_2) is said to be an achievable rate if there exists a sequence of (n, R_0, R_1, R_2) codes whose probability of error

$$\begin{aligned} & \Pr \left((\hat{M}_{01}, \hat{M}_1) \neq (M_0, M_1) \vee (\hat{M}_{02}, \hat{M}_2) \neq (M_0, M_2) \right) \\ &= \Pr \left(\hat{M}_{01} \neq M_0 \vee \hat{M}_1 \neq M_1 \vee \hat{M}_{02} \neq M_0 \vee \hat{M}_2 \neq M_2 \right) \end{aligned}$$

goes to 0 as $n \rightarrow \infty$. We say that such a sequence of codes achieves rate (R_0, R_1, R_2) . The closure of the set of all achievable rates is called the capacity region $\mathcal{C}$ of the channel.

4.1 Properties of the capacity region

Determining the capacity region of a general broadcast channel has been a long-standing open problem. Not even the capacity region in the two-user case is known, let alone the three-or-more-user broadcast channel which is conceptually much harder to solve.

We know the projections $(R_0, 0, R_2)$ and $(R_0, R_1, 0)$ of the capacity region. These settings are called degraded message set. We also know the capacity region for a variety of classes of broadcast channels. However, they are not exhaustive.

Let $\mathcal{C}$ denote the capacity region of a broadcast channel $p(y_1, y_2|x)$. We have the following properties:

- $\mathcal{C}$ is convex. This follows from a time-sharing argument.
- If $(R_0, R_1, R_2) \in \mathcal{C}$, then $(\tilde{R}_0, \tilde{R}_1, \tilde{R}_2) \in \mathcal{C}$ for any $\tilde{R}_0, \tilde{R}_1, \tilde{R}_2 \geq 0$ with $\tilde{R}_0 \leq R_0$, $\tilde{R}_1 \leq R_1$, $\tilde{R}_2 \leq R_2$. This can be done by inserting dummy messages at the transmitter.
- If $(R_0, R_1, R_2) \in \mathcal{C}$, then $(R_0 - \tilde{R}_1 - \tilde{R}_2, R_1 + \tilde{R}_1, R_2 + \tilde{R}_2) \in \mathcal{C}$ for any $\tilde{R}_1, \tilde{R}_2 \geq 0$ with $\tilde{R}_1 + \tilde{R}_2 \leq R_0$. This can be done by moving some parts of the common message to the private messages.
- If $(R_0, R_1, R_2) \in \mathcal{C}$, then $(R_0 + \tilde{R}_0, R_1 - \tilde{R}_0, R_2 - \tilde{R}_0) \in \mathcal{C}$ for any $0 \leq \tilde{R}_0 \leq \min\{R_1, R_2\}$. This can be done by moving some parts of the private messages to the common message.
- The maximum one-sided rates are given by

$$\begin{aligned} \max_{(R_0, R_1, R_2) \in \mathcal{C}} (R_0 + R_1) &= \max_{p(x)} I(X; Y_1), \\ \max_{(R_0, R_1, R_2) \in \mathcal{C}} (R_0 + R_2) &= \max_{p(x)} I(X; Y_2). \end{aligned}$$

This follows from the following reasoning: If $(R_0, R_1, R_2) \in \mathcal{C}$, then $(R_0, R_1, 0) \in \mathcal{C}$, which in turn implies $(0, R_0 + R_1, 0) \in \mathcal{C}$. Hence $\max_{(R_0, R_1, R_2) \in \mathcal{C}} (R_0 + R_1) = \max_{(0, R_1, 0) \in \mathcal{C}} R_1$ which reduces to the maximum rate of the point-to-point channel $p(y_1|x)$. The result for $R_0 + R_2$ follows from a similar argument.

- The maximum common message rate is given by

$$\max_{(R_0, R_1, R_2) \in \mathcal{C}} R_0 = \max_{p(x)} \min \{I(X; Y_1), I(X; Y_2)\}.$$

This follows from the following reasoning: Note that $\max_{(R_0, R_1, R_2) \in \mathcal{C}} R_0 = \max_{(R_0, 0, 0) \in \mathcal{C}} R_0$, since $(R_0, R_1, R_2) \in \mathcal{C}$ implies $(R_0, 0, 0) \in \mathcal{C}$. The achievability follows from random coding and joint typicality decoding, requiring $R_0 < I(X; Y_1)$ and $R_0 < I(X; Y_2)$ for some distribution $p(x)$. The converse follows from a standard application of Fano's inequality similar to that of the point-to-point channel.

- The capacity region of the channel $p(y_1, y_2|x)$ depends only on the marginal channels $p(y_1|x)$ and $p(y_2|x)$, because the two decoders observe only Y_1^n and Y_2^n , respectively: the error events of the two receivers involve only $p(y_1|x)$ and $p(y_2|x)$ separately. In particular, the capacity region of $p(y_1, y_2|x)$ coincides with that of the product channel $p(y_1|x)p(y_2|x)$.

4.2 Time-division inner bound

With $R_0 = R_2 = 0$ or $R_0 = R_1 = 0$, the broadcast channel reduces to the point-to-point channel. Then $(0, R_1, 0)$ and $(0, 0, R_2)$ are achievable for any $R_1 < C_1$ and $R_2 < C_2$, where $C_\ell \triangleq \max_{p(x)} I(X; Y_\ell)$ is the capacity of the point-to-point channel $p(y_\ell|x)$. By time-sharing, any convex combination of these two points is also achievable. This achievable region is called the time-division inner bound. This is mostly the scheme used in current communication systems, e.g. TDMA.




Figure 4.2: Time-division inner bound

The idea of the time-division scheme is very simple: The transmitter only transmits to one of the receivers at a time, ignoring the other. Note that the receivers not only can decode their own messages when it is their turn, but also know when are they able to decode. Thus, some additional information can be encoded in whom is the transmitter sending to. This is the idea

of “randomized time-division”, by which we can convey more information than the vanilla time-division.

4.3 Degraded broadcast channel

When Cover introduced the broadcast channel, he also introduced the degraded broadcast channel. Consider a broadcast channel in which one of the receivers is “stronger” than the other, for example, because it is closer to the transmitter, such that it can always recover the message intended for the “weaker” receiver. Such a setting is called the degraded broadcast channel.

Formally, degradation means that $X \rightarrow Y_1 \rightarrow Y_2$ forms a Markov chain. In such a case, receiver 1 is the stronger receiver, and receiver 2 is the weaker receiver that receives a degraded version Y_2^n of the signal Y_1^n received by receiver 1.




Figure 4.3: A degraded broadcast channel

Definition 4.1. (Degraded broadcast channel) A broadcast channel $p_{Y_1, Y_2 | X} : \mathcal{X} \rightarrow \mathcal{Y}_1 \times \mathcal{Y}_2$ is called (physically) degraded if

$$p_{Y_1, Y_2 | X}(y_1, y_2 | x) = p_{Y_1 | X}(y_1 | x) p_{Y_2 | Y_1}(y_2 | y_1).$$

○

Remark 4.1. The definition above is of a *physically degraded* channel, in which the joint law factors through a fixed channel $p_{Y_2 | Y_1}$. A weaker notion, *stochastically degraded*, requires only the existence of some $p(y_2 | y_1)$ with $p(y_2 | x) = \sum_{y_1} p(y_1 | x) p(y_2 | y_1)$; for the two-receiver setting the two notions give the same capacity region. Throughout, we use the physically degraded definition, which suffices for the notes. ○

For a degraded broadcast channel, the stronger receiver can decode any message that can be decoded by the weaker receiver. Hence, we can assume $R_2 = 0$ without loss of generality, by making the private message intended for the weaker receiver a part of the common message. Thus, we are only interested in the projection $(R_0, R_1, 0)$ of the capacity region.

Proposition 4.2. For a degraded broadcast channel, (R_0, R_1, R_2) is in the capacity region if and only if $(R_0 + R_2, R_1, 0)$ is in the capacity region. ○

Proof. The “if” part is done by moving some parts of the common message to the private message intended for receiver 2.

The “only if” part follows since receiver 1 can decode any message that can be decoded by receiver 2, by locally simulating a version of Y_2^n . Suppose a sequence of (n, R_0, R_1, R_2) codes is given. Let $\Phi : \mathcal{Y}_1^n \rightarrow \mathcal{Y}_2^n$ be a random function such that the probability that $\phi : \mathcal{Y}_1^n \rightarrow \mathcal{Y}_2^n$ is picked is

$$\Pr(\Phi = \phi) = \prod_{y_1^n \in \mathcal{Y}_1^n} \prod_{i=1}^n p(\phi(y_1^n)_i | y_{1i}).$$

Degradation implies that $(\Phi(Y_1^n), M_2)$ and (Y_2^n, M_2) have the same distribution. Thus,

$$\begin{aligned} \Pr(\hat{m}_2(\Phi(Y_1^n)) \neq M_2) &= \Pr(\hat{m}_2(Y_2^n) \neq M_2) \\ &\rightarrow 0 \end{aligned}$$

as $n \rightarrow \infty$. For each n , we can pick a ϕ such that

$$\Pr(\hat{m}_2(\phi(Y_1^n)) \neq M_2) \leq \Pr(\hat{m}_2(\Phi(Y_1^n)) \neq M_2).$$

Now, with the decoder at receiver 1 given by $(\hat{m}_{01}, \hat{m}_1, \hat{m}_2 \circ \phi) : \mathcal{Y}_1^n \rightarrow \mathcal{M}_0 \times \mathcal{M}_1 \times \mathcal{M}_2$, we get a sequence of codes that achieves $(R_0 + R_2, R_1, 0)$. ■

4.4 Superposition coding

Motivated by the degraded broadcast channel, Cover introduced the idea of superposition coding. The idea of superposition coding is to have a nested structure for the codebook: codewords are grouped into well-separated clusters with the “cluster centers”, which are represented using an auxiliary codebook $W^n(m_0)$, carrying the common message m_0 , and codewords $X^n(m_0, m_1)$ within a cluster carry the private message m_1 for one of the receivers.

The idea of superposition coding for the degraded broadcast channel is illustrated in Figure 4.4. The codewords in the same cluster share the same common message M_0 but have different private messages M_1 for the stronger receiver. The codewords in the same cluster can be distinguished by the stronger receiver who knows Y_1^n . The clusters themselves can be distinguished by the weaker receiver who knows Y_2^n , as well as the stronger receiver.

Superposition coding yields an achievable region for the general (not necessarily degraded) broadcast channel. As there is private message for only one receiver, there are two superposition coding regions, one with $R_2 = 0$ (where the private message is for receiver 1) and one with $R_1 = 0$ (where the private message is for receiver 2). It suffices to consider the region with $R_2 = 0$, since the other can be obtained by symmetry.

There are no restrictions on the alphabet $\mathcal{W}$ of the auxiliary codebook $W^n(m_0)$. However, we need a cardinality bound for $\mathcal{W}$ so that the superposition coding region becomes computable.




Figure 4.4: Superposition coding for degraded broadcast channel




Figure 4.5: Codebook generation for superposition coding

Theorem 4.3. (Superposition coding) For a broadcast channel $p_{Y_1, Y_2|X}$, any rate tuple $(R_0, R_1, 0)$ such that

$$\begin{aligned} R_0 &< I(W; Y_2), \\ R_0 + R_1 &< I(X; Y_1), \\ R_1 &< I(X; Y_1|W) \end{aligned}$$

for some distribution $p(w, x)$, is achievable. This region is called the superposition coding region. $\circ$

Proof. We will show the required achievability by random coding and joint typicality decoding.

Fix a distribution $p(w, x)$ on $\mathcal{W} \times \mathcal{X}$. This induces a distribution on (W, X, Y_1, Y_2) by

$$p(w, x, y_1, y_2) = p(w, x) p(y_1, y_2|x).$$

Generate the auxiliary codebook $W_i(m_0)$ and the random codebook $X_i(m_0, m_1)$ ($m_0 \in \mathcal{M}_0$, $m_1 \in \mathcal{M}_1$, $i \in [1 : n]$) according to

$$\left(\prod_{m_0 \in \mathcal{M}_0} \prod_{i=1}^n p_W(w_i(m_0)) \right) \left(\prod_{m_0 \in \mathcal{M}_0} \prod_{m_1 \in \mathcal{M}_1} \prod_{i=1}^n p_{X|W}(x_i(m_0, m_1)|w_i(m_0)) \right).$$

The encoder maps any message pair $(m_0, m_1) \in \mathcal{M}_0 \times \mathcal{M}_1$ to the codeword $X^n(m_0, m_1)$. Decoder 1 maps any received codeword y_1^n to the unique pair $(\hat{m}_{01}, \hat{m}_1)$ with $(W^n(\hat{m}_{01}), X^n(\hat{m}_{01}, \hat{m}_1), y_1^n) \in \mathcal{T}_\epsilon^{(n)}(W, X, Y_1)$, if such a pair exists, and to the default pair $(1, 1)$ otherwise. (The entire error analysis is conditional on the generated codebooks: each decoder is a deterministic function of the *whole* generated cloud/satellite codebook $\mathcal{C}$ and its own output, with kernel $P_{\hat{M}_{01}, \hat{M}_1|Y_1^n, \mathcal{C}}$ at decoder 1 and $P_{\hat{M}_{02}|Y_2^n, \mathcal{C}}$ at decoder 2, exactly as in the random-codebook ensemble laws of Chapters 2 and 3. The conditioning on Y_1^n alone would be legitimate only for a fixed codebook.) Decoder 2 maps any received codeword y_2^n to the unique $\hat{m}_{02}$ with $(W^n(\hat{m}_{02}), y_2^n) \in \mathcal{T}_\epsilon^{(n)}(W, Y_2)$, if such a message exists, and to the default message $\hat{m}_{02} = 1$ otherwise.

The error event $(\hat{M}_{01}, \hat{M}_1) \neq (M_0, M_1) \vee \hat{M}_{02} \neq M_0$ occurs only if at least one of the following five events have occurred:

- $\mathcal{E}_{1'}: (W^n(M_0), X^n(M_0, M_1), Y_1^n) \notin \mathcal{T}_\epsilon^{(n)}(W, X, Y_1).$
- $\mathcal{E}_{2'}: (W^n(M_0), Y_2^n) \notin \mathcal{T}_\epsilon^{(n)}(W, Y_2).$
- $\mathcal{E}_{01}: (W^n(\hat{m}_{01}), X^n(\hat{m}_{01}, \hat{m}_1), Y_1^n) \in \mathcal{T}_\epsilon^{(n)}(W, X, Y_1)$ for some $\hat{m}_{01} \neq M_0$ and $\hat{m}_1$.
- $\mathcal{E}_1: (W^n(M_0), X^n(M_0, \hat{m}_1), Y_1^n) \in \mathcal{T}_\epsilon^{(n)}(W, X, Y_1)$ for some $\hat{m}_1 \neq M_1$.
- $\mathcal{E}_{02}: (W^n(\hat{m}_{02}), Y_2^n) \in \mathcal{T}_\epsilon^{(n)}(W, Y_2)$ for some $\hat{m}_{02} \neq M_0$.

Note that, by symmetry, the conditional probability of each of the above events under $(M_0, M_1) = (m_0, m_1)$ does not depend on the choice of (m_0, m_1) . Also note that $(W_i(\hat{m}_{01}), X_i(\hat{m}_{01}, \hat{m}_1), Y_{1i})$ conditioned on $(M_0, M_1) = (m_0, m_1)$ is independent across $i \in [1 : n]$ and has joint distribution summarized in the following table:

Event	$\hat{m}_{01}$	$\hat{m}_1$	Joint distribution
$\mathcal{E}_{1'}$	$= m_0$	$= m_1$	$p(w, x)p(y_1 x) = p(w, x, y_1)$
$\mathcal{E}_{01}$	$\neq m_0$		$p(w, x)p(y_1)$
$\mathcal{E}_1$	$= m_0$	$\neq m_1$	$p(w, x)p(y_1 w) = p(w)p(x w)p(y_1 w)$

And $(W_i(\hat{m}_{02}), Y_{2i})$ conditioned on $(M_0, M_1) = (m_0, m_1)$ is independent across $i \in [1 : n]$ and has joint distribution summarized in the following table:

Event	$\hat{m}_{02}$	Joint distribution
$\mathcal{E}_{2'}$	$= m_0$	$p(w)p(y_2 w) = p(w, y_2)$
$\mathcal{E}_{02}$	$\neq m_0$	$p(w)p(y_2)$

The law of large numbers gives

$$\begin{aligned} \Pr(\mathcal{E}_{1'}) &= \Pr(\mathcal{E}_{1'} | M_0 = m_0, M_1 = m_1) \\ &\rightarrow 0, \\ \Pr(\mathcal{E}_{2'}) &= \Pr(\mathcal{E}_{2'} | M_0 = m_0, M_1 = m_1) \\ &\rightarrow 0, \end{aligned}$$

as $n \rightarrow \infty$. Also,

$$\begin{aligned} \Pr(\mathcal{E}_{01}) &= \Pr(\mathcal{E}_{01} | M_0 = m_0, M_1 = m_1) \\ &\leq |\mathcal{M}_0 \times \mathcal{M}_1| \cdot 2^{-nD(p(w, x, y_1) \| p(w, x)p(y_1)) + o(n)} \\ &= 2^{-n(I(W; X; Y_1) - R_0 - R_1) + o(n)} \\ &= 2^{-n(I(X; Y_1) - R_0 - R_1) + o(n)} && (W \rightarrow X \rightarrow Y_1 \text{ is Markov}) \\ &\rightarrow 0, && (R_0 + R_1 < I(X; Y_1)) \\ \Pr(\mathcal{E}_1) &= \Pr(\mathcal{E}_1 | M_0 = m_0, M_1 = m_1) \\ &\leq |\mathcal{M}_1| \cdot 2^{-nD(p(w, x, y_1) \| p(w)p(x|w)p(y_1|w)) + o(n)} \\ &= 2^{-n(I(X; Y_1 | W) - R_1) + o(n)} \\ &\rightarrow 0, && (R_1 < I(X; Y_1 | W)) \\ \Pr(\mathcal{E}_{02}) &= \Pr(\mathcal{E}_{02} | M_0 = m_0, M_1 = m_1) \\ &\leq |\mathcal{M}_0| \cdot 2^{-nD(p(w, y_2) \| p(w)p(y_2)) + o(n)} \\ &= 2^{-n(I(W; Y_2) - R_0) + o(n)} \\ &\rightarrow 0, && (R_0 < I(W; Y_2)) \end{aligned}$$

as $n \rightarrow \infty$. Therefore, the union bound gives

$$\Pr((\hat{M}_{01}, \hat{M}_1) \neq (M_0, M_1) \vee \hat{M}_{02} \neq M_0) \rightarrow 0$$

as $n \rightarrow \infty$. ■

Remark 4.2 (Superposition contains time-division). The superposition coding region *contains* the time-division inner bound (in the private-message plane). Indeed, the corner point $(C_1, 0)$ is achieved by superposition coding with $W = \emptyset$ (so that $R_1 < I(X; Y_1)$ and $R_2 = 0$), the corner $(0, C_2)$ is achieved symmetrically, and the superposition region is convex. Hence the convex hull of these two corners — exactly the time-division region — lies inside it, and the containment can be strict. $\circ$

Remark 4.3 (Two facts about superposition coding). 1. Superposition coding achieves the capacity region of degraded broadcast channels [Gallager \(1974\)](#); [Bergmans \(1973\)](#);
2. Superposition coding gives the optimal $(R_0, R_1, 0)$ for *any* broadcast channel (Körner–Marton ’77). $\circ$

It turns out that the superposition coding region is optimal if the broadcast channel is degraded.

Theorem 4.4. (Capacity region of degraded broadcast channel) The capacity region of a degraded broadcast channel $p_{Y_1|X}p_{Y_2|Y_1}$ is the set of rate pairs (R_0, R_1) such that

$$\begin{aligned} R_0 &\leq I(W; Y_2), \\ R_1 &\leq I(X; Y_1|W) \end{aligned}$$

for some distribution $p(w, x)$ with $|\mathcal{W}| \leq \min\{|\mathcal{X}|, |\mathcal{Y}_1|, |\mathcal{Y}_2|\} + 1$. $\circ$

Proof. Achievability: We want to show that any rate pair (R_0, R_1) such that

$$\begin{aligned} R_0 &< I(W; Y_2), \\ R_1 &< I(X; Y_1|W) \end{aligned}$$

for some distribution $p(w, x)$, is achievable. These two inequalities imply

$$\begin{aligned} R_0 + R_1 &< I(W; Y_2) + I(X; Y_1|W) \\ &\leq I(W; Y_1) + I(X; Y_1|W) && (W \rightarrow Y_1 \rightarrow Y_2 \text{ is Markov}) \\ &= I(W, X; Y_1) && (\text{chain rule}) \\ &= I(X; Y_1). && (W \rightarrow X \rightarrow Y_1 \text{ is Markov}) \end{aligned}$$

Therefore, (R_0, R_1) is in the superposition coding region and hence achievable.

Converse: Suppose a sequence of (R_0, R_1) codes is given. Consider

$$\begin{aligned} nR_0 &= \log |\mathcal{M}_0| + o(n) \\ &= H(M_0) + o(n) \\ &= I(M_0; \hat{M}_{02}) + H(M_0|\hat{M}_{02}) + o(n) && (\text{chain rule}) \\ &= I(M_0; \hat{M}_{02}) + o(n) && (\text{Fano's inequality}) \\ &\leq I(M_0; Y_2^n) + o(n), && (M_0 \rightarrow Y_2^n \rightarrow \hat{M}_{02} \text{ is Markov}) \end{aligned}$$

and

$$\begin{aligned}
I(M_0; Y_2^n) &= \sum_{i=1}^n I(M_0; Y_{2i} | Y_2^{i-1}) && \text{(chain rule)} \\
&\leq \sum_{i=1}^n (I(M_0; Y_{2i} | Y_2^{i-1}) + I(Y_{2i}; Y_2^{i-1})) \\
&= \sum_{i=1}^n I(M_0, Y_2^{i-1}; Y_{2i}) && \text{(chain rule)} \\
&\leq \sum_{i=1}^n I(M_0, Y_1^{i-1}; Y_{2i}). && (Y_{2i} \rightarrow (M_0, Y_1^{i-1}) \rightarrow (M_0, Y_2^{i-1}) \text{ is Markov})
\end{aligned}$$

Also,

$$\begin{aligned}
nR_1 &= \log |\mathcal{M}_1| + o(n) \\
&= H(M_1) + o(n) \\
&= H(M_1 | M_0) + o(n) && (M_0, M_1 \text{ are independent}) \\
&= I(M_1; \hat{M}_1 | M_0) + H(M_1 | \hat{M}_1, M_0) + o(n) && \text{(chain rule)} \\
&\leq I(M_1; \hat{M}_1 | M_0) + H(M_1 | \hat{M}_1) + o(n) \\
&= I(M_1; \hat{M}_1 | M_0) + o(n) && \text{(Fano's inequality)} \\
&\leq I(X^n; \hat{M}_1 | M_0) + o(n) && (M_1 \rightarrow (M_0, X^n) \rightarrow \hat{M}_1 \text{ is Markov}) \\
&\leq I(X^n; Y_1^n | M_0) + o(n), && (X^n \rightarrow (M_0, Y_1^n) \rightarrow \hat{M}_1 \text{ is Markov})
\end{aligned}$$

and

$$\begin{aligned}
I(X^n; Y_1^n | M_0) &= \sum_{i=1}^n I(X^n; Y_{1i} | M_0, Y_1^{i-1}) && \text{(chain rule)} \\
&\leq \sum_{i=1}^n I(X_i; Y_{1i} | M_0, Y_1^{i-1}). && (X^{n \setminus i} \rightarrow (X_i, M_0, Y_1^{i-1}) \rightarrow Y_{1i} \text{ is Markov})
\end{aligned}$$

Let $W_i \triangleq (M_0, Y_1^{i-1})$ and let Q be a random variable uniformly distributed on $[1 : n]$. Let $X \triangleq X_Q$, $Y_1 \triangleq Y_{1Q}$, $Y_2 \triangleq Y_{2Q}$, and $W \triangleq (W_Q, Q)$. Then we have

$$\begin{aligned}
R_0 &\leq I(W_Q; Y_{2Q} | Q) + o(1) \\
&\leq I(W_Q; Y_{2Q} | Q) + I(Q; Y_{2Q}) + o(1) \\
&= I(W_Q, Q; Y_{2Q}) + o(1) && \text{(chain rule)} \\
&= I(W; Y_2) + o(1), \\
R_1 &\leq I(X_Q; Y_{1Q} | W_Q, Q) + o(1) \\
&\leq I(X; Y_1 | W) + o(1).
\end{aligned}$$

Note that (W, X, Y_1, Y_2) is indeed distributed according to $p(w, x)p(y_1, y_2 | x)$. Then, since $\mathcal{X}$ is finite, the existence of the required distribution $p(w, x)$ follows by compactness of the finite-dimensional simplex $\mathcal{P}_{\mathcal{X}}$: take a convergent subsequence of the achieving laws.

Cardinality bound: Without loss of generality denote $\mathcal{X} = [1 : |\mathcal{X}|]$, $\mathcal{Y}_1 = [1 : |\mathcal{Y}_1|]$ and $\mathcal{Y}_2 = [1 : |\mathcal{Y}_2|]$. We first show that we can impose the cardinality bound $|\mathcal{W}| \leq |\mathcal{X}| + 1$ without shrinking the capacity region. The image of the set of all distributions on $\mathcal{X}$ under the continuous function

$$f : p(x) \mapsto (\Pr(X = 1), \dots, \Pr(X = |\mathcal{X}| - 1), H(Y_2), I(X; Y_1))$$

is a connected subset of $\mathbb{R}^{|\mathcal{X}|+1}$. Hence, by the Fenchel–Bunt extension of Carathéodory’s theorem, any joint distribution $p(x, w)$ on $\mathcal{X} \times \mathcal{W}$ can be replaced with another joint distribution such that $|\mathcal{W}| \leq |\mathcal{X}| + 1$ while preserving $\sum_{w \in \mathcal{W}} p(w) f(p(x|w))$, or in other words, $p(x)$, $H(Y_2|W)$, $I(X; Y_1|W)$ are preserved. Note that such replacement also preserves $p(y_2) = \sum_{x \in \mathcal{X}} p(x) p(y_2|x)$ and $I(W; Y_2) = H(Y_2) - H(Y_2|W)$. This gives the cardinality bound $|\mathcal{W}| \leq |\mathcal{X}| + 1$. Similarly, by considering

$$\begin{aligned} p(x) &\mapsto (\Pr(Y_1 = 1), \dots, \Pr(Y_1 = |\mathcal{Y}_1| - 1), H(Y_2), I(X; Y_1)), \\ p(x) &\mapsto (\Pr(Y_2 = 1), \dots, \Pr(Y_2 = |\mathcal{Y}_2| - 1), H(Y_2), I(X; Y_1)), \end{aligned}$$

we can get the cardinality bounds $|\mathcal{W}| \leq |\mathcal{Y}_1| + 1$ and $|\mathcal{W}| \leq |\mathcal{Y}_2| + 1$. ■

4.5 Upper concave envelope and sub-additivity

A recurring computation in this chapter is the maximization of a weighted sum rate, or more generally of a linear combination of information functionals, over an auxiliary random variable. The tool that organizes such computations is the *upper concave envelope*, together with its sub-additivity. We develop this toolbox here; it is used below for the Gaussian degraded broadcast channel, for the sub-additivity form of the *UV* bound, and for the Gaussian vector broadcast channel.

The systematic use of upper concave envelopes as a replacement for auxiliary random variables in the representation of inner and outer regions is developed in [Nair \(2013\)](#): besides simplifying boundary-point computations, the envelope representation leads to factorisation-type inequalities of the information quantities involved, and in several cases provides a direct route to proving optimality of an achievable region; see [Liu et al. \(2014\)](#) for the interference-channel application. This section presents the two elementary facts from that programme that the rest of this chapter needs.

Definition 4.5. (Upper concave envelope) Let $f : \mathcal{P}_{\mathcal{X}} \rightarrow \mathbb{R}$ be a functional on the set $\mathcal{P}_{\mathcal{X}}$ of probability distributions on $\mathcal{X}$. The upper concave envelope of f is the functional

$$\mathfrak{C}_{\mathcal{X}}[f](p_X) \triangleq \sup \left\{ \sum_w \lambda_w f(p_w) : \lambda_w \geq 0, \sum_w \lambda_w = 1, \sum_w \lambda_w p_w = p_X \right\},$$

where the supremum is over all finite mixtures of laws $p_w \in \mathcal{P}_{\mathcal{X}}$ whose mixture equals p_X . We write $\mathfrak{C}[f]$ when the variable X is clear from context. (Here and below, a “concave” functional is one concave in the distribution p_X . This mixture-supremum form is the operational definition; Lemma 4.7 below gives the equivalent auxiliary-random-variable representation.) ○

Lemma 4.6. The upper concave envelope $\mathfrak{C}_X[f]$ is concave, dominates f pointwise, and is the smallest concave function dominating f . $\circ$

Proof. Taking the trivial mixture ($\lambda = 1, p_w = p_X$) shows $\mathfrak{C}_X[f](p_X) \geq f(p_X)$. For concavity, given p_{X_1}, p_{X_2} and $p_X = \alpha p_{X_1} + \bar{\alpha} p_{X_2}$, for any $\epsilon > 0$ pick finite mixtures $\{(\lambda_w^1, p_w^1)\}$ and $\{(\lambda_w^2, p_w^2)\}$ representing p_{X_1} and p_{X_2} with $\sum_w \lambda_w^1 f(p_w^1) \geq \mathfrak{C}_X[f](p_{X_1}) - \epsilon$ and $\sum_w \lambda_w^2 f(p_w^2) \geq \mathfrak{C}_X[f](p_{X_2}) - \epsilon$. Concatenating the two mixtures with weights $\alpha \lambda_w^1$ and $\bar{\alpha} \lambda_w^2$ represents p_X , so

$$\begin{aligned} \mathfrak{C}_X[f](p_X) &\geq \alpha \sum_w \lambda_w^1 f(p_w^1) + \bar{\alpha} \sum_w \lambda_w^2 f(p_w^2) \\ &\geq \alpha \mathfrak{C}_X[f](p_{X_1}) + \bar{\alpha} \mathfrak{C}_X[f](p_{X_2}) - 2\epsilon. \end{aligned}$$

Letting $\epsilon \rightarrow 0$ proves concavity. For minimality, let g be any concave function with $g \geq f$ pointwise and let (λ_w, p_w) be any mixture representing p_X . Jensen's inequality gives

$$\sum_w \lambda_w f(p_w) \leq \sum_w \lambda_w g(p_w) \leq g\left(\sum_w \lambda_w p_w\right) = g(p_X),$$

so taking the supremum over all representations gives $\mathfrak{C}_X[f](p_X) \leq g(p_X)$. Thus $\mathfrak{C}_X[f]$ is the smallest concave majorant of f . By ordinary Carathéodory's theorem on the simplex, at most $|X|$ mixture points suffice for each fixed input law p_X . $\blacksquare$

Lemma 4.7. (Envelope as an auxiliary-random-variable supremum) Let f be a functional on $\mathcal{P}_X$, extended to conditional laws by averaging,

$$f(X|W) \triangleq \sum_w p_W(w) f(p_{X|W=w}).$$

Then

$$\mathfrak{C}_X[f](p_X) = \sup_{p_{X|W}} f(X|W),$$

where the supremum is over all channels $p_{X|W}$ (equivalently, over all p_W and $p_{X|W}$ with $\sum_w p_W(w) p_{X|W}(x|w) = p_X(x)$), with $|\mathcal{W}|$ unbounded. $\circ$

Proof. A channel $p_{X|W}$ together with p_W is exactly a mixture representation $\lambda_w = p_W(w)$, $p_w = p_{X|W=w}$ of the law $p_X = \sum_w p_W(w) p_{X|W=w}$, and conversely. The two suprema are therefore the same functional, and concavity and minimality were established in the preceding lemma. $\blacksquare$

4.5.1 Sub-additivity of the envelope in a degraded channel

Let $X \rightarrow Y_1 \rightarrow Y_2$ be a (physically) degraded channel, i.e. $p(y_1, y_2|x) = p(y_1|x)p(y_2|y_1)$. For $\lambda \geq 1$ consider the functional

$$f_\lambda(p_X) \triangleq I(X; Y_1) - \lambda I(X; Y_2),$$

whose upper concave envelope is $\mathfrak{C}_X[f_\lambda](p_X) = \sup_{p_{X|W}} [I(X; Y_1|W) - \lambda I(X; Y_2|W)]$.

Lemma 4.8. (Sub-additivity of the upper concave envelope) For $\lambda \geq 1$, the mapping

$$p_X \mapsto \mathfrak{C}_X[I(X; Y_1) - \lambda I(X; Y_2)]$$

is sub-additive: for the product of two independent degraded channels, (X_1, Y_{11}, Y_{21}) and (X_2, Y_{12}, Y_{22}) , and any joint input law $p_{X_1 X_2}$,

$$\begin{aligned} \mathfrak{C}_{X_1 X_2}[I(X_1 X_2; Y_{11} Y_{12}) - \lambda I(X_1 X_2; Y_{21} Y_{22})] &\leq \mathfrak{C}_{X_1}[I(X_1; Y_{11}) - \lambda I(X_1; Y_{21})] \\ &\quad + \mathfrak{C}_{X_2}[I(X_2; Y_{12}) - \lambda I(X_2; Y_{22})]. \end{aligned}$$

○

Proof sketch (cited factorization lemma). The following factorization lemma is standard in the proof of the degraded Gaussian BC and is due to Bergmans 1974; we state it in the form needed here and refer to Bergmans (1974) (see also the rotation-trick account of §1.13 of Chapter 1) for the chain-rule proof. Its proof uses the auxiliary-variable representation of the envelope (Lemma 4.7), successive conditioning on the coordinates of the product channel, and the non-positive remainder available for $\lambda \geq 1$. The displayed computation below is that derivation in outline; we import it as a lemma rather than claiming the three-line manipulation alone proves the statement. Writing the envelope as a supremum over $p_{X_1 X_2 | W}$ and expanding the mutual informations on the product channel, the Markov structure of the degraded product channel (the “ d -separation” graph $W \rightarrow (X_1, X_2) \rightarrow Y_1 \rightarrow Y_2$ on each coordinate) gives the chain-rule decomposition

$$\begin{aligned} &I(X_1, X_2; Y_{11}, Y_{12} | W) - \lambda I(X_1, X_2; Y_{21}, Y_{22} | W) \\ &= I(X_1; Y_{11} | W) + I(X_2; Y_{12} | W, Y_{11}) - \lambda I(X_1; Y_{21} | W) - \lambda I(X_2; Y_{22} | W, Y_{11}) - \lambda I(Y_{11}; Y_{22} | W, Y_{21}) \\ &\leq \mathfrak{C}_{X_1}[I(X_1; Y_{11}) - \lambda I(X_1; Y_{21})] + \mathfrak{C}_{X_2}[I(X_2; Y_{12}) - \lambda I(X_2; Y_{22})], \end{aligned}$$

where in the last step the first two terms are bounded by the envelope over X_1 and the remaining terms (using $\lambda \geq 1$ to absorb the final non-positive cross term $-\lambda I(Y_{11}; Y_{22} | W, Y_{21})$, together with $I(X_2; Y_{22} | W, Y_{11}) \geq I(X_2; Y_{22} | W, Y_{11}, Y_{21})$) are bounded by the envelope over X_2 . This is the standard “single-letterization” of the envelope; a fully rigorous treatment appears in the Gaussian optimality proofs (Bergmans, and the rotation-trick account of §1.13 of Chapter 1). ■

Remark 4.4. The hypothesis $\lambda \geq 1$ is essential: it is what makes the cross term $-\lambda I(Y_{11}; Y_{22} | W, Y_{21})$ non-positive so that it can be dropped. (The weaker requirement $\lambda \geq 0$ does not suffice for the last inequality.) ○

Corollary 4.9. (Weighted sum via the envelope) For a degraded broadcast channel, the weighted sum rate satisfies

$$\max_{(R_0, R_1) \in \mathcal{C}} (\lambda R_0 + R_1) = \max_{p_X} \left\{ \lambda I(X; Y_2) + \mathfrak{C}_X[I(X; Y_1) - \lambda I(X; Y_2)] \right\}.$$

○

Proof. By the capacity region of the degraded BC (Theorem above), $\lambda R_0 + R_1 \leq \lambda I(W; Y_2) + I(X; Y_1|W)$. Using $W \rightarrow X \rightarrow Y_2$,

$$\begin{aligned}\lambda I(W; Y_2) + I(X; Y_1|W) &= \lambda I(X; Y_2) - \lambda I(X; Y_2|W) + I(X; Y_1|W) \\ &= \lambda I(X; Y_2) + [I(X; Y_1|W) - \lambda I(X; Y_2|W)].\end{aligned}$$

Maximizing first over $p(x)$ and then over $p(w|x)$ (equivalently $p_{X|W}$) and using Lemma 4.7 gives the claim. ■

4.6 AWGN degraded broadcast channel

Consider the scalar degraded additive white Gaussian noise (AWGN) broadcast channel

$$\begin{aligned}Y_1 &= X + Z_1, \\ Y_2 &= X + Z_1 + Z_2 = Y_1 + Z_2,\end{aligned}$$

where $Z_1 \sim \mathcal{N}(0, N_1)$ and $Z_2 \sim \mathcal{N}(0, N_2)$ are independent Gaussians and the input satisfies $E[X^2] \leq P$. Thus Y_2 is a degraded version of Y_1 , and receiver 1 is stronger. By the degraded-BC result, the capacity region is

$$\left\{ (R_1, R_2) : R_1 \leq I(X; Y_1|W), R_2 \leq I(W; Y_2) \text{ for some } p(w, x), E[X^2] \leq P \right\}.$$

Theorem 4.10. (AWGN degraded BC capacity) The capacity region of the scalar AWGN degraded broadcast channel is the set of rate pairs (R_1, R_2) such that

$$\begin{aligned}R_1 &\leq \frac{1}{2} \log \left(1 + \frac{\alpha P}{N_1} \right), \\ R_2 &\leq \frac{1}{2} \log \left(1 + \frac{(1 - \alpha)P}{\alpha P + N_1 + N_2} \right)\end{aligned}$$

for some $0 \leq \alpha \leq 1$. ○

Proof. Achievability: Superposition coding with $X = V + W$, where $W \sim \mathcal{N}(0, (1 - \alpha)P)$ is the cloud center (decoded by both receivers) and $V \sim \mathcal{N}(0, \alpha P)$ is the satellite (decoded only by receiver 1), $V \perp W$. Receiver 1 decodes V after subtracting W , and receiver 2 decodes only W , treating V as noise:

$$\begin{aligned}I(X; Y_1|W) &= \frac{1}{2} \log \left(1 + \frac{\alpha P}{N_1} \right), \\ I(W; Y_2) &= \frac{1}{2} \log \left(1 + \frac{(1 - \alpha)P}{\alpha P + N_1 + N_2} \right).\end{aligned}$$

Converse (Bergmans 1974): Suppose (R_1, R_2) is achievable. By the degraded-BC capacity region, there is $p(w, x)$ with $E[X^2] \leq P$, $R_1 \leq I(X; Y_1|W)$, $R_2 \leq I(W; Y_2)$. We work in *natural*

units (nats) in the derivation and convert to bits at the end: below h denotes natural-log differential entropy and $\ln$ the natural logarithm. We have

$$\begin{aligned} R_1 &\leq I(X; Y_1|W) = h(Y_1|W) - h(Z_1) = h(Y_1|W) - \frac{1}{2} \ln(2\pi e N_1), \\ R_2 &\leq I(W; Y_2) = h(Y_2) - h(Y_2|W) \leq \frac{1}{2} \ln(2\pi e(P + N_1 + N_2)) - h(Y_2|W), \end{aligned}$$

where the last inequality uses $h(Y_2) \leq \frac{1}{2} \ln(2\pi e(P + N_1 + N_2))$, since Gaussian maximizes differential entropy subject to $E[X^2] \leq P$. The key step is the conditional entropy power inequality (EPI) of Bergmans 1974. Since $Y_2 = Y_1 + Z_2$ with Z_2 Gaussian and independent of (X, Z_1, W) , the conditional EPI gives

$$e^{2h(Y_2|W)} = e^{2h(Y_1+Z_2|W)} \geq e^{2h(Y_1|W)} + e^{2h(Z_2)} = e^{2h(Y_1|W)} + 2\pi e N_2.$$

(The conditional entropies here are finite — $h(Z_2)$ is finite and $h(Y_1|W) \leq h(Y_1) \leq \frac{1}{2} \ln(2\pi e(P + N_1))$ — so the conditional EPI applies directly, without the regularized extension.) Now $h(Y_1|W)$ ranges over $[\frac{1}{2} \ln(2\pi e N_1), \frac{1}{2} \ln(2\pi e(P + N_1))]$ (entropy cannot decrease by adding X , and is maximized by Gaussian input), so there is a unique $\alpha \in [0, 1]$ with

$$h(Y_1|W) = \frac{1}{2} \ln(2\pi e(\alpha P + N_1)).$$

Then $R_1 \leq \frac{1}{2} \ln(1 + \frac{\alpha P}{N_1})$ (nats), and the conditional EPI gives

$$e^{2h(Y_2|W)} \geq 2\pi e(\alpha P + N_1) + 2\pi e N_2 = 2\pi e(\alpha P + N_1 + N_2),$$

so that $h(Y_2|W) \geq \frac{1}{2} \ln(2\pi e(\alpha P + N_1 + N_2))$ and hence

$$R_2 \leq \frac{1}{2} \ln(2\pi e(P + N_1 + N_2)) - \frac{1}{2} \ln(2\pi e(\alpha P + N_1 + N_2)) = \frac{1}{2} \ln\left(1 + \frac{(1-\alpha)P}{\alpha P + N_1 + N_2}\right)$$

in nats. Converting nats to bits (dividing by $\ln 2$, which replaces each $\ln$ by $\log_2$) yields the stated bit-rate bounds $R_1 \leq \frac{1}{2} \log_2(1 + \frac{\alpha P}{N_1})$ and $R_2 \leq \frac{1}{2} \log_2(1 + \frac{(1-\alpha)P}{\alpha P + N_1 + N_2})$. The Gaussian superposition codebook attains both bounds with equality. (The sub-additivity / rotation-trick route to Gaussian optimality is the alternative proof described in §1.13 of Chapter 1; the EPI proof above is the classical one.) ■

4.7 Degraded broadcast channel with K receivers

The superposition-coding story extends verbatim to a chain of degraded receivers

$$X \rightarrow Y_1 \rightarrow Y_2 \rightarrow \cdots \rightarrow Y_K.$$

Theorem 4.11. (K -receiver degraded BC) For the degraded chain above, superposition coding achieves the capacity region consisting of all *nonnegative* rate tuples $(R_1, \dots, R_K)$ satisfying

$$R_K \leq I(W_K; Y_K),$$

$$\begin{aligned}
R_{K-1} &\leq I(W_{K-1}; Y_{K-1} | W_K), \\
&\vdots \\
R_1 &\leq I(X; Y_1 | W_2, \dots, W_K)
\end{aligned}$$

for some layered law

$$p(w_K) p(w_{K-1} | w_K) \cdots p(w_2 | w_3) p(x | w_2),$$

equivalently for some $p(w_2, \dots, w_K, x)$ in which the Markov chain

$$W_K \rightarrow W_{K-1} \rightarrow \cdots \rightarrow W_2 \rightarrow X \rightarrow Y_1 \rightarrow \cdots \rightarrow Y_K$$

holds. Here R_j is the rate of the j -th *layered* message, carried by the auxiliary W_j and decoded by receivers $1, \dots, j$ (the conventional degraded-BC formulation): W_K is the outermost cloud center decoded by everyone, and X plays the role of W_1 . Boundary rate tuples are included by closure of the achievable region. $\circ$

Proof sketch. Achievability is the natural K -level superposition code generated according to the layered factorization above. For the converse, single-letterize with the auxiliaries $W_{li} = (M_i, \dots, M_K, Y_{li}^{i-1})$ (the hint $U_i = M_i^K, Y_{li}^{i-1}$), and use the degradedness $X \rightarrow Y_i \rightarrow Y_{i+1}$ at each level; the standard support-lemma cardinality bounds for the layered chain are as in the two-receiver case (see also Gallager (1974); Bergmans (1973)). $\blacksquare$

Remark 4.5 (More than two receivers). For the less-noisy and more-capable orders the situation is richer. For *three* receivers with a less-noisy chain $Y_1 \preceq Y_2 \preceq Y_3$ (in the less-noisy order), superposition coding is still optimal: the capacity region is

$$R_3 \leq I(W_3; Y_3), \quad R_2 \leq I(W_2; Y_2 | W_3), \quad R_1 \leq I(X; Y_1 | W_2, W_3),$$

with auxiliaries $W_{3i} = (M_3, Y_2^{i-1})$ and $W_{2i} = (M_2, M_3, Y_2^{i-1})$ in the converse. For *four or more* receivers, the optimality of superposition coding even for the *less-noisy* order is open (the physically degraded class is settled by the theorem above for every K). For the *more-capable* order, superposition coding is optimal for two receivers (El Gamal 1979) but is strictly suboptimal already for three receivers (Xia–Gohari et al.). $\circ$

4.8 Less noisy and more capable broadcast channels

We have shown that superposition coding is optimal for a broadcast channel $p_{Y_1, Y_2 | X}$ if Y_2 is a degraded version of Y_1 . In fact, superposition coding can be shown to be optimal for larger classes of broadcast channels, namely less noisy broadcast channels and more capable broadcast channels.

Definition 4.12. (Less noisy channel) A channel $p_{Y_1|X}$ is called less noisy than another channel $p_{Y_2|X}$ if

$$I(U; Y_1) \geq I(U; Y_2)$$

for any distribution $p(u, x)$, where $(U, X, Y_1, Y_2) \sim p(u, x)p(y_1|x)p(y_2|x)$. $\circ$

Remark 4.6. It is immediate that $p_{Y_1|X}$ is less noisy than $p_{Y_2|X}$ if and only if the *conditional* form holds: for any random variables U, V, W such that $(U, V, W) \rightarrow X \rightarrow (Y_1, Y_2)$ forms a Markov chain,

$$I(U; Y_1|W) \geq I(U; Y_2|W).$$

This is obtained by applying the definition with $p(u|w)$ for each fixed w ; the version with an additional V follows by the chain rule $I(U; Y_j, V|W) = I(U; V|W) + I(U; Y_j|V, W)$. $\circ$

Similarly to the case of the degraded broadcast channel, for a less noisy broadcast channel, we can assume $R_2 = 0$ without loss of generality. Thus, we are only interested in the projection $(R_0, R_1, 0)$ of the capacity region.

Proposition 4.13. For a broadcast channel $p_{Y_1, Y_2|X}$ such that $p_{Y_1|X}$ is less noisy than $p_{Y_2|X}$, (R_0, R_1, R_2) is in the capacity region if and only if $(R_0 + R_2, R_1, 0)$ is in the capacity region. $\circ$

Proof. The “if” part is done by moving some parts of the common message to the private message intended for receiver 2.

Now we show the “only if” part. Suppose a sequence of (n, R_0, R_1, R_2) codes is given. A standard application of Fano’s inequality gives

$$\begin{aligned} nR_1 &\leq I(X^n; Y_1^n | M_0, M_2) + o(n), \\ n(R_0 + R_2) &\leq I(M_0, M_2; Y_2^n) + o(n), \\ n(R_0 + R_1 + R_2) &\leq I(X^n; Y_1^n | M_0, M_2) + I(M_0, M_2; Y_2^n) + o(n). \end{aligned}$$

By the less-noisy replacement lemma (Lemma 4.14 below), we have

$$\begin{aligned} I(M_0, M_2; Y_2^n) &= \sum_{i=1}^n I(M_0, M_2; Y_{2i} | Y_2^{i-1}) && \text{(chain rule)} \\ &\leq \sum_{i=1}^n I(M_0, M_2, Y_2^{i-1}; Y_{2i}) \\ &\leq \sum_{i=1}^n I(M_0, M_2, Y_1^{i-1}; Y_{2i}) \\ &= \sum_{i=1}^n I(W_i; Y_{2i}), \end{aligned}$$

where $W_i \triangleq (M_0, M_2, Y_1^{i-1})$. Also,

$$\begin{aligned}
 I(X^n; Y_1^n | M_0, M_2) &= \sum_{i=1}^n I(X^n; Y_{1i} | M_0, M_2, Y_1^{i-1}) \quad (\text{chain rule}) \\
 &\leq \sum_{i=1}^n I(X_i; Y_{1i} | M_0, M_2, Y_1^{i-1}) \quad (X^{n \setminus i} \rightarrow (X_i, M_0, M_2, Y_1^{i-1}) \rightarrow Y_{1i} \text{ is Markov}) \\
 &= \sum_{i=1}^n I(X_i; Y_{1i} | W_i).
 \end{aligned}$$

Note that we also have

$$\begin{aligned}
 &I(X^n; Y_1^n | M_0, M_2) + I(M_0, M_2; Y_2^n) \\
 &\leq \sum_{i=1}^n (I(X_i; Y_{1i} | W_i) + I(W_i; Y_{2i})) \\
 &\leq \sum_{i=1}^n (I(X_i; Y_{1i} | W_i) + I(W_i; Y_{1i})) \quad (\text{less noisy condition}) \\
 &= \sum_{i=1}^n I(W_i, X_i; Y_{1i}) \quad (\text{chain rule}) \\
 &= \sum_{i=1}^n I(X_i; Y_{1i}). \quad (W_i \rightarrow X_i \rightarrow Y_{1i} \text{ is Markov})
 \end{aligned}$$

Let Q be a random variable uniformly distributed on $[1 : n]$. Let $X \triangleq X_Q$, $Y_1 \triangleq Y_{1Q}$, $Y_2 \triangleq Y_{2Q}$, and $W \triangleq (W_Q, Q)$. Then we have

$$\begin{aligned}
 R_1 &\leq I(X_Q; Y_{1Q} | W_Q, Q) + o(1) \\
 &= I(X; Y_1 | W) + o(1), \\
 R_0 + R_2 &\leq I(W_Q; Y_{2Q} | Q) + o(1) \\
 &\leq I(W_Q; Y_{2Q} | Q) + I(Q; Y_{2Q}) + o(1) \\
 &= I(W_Q, Q; Y_{2Q}) + o(1) \quad (\text{chain rule}) \\
 &= I(W; Y_2) + o(1), \\
 R_0 + R_1 + R_2 &\leq I(X_Q; Y_{1Q} | Q) + o(1) \\
 &\leq I(X_Q; Y_{1Q} | Q) + I(Q; Y_{1Q}) + o(1) \\
 &= I(X_Q, Q; Y_{1Q}) + o(1) \quad (\text{chain rule}) \\
 &= I(X_Q; Y_{1Q}) + o(1) \quad (Q \rightarrow X_Q \rightarrow Y_{1Q} \text{ is Markov}) \\
 &= I(X; Y_1) + o(1).
 \end{aligned}$$

Hence, for every fixed $\epsilon > 0$, the tuple $(R_0 + R_2 - \epsilon, R_1 - \epsilon, 0)$ satisfies the superposition-coding inequalities strictly and is achievable. Taking $\epsilon \rightarrow 0$ and using closure of the capacity region, $(R_0 + R_2, R_1, 0)$ belongs to the capacity region (closure, not exact fixed-rate achievability of the limiting boundary tuple). ■

Lemma 4.14. (Less-noisy replacement) Suppose $p_{Y_1|X}$ is less noisy than $p_{Y_2|X}$. Then for any (n, R_0, R_1, R_2) code and each $i \in [1 : n]$,

$$I(M_0, M_2, Y_2^{i-1}; Y_{2i}) \leq I(M_0, M_2, Y_1^{i-1}; Y_{2i}).$$

○

Proof interface. The displayed inequality is the standard “replace past Y_2 by past Y_1 ” step of the less-noisy converse, due to Körner–Marton 1977a; it is presented in full in the textbook treatment of El Gamal and Kim (2012) (Thm. 5.3). We import it here as a known less-noisy information inequality rather than deriving it from scratch. What follows directly from the definition is the *end-to-end* tensorization: for any $U \rightarrow X^n \rightarrow (Y_1^n, Y_2^n)$ one has $I(U; Y_1^n) \geq I(U; Y_2^n)$. Indeed, writing $U_i = (U, Y_1^{i-1}, Y_{2,i+1}^n)$, the Csiszár sum identity (the same telescoping used in the UV bound below) gives

$$I(U; Y_1^n) - I(U; Y_2^n) = \sum_{i=1}^n [I(U_i; Y_{1i}) - I(U_i; Y_{2i})] \geq 0,$$

since the channel is memoryless, so $U_i \rightarrow X_i \rightarrow (Y_{1i}, Y_{2i})$ and each summand is non-negative by less-noisy. This end-to-end form does *not* by itself yield the local replacement with output Y_{2i} : the local step additionally uses the image-of-a-set argument of El Gamal and Kim (2012) (Thm. 5.3), in which the past output sequence is replaced coordinate by coordinate while the future inputs are conditioned on; the induction is carried out there in full. ■

Theorem 4.15. (Capacity region of less noisy broadcast channel (Körner–Marton 1977a))

The capacity region of a broadcast channel $p_{Y_1, Y_2|X}$ such that $p_{Y_1|X}$ is less noisy than $p_{Y_2|X}$ is the set of rate pairs (R_0, R_1) such that

$$\begin{aligned} R_0 &\leq I(W; Y_2), \\ R_1 &\leq I(X; Y_1|W) \end{aligned}$$

for some distribution $p(w, x)$ with $|\mathcal{W}| \leq \min\{|\mathcal{X}|, |\mathcal{Y}_2|\} + 1$.

○

Remark 4.7 (Beyond two receivers). The two-receiver results above extend, with genuine difficulty, to three and more receivers. The obstruction is the converse: the single-letterisation of the two-receiver converses rests on a Csiszár-sum telescoping between *two* output sequences, and no direct three-sequence generalisation is known. For the class of *less noisy* receivers this obstruction was bypassed by a new information inequality, and the capacity region of the three-receiver less noisy broadcast channel is known Nair and Wang (2011); the three-receiver more capable capacity region is likewise known Nair and Xia (2012). On the other hand, for three receivers with two degraded message sets, superposition coding is strictly suboptimal in general Nair and Yazdanpanah (2017), whereas it is exactly optimal for the class of three-receiver channels with degraded message sets treated in Nair and El Gamal (2009). These results make good optional reading: they show the two-receiver tools — the sub-additivity step, the Csiszár sum identity, superposition coding — each finding their precise limit. ○

Proof. The proof is the same as that of the capacity of the degraded broadcast channel, except that:

- In the proof of achievability, we lack the Markov chain $W \rightarrow Y_1 \rightarrow Y_2$, but the less noisy condition still gives $I(W; Y_2) \leq I(W; Y_1)$, which is all that is needed (so $R_0 + R_1 < I(W; Y_2) + I(X; Y_1|W) \leq I(X; Y_1)$).
- In the proof of converse, the step $I(M_0, Y_2^{i-1}; Y_{2i}) \leq I(M_0, Y_1^{i-1}; Y_{2i})$ (valid by degradation) is replaced by the less-noisy replacement Lemma 4.14.
- In the proof of cardinality bound for $\mathcal{W}$, we can no longer obtain a cardinality bound by considering

$$p(x) \mapsto (\Pr(Y_1 = 1), \dots, \Pr(Y_1 = |\mathcal{Y}_1| - 1), H(Y_2), I(X; Y_1)),$$

since preserving $p(y_1)$ no longer preserves $H(Y_2)$, as we no longer have degradation. Hence only $|\mathcal{W}| \leq \min\{|\mathcal{X}|, |\mathcal{Y}_2|\} + 1$ survives. ■

Definition 4.16. (More capable channel) A channel $p_{Y_1|X}$ is called more capable than another channel $p_{Y_2|X}$ if

$$I(X; Y_1) \geq I(X; Y_2)$$

for any distribution $p(x)$, where $(X, Y_1, Y_2) \sim p(x)p(y_1|x)p(y_2|x)$. ○

Remark 4.8. It is immediate that $p_{Y_1|X}$ is more capable than $p_{Y_2|X}$ if and only if the conditional form holds: for any random variables U such that $U \rightarrow X \rightarrow (Y_1, Y_2)$ forms a Markov chain,

$$I(X; Y_1|U) \geq I(X; Y_2|U).$$

This is obtained by applying the definition to $P_{X|U=u}$ for each fixed u : both sides are averages over U of the one-letter quantities $I(X; Y_j | U = u)$ evaluated under $p(x|u)$. If an additional variable V is present, apply the same argument conditionally on $(U, V) = (u, v)$: for any $(U, V) \rightarrow X \rightarrow (Y_1, Y_2)$ one has $I(X; Y_1|U, V) \geq I(X; Y_2|U, V)$, using $P_{X|U=u, V=v}$ in place of $p(x)$. ○

Remark 4.9 (Chain of inclusions). The three classes satisfy

$$\text{Degraded} \subsetneq \text{Less noisy} \subsetneq \text{More capable}.$$

○

Similarly to the cases of the degraded and less noisy broadcast channels, for a more capable broadcast channel, we can assume $R_2 = 0$ without loss of generality. Thus, we are only interested in the projection $(R_0, R_1, 0)$ of the capacity region.

Proposition 4.17. For a broadcast channel $p_{Y_1, Y_2|X}$ such that $p_{Y_1|X}$ is more capable than $p_{Y_2|X}$, (R_0, R_1, R_2) is in the capacity region if and only if $(R_0 + R_2, R_1, 0)$ is in the capacity region. $\circ$

Proof. The “if” part is done by moving some parts of the common message to the private message intended for receiver 2.

Now we show the “only if” part. Suppose a sequence of (n, R_0, R_1, R_2) codes is given. We will show that $(R_0 + R_2, R_1, 0)$ is in the superposition coding region and hence in the capacity region. It suffices to show instead that

$$\begin{aligned} R_0 + R_2 &\leq I(W; Y_2), \\ R_0 + R_1 + R_2 &\leq I(W; Y_2) + I(X; Y_1|W), \\ R_0 + R_1 + R_2 &\leq I(X; Y_1) \end{aligned}$$

for some distribution $p(w, x)$. To show that this suffices, observe that the superposition coding region is convex, and that for every point (R_0, R_1, R_2) in this alternative region, $(R_0 + R_2, R_1)$ is a convex combination of the points

- $(0, 0)$,
- $(I(W; Y_2), 0)$,
- $(I(W; Y_2), \min\{I(X; Y_1), I(W; Y_2) + I(X; Y_1|W)\} - I(W; Y_2))$,
- $(0, \min\{I(X; Y_1), I(W; Y_2) + I(X; Y_1|W)\})$,

which can all be shown to be in the superposition region.

A standard application of Fano’s inequality gives

$$\begin{aligned} n(R_0 + R_2) &\leq I(M_0, M_2; Y_2^n) + o(n), \\ n(R_0 + R_1 + R_2) &\leq I(M_0, M_2; Y_2^n) + I(X^n; Y_1^n | M_0, M_2) + o(n), \\ n(R_0 + R_1 + R_2) &\leq I(M_0, M_1; Y_1^n) + I(X^n; Y_2^n | M_0, M_1) + o(n). \end{aligned}$$

Then,

$$\begin{aligned} I(M_0, M_2; Y_2^n) &= \sum_{i=1}^n I(M_0, M_2; Y_{2i} | Y_{2(i+1)}^n) && \text{(chain rule)} \\ &\leq \sum_{i=1}^n \left(I(M_0, M_2; Y_{2i} | Y_{2(i+1)}^n) + I(Y_{2i}; Y_{2(i+1)}^n) \right) \\ &= \sum_{i=1}^n I(M_0, M_2, Y_{2(i+1)}^n; Y_{2i}) && \text{(chain rule)} \\ &\leq \sum_{i=1}^n I(M_0, M_2, Y_1^{i-1}, Y_{2(i+1)}^n; Y_{2i}) \\ &= \sum_{i=1}^n I(W_i; Y_{2i}), \end{aligned}$$

where $W_i \triangleq (M_0, M_2, Y_1^{i-1}, Y_{2(i+1)}^n)$. Also, $X^{n \setminus i} \rightarrow (X_i, M_0, M_2, Y_1^{i-1}) \rightarrow Y_{1i}$ is Markov, since the channel is memoryless,

$$\begin{aligned}
& I(M_0, M_2; Y_2^n) + I(X^n; Y_1^n | M_0, M_2) \\
&= \sum_{i=1}^n I(M_0, M_2; Y_{2i} | Y_{2(i+1)}^n) + \sum_{i=1}^n I(X^n; Y_{1i} | M_0, M_2, Y_1^{i-1}) \quad (\text{chain rule}) \\
&= \sum_{i=1}^n I(M_0, M_2; Y_{2i} | Y_{2(i+1)}^n) + \sum_{i=1}^n I(X_i; Y_{1i} | M_0, M_2, Y_1^{i-1}) \quad (\text{Markov}) \\
&\leq \sum_{i=1}^n I(M_0, M_2; Y_{2i} | Y_{2(i+1)}^n) + \sum_{i=1}^n I(X_i, Y_{2(i+1)}^n; Y_{1i} | M_0, M_2, Y_1^{i-1}) \\
&\leq \sum_{i=1}^n I(M_0, M_2; Y_{2i} | Y_{2(i+1)}^n) + \sum_{i=1}^n I(Y_{2(i+1)}^n; Y_{2i}) \\
&\quad + \sum_{i=1}^n I(X_i, Y_{2(i+1)}^n; Y_{1i} | M_0, M_2, Y_1^{i-1}) \\
&= \sum_{i=1}^n I(M_0, M_2, Y_{2(i+1)}^n; Y_{2i}) \\
&\quad + \sum_{i=1}^n I(X_i, Y_{2(i+1)}^n; Y_{1i} | M_0, M_2, Y_1^{i-1}) \quad (\text{chain rule}) \\
&= \sum_{i=1}^n \left(I(M_0, M_2, Y_1^{i-1}, Y_{2(i+1)}^n; Y_{2i}) - I(Y_1^{i-1}; Y_{2i} | M_0, M_2, Y_{2(i+1)}^n) \right. \\
&\quad \left. + I(X_i; Y_{1i} | M_0, M_2, Y_1^{i-1}, Y_{2(i+1)}^n) + I(Y_{2(i+1)}^n; Y_{1i} | M_0, M_2, Y_1^{i-1}) \right) \quad (\text{chain rule}) \\
&= \sum_{i=1}^n (I(W_i; Y_{2i}) + I(X_i; Y_{1i} | W_i)) \\
&\quad + \sum_{i=1}^n I(Y_{2(i+1)}^n; Y_{1i} | M_0, M_2, Y_1^{i-1}) - \sum_{i=1}^n I(Y_1^{i-1}; Y_{2i} | M_0, M_2, Y_{2(i+1)}^n) \\
&= \sum_{i=1}^n (I(W_i; Y_{2i}) + I(X_i; Y_{1i} | W_i)) \\
&\quad + \sum_{i=1}^n I(Y_1^i; Y_{2(i+1)}^n | M_0, M_2) - \sum_{i=1}^n I(Y_1^{i-1}; Y_{2i} | M_0, M_2) \\
&\quad \quad \quad (\text{Csiszár sum identity}) \\
&= \sum_{i=1}^n (I(W_i; Y_{2i}) + I(X_i; Y_{1i} | W_i)) .
\end{aligned}$$

Similarly,

$$I(M_0, M_1; Y_1^n) + I(X^n; Y_2^n | M_0, M_1) \leq \sum_{i=1}^n (I(V_i; Y_{1i}) + I(X_i; Y_{2i} | V_i)),$$

where $V_i \triangleq (M_0, M_1, Y_1^{i-1}, Y_{2(i+1)}^n)$. Let Q be a random variable uniformly distributed on $[1 : n]$ and let $X \triangleq X_Q$, $Y_1 \triangleq Y_{1Q}$, $Y_2 \triangleq Y_{2Q}$, $W \triangleq (W_Q, Q)$, and $V \triangleq (V_Q, Q)$. Then,

$$\begin{aligned} R_0 + R_2 &\leq I(W_Q; Y_{2Q} | Q) + o(1) \\ &\leq I(W_Q; Y_{2Q} | Q) + I(Q; Y_{2Q}) + o(1) \\ &= I(W_Q, Q; Y_{2Q}) + o(1) && \text{(chain rule)} \\ &= I(W; Y_2) + o(1), \\ R_0 + R_1 + R_2 &\leq I(W_Q; Y_{2Q} | Q) + I(X_Q; Y_{1Q} | W_Q, Q) + o(1) \\ &\leq I(W_Q; Y_{2Q} | Q) + I(Q; Y_{2Q}) + I(X_Q; Y_{1Q} | W_Q, Q) + o(1) \\ &= I(W_Q, Q; Y_{2Q}) + I(X_Q; Y_{1Q} | W_Q, Q) + o(1) && \text{(chain rule)} \\ &= I(W; Y_2) + I(X; Y_1 | W) + o(1), \\ R_0 + R_1 + R_2 &\leq I(V_Q; Y_{1Q} | Q) + I(X_Q; Y_{2Q} | V_Q, Q) + o(1) \\ &\leq I(V_Q; Y_{1Q} | Q) + I(Q; Y_{1Q}) + I(X_Q; Y_{2Q} | V_Q, Q) + o(1) \\ &= I(V_Q, Q; Y_{1Q}) + I(X_Q; Y_{2Q} | V_Q, Q) + o(1) && \text{(chain rule)} \\ &= I(V; Y_1) + I(X; Y_2 | V) + o(1) \\ &\leq I(V; Y_1) + I(X; Y_1 | V) + o(1) && \text{(more capable condition)} \\ &= I(V, X; Y_1) + o(1) && \text{(chain rule)} \\ &= I(X; Y_1) + o(1). && (V \rightarrow X \rightarrow Y_1 \text{ is Markov}) \end{aligned}$$

The displayed bounds hold up to $o(1)$: for every fixed $\epsilon > 0$ the tuple $(R_0 + R_2 - \epsilon, R_1 - \epsilon, 0)$ satisfies the alternative region's inequalities strictly and hence lies in the superposition coding region, and taking $\epsilon \rightarrow 0$ with closure of the capacity region gives $(R_0 + R_2, R_1, 0)$ in the capacity region (closure, not exact fixed-rate achievability of the limiting boundary tuple). ■

Theorem 4.18. (*Capacity region of more capable broadcast channel (El Gamal 1979)*)

The capacity region of a broadcast channel $p_{Y_1, Y_2 | X}$ such that $p_{Y_1 | X}$ is more capable than $p_{Y_2 | X}$ is the set of rate pairs (R_0, R_1) such that

$$\begin{aligned} R_0 &\leq I(W; Y_2), \\ R_0 + R_1 &\leq I(X; Y_1), \\ R_1 &\leq I(X; Y_1 | W) \end{aligned}$$

for some distribution $p(w, x)$ with $|\mathcal{W}| \leq \min\{|\mathcal{X}|, |\mathcal{Y}_1| \cdot |\mathcal{Y}_2|\} + 2$. ○

Proof. Achievability: Achievability follows from superposition coding.

Converse: We shall show instead that

$$\begin{aligned} R_0 &\leq I(W; Y_2), \\ R_0 + R_1 &\leq I(W; Y_2) + I(X; Y_1|W), \\ R_0 + R_1 &\leq I(X; Y_1) \end{aligned}$$

for some distribution $p(w, x)$. To show that this suffices, observe that the original region is convex, and that every point in this alternative region is a convex combination of the points

- $(0, 0)$,
- $(I(W; Y_2), 0)$,
- $(I(W; Y_2), \min\{I(X; Y_1), I(W; Y_2) + I(X; Y_1|W)\} - I(W; Y_2))$,
- $(0, \min\{I(X; Y_1), I(W; Y_2) + I(X; Y_1|W)\})$,

which can all be shown to be in the original region.

Suppose a sequence of (R_0, R_1) codes is given. A standard application of Fano's inequality gives

$$\begin{aligned} nR_0 &\leq I(M_0; Y_2^n) + o(n), \\ n(R_0 + R_1) &\leq I(M_0; Y_2^n) + I(X^n; Y_1^n|M_0) + o(n), \\ n(R_0 + R_1) &\leq I(M_1; Y_1^n) + I(X^n; Y_2^n|M_1) + o(n). \end{aligned}$$

The telescoping is the same as in the previous proposition, with the identifications $W_i \triangleq (M_0, Y_1^{i-1}, Y_{2(i+1)}^n)$ and $V_i \triangleq (M_1, Y_1^{i-1}, Y_{2(i+1)}^n)$; it yields

$$\begin{aligned} R_0 &\leq I(W; Y_2) + o(1), \\ R_0 + R_1 &\leq I(W; Y_2) + I(X; Y_1|W) + o(1), \\ R_0 + R_1 &\leq I(V; Y_1) + I(X; Y_2|V) + o(1) \leq I(X; Y_1) + o(1), \end{aligned}$$

where the last inequality uses the more capable condition $I(X; Y_2|V) \leq I(X; Y_1|V)$ and the chain rule.

Cardinality bound: Without loss of generality denote $\mathcal{X} = [1 : |\mathcal{X}|]$, $\mathcal{Y}_1 = [1 : |\mathcal{Y}_1|]$ and $\mathcal{Y}_2 = [1 : |\mathcal{Y}_2|]$. Note that

$$\begin{aligned} I(W; Y_2) &= H(Y_2) - H(Y_2|W), \\ I(X; Y_1) &= I(W, X; Y_1) && (W \rightarrow X \rightarrow Y_1 \text{ is Markov}) \\ &= I(X; Y_1|W) + I(W; Y_1) && (\text{chain rule}) \\ &= I(X; Y_1|W) + H(Y_1) - H(Y_1|W). \end{aligned}$$

Thus, to obtain the required cardinality bounds, it suffices to show that the joint distribution $p(w, x)$ can be replaced such that the cardinality bounds are satisfied while $H(Y_1)$, $H(Y_2)$, $H(Y_1|W)$, $H(Y_2|W)$, $I(X; Y_1|W)$ are preserved. This is done by a standard application of the

Fenchel–Bunt extension of Carathéodory’s theorem to the images of the set of all distributions on $\mathcal{X}$ under the continuous functions

$$\begin{aligned} p(x) &\mapsto (\Pr(X=1), \dots, \Pr(X=|\mathcal{X}|-1), H(Y_1), H(Y_2), I(X; Y_1)), \\ p(x) &\mapsto \left(\{\Pr(Y_1 = y_1, Y_2 = y_2)\}_{(y_1, y_2) \in (\mathcal{Y}_1 \times \mathcal{Y}_2) \setminus (1,1)}, H(Y_1), H(Y_2), I(X; Y_1) \right). \end{aligned}$$

■

Example 4.19. (BEC/BSC thresholds) Consider the broadcast channel specified by

$$\begin{aligned} p_{Y_1|X} &= \text{BEC}(\epsilon), \\ p_{Y_2|X} &= \text{BSC}(p), \end{aligned}$$

where $0 \leq \epsilon \leq 1$ and $0 \leq p \leq \frac{1}{2}$, with $X \sim \text{Bern}(1/2)$. Then we have the following:

- $\epsilon \leq 2p$ if and only if Y_2 is a degraded version of Y_1 .
- $\epsilon \leq 4p(1-p)$ if and only if Y_1 is less noisy than Y_2 .
- $\epsilon \leq H_2(p)$ if and only if Y_1 is more capable than Y_2 .

In particular the containments degraded $\subset$ less noisy $\subset$ more capable are strict (there are parameter ranges where Y_1 is less noisy but not degraded, etc.).

4.9 UV outer bound

We now turn from the classes where superposition coding is optimal to a general two-receiver broadcast channel. The first goal is an *outer bound* on the private-message capacity region $(0, R_1, R_2)$.

As emphasized in §1.13 of Chapter 1, the engine of every converse here is *sub-additivity*. Throughout this chapter the word denotes two distinct but related mechanisms, flagged once for clarity: (i) the *sub-additivity step* $F^{(n)} \leq \sum_i F^{(1)}$, realised by the Csiszár–Körner sum identity (a submodularity-style telescoping of conditional mutual informations) that single-letterizes an n -letter functional; and (ii) *functional sub-additivity* $F(p_{X_1 X_2}) \leq F(p_{X_1}) + F(p_{X_2})$ of a concave-envelope object, which supplies the convexity needed for the min–max theorem (Corollary 9.16). The UV bound is the lower-set

$$\{(R_1, R_2) : R_1 \leq T_1, R_2 \leq T_2, R_1 + R_2 \leq \min\{T_3, T_4\}\}$$

with

$$T_1 = I(U; Y_1), \quad T_2 = I(V; Y_2), \quad T_3 = I(U; Y_1) + I(X; Y_2|U), \quad T_4 = I(V; Y_2) + I(X; Y_1|V).$$

The sub-additive functional driving the converse is

$$F^{(n)} \triangleq I(M_1; Y_1^n) + I(X^n; Y_2^n | M_1),$$

and the sub-additivity step is the Körner–Marton telescoping inequality

$$I(M_1; Y_1^n) + I(X^n; Y_2^n | M_1) \leq \sum_{i=1}^n (I(U_i; Y_{1i}) + I(X_i; Y_{2i} | U_i)),$$

with the enlarged auxiliaries $U_i = (M_1, Y_1^{i-1}, Y_{2,i+1}^n)$. The derivation below first adds nonnegative terms (conditioning enlarges mutual information) and then cancels the added past/future cross terms exactly via the Csiszár sum identity

$$\sum_i I(Y_1^{i-1}; Y_{2i} | M_1, Y_{2,i+1}^n) = \sum_i I(Y_{2,i+1}^n; Y_{1i} | M_1, Y_1^{i-1});$$

the cancellation is exact, so equality in fact holds for the enlarged auxiliaries, but only the inequality is used in the converse. This collapses the n -letter quantity to a sum of single-letter quantities, which is then convexified by the time-sharing variable Q .

Theorem 4.20. (UV outer bound (Körner–Marton ’79, El Gamal ’79)) If $(0, R_1, R_2)$ is achievable for the broadcast channel $p_{Y_1, Y_2 | X}$, then

$$\begin{aligned} R_1 &\leq I(U; Y_1), \\ R_2 &\leq I(V; Y_2), \\ R_1 + R_2 &\leq \min\{I(U; Y_1) + I(X; Y_2 | U), I(V; Y_2) + I(X; Y_1 | V)\} \end{aligned}$$

for some distribution $p(u, v, x)$ with $|\mathcal{U}| \leq \min\{|\mathcal{X}|, |\mathcal{Y}_1|\} + 1$ and $|\mathcal{V}| \leq \min\{|\mathcal{X}|, |\mathcal{Y}_2|\} + 1$. $\circ$

Proof. Suppose a sequence of $(0, R_1, R_2)$ codes is given. A standard application of Fano’s inequality gives

$$\begin{aligned} nR_1 &\leq I(M_1; Y_1^n) + o(n), \\ nR_2 &\leq I(M_2; Y_2^n) + o(n), \\ n(R_1 + R_2) &\leq I(M_1; Y_1^n) + I(X^n; Y_2^n | M_1) + o(n), \\ n(R_1 + R_2) &\leq I(M_2; Y_2^n) + I(X^n; Y_1^n | M_2) + o(n). \end{aligned}$$

Then,

$$\begin{aligned} I(M_1; Y_1^n) &= \sum_{i=1}^n I(M_1; Y_{1i} | Y_1^{i-1}) && \text{(chain rule)} \\ &\leq \sum_{i=1}^n (I(M_1; Y_{1i} | Y_1^{i-1}) + I(Y_{1i}; Y_1^{i-1})) \\ &= \sum_{i=1}^n I(M_1, Y_1^{i-1}; Y_{1i}) && \text{(chain rule)} \\ &\leq \sum_{i=1}^n I(M_1, Y_1^{i-1}, Y_{2,i+1}^n; Y_{1i}) \end{aligned}$$

$$= \sum_{i=1}^n I(U_i; Y_{1i}),$$

where $U_i \triangleq (M_1, Y_1^{i-1}, Y_{2(i+1)}^n)$. Similarly,

$$I(M_2; Y_2^n) \leq \sum_{i=1}^n I(V_i; Y_{2i}),$$

where $V_i \triangleq (M_2, Y_2^{i-1}, Y_{1(i+1)}^n)$. Now, $X^n \rightarrow (X_i, M_1, Y_{2(i+1)}^n) \rightarrow Y_{2i}$ is Markov, since the channel is memoryless,

$$\begin{aligned}
& I(M_1; Y_1^n) + I(X^n; Y_2^n | M_1) \\
&= \sum_{i=1}^n \left(I(M_1; Y_{1i} | Y_1^{i-1}) + I(X^n; Y_{2i} | M_1, Y_{2(i+1)}^n) \right) && \text{(chain rule)} \\
&= \sum_{i=1}^n \left(I(M_1; Y_{1i} | Y_1^{i-1}) + I(X_i; Y_{2i} | M_1, Y_{2(i+1)}^n) \right) && \text{(Markov)} \\
&\leq \sum_{i=1}^n \left(I(M_1; Y_{1i} | Y_1^{i-1}) + I(X_i, Y_1^{i-1}; Y_{2i} | M_1, Y_{2(i+1)}^n) \right) \\
&\leq \sum_{i=1}^n I(M_1; Y_{1i} | Y_1^{i-1}) + \sum_{i=1}^n I(Y_1^{i-1}; Y_{1i}) \\
&\quad + \sum_{i=1}^n I(X_i, Y_1^{i-1}; Y_{2i} | M_1, Y_{2(i+1)}^n) \\
&= \sum_{i=1}^n I(M_1, Y_1^{i-1}; Y_{1i}) + \sum_{i=1}^n I(X_i, Y_1^{i-1}; Y_{2i} | M_1, Y_{2(i+1)}^n) && \text{(chain rule)} \\
&= \sum_{i=1}^n \left(I(M_1, Y_1^{i-1}, Y_{2(i+1)}^n; Y_{1i}) - I(Y_{2(i+1)}^n; Y_{1i} | M_1, Y_1^{i-1}) \right. && \text{(chain rule)} \\
&\quad \left. + I(X_i; Y_{2i} | M_1, Y_1^{i-1}, Y_{2(i+1)}^n) + I(Y_1^{i-1}; Y_{2i} | M_1, Y_{2(i+1)}^n) \right) \\
&= \sum_{i=1}^n (I(U_i; Y_{1i}) + I(X_i; Y_{2i} | U_i)) \\
&\quad + \sum_{i=1}^n I(Y_1^{i-1}; Y_{2i} | M_1, Y_{2(i+1)}^n) - \sum_{i=1}^n I(Y_{2(i+1)}^n; Y_{1i} | M_1, Y_1^{i-1}) \\
&= \sum_{i=1}^n (I(U_i; Y_{1i}) + I(X_i; Y_{2i} | U_i)) \\
&\quad + \sum_{i=1}^n I(Y_1^{i-1}; Y_{2i} | M_1) - \sum_{i=1}^n I(Y_1^i; Y_{2(i+1)}^n | M_1) \\
&\quad \text{(Csiszár sum identity)}
\end{aligned}$$

$$= \sum_{i=1}^n (I(U_i; Y_{1i}) + I(X_i; Y_{2i}|U_i)).$$

Similarly,

$$I(M_2; Y_2^n) + I(X^n; Y_1^n|M_2) \leq \sum_{i=1}^n (I(V_i; Y_{2i}) + I(X_i; Y_{1i}|V_i)).$$

Let Q be a random variable uniformly distributed on $[1 : n]$ and let $X \triangleq X_Q$, $Y_1 \triangleq Y_{1Q}$, $Y_2 \triangleq Y_{2Q}$, $U \triangleq (U_Q, Q)$, and $V \triangleq (V_Q, Q)$. Note that

$$\begin{aligned} I(U_Q; Y_{1Q}|Q) &\leq I(U_Q; Y_{1Q}|Q) + I(Q; Y_{1Q}) \\ &= I(U_Q, Q; Y_{1Q}) && \text{(chain rule)} \\ &= I(U; Y_1), \end{aligned}$$

and similarly

$$I(V_Q; Y_{2Q}|Q) \leq I(V; Y_2).$$

We have

$$\begin{aligned} R_1 &\leq I(U_Q; Y_{1Q}|Q) + o(1) \\ &\leq I(U; Y_1) + o(1), \\ R_2 &\leq I(V_Q; Y_{2Q}|Q) + o(1) \\ &\leq I(V; Y_2) + o(1), \\ R_1 + R_2 &\leq I(U_Q; Y_{1Q}|Q) + I(X_Q; Y_{2Q}|U_Q, Q) + o(1) \\ &\leq I(U; Y_1) + I(X; Y_2|U) + o(1), \\ R_1 + R_2 &\leq I(V_Q; Y_{2Q}|Q) + I(X_Q; Y_{1Q}|V_Q, Q) + o(1) \\ &\leq I(V; Y_2) + I(X; Y_1|V) + o(1). \end{aligned}$$

This shows the outer bound.

Now we bound the cardinalities of $\mathcal{U}$ and $\mathcal{V}$. Without loss of generality denote $\mathcal{X} = [1 : |\mathcal{X}|]$, $\mathcal{Y}_1 = [1 : |\mathcal{Y}_1|]$ and $\mathcal{Y}_2 = [1 : |\mathcal{Y}_2|]$. Note that

$$\begin{aligned} I(U; Y_1) &= H(Y_1) - H(Y_1|U), \\ I(V; Y_2) &= H(Y_2) - H(Y_2|V). \end{aligned}$$

Thus, to obtain the required cardinality bounds, it suffices to show that the joint distribution $p(u, v, x)$ can be replaced such that the cardinality bounds are satisfied while $H(Y_1)$, $H(Y_2)$, $H(Y_1|U)$, $H(Y_2|V)$, $I(X; Y_2|U)$, $I(X; Y_1|V)$ are preserved. The cardinality reduction is the staged support-lemma (Fenchel–Eggleston–Carathéodory) construction of Nair–El Gamal 2007 (see also the textbook treatment in El Gamal and Kim (2012)); the verification that the staged replacements preserve the six functionals simultaneously is carried out there in full, and we

sketch its structure here. It proceeds by applying the support lemma to the images of the set of all distributions on $\mathcal{X}$ under the continuous functions

$$\begin{aligned} p(x) &\mapsto (\Pr(X = 1), \dots, \Pr(X = |\mathcal{X}| - 1), I(X; Y_1), I(X; Y_2)), \\ p(x) &\mapsto (\Pr(Y_1 = 1), \dots, \Pr(Y_1 = |\mathcal{Y}_1| - 1), H(Y_1), I(X; Y_2)), \\ p(x) &\mapsto (\Pr(Y_2 = 1), \dots, \Pr(Y_2 = |\mathcal{Y}_2| - 1), H(Y_2), I(X; Y_1)). \end{aligned}$$

■

Remark 4.10 (Min–max form of the UV bound). The lower-set defining the UV bound includes the two sum-rate constraints $R_1 + R_2 \leq \min\{T_3, T_4\}$ and is *not* simply $\{(R_1, R_2) : R_i \leq T_i\}$; functional sub-additivity of the envelope alone does not automatically verify the lower-set convexity demanded by the min–max corollary (Corollary 9.16). The supporting-hyperplane form

$$\sup_p \min_{\lambda \in [0,1]} \sum_i \lambda_i T_i(p) = \min_{\lambda \in [0,1]} \sup_p \sum_i \lambda_i T_i(p),$$

with (λ_1, λ_2) the simplex weights, is established in the treatment of Nair–El Gamal 2007, where the sub-additivity of the upper concave envelope (Lemma 4.8) supplies the convexity of the two sum functionals needed to justify the interchange; we import it here as a cited computational consequence rather than re-deriving the convexity verification. The left side is the quantity bounded in the converse and the right side is the computable weighted-sum bound. ○

Remark 4.11 (Historical note). Körner–Marton ’79 formulated this upper bound in the two “one-sided” forms

$$\begin{aligned} \mathcal{R}_1 &: \{R_1 \leq I(U; Y_1), R_2 \leq I(X; Y_2), R_1 + R_2 \leq I(U; Y_1) + I(X; Y_2|U)\}, \\ \mathcal{R}_2 &: \{R_1 \leq I(X; Y_1), R_2 \leq I(V; Y_2), R_1 + R_2 \leq I(V; Y_1) + I(X; Y_2|V)\}, \end{aligned}$$

each being a union over the corresponding auxiliary. The intersection $\mathcal{R}_1 \cap \mathcal{R}_2$ is *not* equal to the UV outer bound in general: there are channels for which $\mathcal{R} \subsetneq \mathcal{R}_1 \cap \mathcal{R}_2$. The UV bound (Theorem 4.20) is implicit in the arguments of the old papers, but was only made explicit later. It gives the capacity region for more-capable channels — after the standard common-to-private-message reduction $(R_0 + R_2, R_1, 0)$ of the more-capable class treated above, which is exactly the setting in which the private-message UV region applies — and for the Gaussian vector BC with private and common messages (see below). ○

4.10 UVW outer bound

We now add a common message. The resulting outer bound, due to Nair–El Gamal (2007), is the best known single-letter outer bound for the general two-receiver broadcast channel. We state the computable form established by Nair 2011. The argument below derives the classical constraints from the code and explains their single-letterisation. The computable form is then

invoked as the separate cited result, Bound 3 of that note, with its own auxiliaries. The note [Nair \(2011\)](#) shows that the stated region is equivalent to the “New-Jersey” outer bound of Liang–Kramer–Shamai (a twelve-inequality bound with five structured auxiliary random variables), and that the structural restrictions of that bound are red herrings; we do not reproduce the New-Jersey bound here.

Theorem 4.21. (*UVW outer bound (Nair–El Gamal; computable form of [Nair \(2011\)](#))*) If (R_0, R_1, R_2) is achievable for the broadcast channel $p_{Y_1, Y_2|X}$, then

$$\begin{aligned} R_0 &\leq \min\{I(W; Y_1), I(W; Y_2)\}, \\ R_0 + R_1 &\leq I(U; Y_1|W) + \min\{I(W; Y_1), I(W; Y_2)\}, \\ R_0 + R_2 &\leq I(V; Y_2|W) + \min\{I(W; Y_1), I(W; Y_2)\}, \\ R_0 + R_1 + R_2 &\leq \min\{I(W; Y_1), I(W; Y_2)\} + I(X; Y_1|V, W) + I(V; Y_2|W), \\ R_0 + R_1 + R_2 &\leq \min\{I(W; Y_1), I(W; Y_2)\} + I(U; Y_1|W) + I(X; Y_2|U, W) \end{aligned}$$

for some distribution $p(u, v, w)p(x|u, v, w)$ with $(U, V, W) \rightarrow X \rightarrow (Y_1, Y_2)$ forming a Markov chain. Moreover, it suffices to consider $|\mathcal{W}| \leq |\mathcal{X}| + 5$ and $|\mathcal{U}|, |\mathcal{V}| \leq |\mathcal{X}| + 1$. $\circ$

Remark 4.12 (Classical form). The proof below establishes, en route, the classical constraints (with the classical-code auxiliaries used in the following derivation, which moreover satisfy $U \perp\!\!\!\perp V$ and $X = f(W, U, V)$):

$$R_0 \leq \min\{I(W; Y_1), I(W; Y_2)\}, \quad (4.1)$$

$$R_0 + R_1 \leq I(W, U; Y_1), \quad (4.2)$$

$$R_0 + R_2 \leq I(W, V; Y_2), \quad (4.3)$$

$$R_0 + R_1 + R_2 \leq I(W, U; Y_1) + I(V; Y_2|W, U), \quad (4.4)$$

$$R_0 + R_1 + R_2 \leq I(W, V; Y_2) + I(U; Y_1|W, V), \quad (4.5)$$

for some $p(u)p(v)p(w|u, v)$ and deterministic $x = f(w, u, v)$, with $|\mathcal{U}| \leq |\mathcal{X}| + 2$, $|\mathcal{V}| \leq |\mathcal{X}| + 2$. This is the form in which the UVW bound is usually stated (see, e.g., the textbook treatment); the semi-deterministic converse of §4.12 uses it directly. $\circ$

Proof. *Scope of the proof.* We give the structural single-letterization that motivates the UVW bound: the Fano inequalities, the Csiszár-sum telescoping, and the assembly of the classical constraints are derived in full below. The auxiliary identification ($U = M_1$, $V = M_2$, deterministic $X = f(W, U, V)$) is read off the code; the conversion from the classical to the computable form is *cited* (Bound 3 of [Nair \(2011\)](#), with its own auxiliaries — it is not partially re-derived here), and the simultaneous cardinality reductions likewise rest on the support argument of that note; the full verification is in Nair–El Gamal 2007 and Nair’s computable-form note [Nair \(2011\)](#).

Suppose a sequence of (R_0, R_1, R_2) codes is given. A standard application of Fano’s inequality gives

$$nR_0 \leq I(M_0; Y_1^n) + o(n),$$

$$\begin{aligned}
nR_0 &\leq I(M_0; Y_2^n) + o(n), \\
n(R_0 + R_1) &\leq I(M_0, M_1; Y_1^n) + o(n), \\
n(R_0 + R_2) &\leq I(M_0, M_2; Y_2^n) + o(n), \\
n(R_0 + R_1 + R_2) &\leq I(M_0, M_1; Y_1^n) + I(M_2; Y_2^n | M_0, M_1) + o(n), \\
n(R_0 + R_1 + R_2) &\leq I(M_0, M_2; Y_2^n) + I(M_1; Y_1^n | M_0, M_2) + o(n).
\end{aligned}$$

Define the auxiliary random variables $W_i \triangleq (M_0, Y_1^{i-1}, Y_{2(i+1)}^n)$, $U_i \triangleq M_1$, $V_i \triangleq M_2$. Let Q be a random variable uniformly distributed on $[1 : n]$ and let $X \triangleq X_Q$, $Y_1 \triangleq Y_{1Q}$, $Y_2 \triangleq Y_{2Q}$, $W \triangleq (W_Q, Q)$, $U \triangleq U_Q$, and $V \triangleq V_Q$. Then,

$$\begin{aligned}
I(M_0; Y_1^n) &= \sum_{i=1}^n I(M_0; Y_{1i} | Y_1^{i-1}) && \text{(chain rule)} \\
&\leq \sum_{i=1}^n (I(M_0; Y_{1i} | Y_1^{i-1}) + I(Y_{1i}; Y_1^{i-1})) \\
&= \sum_{i=1}^n I(M_0, Y_1^{i-1}; Y_{1i}) && \text{(chain rule)} \\
&\leq \sum_{i=1}^n I(M_0, Y_1^{i-1}, Y_{2(i+1)}^n; Y_{1i}) \\
&= \sum_{i=1}^n I(W_i; Y_{1i}) \\
&= nI(W_Q; Y_{1Q} | Q) \\
&\leq n(I(W_Q; Y_{1Q} | Q) + I(Q; Y_{1Q})) \\
&= nI(W_Q, Q; Y_{1Q}) && \text{(chain rule)} \\
&= nI(W; Y_1).
\end{aligned}$$

Similarly,

$$I(M_0; Y_2^n) \leq nI(W; Y_2).$$

Also,

$$\begin{aligned}
I(M_0, M_1; Y_1^n) &= \sum_{i=1}^n I(M_0, M_1; Y_{1i} | Y_1^{i-1}) && \text{(chain rule)} \\
&\leq \sum_{i=1}^n (I(M_0, M_1; Y_{1i} | Y_1^{i-1}) + I(Y_{1i}; Y_1^{i-1})) \\
&= \sum_{i=1}^n I(M_0, M_1, Y_1^{i-1}; Y_{1i}) && \text{(chain rule)} \\
&\leq \sum_{i=1}^n I(M_0, M_1, Y_1^{i-1}, Y_{2(i+1)}^n; Y_{1i})
\end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^n I(W_i, U_i; Y_{1i}) \\
&= nI(W_Q, U_Q; Y_{1Q}|Q) \\
&\leq n(I(W_Q, U_Q; Y_{1Q}|Q) + I(Q; Y_{1Q})) \\
&= nI(W_Q, U_Q, Q; Y_{1Q}) \quad (\text{chain rule}) \\
&\leq nI(W, U; Y_1).
\end{aligned}$$

Similarly,

$$I(M_0, M_2; Y_2^n) \leq nI(W, V; Y_2).$$

The classical form above uses the code auxiliaries $U = M_1$, $V = M_2$ and the code-history cloud W ; those are the auxiliaries of the classical constraints (4.1)–(4.5) derived next. The *computable* form stated in Theorem 4.21 is a separate cited result: it is Bound 3 of Nair’s note [Nair \(2011\)](#), proved there to be equivalent to the New-Jersey outer bound of Liang–Kramer–Shamai (2008). That equivalence introduces its *own* auxiliaries (U, V, W) — not the code identifications above — and supplies all five inequalities simultaneously; it is not inferred from the classical-code auxiliaries. In particular the standalone private constraints of the New-Jersey form are statements about those constructed auxiliaries, and fail for the naive identification: for a noiseless $Y_1 = Y_2 = X$ with M_0, M_2 constant and each bit of M_1 repeated twice, the code-history cloud contains the other copy of every bit, so $I(M_1; Y_{1i} | W_i) = 0$ while $R_1 = \frac{1}{2}$ is reliably communicated. We import Bound 3 rather than re-derive the equivalence. Also,

$$\begin{aligned}
&I(M_0, M_1; Y_1^n) + I(M_2; Y_2^n | M_0, M_1) \\
&= \sum_{i=1}^n \left(I(M_0, M_1; Y_{1i} | Y_1^{i-1}) + I(M_2; Y_{2i} | M_0, M_1, Y_{2(i+1)}^n) \right) \quad (\text{chain rule}) \\
&\leq \sum_{i=1}^n \left(I(M_0, M_1; Y_{1i} | Y_1^{i-1}) + I(M_2, Y_1^{i-1}; Y_{2i} | M_0, M_1, Y_{2(i+1)}^n) \right) \\
&\leq \sum_{i=1}^n \left(I(M_0, M_1; Y_{1i} | Y_1^{i-1}) + I(Y_1^{i-1}; Y_{1i}) + I(M_2, Y_1^{i-1}; Y_{2i} | M_0, M_1, Y_{2(i+1)}^n) \right) \\
&= \sum_{i=1}^n \left(I(M_0, M_1, Y_1^{i-1}; Y_{1i}) + I(M_2, Y_1^{i-1}; Y_{2i} | M_0, M_1, Y_{2(i+1)}^n) \right) \quad (\text{chain rule}) \\
&= \sum_{i=1}^n \left(I(M_0, M_1, Y_1^{i-1}, Y_{2(i+1)}^n; Y_{1i}) - I(Y_{2(i+1)}^n; Y_{1i} | M_0, M_1, Y_1^{i-1}) \right. \\
&\quad \left. + I(M_2; Y_{2i} | M_0, M_1, Y_1^{i-1}, Y_{2(i+1)}^n) + I(Y_1^{i-1}; Y_{2i} | M_0, M_1, Y_{2(i+1)}^n) \right) \quad (\text{chain rule}) \\
&= \sum_{i=1}^n (I(W_i, U_i; Y_{1i}) + I(V_i; Y_{2i} | W_i, U_i)) \\
&\quad + \sum_{i=1}^n \left(I(Y_1^{i-1}; Y_{2i} | M_0, M_1, Y_{2(i+1)}^n) - I(Y_{2(i+1)}^n; Y_{1i} | M_0, M_1, Y_1^{i-1}) \right)
\end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^n (I(W_i, U_i; Y_{1i}) + I(V_i; Y_{2i}|W_i, U_i)) \\
&\quad + \sum_{i=1}^n I(Y_1^{i-1}; Y_{2i}^n | M_0, M_1) - \sum_{i=1}^n I(Y_1^i; Y_{2(i+1)}^n | M_0, M_1) \\
&\quad \text{(Csiszár sum identity)} \\
&= \sum_{i=1}^n (I(W_i, U_i; Y_{1i}) + I(V_i; Y_{2i}|W_i, U_i)) \\
&= n (I(W_Q, U_Q; Y_{1Q}|Q) + I(V_Q; Y_{2Q}|W_Q, U_Q, Q)) \\
&\leq n (I(W_Q, U_Q; Y_{1Q}|Q) + I(Q; Y_{1Q}) + I(V_Q; Y_{2Q}|W_Q, U_Q, Q)) \\
&= n (I(W_Q, U_Q, Q; Y_{1Q}) + I(V_Q; Y_{2Q}|W_Q, U_Q, Q)) \quad \text{(chain rule)} \\
&\leq n (I(W, U; Y_1) + I(V; Y_2|W, U)).
\end{aligned}$$

Similarly,

$$I(M_0, M_2; Y_2^n) + I(M_1; Y_1^n | M_0, M_2) \leq n (I(W, V; Y_2) + I(U; Y_1 | W, V)).$$

Assembly. The five classical constraints (4.1)–(4.5) now follow from the chains above: the first three are the Fano chains, and the two sum constraints are the two Csiszár-telescoping chains just displayed. The computable form of Theorem 4.21 is then *imported* as Bound 3 of Nair (2011) with its own auxiliaries, as stated above; the five computable inequalities are not assembled from the classical-code auxiliaries by summing private bounds (a route that fails, as the counterexample shows).

In the imported computable form one may take X to be a function of (W, U, V) without shrinking the region; the displayed X -based sum terms then follow from the two vanishing terms, which have *separate reasons*: comparing the two expansions of $I(U, X; Y_1 | V, W)$,

$$I(X; Y_1 | V, W) = I(U, X; Y_1 | V, W) = I(U; Y_1 | V, W) + I(X; Y_1 | U, V, W),$$

we have $I(U; Y_1 | X, V, W) = 0$ from *channel Markovity* (given the input X , the output Y_1 is independent of the auxiliaries through the memoryless broadcast kernel), and $I(X; Y_1 | U, V, W) = 0$ because $X = f(W, U, V)$ is a *deterministic function* of the conditioning variables. Hence $I(U; Y_1 | V, W) = I(X; Y_1 | V, W)$; the fifth term is symmetric. Finally, for the classical-code auxiliaries, $M_1 \perp\!\!\!\perp M_2$ and W_i a function of $(M_0, Y_1^{i-1}, Y_{2,i+1}^n)$ give the law $p(u)p(v)p(w|u, v)$ by ordinary factorisation (no graphical independence argument is needed here, and no claim that $U \perp\!\!\!\perp V | W$ is made), so the classical form is a special case of the theorem's $p(u, v, w)p(x|u, v, w)$ with $x = f(w, u, v)$ deterministic.

Cardinality bounds. By the Fenchel–Eggleston–Carathéodory (support) lemma it suffices to consider $|\mathcal{W}| \leq |\mathcal{X}| + 5$: one must preserve the $|\mathcal{X}| - 1$ components of $p(x)$ together with the six functionals $H(Y_1|W), H(Y_2|W), H(Y_1|U, W), H(Y_2|U, W), H(Y_1|V, W), H(Y_2|V, W)$. Conditioned on W , the support lemma gives $|\mathcal{U}| \leq |\mathcal{X}| + 1$ preserving $p(x|w)$ and the two quantities $H(Y_1|U, W = w), H(Y_2|U, W = w)$, and similarly $|\mathcal{V}| \leq |\mathcal{X}| + 1$. This is the argument of Nair (2011); it makes the outer bound computable. ■

Remark 4.13 (Sub-additivity and the min–max bridge). The two sum-rate functionals

$$I(M_0, M_1; Y_1^n) + I(M_2; Y_2^n | M_0, M_1) \quad \text{and} \quad I(M_0, M_2; Y_2^n) + I(M_1; Y_1^n | M_0, M_2)$$

are each *sub-additive* in the sense of §1.13 of Chapter 1: the Csiszár sum identity collapses each to a sum of single-letter terms $\sum_i [I(W_i, U_i; Y_{1i}) + I(V_i; Y_{2i} | W_i, U_i)]$ and $\sum_i [I(W_i, V_i; Y_{2i}) + I(U_i; Y_{1i} | W_i, V_i)]$. The minimum of the two in the sum-rate inequality is the “ $\sup_p \min_\lambda$ ” side; by the min–max theorem (Theorem 9.15 and Corollary 9.16 of the appendix, and the abstract template of §1.13) it equals the computable “ $\min_\lambda \sup_p$ ” side, i.e. the union over $p(u, v, w)$ is converted to a $\min_{\alpha \in [0,1]}$ of weighted-sum bounds — which is how the UVW bound (and the Marton sum rate, §1.13-style) is evaluated in practice. $\circ$

Remark 4.14 (Reduction to the UV bound). With no common message ($R_0 = 0$), the UVW bound of Theorem 4.21 reduces *exactly* to the UV bound of Theorem 4.20: given any (U, V, W, X) admissible for the UVW bound, set $U' \triangleq (U, W)$ and $V' \triangleq (V, W)$. Then the UVW constraints at $R_0 = 0$ imply

$$\begin{aligned} R_1 &\leq I(U; Y_1 | W) + \min\{I(W; Y_1), I(W; Y_2)\} \leq I(U; Y_1 | W) + I(W; Y_1) = I(U'; Y_1), \\ R_2 &\leq I(V; Y_2 | W) + \min\{I(W; Y_1), I(W; Y_2)\} \leq I(V; Y_2 | W) + I(W; Y_2) = I(V'; Y_2), \\ R_1 + R_2 &\leq \min\{I(W; Y_1), I(W; Y_2)\} + I(X; Y_1 | V, W) + I(V; Y_2 | W) \\ &\leq I(W; Y_2) + I(X; Y_1 | V, W) + I(V; Y_2 | W) = I(X; Y_1 | V') + I(V'; Y_2), \\ R_1 + R_2 &\leq \min\{I(W; Y_1), I(W; Y_2)\} + I(U; Y_1 | W) + I(X; Y_2 | U, W) \\ &\leq I(W; Y_1) + I(U; Y_1 | W) + I(X; Y_2 | U, W) = I(U'; Y_1) + I(X; Y_2 | U'), \end{aligned}$$

where the equalities use $I(V'; Y_2) = I(W; Y_2) + I(V; Y_2 | W)$ and $I(U'; Y_1) = I(W; Y_1) + I(U; Y_1 | W)$ (chain rule). Hence every point of the $R_0 = 0$ section of the UVW region lies in the UV region, and the reverse containment is trivial (take $W = \emptyset$). This equivalence is proved in Nair (2011). $\circ$

4.11 Marton’s inner bound

We now discuss the best known inner bound for the broadcast channel, due to Marton (1979). It is obtained by superposition coding together with *random binning* (multicoding) to correlate the two private-message codewords through a common cloud center.

Theorem 4.22. (Marton’s inner bound) For a broadcast channel $p_{Y_1, Y_2 | X}$, any rate tuple (R_0, R_1, R_2) such that

$$\begin{aligned} R_0 + R_1 &< I(W, U; Y_1), \\ R_0 + R_2 &< I(W, V; Y_2), \\ R_0 + R_1 + R_2 &< \min\{I(W; Y_1), I(W; Y_2)\} + I(U; Y_1 | W) + I(V; Y_2 | W) - I(U; V | W), \\ 2R_0 + R_1 + R_2 &< I(W; Y_1) + I(W; Y_2) + I(U; Y_1 | W) + I(V; Y_2 | W) - I(U; V | W), \end{aligned}$$

for some distribution $p(w, u, v)$ and deterministic function $x = f(w, u, v)$, is achievable. This region is called Marton's achievable region. $\circ$

Lemma 4.23. (Mutual covering lemma) Let $p(w, u, v)$ be a distribution on $\mathcal{W} \times \mathcal{U} \times \mathcal{V}$. Let $W^n, U^n(\ell_1)$ ($\ell_1 \in [1 : 2^{nL_1}]$) and $V^n(\ell_2)$ ($\ell_2 \in [1 : 2^{nL_2}]$) be random variables distributed according to

$$\left(\prod_{i=1}^n p_W(w_i) \right) \left(\prod_{\ell_1} \prod_{i=1}^n p_{U|W}(u_i(\ell_1)|w_i) \right) \left(\prod_{\ell_2} \prod_{i=1}^n p_{V|W}(v_i(\ell_2)|w_i) \right).$$

Suppose that $L_1, L_2 > 0$ and $L_1 + L_2 > I(U; V|W)$. Then, for $\epsilon > 0$ small enough,

$$\Pr \left((W^n, U^n(\ell_1), V^n(\ell_2)) \notin \mathcal{T}_\epsilon^{(n)}(W, U, V) \text{ for all } \ell_1, \ell_2 \right) \rightarrow 0$$

as $n \rightarrow \infty$. $\circ$

Proof. It suffices to show the case where W is constant. The general case can be shown by conditioning on $W^n \in \mathcal{T}_\epsilon^{(n)}(W)$ and using the law of large numbers.

Let

$$A(\ell_1, \ell_2) \triangleq \begin{cases} 1, & \text{if } (U^n(\ell_1), V^n(\ell_2)) \in \mathcal{T}_\epsilon^{(n)}(U, V), \\ 0, & \text{otherwise,} \end{cases}$$

$$A \triangleq \sum_{\ell_1, \ell_2} A(\ell_1, \ell_2).$$

Note that $A(\ell_1, \ell_2)$ and $A(\ell'_1, \ell'_2)$ are independent if $\ell_1 \neq \ell'_1$ and $\ell_2 \neq \ell'_2$. Consider

$$\begin{aligned} \mathbb{E}[A] &= \sum_{\ell_1, \ell_2} \mathbb{E}[A(\ell_1, \ell_2)] \\ &= 2^{n(L_1+L_2)} \cdot 2^{-n(I(U;V)+O(\epsilon))+o(n)}, \\ \mathbb{E}[A^2] &= \sum_{\ell_1, \ell_2} \mathbb{E}[A(\ell_1, \ell_2)] + \sum_{\substack{\ell_1, \ell_2, \ell'_2 \\ \ell'_2 \neq \ell_2}} \mathbb{E}[A(\ell_1, \ell_2)A(\ell_1, \ell'_2)] + \sum_{\substack{\ell_1, \ell_2, \ell'_1 \\ \ell'_1 \neq \ell_1}} \mathbb{E}[A(\ell_1, \ell_2)A(\ell'_1, \ell_2)] \\ &\quad + \sum_{\substack{\ell_1, \ell_2, \ell'_1, \ell'_2 \\ \ell'_1 \neq \ell_1, \ell'_2 \neq \ell_2}} \mathbb{E}[A(\ell_1, \ell_2)A(\ell'_1, \ell'_2)] \\ &\leq \mathbb{E}[A] + 2^{n(L_1+2L_2)} \mathbb{E}[A(1, 1)A(1, 2)] + 2^{n(2L_1+L_2)} \mathbb{E}[A(1, 1)A(2, 1)] + \mathbb{E}[A]^2 \\ &\leq \mathbb{E}[A] + 2^{n(L_1+2L_2)} \cdot 2^{-n(2I(U;V)+O(\epsilon))+o(n)} + 2^{n(2L_1+L_2)} \cdot 2^{-n(2I(U;V)+O(\epsilon))+o(n)} + \mathbb{E}[A]^2. \end{aligned}$$

Then, by Chebyshev's inequality, we have

$$\begin{aligned} \Pr(A = 0) &\leq \Pr(|A - \mathbb{E}[A]| \geq \mathbb{E}[A]) \\ &\leq \frac{\mathbb{E}[A^2] - \mathbb{E}[A]^2}{\mathbb{E}[A]^2} \end{aligned}$$

$$\leq 2^{-n(L_1+L_2-I(U;V)+O(\epsilon))+o(n)} + 2^{-n(L_1+O(\epsilon))+o(n)} + 2^{-n(L_2+O(\epsilon))+o(n)}.$$

Taking $\epsilon > 0$ small enough and then letting $n \rightarrow \infty$ gives $\Pr(A = 0) \rightarrow 0$. $\blacksquare$

Proof of Marton's inner bound. We will first show the achievability of a rate-split region by random coding, random binning and joint typicality decoding. The required achievable region will then be obtained using Fourier–Motzkin elimination.

Rate splitting: Each private message M_j ($j = 1, 2$) is split into two independent parts, namely the common part $M_{0j} \in [1 : 2^{nR_{0j}}]$ and the private part $M_{jj} \in [1 : 2^{nR_{jj}}]$. Thus, the private message rates are split as

$$\begin{aligned} R_1 &= R_{01} + R_{11}, \\ R_2 &= R_{02} + R_{22}. \end{aligned}$$

Define the combined common message $\tilde{M}_0$ and the new private messages $\tilde{M}_1, \tilde{M}_2$ as

$$\begin{aligned} \tilde{M}_0 &\triangleq (M_0, M_{01}, M_{02}), \\ \tilde{M}_1 &\triangleq M_{11}, \\ \tilde{M}_2 &\triangleq M_{22}, \end{aligned}$$

with rates

$$\begin{aligned} \tilde{R}_0 &\triangleq R_0 + R_{01} + R_{02}, \\ \tilde{R}_1 &\triangleq R_{11}, \\ \tilde{R}_2 &\triangleq R_{22}, \end{aligned}$$

respectively. We want to show the achievability of $(\tilde{R}_0, \tilde{R}_1, \tilde{R}_2)$ such that

$$\begin{aligned} \tilde{R}_0 + \tilde{R}_1 + L_1 &< I(W, U; Y_1), \\ \tilde{R}_1 + L_1 &< I(U; Y_1|W), \\ \tilde{R}_0 + \tilde{R}_2 + L_2 &< I(W, V; Y_2), \\ \tilde{R}_2 + L_2 &< I(V; Y_2|W), \\ L_1 + L_2 &> I(U; V|W), \\ L_1 &> 0, \\ L_2 &> 0, \end{aligned}$$

for some L_1, L_2 .

Random codebook generation and random binning:

- Generate $W^n(\tilde{m}_0)$ according to

$$\prod_{\tilde{m}_0} \prod_{i=1}^n p_W(w_i(\tilde{m}_0)).$$

- Generate $U^n(\tilde{m}_0, \tilde{m}_1, \ell_1)$, where $\ell_1 \in [1 : 2^{nL_1}]$, according to

$$\prod_{\tilde{m}_0, \tilde{m}_1, \ell_1} \prod_{i=1}^n p_{U|W}(u_i(\tilde{m}_0, \tilde{m}_1, \ell_1) | w_i(\tilde{m}_0)).$$

- Generate $V^n(\tilde{m}_0, \tilde{m}_2, \ell_2)$, where $\ell_2 \in [1 : 2^{nL_2}]$, according to

$$\prod_{\tilde{m}_0, \tilde{m}_2, \ell_2} \prod_{i=1}^n p_{V|W}(v_i(\tilde{m}_0, \tilde{m}_2, \ell_2) | w_i(\tilde{m}_0)).$$

- For each $(\tilde{m}_0, \tilde{m}_1, \tilde{m}_2)$, find a pair (ℓ_1, ℓ_2) such that

$$(W^n(\tilde{m}_0), U^n(\tilde{m}_0, \tilde{m}_1, \ell_1), V^n(\tilde{m}_0, \tilde{m}_2, \ell_2)) \in \mathcal{T}_\epsilon^{(n)}(W, U, V),$$

if such pair exists, or choose $(\ell_1, \ell_2) = (1, 1)$ otherwise. Generate $X^n(\tilde{m}_0, \tilde{m}_1, \tilde{m}_2)$ according to

$$\prod_{\tilde{m}_0, \tilde{m}_1, \tilde{m}_2} \prod_{i=1}^n p_{X|W, U, V}(x_i | w_i(\tilde{m}_0), u_i(\tilde{m}_0, \tilde{m}_1, \ell_1), v_i(\tilde{m}_0, \tilde{m}_2, \ell_2)).$$

Encoding: The encoder maps any message tuple $(\tilde{m}_0, \tilde{m}_1, \tilde{m}_2)$ to the codeword $X^n(\tilde{m}_0, \tilde{m}_1, \tilde{m}_2)$.

Decoding: Decoder j ($j = 1, 2$) maps any received codeword y_j^n to the unique pair $(\hat{m}_{0j}, \hat{m}_j)$ such that

$$(W^n(\hat{m}_{0j}), U_j^n(\hat{m}_{0j}, \hat{m}_j, \ell_j), y_j^n) \in \mathcal{T}_\epsilon^{(n)}(W, U_j, Y_j)$$

for some ℓ_j , if such pair exists, or some predefined default message pair $(\hat{m}_{0j}, \hat{m}_j) = (1, 1)$ otherwise, where $U_1 \triangleq U$ and $U_2 \triangleq V$.

Probability of error analysis: The error event $(\hat{M}_{01}, \hat{M}_1) \neq (\tilde{M}_0, \tilde{M}_1) \vee (\hat{M}_{02}, \hat{M}_2) \neq (\tilde{M}_0, \tilde{M}_2)$ occurs only if at least one of the following events have occurred:

- $\mathcal{E}_{\text{enc}}: (W^n(\tilde{M}_0), U^n(\tilde{M}_0, \tilde{M}_1, \ell_1), V^n(\tilde{M}_0, \tilde{M}_2, \ell_2)) \notin \mathcal{T}_\epsilon^{(n)}(W, U, V)$ for all ℓ_1, ℓ_2 .
- $\mathcal{E}_{j'}: (W^n(\tilde{M}_0), U_j^n(\tilde{M}_0, \tilde{M}_j, \ell_j), Y_j^n) \notin \mathcal{T}_\epsilon^{(n)}(W, U_j, Y_j)$ for all ℓ_j .
- $\mathcal{E}_{0j}: (W^n(\hat{m}_{0j}), U_j^n(\hat{m}_{0j}, \hat{m}_j, \ell_j), Y_j^n) \in \mathcal{T}_\epsilon^{(n)}(W, U_j, Y_j)$ for some $\hat{m}_{0j} \neq \tilde{M}_0$, $\hat{m}_j$ and ℓ_j .
- $\mathcal{E}_j: (W^n(\tilde{M}_0), U_j^n(\tilde{M}_0, \hat{m}_j, \ell_j), Y_j^n) \in \mathcal{T}_\epsilon^{(n)}(W, U_j, Y_j)$ for some $\hat{m}_j \neq \tilde{M}_j$ and ℓ_j .

By the mutual covering lemma, the probability of unsuccessful random binning in the codebook generation step is

$$\Pr(\mathcal{E}_{\text{enc}}) \rightarrow 0$$

as $n \rightarrow \infty$.

Note that, by symmetry, the conditional probability of each of the above events under $(\tilde{M}_0, \tilde{M}_j) = (\tilde{m}_0, \tilde{m}_j)$ does not depend on the choice of $(\tilde{m}_0, \tilde{m}_j)$. Also note that $(W^n(\hat{m}_{0j}), U_j^n(\hat{m}_{0j}, \hat{m}_j, \ell_j), Y_j^n)$ conditioned on $(\tilde{M}_0, \tilde{M}_j) = (\tilde{m}_0, \tilde{m}_j)$ is independent across $i \in [1 : n]$ and has joint distribution summarized in the following table:

Event	$\hat{m}_{0j}$	$\hat{m}_j$	Joint distribution
$\mathcal{E}_{j'}$	$= \tilde{m}_0$	$= \tilde{m}_j$	$p(w, u_j, y_j)$
$\mathcal{E}_{0j}$	$\neq \tilde{m}_0$		$p(w, u_j)p(y_j)$
$\mathcal{E}_j$	$= \tilde{m}_0$	$\neq \tilde{m}_j$	$p(w, u_j)p(y_j w)$

The law of large numbers gives

$$\Pr(\mathcal{E}_{j'}) \rightarrow 0$$

as $n \rightarrow \infty$. Also,

$$\begin{aligned}
\Pr(\mathcal{E}_{0j}) &\leq 2^{-n(D(p(w, u_j, y_j) \| p(w, u_j)p(y_j)) - \tilde{R}_0 - \tilde{R}_j - L_j) + o(n)} \\
&= 2^{-n(I(W, U_j; Y_j) - \tilde{R}_0 - \tilde{R}_j - L_j) + o(n)} \\
&\rightarrow 0, & (\tilde{R}_0 + \tilde{R}_j + L_j < I(W, U_j; Y_j)) \\
\Pr(\mathcal{E}_j) &\leq 2^{-n(D(p(w, u_j, y_j) \| p(w, u_j)p(y_j|w)) - \tilde{R}_j - L_j) + o(n)} \\
&= 2^{-n(I(U_j; Y_j|W) - \tilde{R}_j - L_j) + o(n)} \\
&\rightarrow 0, & (\tilde{R}_j + L_j < I(U_j; Y_j|W))
\end{aligned}$$

as $n \rightarrow \infty$. Therefore, the union bound gives

$$\Pr((\hat{M}_{01}, \hat{M}_1) \neq (\tilde{M}_0, \tilde{M}_1) \vee (\hat{M}_{02}, \hat{M}_2) \neq (\tilde{M}_0, \tilde{M}_2)) \rightarrow 0$$

as $n \rightarrow \infty$.

Fourier–Motzkin elimination: We have shown the achievability of $(\tilde{R}_0, \tilde{R}_1, \tilde{R}_2)$ such that

$$\begin{aligned}
\tilde{R}_0 + \tilde{R}_1 + L_1 &< I(W, U; Y_1), \\
\tilde{R}_1 + L_1 &< I(U; Y_1|W), \\
\tilde{R}_0 + \tilde{R}_2 + L_2 &< I(W, V; Y_2), \\
\tilde{R}_2 + L_2 &< I(V; Y_2|W), \\
L_1 + L_2 &> I(U; V|W), \\
L_1 &> 0, \\
L_2 &> 0,
\end{aligned}$$

for some L_1, L_2 . We will use Fourier–Motzkin elimination to reduce this region to the required form. The elimination of the binning rates L_1, L_2 and of the split rates R_{01}, R_{02} below is routine

Fourier–Motzkin elimination; its output — the simplified four-inequality common-message form of Proposition 4.24 — is the standard simplification of the Marton region (see Chong et al. (2008) and the Marton-region evaluation of Anantharam et al. (2019)). These inequalities can be rewritten as

$$\begin{aligned} L_1 &< \min\{I(W; Y_1), \tilde{R}_0\} + I(U; Y_1|W) - \tilde{R}_0 - \tilde{R}_1, \\ L_2 &< \min\{I(W; Y_2), \tilde{R}_0\} + I(V; Y_2|W) - \tilde{R}_0 - \tilde{R}_2, \\ L_1 + L_2 &> I(U; V|W), \\ L_1 &> 0, \\ L_2 &> 0, \end{aligned}$$

for some L_1, L_2 . To eliminate L_1, L_2 , we notice that the necessary and sufficient conditions for such L_1, L_2 to exist are given by

$$\begin{aligned} 0 &< \min\{I(W; Y_1), \tilde{R}_0\} + I(U; Y_1|W) - \tilde{R}_0 - \tilde{R}_1, \\ 0 &< \min\{I(W; Y_2), \tilde{R}_0\} + I(V; Y_2|W) - \tilde{R}_0 - \tilde{R}_2, \\ I(U; V|W) &< \min\{I(W; Y_1), \tilde{R}_0\} + I(U; Y_1|W) - \tilde{R}_0 - \tilde{R}_1 \\ &\quad + \min\{I(W; Y_2), \tilde{R}_0\} + I(V; Y_2|W) - \tilde{R}_0 - \tilde{R}_2. \end{aligned}$$

This is equivalent to

$$\begin{aligned} \tilde{R}_0 + \tilde{R}_1 &< I(W, U; Y_1), \\ \tilde{R}_1 &< I(U; Y_1|W), \\ \tilde{R}_0 + \tilde{R}_2 &< I(W, V; Y_2), \\ \tilde{R}_2 &< I(V; Y_2|W), \\ \tilde{R}_1 + \tilde{R}_2 &< I(U; Y_1|W) + I(V; Y_2|W) - I(U; V|W), \\ \tilde{R}_0 + \tilde{R}_1 + \tilde{R}_2 &< \min\{I(W; Y_1), I(W; Y_2)\} + I(U; Y_1|W) + I(V; Y_2|W) - I(U; V|W), \\ 2\tilde{R}_0 + \tilde{R}_1 + \tilde{R}_2 &< I(W; Y_1) + I(W; Y_2) + I(U; Y_1|W) + I(V; Y_2|W) - I(U; V|W). \end{aligned}$$

Note that

$$\begin{aligned} \tilde{R}_0 &= R_0 + R_{01} + R_{02}, \\ \tilde{R}_1 &= R_1 - R_{01}, \\ \tilde{R}_2 &= R_2 - R_{02}. \end{aligned}$$

Hence, (R_0, R_1, R_2) is achievable if the set of R_{01}, R_{02} such that

$$\begin{aligned} R_{01} &< \min\{I(W, V; Y_2) - R_0 - R_2, R_1\}, \\ R_{02} &< \min\{I(W, U; Y_1) - R_0 - R_1, R_2\}, \\ R_{01} + R_{02} &< I(W; Y_1) + I(W; Y_2) + I(U; Y_1|W) + I(V; Y_2|W) - I(U; V|W) - 2R_0 - R_1 - R_2, \\ R_{01} &> \max\{R_1 - I(U; Y_1|W), 0\}, \end{aligned}$$

$$R_{02} > \max\{R_2 - I(V; Y_2|W), 0\},$$

$$R_{01} + R_{02} > R_1 + R_2 - I(U; Y_1|W) - I(V; Y_2|W) + I(U; V|W),$$

is nonempty. It can be verified that this set is nonempty if the inequalities of Marton's achievable region and

$$I(U; Y_1|W) + I(V; Y_2|W) > I(U; V|W)$$

are satisfied. If we have

$$I(U; Y_1|W) + I(V; Y_2|W) \leq I(U; V|W)$$

instead, then one of the inequalities of Marton's achievable region gives

$$\begin{aligned} R_0 + R_1 + R_2 &< \min\{I(W; Y_1), I(W; Y_2)\} \\ &\quad + I(U; Y_1|W) + I(V; Y_2|W) - I(U; V|W) \\ &\leq \min\{I(W; Y_1), I(W; Y_2)\} \\ &\leq \min\{I(X; Y_1), I(X; Y_2)\}, \end{aligned} \quad (W \rightarrow X \rightarrow (Y_1, Y_2) \text{ is Markov})$$

and we already know that this implies that (R_0, R_1, R_2) is achievable. $\blacksquare$

4.11.1 Simplification of Marton's region

Proposition 4.24. Marton's achievable region (Theorem 4.22) is the same as the set of rate tuples (R_0, R_1, R_2) such that

$$\begin{aligned} R_0 &< \min\{I(W; Y_1), I(W; Y_2)\}, \\ R_0 + R_1 &< I(W, U; Y_1), \\ R_0 + R_2 &< I(W, V; Y_2), \\ R_0 + R_1 + R_2 &< \min\{I(W; Y_1), I(W; Y_2)\} + I(U; Y_1|W) + I(V; Y_2|W) - I(U; V|W), \end{aligned}$$

for some distribution $p(w, u, v)$ and deterministic $x = f(w, u, v)$, where $|\mathcal{W}| \leq |\mathcal{X}| + 4$, $|\mathcal{U}| \leq |\mathcal{X}|$ and $|\mathcal{V}| \leq |\mathcal{X}|$. $\circ$

Proof. ($\subseteq$): If (R_0, R_1, R_2) satisfies the four simplified inequalities, then

$$\begin{aligned} 2R_0 + R_1 + R_2 &= R_0 + (R_0 + R_1 + R_2) \\ &< \min\{I(W; Y_1), I(W; Y_2)\} + \min\{I(W; Y_1), I(W; Y_2)\} \\ &\quad + I(U; Y_1|W) + I(V; Y_2|W) - I(U; V|W) \\ &\leq I(W; Y_1) + I(W; Y_2) + I(U; Y_1|W) + I(V; Y_2|W) - I(U; V|W), \end{aligned}$$

and the remaining three inequalities coincide, so the point is in Marton's region.

($\supseteq$): We must show that the inequality $2R_0 + R_1 + R_2 < I(W; Y_1) + I(W; Y_2) + I(U; Y_1|W) + I(V; Y_2|W) - I(U; V|W)$ can be replaced by $R_0 < \min\{I(W; Y_1), I(W; Y_2)\}$. This is the standard

reduction of El Gamal–Kim (Thm. 8.2 of [El Gamal and Kim \(2012\)](#); see also [Chong et al. \(2008\)](#)); we sketch it. Marton’s region is convex, so it suffices to check the extreme points. The argument proceeds by the following two observations.

(1) *Move V into the cloud when receiver 1 decodes it faster.* If $I(V; Y_1|W) \geq I(V; Y_2|W)$, then receiver 1 can decode V at least as well as receiver 2, so the private message M_2 can be treated as common. Replacing W by (V, W) and setting $V = \emptyset$ does not shrink any of the three surviving constraints (one checks $I(V, W; Y_1)$ and $I(V, W; Y_2)$ dominate the old common terms), and reduces the region to superposition coding for receiver 1.

(2) *Rebalance by time-sharing.* It remains to consider $I(V; Y_1|W) < I(V; Y_2|W)$. If $I(W; Y_1) > I(W; Y_2)$ and V is non-trivial, time-share the two identifications

Probability	Identification	V	U
ϵ	$W \rightarrow (V, W)$	$\emptyset$	U
$1 - \epsilon$	$W \rightarrow W$	V	U

with the chain-rule expansions

$$I(V, W; Y_1) = I(W; Y_1) + I(V; Y_1|W),$$

$$I(V, W; Y_2) = I(W; Y_2) + I(V; Y_2|W),$$

and fine-tune ϵ to improve the region until R_0 is constrained only by $\min\{I(W; Y_1), I(W; Y_2)\}$. Repeating this, one reaches one of the three canonical cases — $I(W; Y_1) > I(W; Y_2)$ with $U = X, V = \emptyset$; $I(W; Y_2) > I(W; Y_1)$ with $V = X, U = \emptyset$; or $I(W; Y_1) = I(W; Y_2)$ with arbitrary U, V — in each of which $2R_0 + R_1 + R_2 < \dots$ is implied by $R_0 < \min\{I(W; Y_1), I(W; Y_2)\}$ plus the sum-rate inequality. The cardinality bounds follow by the perturbation method (see the sum-rate section below). ■

4.12 Deterministic and semi-deterministic broadcast channels

- Definition 4.25.**
1. A broadcast channel $p(y_1, y_2|x)$ is *semi-deterministic* if one of the outputs is a deterministic function of the input, say $Y_2 = f(X)$.
 2. A broadcast channel is *fully deterministic* if both outputs are deterministic functions of the input, $Y_1 = f_1(X)$ and $Y_2 = f_2(X)$.
 3. A pair of *reversely degraded* broadcast channels $p(y_{11}, y_{21}|x_1)$ and $p(y_{12}, y_{22}|x_2)$ is such that $X_1 \rightarrow Y_{11} \rightarrow Y_{21}$ and $X_2 \rightarrow Y_{22} \rightarrow Y_{12}$. The product is a broadcast channel with $X = (X_1, X_2)$, $Y_1 = (Y_{11}, Y_{12})$, $Y_2 = (Y_{21}, Y_{22})$.

○

Theorem 4.26. (Capacity region of the semi-deterministic BC) For the semi-deterministic broadcast channel with $Y_2 = f(X)$, the capacity region is the set of rate triples (R_0, R_1, R_2) such

that

$$\begin{aligned} R_0 &\leq \min\{I(W; Y_1), I(W; Y_2)\}, \\ R_0 + R_1 &\leq I(W, U; Y_1), \\ R_0 + R_2 &\leq H(Y_2), \\ R_0 + R_1 + R_2 &\leq I(W, U; Y_1) + H(Y_2|W, U) \end{aligned}$$

for some distribution $p(w, u, x)$ with $(W, U) \rightarrow X \rightarrow (Y_1, Y_2)$ forming a Markov chain. The auxiliary pair need *not* determine the input: a residual stochastic law $p(x | w, u)$ is permitted, and is essential – with a deterministic $x = f(w, u)$ the term $H(Y_2|W, U)$ would vanish identically (since $Y_2 = f_2(X)$ is then a function of (W, U)) and the region would exclude the private message of a noiseless $Y_2 = X$ (for which $(R_0, R_1, R_2) = (0, 0, 1)$ is achievable). $\circ$

Proof. Achievability: By Marton’s inner bound with $V = Y_2$ (valid because $Y_2 = f(X)$ is a deterministic function of X). Then

$$\begin{aligned} I(W, V; Y_2) &= I(W, Y_2; Y_2) = H(Y_2), \\ I(V; Y_2|W) - I(U; V|W) &= I(Y_2; Y_2|W) - I(U; Y_2|W) = H(Y_2|W) - I(U; Y_2|W) = H(Y_2|W, U). \end{aligned}$$

Marton’s (simplified) region then reads

$$\begin{aligned} R_0 &< \min\{I(W; Y_1), I(W; Y_2)\}, \\ R_0 + R_1 &< I(W, U; Y_1), \\ R_0 + R_2 &< H(Y_2), \\ R_0 + R_1 + R_2 &< \min\{I(W; Y_1), I(W; Y_2)\} + I(U; Y_1|W) + H(Y_2|W, U). \end{aligned}$$

The simplification of the sum-rate term to $I(W, U; Y_1) + H(Y_2|W, U)$ is the standard auxiliary-reduction step of the semi-deterministic capacity theorem (Marton 1979; the textbook treatment El Gamal and Kim (2012), Thm. 8.4). We cite that reduction rather than state a one-sided “rebalancing” lemma: a coarsening $\tilde{W}$ of W that keeps the satellite U fixed cannot in general preserve $I(\tilde{W}, U; Y_1) = I(W, U; Y_1)$ while lowering $I(\tilde{W}; Y_1)$ to $I(W; Y_2)$ (take binary X with $W = X$, constant U , $Y_1 = X$, Y_2 constant: the demanded $I(\tilde{W}; Y_1) = 0$ contradicts the preserved $I(\tilde{W}, U; Y_1) = 1$). The valid reduction may jointly redefine the auxiliaries, and with it the achievable region is exactly the region in the theorem.

Converse: By the classical form of the UVW outer bound (the constraints (4.1)–(4.5) derived in the proof of Theorem 4.21), with $p(u)p(v)p(w|u, v)$ and $x = f(w, u, v)$:

$$\begin{aligned} R_0 &\leq \min\{I(W; Y_1), I(W; Y_2)\}, \\ R_0 + R_1 &\leq I(W, U; Y_1), \\ R_0 + R_2 &\leq I(W, V; Y_2) \leq H(Y_2), \\ R_0 + R_1 + R_2 &\leq I(W, U; Y_1) + I(V; Y_2|W, U) \leq I(W, U; Y_1) + H(Y_2|W, U), \end{aligned}$$

where the last two lines use $I(W, V; Y_2) \leq H(Y_2)$ and $I(V; Y_2|W, U) \leq H(Y_2|W, U)$ (both immediate since $Y_2 = f(X)$). This is exactly the region in the theorem. $\blacksquare$

Corollary 4.27. (Fully deterministic BC) For the fully deterministic broadcast channel $Y_1 = f_1(X)$, $Y_2 = f_2(X)$, the capacity region is the set of rate triples (R_0, R_1, R_2) such that

$$\begin{aligned} R_0 &\leq \min\{I(W; Y_1), I(W; Y_2)\}, \\ R_0 + R_1 &\leq H(Y_1), \\ R_0 + R_2 &\leq H(Y_2), \\ R_0 + R_1 + R_2 &\leq \min\{I(W; Y_1), I(W; Y_2)\} + H(Y_1, Y_2|W) \end{aligned}$$

for some $p(w)$. For $R_0 = 0$ this reduces to the Slepian–Wolf-type region

$$R_1 \leq H(Y_1), \quad R_2 \leq H(Y_2), \quad R_1 + R_2 \leq H(Y_1, Y_2).$$

○

Remark 4.15. For the reversely-degraded product channel, Marton’s inner bound matches the UVW outer bound, giving the capacity region. By contrast, for the product of two more-capable (or two semi-deterministic) channels, Marton’s inner bound can be strictly contained in the UVW outer bound. ○

4.13 Marton’s sum rate

We now consider the optimality of the sum rate $R_0 + R_1 + R_2$. For a broadcast channel $\mathbf{q}(y_1, y_2|x)$, define the Marton sum rate

$$SR_M(\mathbf{q}) \triangleq \max_{p(w,u,v,x)} \left\{ \min\{I(W; Y_1), I(W; Y_2)\} + I(U; Y_1|W) + I(V; Y_2|W) - I(U; V|W) \right\}, \quad (4.6)$$

over all $p(w, u, v, x)$ with $(W, U, V) \rightarrow X \rightarrow (Y_1, Y_2)$.

By the min–max theorem of Tsernys (Theorem 9.15 and Corollary 9.16 of the appendix),

$$\begin{aligned} SR_M &= \max_{p(w,u,v,x)} \left\{ \min_{\alpha \in [0,1]} \left\{ \alpha I(W; Y_1) + \bar{\alpha} I(W; Y_2) \right\} + I(U; Y_1|W) + I(V; Y_2|W) - I(U; V|W) \right\} \\ &= \min_{\alpha \in [0,1]} \left\{ \max_{p(w,u,v,x)} \left\{ \alpha I(W; Y_1) + \bar{\alpha} I(W; Y_2) + I(U; Y_1|W) + I(V; Y_2|W) - I(U; V|W) \right\} \right\} \\ &= \min_{\alpha \in [0,1]} \alpha - SR_M, \end{aligned}$$

where $\bar{\alpha} \triangleq 1 - \alpha$ and

$$\alpha - SR_M \triangleq \max_{p(w,u,v,x)} \left\{ \alpha I(W; Y_1) + \bar{\alpha} I(W; Y_2) + I(U; Y_1|W) + I(V; Y_2|W) - I(U; V|W) \right\}.$$

The functional $\alpha - SR_M$ is convex in α , and the extreme values $0 - SR_M$ and $1 - SR_M$ reduce to the sum-rate constraints of the UVW outer bound. The single-letter Marton region is optimal (gives capacity) if and only if

$$SR_M(\mathbf{q} \otimes \mathbf{q}) = 2 SR_M(\mathbf{q}) \quad \text{for every } \mathbf{q},$$

equivalently, if and only if SR_M equals the sum capacity. This is open in general; there exist channels q_1, q_2 (products of more-capable/semi-deterministic components) with $SR_M(q_1 \otimes q_2) > SR_M(q_1) + SR_M(q_2)$.

4.13.1 Cardinality bound via the perturbation method

Lemma 4.28. To evaluate $SR_M(q)$, it suffices to consider auxiliary random variables with $|\mathcal{W}|, |\mathcal{U}|, |\mathcal{V}| \leq |\mathcal{X}|$. $\circ$

Proof. By Carathéodory's theorem, $|\mathcal{W}| \leq |\mathcal{X}|$ suffices, since we need to preserve the $|\mathcal{X}| - 1$ components of $p(x)$ together with the single functional

$$-\alpha H(Y_1|W) - \bar{\alpha} H(Y_2|W) + I(U; Y_1|W) + I(V; Y_2|W) - I(U; V|W).$$

It remains to bound $|\mathcal{U}|$ and $|\mathcal{V}|$. Consider the simplified sum rate $SR_M(q) \triangleq \sup_{p(u,v,x)} I(U; Y_1) + I(V; Y_2) - I(U; V)$ (the case $W = \emptyset$), and define $SR_A(q)$ to be the same supremum restricted to $|\mathcal{U}|, |\mathcal{V}| \leq A$. We show $SR_A = SR_{|\mathcal{X}|}$ for every $A > |\mathcal{X}|$; letting $A \rightarrow \infty$ then gives $SR_M = SR_{|\mathcal{X}|}$.

Since $|\mathcal{U}|, |\mathcal{V}| \leq A$ are finite, the objective is continuous on the compact simplex of laws $p(u, v, x)$, so a maximizer $p^*(u, v, x)$ of SR_A exists. Let $L(u)$ be a non-zero solution of

$$\sum_u p^*(u, x) L(u) = 0, \quad \forall x \in \mathcal{X}. \quad (4.7)$$

Such an L exists as long as $|\mathcal{U}| > |\mathcal{X}|$: this is a homogeneous system of $|\mathcal{X}|$ linear equations in $|\mathcal{U}|$ unknowns. Define the perturbed distribution

$$p_\epsilon(u, v, x) \triangleq p^*(u, v, x) (1 + \epsilon L(u)).$$

By the defining equations, $p_\epsilon(x) = p^*(x)$, so the perturbation preserves the input law and hence the laws of (Y_1, Y_2) . A short computation gives

$$\begin{aligned} p_\epsilon(v, y_2) &= p^*(v, y_2) (1 + \epsilon E[L|V, Y_2]), \\ p_\epsilon(u, v) &= p^*(u, v) (1 + \epsilon L(u)), \\ p_\epsilon(u, y_1) &= p^*(u, y_1) (1 + \epsilon L(u)). \end{aligned}$$

Write

$$\begin{aligned} f(\epsilon) &:= I_{p_\epsilon}(U; Y_1) + I_{p_\epsilon}(V; Y_2) - I_{p_\epsilon}(U; V) \\ &= H_{p_\epsilon}(Y_1) + H_{p_\epsilon}(Y_2) + H_{p_\epsilon}(U, V) - H_{p_\epsilon}(U, Y_1) - H_{p_\epsilon}(V, Y_2). \end{aligned}$$

Since p^* is a maximizer, $f'(0) = 0$ and $f''(0) \leq 0$. For the ϵ -dependent part

$$g(\epsilon) := H_{p_\epsilon}(U, V) - H_{p_\epsilon}(U, Y_1) - H_{p_\epsilon}(V, Y_2),$$

a direct calculation (the same as in Note 4.16 below, carried out at a general ϵ rather than at $\epsilon = 0$) gives

$$g''(\epsilon) = E_{p_\epsilon} \left[E_\epsilon [L_\epsilon(U) | V, Y_2]^2 \right] \geq 0 \quad \text{for every feasible } \epsilon,$$

where

$$L_\epsilon(u) \triangleq \frac{L(u)}{1 + \epsilon L(u)}$$

is the *score* of the perturbation at ϵ — the derivative $\frac{d}{d\epsilon} \log p_\epsilon(u, v, x)$ — and $E_\epsilon[\cdot | V, Y_2]$ is the conditional expectation under p_ϵ . (At $\epsilon = 0$ the score reduces to L and the formula reproduces $E[E[L|V, Y_2]^2]$.) The second derivative of g along the perturbation line is thus an average of squares *at every point* of the feasible interval, not merely at $\epsilon = 0$. This global sign is the key step of the perturbation method of [Anantharam et al. \(2019\)](#); a local second-derivative condition at $\epsilon = 0$ alone would not suffice. The entropies $H_{p_\epsilon}(Y_1)$ and $H_{p_\epsilon}(Y_2)$ are constant in ϵ (the perturbation preserves the input law and hence the output laws), so $f(\epsilon) = g(\epsilon) + \text{constant}$ and g is *convex* on the entire feasible interval $[\epsilon_{\min}, \epsilon_{\max}]$. Since $\epsilon = 0$ is an interior point and a maximizer of the convex function g , g — and hence f — is constant on the whole feasible interval, not merely locally. Choosing $\epsilon_{\max} > 0$ as the largest value with $1 + \epsilon_{\max} L(u) \geq 0$ for all u kills one atom ($p_{\epsilon_{\max}}(u_{\max}) = 0$), reducing $|\mathcal{U}|$ by one without changing the value. Iterating gives $|\mathcal{U}| \leq |\mathcal{X}|$; the same argument applies to $|\mathcal{V}|$. ■

Remark 4.16. (Calculus on a perturbed entropy.) For a distribution $p(a)$ and a perturbation $p(a)(1 + \epsilon L(a))$ with $\sum_a p(a)L(a) = 0$, let

$$\begin{aligned} r(\epsilon) &\triangleq - \sum_a p(a)(1 + \epsilon L(a)) \log(p(a)(1 + \epsilon L(a))) \\ &= - \sum_a p(a)(1 + \epsilon L(a)) \log p(a) - \sum_a p(a)(1 + \epsilon L(a)) \log(1 + \epsilon L(a)). \end{aligned}$$

Differentiating twice and setting $\epsilon = 0$ gives

$$\begin{aligned} r'(0) &= - \sum_a p(a)L(a) \log p(a), \\ r''(0) &= - \sum_a p(a)L(a)^2 = -E[L^2(A)]. \end{aligned}$$

The sign is the hallmark of entropy concavity: H is concave in the distribution, so the second derivative of the perturbed entropy along any direction is non-positive. ○

Remark 4.17 (Further reading on Marton's inner bound). The evaluation of Marton's region — computing the auxiliary random variables that attain its boundary — is a non-trivial optimisation problem in its own right, with a substantial literature: an information inequality that reduces the evaluation for binary-input channels is given in [Nair et al. \(2010\)](#); improved cardinality bounds on the auxiliaries in [Anantharam et al. \(2013\)](#); a conjecture on the structure of the optimisers in [Gohari et al. \(2025\)](#); and a local-tensorisation step for product channels with a binary component in [Nair \(2020\)](#). These make good optional reading after the present chapter: they exercise exactly the tools developed here — the support lemma, the perturbation method, and the min-max bridge — on a concrete computational problem. ○

4.14 Functional representation and multi-letter Marton

Lemma 4.29. (Functional representation lemma) Given $p(x, y)$ on finite alphabets, there exists a random variable Q independent of X and a function g such that $Y = g(X, Q)$. $\circ$

Proof. For each $x \in \mathcal{X}$, let $F_x(y) \triangleq \Pr(Y \leq y \mid X = x)$ be the conditional cdf of Y given $X = x$, and let

$$g(x, q) \triangleq \inf\{y : F_x(y) \geq q\}$$

be the generalized inverse (right-continuous). Let $Q \sim \text{Unif}[0, 1]$ be independent of X . The standard inverse-cdf identity gives $\Pr(g(x, Q) \leq y) = F_x(y)$ for every x, y , i.e. $g(x, Q) \sim p_{Y|X=x}$. For finite alphabets the finitely many breakpoints of the cdfs F_x discretize Q to a finite alphabet, so the lemma holds with Q taking values in a finite set. $\blacksquare$

Remark 4.18. Using the functional representation lemma, in Marton's inner bound we may, without loss of generality, take x to be a deterministic function $x = f(w, u, v)$ rather than a general channel $p(x|w, u, v)$. Indeed, write $X = g(W, U, V, Q)$ with $Q \perp (W, U, V)$, set $\tilde{W} = W$, $\tilde{U} = (U, Q)$, $\tilde{V} = V$; then $I(\tilde{U}; Y_1|W) = I(U, Q; Y_1|W) \geq I(U; Y_1|W)$ and $I(\tilde{U}; V|W) = I(U, Q; V|W) = I(U; V|W)$ (since $Q \perp (U, V, W)$), so the region is not shrunk. $\circ$

Consider the multi-letter (blocklength- k) Marton region, obtained by regarding Y_1^k, Y_2^k as the outputs of a k -fold product channel:

$$\begin{aligned} R_0 &< \frac{1}{k} \min\{I(W; Y_1^k), I(W; Y_2^k)\}, \\ R_0 + R_1 &< \frac{1}{k} (\min\{I(W; Y_1^k), I(W; Y_2^k)\} + I(U; Y_1^k|W)), \\ R_0 + R_2 &< \frac{1}{k} (\min\{I(W; Y_1^k), I(W; Y_2^k)\} + I(V; Y_2^k|W)), \\ R_0 + R_1 + R_2 &< \frac{1}{k} (\min\{I(W; Y_1^k), I(W; Y_2^k)\} + I(U; Y_1^k|W) + I(V; Y_2^k|W) - I(U; V|W)). \end{aligned}$$

The containment of the capacity region in the multi-letter Marton region is the standard multi-letter Marton outer bound, due to Marton 1979 (see also the multi-letter Marton discussion in El Gamal and Kim (2012)); we state it here as a cited inclusion rather than re-deriving it. Its starting point is the identification $W = M_0, U = M_1, V = M_2$, under which $I(M_1; M_2|M_0) = 0$ by message independence and Fano's inequality gives the three individual bounds

$$\begin{aligned} kR_0 &\leq \min\{I(M_0; Y_1^k), I(M_0; Y_2^k)\}, \\ kR_1 &\leq I(M_1; Y_1^k|M_0), \\ kR_2 &\leq I(M_2; Y_2^k|M_0); \end{aligned}$$

the fourth (sum-rate) inequality of the region is the additional, genuinely multi-letter step of the cited bound and does not follow from these three alone. Consequently the single-letter Marton region is capacity if and only if $SR_M(\mathbf{q} \otimes \mathbf{q}) = 2SR_M(\mathbf{q})$ for every channel $\mathbf{q}$ (testing

optimality reduces to checking the equality of Marton's sum rate on a single channel with that on its two-fold product; the equivalence rests on the cited multi-letter inclusion and is standard — see [Marton \(1979\)](#) and [El Gamal and Kim \(2012\)](#)).

4.15 Equivalence of average- and maximal-error capacity regions

Let P_{avg} denote the probability of *average* error (averaged over the uniformly chosen messages) and let P_{max} denote the probability of *maximal* error (supremum over messages).

Theorem 4.30. For any broadcast channel, the capacity region with respect to P_{max} equals the capacity region with respect to P_{avg} . $\circ$

Citation (Willems 1990). Maximal error trivially dominates average error, so $\mathcal{C}_{\text{max}} \subseteq \mathcal{C}_{\text{avg}}$. The reverse containment is a theorem of Willems 1990: for the broadcast channel the two capacity regions coincide. The naive argument — by Markov's inequality at most half the message triples have error above $2\epsilon_n$, so delete the bad half and re-index — is *not* a proof: an arbitrary half-density subset of the Cartesian triple-message set need not contain a large Cartesian product $\mathcal{M}'_0 \times \mathcal{M}'_1 \times \mathcal{M}'_2$, so the deletion need not leave a valid broadcast code at rates $R_j - 1/n$. The correct proof uses the single-encoder structure of the BC through a more delicate expurgation/permutation argument that preserves the product message structure; see [Willems \(1990\)](#). (The equivalence is specific to the broadcast channel: for the MAC and the interference channel, maximal-error and average-error capacity regions can differ, so this statement must not be imported there.) $\blacksquare$

4.16 Gaussian vector broadcast channel

Consider the MIMO Gaussian broadcast channel

$$\begin{aligned} \mathbf{Y}_1 &= A_1 \mathbf{X} + \mathbf{Z}_1, \\ \mathbf{Y}_2 &= A_2 \mathbf{X} + \mathbf{Z}_2, \end{aligned}$$

with channel gain matrices A_1, A_2 of compatible dimensions, independent Gaussian noise vectors $\mathbf{Z}_1, \mathbf{Z}_2$ with positive-definite covariance matrices $K_{Z_1}, K_{Z_2} \succ 0$, and a second-moment (covariance) input constraint $\mathbb{E}[\mathbf{X} \mathbf{X}^T] \preceq K$.

Theorem 4.31. (Gaussian vector BC)

1. When $R_0 = 0$, Marton's inner bound matches the UV outer bound, and the capacity region is the dirty-paper-coding region of Weingarten–Steinberg–Shamai [Weingarten et al. \(2006\)](#) (private messages).
2. For general R_0 , Marton's inner bound matches the UVW outer bound, giving the common-plus-private-message capacity region of Geng–Nair [2014](#).

$\circ$

Proof sketch (cited deep theorem). This is a deep theorem; the following is a motivational sketch, not a derivation. The full proof supplies (i) the dirty-paper-coding covariance-split construction achieving the region, and (ii) the Gaussian-optimality upgrade of the outer bounds, via the perturbation/concave-envelope machinery; both are in the cited sources. The sketch: both cases are shown by checking that the extreme points of the outer bound lie in Marton's inner bound; the second case additionally uses the sub-additivity of the upper concave envelope. The insight comes from dirty-paper coding. Suppose U and X are jointly Gaussian; a Gaussian V is then constructed (the covariance-split of dirty-paper coding, computed by linear algebra in Weingarten et al. (2006)) with $(U, V) \rightarrow X$ and

$$I(U; Y_1) + I(X; Y_2|U) = I(U; Y_1) + I(V; Y_2) - I(U; V).$$

In general (non-Gaussian U, X) this is an inequality: writing $X = f(U, V)$,

$$\begin{aligned} I(X; Y_2|U) &\geq I(V; Y_2|U) && \text{(DPI, since } X = f(U, V)\text{)} \\ &= I(V; Y_2, U) - I(U; V) \\ &= I(V; Y_2) + I(V; U|Y_2) - I(U; V) \\ &\geq I(V; Y_2) - I(U; V). \end{aligned}$$

Gaussian optimality upgrades this inequality to equality at the maximizer, which is exactly what is needed for the outer and inner bounds to meet; this upgrade is a theorem of the cited sources, not an immediate consequence of the one-line inequality. The chain of tools is the template of §1.13 in Chapter 1: sub-additivity sets up the gap, the rotation trick forces the gap to vanish at the maximizer, the double-Markovity lemma (Lemma 1.56 of Chapter 1) turns the vanishing gap into conditional independence, and Bernstein's theorem (Lemma 1.60 of Chapter 1) upgrades independence to Gaussianity. ■

Remark 4.19 (Further reading on the Gaussian broadcast channel). The equality of the private- and common-message capacity regions with the UV/UVW bounds is due to Geng and Nair (2014). A key technical component — the optimality of Gaussian inputs for the weighted-sum functionals — is obtained there by the perturbation/concave-envelope machinery of Chapter 1 (§1.13), together with the local-optimality results for log-determinant optimisation problems treated by information-theoretic methods in Lau et al. (2022). For the two-receiver *scalar* Gaussian broadcast channel these computations reduce to the closed forms derived earlier in this chapter (§4.6). ○

5 Interference channel

An interference channel models two sender–receiver pairs that share the same transmission medium. Each sender wishes to convey its own message to its own receiver, but the signal radiated by one sender reaches *both* receivers, so the two transmissions interfere with one another. The setting is depicted below.




Figure 5.1: An interference channel

5.1 Model and Definitions

Definition 5.1. (Discrete-memoryless interference channel (DM-IC)) A two-user discrete-memoryless interference channel is a tuple $(\mathcal{X}_1 \times \mathcal{X}_2, W(y_1, y_2 | x_1, x_2), \mathcal{Y}_1 \times \mathcal{Y}_2)$, where $\mathcal{X}_j$ is the input alphabet of sender j ($j = 1, 2$), $\mathcal{Y}_j$ the output alphabet of receiver j , and W a conditional distribution on $\mathcal{Y}_1 \times \mathcal{Y}_2$ given $\mathcal{X}_1 \times \mathcal{X}_2$. The channel is memoryless:

$$W^{\otimes n}(y_1^n, y_2^n | x_1^n, x_2^n) = \prod_{i=1}^n W(y_{1i}, y_{2i} | x_{1i}, x_{2i}).$$

The two senders do not cooperate and there is no feedback: sender j observes only its own message M_j , and receiver j observes only Y_j^n . $\circ$

Definition 5.2. (Code, probability of error, achievable pair, capacity region) The rates are nonnegative, $R_1, R_2 \geq 0$. A $(2^{nR_1}, 2^{nR_2}, n)$ code for the DM-IC consists of

- two message sets $[1 : 2^{nR_1}]$, $[1 : 2^{nR_2}]$, with M_1, M_2 independent and uniform;
- two encoders: $x_1^n(m_1)$ for $m_1 \in [1 : 2^{nR_1}]$ and $x_2^n(m_2)$ for $m_2 \in [1 : 2^{nR_2}]$;
- two decoders: $\hat{m}_1(y_1^n)$ and $\hat{m}_2(y_2^n)$.

The (average) probability of error is

$$P_e^{(n)} = \Pr \{(\hat{M}_1, \hat{M}_2) \neq (M_1, M_2)\}.$$

Throughout, message-set sizes are read with the ceiling convention (2^{nR_j} stands for $\lceil 2^{nR_j} \rceil$), which changes rates by $O(1/n)$ and hence not their limits. A rate pair (R_1, R_2) is *achievable* if there exists a sequence of $(2^{nR_1}, 2^{nR_2}, n)$ codes with $P_e^{(n)} \rightarrow 0$ as $n \rightarrow \infty$. The *capacity region* $\mathcal{C}$ is the closure of the set of achievable rate pairs. $\circ$

A useful first observation is that the capacity region depends only on the *marginal* channel laws, not on the joint law of (Y_1, Y_2) given (X_1, X_2) .

Lemma 5.3. The capacity region of a DM-IC (with no feedback) depends only on the marginal distributions $W(y_1|x_1, x_2)$ and $W(y_2|x_1, x_2)$. $\circ$

Proof. Decoder j observes only Y_j^n ; hence the error probability, and therefore achievability, is a function of the two marginal laws $W(y_1|x_1, x_2)$ and $W(y_2|x_1, x_2)$ alone. Formally, the joint law of (Y_1^n, Y_2^n) given the inputs never enters the expression for $P_e^{(n)}$: each decoder j only needs the conditional law of Y_j^n given (X_1^n, X_2^n) , which is determined by the j -th marginal of W . $\blacksquare$

Before turning to the best-known inner bound and the various converses, we record the organising principle that will run through the entire chapter. Recall from §1.13 of Chapter 1 that *many single-letterisation arguments in multi-user information theory amount to proving sub-additivity of an associated functional*: a functional $F^{(n)}$ on n -letter laws is sub-additive if $F^{(n)}(p_{X^n Y^n}) \leq \sum_{i=1}^n F(p_{X_i Y_i})$, in which case a Fano-type multi-letter bound $nR \leq F^{(n)} + o(n)$ collapses to the single-letter bound $R \leq \sup_p F(p)$. We shall name the sub-additive functional at the start of each converse below, and the Gaussian converses will additionally invoke Gaussian optimality (Lemma 1.59, proved via the rotation trick in Chapter 1).

5.2 Some Simple Inner and Outer Bounds

5.2.1 Time-division inner bound

The maximum achievable individual rates, when one user transmits alone and the other is silent, are

$$\begin{aligned} C_1 &= \max_{p(x_1), x_2 \in \mathcal{X}_2} I(X_1; Y_1 | X_2 = x_2), \\ C_2 &= \max_{p(x_2), x_1 \in \mathcal{X}_1} I(X_2; Y_2 | X_1 = x_1). \end{aligned} \tag{5.1}$$

Here x_2 is the “best silent symbol” for user 2: by choosing it, user 2 removes the interference it would otherwise cause at receiver 1. By *time-sharing* between the two point-to-point links, we obtain the following inner bound.

Proposition 5.4. (Time-division inner bound) All rate pairs (R_1, R_2) with

$$R_1 \leq \alpha C_1, \quad R_2 \leq \bar{\alpha} C_2, \quad (5.2)$$

for some $\alpha \in [0, 1]$ (and $\bar{\alpha} = 1 - \alpha$), are achievable. $\circ$

5.2.2 Treating interference as noise (TIN)

A better inner bound is obtained by using independent point-to-point codes for the two pairs, time-sharing, and letting each receiver *treat the other user's signal as noise*.

Proposition 5.5. (TIN inner bound) All rate pairs (R_1, R_2) satisfying

$$\begin{aligned} R_1 &\leq I(X_1; Y_1|Q), \\ R_2 &\leq I(X_2; Y_2|Q), \end{aligned} \quad (5.3)$$

for some $p(q)p(x_1|q)p(x_2|q)$, are achievable. $\circ$

Example 5.6. When the channel has no interference, i.e. $W(y_1, y_2|x_1, x_2) = W(y_1|x_1)W(y_2|x_2)$, TIN reduces to the rectangle $R_1 \leq C_1, R_2 \leq C_2$ (taking X_j the point-to-point capacity-achieving input and $Q = \emptyset$), which is the full capacity region and is strictly larger than the time-division bound whenever both $C_1, C_2 > 0$.

5.2.3 Simultaneous-decoding inner bound

At the opposite extreme, each receiver decodes *both* messages. The channel to each receiver is then a two-user multiple access channel, and the two decoders must simultaneously succeed.

Proposition 5.7. (Simultaneous-decoding inner bound) All rate pairs (R_1, R_2) satisfying

$$\begin{aligned} R_1 &\leq \min\{I(X_1; Y_1|X_2, Q), I(X_1; Y_2|X_2, Q)\}, \\ R_2 &\leq \min\{I(X_2; Y_1|X_1, Q), I(X_2; Y_2|X_1, Q)\}, \\ R_1 + R_2 &\leq \min\{I(X_1, X_2; Y_1|Q), I(X_1, X_2; Y_2|Q)\}, \end{aligned} \quad (5.4)$$

for some $p(q)p(x_1|q)p(x_2|q)$, are achievable. $\circ$

Proof. This is the DM-MAC achievability (simultaneous decoding) applied to each receiver, together with coded time-sharing through Q . Each receiver runs the MAC decoder for both messages; the three constraints per receiver are exactly the MAC region constraints, and intersecting the two receivers' MAC regions (taking the min) ensures both decoders succeed simultaneously. $\blacksquare$

5.3 Strong and Very Strong Interference

The interference channel is at its most tractable when the interference is so strong that the “cross” link carries more information than the “direct” link. There are two classical notions.

5.3.1 Very strong interference

Definition 5.8. (Very strong interference (Carleial 1975)) A DM-IC is said to have *very strong interference* if

$$\begin{aligned} I(X_1; Y_2) &\geq I(X_1; Y_1|X_2), \\ I(X_2; Y_1) &\geq I(X_2; Y_2|X_1), \end{aligned} \tag{5.5}$$

for all product input distributions $p(x_1)p(x_2)$. ○

Proposition 5.9. The capacity region of a DM-IC with very strong interference is

$$\begin{aligned} R_1 &\leq I(X_1; Y_1|X_2, Q), \\ R_2 &\leq I(X_2; Y_2|X_1, Q), \end{aligned} \tag{5.6}$$

for some $p(q)p(x_1|q)p(x_2|q)$. ○

Proof. Achievability. Use simultaneous decoding. Under very strong interference, the cross terms dominate the direct terms:

$$I(X_1; Y_2|X_2, Q) \geq I(X_1; Y_2|Q) \geq I(X_1; Y_1|X_2, Q),$$

where the first inequality uses the independence identity $I(X_1; Y_2|X_2, Q) = I(X_1; Y_2|Q) + I(X_1; X_2|Y_2, Q) \geq I(X_1; Y_2|Q)$ (since $X_1 \perp\!\!\!\perp X_2 \mid Q$) and the second is the very-strong condition averaged over Q ; symmetrically $I(X_2; Y_1|X_1, Q) \geq I(X_2; Y_2|X_1, Q)$. Hence in (5.4) the min’s reduce to the direct terms,

$$R_1 \leq I(X_1; Y_1|X_2, Q), \quad R_2 \leq I(X_2; Y_2|X_1, Q),$$

and the sum constraint is implied: since $I(X_2; Y_1|Q) \geq I(X_2; Y_2|X_1, Q)$,

$$I(X_1, X_2; Y_1|Q) = I(X_2; Y_1|Q) + I(X_1; Y_1|X_2, Q) \geq R_2 + R_1,$$

and symmetrically for Y_2 . Thus (5.4) collapses to the rectangle above.

Converse. The individual cut-set bounds give

$$nR_1 \leq I(X_1^n; Y_1^n|X_2^n) + n\varepsilon_n, \quad nR_2 \leq I(X_2^n; Y_2^n|X_1^n) + n\varepsilon_n$$

(equivalently $I(M_j; Y_j^n) \leq I(X_j^n; Y_j^n) \leq I(X_j^n; Y_j^n | X_j^n)$ by the data-processing inequality and conditioning). The functional $F = P_{X_1, X_2} \mapsto I(X_1; Y_1 | X_2)$ is sub-additive over the memoryless channel (see §1.13):

$$I(X_1^n; Y_1^n | X_2^n) \leq \sum_{i=1}^n I(X_{1i}; Y_{1i} | X_{2i}),$$

so after time-sharing with $Q \sim \text{Unif}[1 : n]$ we obtain $R_1 \leq I(X_1; Y_1 | X_2, Q)$ and $R_2 \leq I(X_2; Y_2 | X_1, Q)$. ■

5.3.2 Strong interference

Definition 5.10. (Strong interference) A DM-IC is said to have *strong interference* if

$$\begin{aligned} I(X_1; Y_2 | X_2) &\geq I(X_1; Y_1 | X_2), \\ I(X_2; Y_1 | X_1) &\geq I(X_2; Y_2 | X_1), \end{aligned} \tag{5.7}$$

for all product input distributions $p(x_1)p(x_2)$. ○

Very strong interference implies strong interference. Indeed, since $X_1 \perp\!\!\!\perp X_2$,

$$I(X_1; Y_2 | X_2) = I(X_1; Y_2) + I(X_1; X_2 | Y_2) \geq I(X_1; Y_2) \geq I(X_1; Y_1 | X_2),$$

and symmetrically

$$I(X_2; Y_1 | X_1) = I(X_2; Y_1) + I(X_2; X_1 | Y_1) \geq I(X_2; Y_1) \geq I(X_2; Y_2 | X_1),$$

where the first inequality in each line is the identity $I(A; B | C) = I(A; B) + I(A; C | B)$ with C the *other* input (independent of A). Note the second line couples X_2 to the *other* input X_1 : $I(X_2; Y_1 | X_1) \geq I(X_2; Y_1)$; the superficially analogous statement $I(X_2; Y_1 | X_2) \geq I(X_2; Y_1)$ is false (its left-hand side is 0), so the conditioning variable matters. The converse implication (strong $\Rightarrow$ very strong) is false in general; for instance the Gaussian IC with $a^2 = b^2 = 1$ is strong but not very strong (see §5.5).

To prove the strong-interference capacity theorem we need two lemmas. The first says that the product of two strong-interference channels is again strong.

Lemma 5.11. (Strong interference extends to arbitrary input laws) Suppose a channel satisfies

$$I(X_1; Y_2 | X_2) \geq I(X_1; Y_1 | X_2)$$

for every product law $p(x_1)p(x_2)$. Then the same inequality holds for every joint law $p(x_1, x_2)$. Consequently, for every random variable U with

$$U \rightarrow (X_1, X_2) \rightarrow (Y_1, Y_2),$$

one has

$$I(X_1; Y_2 | X_2, U) \geq I(X_1; Y_1 | X_2, U).$$

○

Proof. Fix x_2 of positive probability. Under the conditional law given $X_2 = x_2$,

$$P_{X_1, X_2 | X_2 = x_2}(x_1, x'_2) = P_{X_1 | X_2 = x_2}(x_1) \mathbf{1}\{x'_2 = x_2\},$$

a product law, namely $P_{X_1 | X_2 = x_2} \delta_{x_2}$. The assumed strong-interference inequality therefore gives

$$I(X_1; Y_2 | X_2 = x_2) \geq I(X_1; Y_1 | X_2 = x_2).$$

Averaging over X_2 proves the arbitrary-joint-law statement. Condition next on every $U = u$ of positive probability and apply that statement to the joint conditional law $P_{X_1, X_2 | U = u}$; averaging in u gives the final conditional inequality. ■

Lemma 5.12. (Tensorization of strong interference) If two channels W_1 and W_2 satisfy the strong-interference condition, then $W_1 \otimes W_2$ also satisfies it. Hence for every block-product input law $P_{X_1^n} P_{X_2^n}$,

$$I(X_1^n; Y_2^n | X_2^n) \geq I(X_1^n; Y_1^n | X_2^n).$$

○

Proof. It is enough to prove the two-coordinate product statement. Let

$$P_{X_{11}, X_{12}, X_{21}, X_{22}} = P_{X_{11}, X_{12}} P_{X_{21}, X_{22}}.$$

The standard two-coordinate chain-rule identity is

$$\begin{aligned} & I(X_{11}, X_{12}; Y_{21}, Y_{22} | X_{21}, X_{22}) - I(X_{11}, X_{12}; Y_{11}, Y_{12} | X_{21}, X_{22}) \\ &= \left[I(X_{11}; Y_{21} | X_{21}, X_{22}, Y_{12}) - I(X_{11}; Y_{11} | X_{21}, X_{22}, Y_{12}) \right] \\ &+ \left[I(X_{12}; Y_{22} | X_{21}, X_{22}, Y_{21}) - I(X_{12}; Y_{12} | X_{21}, X_{22}, Y_{21}) \right]. \end{aligned}$$

The first bracket is nonnegative by Lemma 5.11, applied to the first channel conditionally on $U = (X_{22}, Y_{12})$. The second bracket is nonnegative by the same lemma applied to the second channel conditionally on $U' = (X_{21}, Y_{21})$. Thus the product channel is strong. Iteration proves the blocklength- n statement. No independence between X_{11} and X_{12} , or between X_{21} and X_{22} , is used or required: conditioning on a point mass in the interfering input is precisely what reduces each bracket to the one-letter condition. ■

A direct consequence is the n -letter strong-interference inequality.

Lemma 5.13. For a DM-IC with strong interference,

$$I(X_1^n; Y_2^n | X_2^n) \geq I(X_1^n; Y_1^n | X_2^n), \quad (5.8)$$

for all $(X_1^n, X_2^n) \sim p(x_1^n)p(x_2^n)$ and all $n \geq 1$. ○

Proof. By induction on n using the previous lemma: $W^{\otimes n}$ is a strong-interference channel, so the strong-interference condition holds at block length n for product inputs $p(x_1^n)p(x_2^n)$, which is exactly the displayed inequality. ■

We can now state and prove the capacity region under strong interference.

Theorem 5.14. (Capacity region under strong interference) The capacity region of a DM-IC with strong interference is the set of *nonnegative* rate pairs (R_1, R_2) satisfying

$$\begin{aligned} R_1 &\leq I(X_1; Y_1 | X_2, Q), \\ R_2 &\leq I(X_2; Y_2 | X_1, Q), \\ R_1 + R_2 &\leq \min\{I(X_1, X_2; Y_1 | Q), I(X_1, X_2; Y_2 | Q)\}, \end{aligned} \quad (5.9)$$

for some $p(q)p(x_1|q)p(x_2|q)$ with $|Q| \leq 4$; boundary points are included by closure of the achievable region. $\circ$

Remark 5.1. The Gaussian case was proved by Sato (1981); the DM-IC case by El Gamal and Costa (1987). The cardinality bound $|Q| \leq 4$ follows from the Fenchel–Eggleston–Carathéodory (support) lemma: the region is described by three functionals of $(p(x_1), p(x_2))$, so three points of the underlying simplex suffice, plus one more to preserve the convex-hull (time-sharing) structure. $\circ$

Proof. Achievability. Use simultaneous decoding and apply the strong-interference condition to drop the inactive constraints. In (5.4), under strong interference $I(X_1; Y_2 | X_2, Q) \geq I(X_1; Y_1 | X_2, Q)$ and $I(X_2; Y_1 | X_1, Q) \geq I(X_2; Y_2 | X_1, Q)$, so

$$R_1 \leq I(X_1; Y_1 | X_2, Q), \quad R_2 \leq I(X_2; Y_2 | X_1, Q),$$

and the sum constraint $R_1 + R_2 \leq \min\{I(X_1, X_2; Y_1 | Q), I(X_1, X_2; Y_2 | Q)\}$ remains. This is exactly (5.9).

Converse. The first two inequalities are cut-set bounds. For the third, it suffices by symmetry to show $R_1 + R_2 \leq I(X_1, X_2; Y_2 | Q)$. We proceed by the sub-additivity template. Fano gives

$$\begin{aligned} n(R_1 + R_2) &= H(M_1) + H(M_2) \\ &\leq I(M_1; Y_1^n) + I(M_2; Y_2^n) + 2n\epsilon_n \\ &\leq I(X_1^n; Y_1^n) + I(X_2^n; Y_2^n) + 2n\epsilon_n \\ &\leq I(X_1^n; Y_1^n | X_2^n) + I(X_2^n; Y_2^n) + 2n\epsilon_n \\ &\leq I(X_1^n; Y_2^n | X_2^n) + I(X_2^n; Y_2^n) + 2n\epsilon_n \\ &= I(X_1^n, X_2^n; Y_2^n) + 2n\epsilon_n \\ &\leq \sum_{i=1}^n I(X_{1i}, X_{2i}; Y_{2i}) + 2n\epsilon_n \\ &= n I(X_1, X_2; Y_2 | Q) + 2n\epsilon_n, \end{aligned} \quad (5.10)$$

where the steps are:

- (a) M_1, M_2 are uniform over $[1 : 2^{nR_1}]$, $[1 : 2^{nR_2}]$;

- (b) Fano's inequality (two terms, hence $2n\varepsilon_n$);
- (c) the data-processing inequality via $M_j \rightarrow X_j^n \rightarrow Y_j^n$;
- (d) the independence identity $I(X_1^n; Y_1^n | X_2^n) = I(X_1^n; Y_1^n) + I(X_1^n; X_2^n | Y_1^n) \geq I(X_1^n; Y_1^n)$ (valid since $X_1^n \perp\!\!\!\perp X_2^n$);
- (e) Lemma 5.13 (strong interference at block length n);
- (f) the chain rule $I(X_1; Y_2 | X_2) + I(X_2; Y_2) = I(X_1, X_2; Y_2)$;
- (g) *sub-additivity of the functional* $F = P_{X_1, X_2} \mapsto I(X_1, X_2; Y_2)$ over the memoryless channel (§1.13): $I(X_1^n, X_2^n; Y_2^n) \leq \sum_i I(X_{1i}, X_{2i}; Y_{2i})$;
- (h) time-sharing with $Q \sim \text{Unif}[1 : n]$ and $(X_1, X_2) = (X_{1Q}, X_{2Q})$.

Dividing by n and letting $n \rightarrow \infty$ gives $R_1 + R_2 \leq I(X_1, X_2; Y_2 | Q)$; by symmetry the same bound holds with Y_1 , giving the min in (5.9). ■

5.4 Han–Kobayashi Inner Bound

We now discuss the best-known general inner bound for the interference channel, due to Han and Kobayashi (1981). The key ideas are *rate splitting* (each message is split into a “common” part, decoded by both receivers, and a “private” part, decoded only by its own receiver), *superposition coding* (the private part is a satellite of the common part), and *simultaneous non-unique decoding*.

Theorem 5.15. (Han–Kobayashi inner bound) For a DM-IC $W(y_1, y_2 | x_1, x_2)$, a rate pair (R_1, R_2) is achievable if

$$\begin{aligned}
 R_1 &< I(X_1; Y_1 | U_2, Q), \\
 R_2 &< I(X_2; Y_2 | U_1, Q), \\
 R_1 + R_2 &< I(X_1, U_2; Y_1 | Q) + I(X_2; Y_2 | U_1, U_2, Q), \\
 R_1 + R_2 &< I(X_2, U_1; Y_2 | Q) + I(X_1; Y_1 | U_1, U_2, Q), \\
 R_1 + R_2 &< I(X_1, U_2; Y_1 | U_1, Q) + I(X_2, U_1; Y_2 | U_2, Q), \\
 2R_1 + R_2 &< I(X_1, U_2; Y_1 | Q) + I(X_1; Y_1 | U_1, U_2, Q) + I(X_2, U_1; Y_2 | U_2, Q), \\
 R_1 + 2R_2 &< I(X_2, U_1; Y_2 | Q) + I(X_2; Y_2 | U_1, U_2, Q) + I(X_1, U_2; Y_1 | U_1, Q),
 \end{aligned} \tag{5.11}$$

for some $p(q)p(u_1, x_1 | q)p(u_2, x_2 | q)$ and $|Q| \leq 7$. ○

Proof. Codebook generation. Fix $p(q)p(u_1, x_1 | q)p(u_2, x_2 | q)$. Generate a time-sharing sequence $q^n \sim \prod_i p(q_i)$. For $j = 1, 2$, randomly and conditionally independently generate $2^{nR_{j0}}$ sequences $u_j^n(m_{j0})$, $m_{j0} \in [1 : 2^{nR_{j0}}]$, each according to $\prod_i p(u_{ji} | q_i)$. Then, for each m_{j0} , randomly and conditionally independently generate $2^{nR_{j1}}$ sequences $x_j^n(m_{j0}, m_{j1})$, $m_{j1} \in [1 : 2^{nR_{j1}}]$, each according to $\prod_i p(x_{ji} | u_{ji}(m_{j0}), q_i)$. (With time-sharing one may equally use block-dependent rates $R_{j0}(q_i), R_{j1}(q_i)$; the exponent above is the averaged, rate-splitting form $R_j = R_{j0} + R_{j1}$.)

Encoding. To send $m_j = (m_{j0}, m_{j1})$, sender j transmits $x_j^n(m_{j0}, m_{j1})$.

Decoding. Each decoder uses simultaneous non-unique decoding. Decoder 1 finds $(\hat{m}_{10}, \hat{m}_{11})$ such that

$$(q^n, u_1^n(\hat{m}_{10}), u_2^n(m_{20}), x_1^n(\hat{m}_{10}, \hat{m}_{11}), y_1^n) \in \mathcal{T}_\epsilon^{(n)}$$

for *some* m_{20} (decoded non-uniquely); otherwise it declares an error. Decoder 2 proceeds symmetrically.

Error analysis at decoder 1. Assume $(m_{10}, m_{11}, m_{20}) = (1, 1, 1)$ is sent. The four error events and the constraints under which each has vanishing probability are given in Table 5.1.

m_{10}	m_{11}	m_{20}	constraint at decoder 1 (Y_1^n)
×	?	×	$R_{10} + R_{11} + R_{20} < I(U_2, X_1; Y_1 Q)$
✓	×	×	$R_{11} + R_{20} < I(U_2, X_1; Y_1 U_1, Q)$
×	?	✓	$R_{10} + R_{11} < I(X_1; Y_1 U_2, Q)$
✓	×	✓	$R_{11} < I(X_1; Y_1 U_1, U_2, Q)$

Table 5.1: Decoder-1 error events. “✓” = correct, “×

Each row is controlled by the packing lemma (rows 1–3) or by the law of large numbers / packing lemma (row 4). Hence $P[E_1] \rightarrow 0$ whenever the four constraints hold.

Error analysis at decoder 2. Decoder 2 recovers (m_{20}, m_{21}) and non-uniquely decodes m_{10} . The symmetric constraints are

$$\begin{aligned} R_{20} + R_{21} + R_{10} &< I(U_1, X_2; Y_2|Q), \\ R_{21} + R_{10} &< I(U_1, X_2; Y_2|U_2, Q), \\ R_{20} + R_{21} &< I(X_2; Y_2|U_1, Q), \\ R_{21} &< I(X_2; Y_2|U_2, U_1, Q). \end{aligned} \tag{5.12}$$

(Note that decoder 2 tests typicality of (u_2, x_2, u_1, y_2) ; the roles of the two users are interchanged relative to Table 5.1, and the private rate of user 2 is R_{21} .)

Fourier–Motzkin elimination. Substituting $R_{10} + R_{11} = R_1$ and $R_{20} + R_{21} = R_2$ and eliminating the four split rates $R_{10}, R_{11}, R_{20}, R_{21}$ (equivalently the two common rates R_{10}, R_{20}) yields nine inequalities: the seven in (5.11) together with

$$\begin{aligned} R_1 &< I(X_1; Y_1|U_1, U_2, Q) + I(X_2, U_1; Y_2|U_2, Q), \\ R_2 &< I(X_1, U_2; Y_1|U_1, Q) + I(X_2; Y_2|U_1, U_2, Q). \end{aligned} \tag{5.13}$$

These two are redundant (each is implied by the seven), so the nine-inequality region equals the seven-inequality region (5.11). The elimination is routine Fourier–Motzkin, and the redundancy of the two extra inequalities together with the cardinality bound $|Q| \leq 7$ is the standard simplification of Chong–Motani–Garg–El Gamal 2008; we follow their reduction and omit the

algebra. All inequalities are strict, and boundary points of the region follow by closure of the capacity region (strict margins, then limits), as in Chapters 2–3. ■

Remark 5.2 (Extreme specialisations). • Setting $U_1 = U_2 = \emptyset$ reduces (5.11) to the TIN inner bound (5.3).

- Setting $U_1 = X_1, U_2 = X_2$ reduces (5.11) to the MAC (simultaneous decoding) region: each user's entire message is common, so each receiver decodes both messages.

○

Remark 5.3 (Further reading: how good is Han–Kobayashi?). For many years it was conjectured that the Han–Kobayashi region is the capacity region of the interference channel; the Gaussian weak-interference sum capacity of §5.8 was a key step toward it, and the region is known to be optimal for strong interference (§5.3.2) and for the injective deterministic channel (§5.6). It is now known that Han–Kobayashi coding is *not* optimal in general: [Nair et al. \(2015\)](#) exhibits a channel for which the capacity region strictly contains the Han–Kobayashi region (roughly, Han–Kobayashi coding is limited to a two-letter style of superposition, while optimal codes can exploit the channel's structure more flexibly). The auxiliary-receiver bounds of §5.7.2 below are the current state of the art on the converse side. ○

5.5 The Scalar Gaussian Interference Channel

The most studied instance of the interference channel is the *scalar Gaussian IC*

$$Y_1 = X_1 + bX_2 + Z_1, \quad Y_2 = aX_1 + X_2 + Z_2,$$

where $a, b \geq 0$ are the cross-channel gains (from sender 1 to receiver 2, and from sender 2 to receiver 1, respectively) and $Z_1, Z_2 \sim \mathcal{N}(0, 1)$ are independent standard Gaussians, independent of the inputs. The inputs satisfy the average power constraints $\mathbb{E}[X_1^2] \leq P_1$ and $\mathbb{E}[X_2^2] \leq P_2$; as in the Gaussian MAC chapter these are read at code level, i.e. every codeword satisfies the time-average second-moment constraint $\frac{1}{n} \sum_{i=1}^n X_{ji}^2 \leq P_j$ (the converses below use only time-averaged conditional variances and Jensen's inequality, so any convention implying the averaged constraint suffices). Throughout this chapter log is base 2 (bits), so $\frac{1}{2} \log(\cdot)$ denotes $\frac{1}{2} \log_2(\cdot)$; by Gaussian maximality (Lemma 1.59) a Gaussian input maximises $h(X + Z)$ under a power constraint.

5.5.1 Strong interference

Proposition 5.16. If $a^2 \geq 1$ and $b^2 \geq 1$, then the Gaussian IC satisfies the strong-interference condition. ○

Proof. By symmetry it suffices to show $I(X_1; Y_2|X_2) \geq I(X_1; Y_1|X_2)$. We have

$$I(X_1; Y_2|X_2) = I(X_1; aX_1 + X_2 + Z_2|X_2) = I(X_1; aX_1 + Z_2|X_2) = I(X_1; aX_1 + Z_2) = I(X_1; X_1 + \frac{Z_2}{a}),$$




Figure 5.2: A scalar Gaussian interference channel

where the third equality uses $X_2 \perp\!\!\!\perp (X_1, Z_2)$, and

$$I(X_1; Y_1 | X_2) = I(X_1; X_1 + bX_2 + Z_1 | X_2) = I(X_1; X_1 + Z_1 | X_2) = I(X_1; X_1 + Z_1).$$

Since Z_2/a has variance $1/a^2 \leq 1 = \text{Var}(Z_1)$ (as $a^2 \geq 1$), the channel $X_1 \rightarrow X_1 + Z_2/a$ is less noisy than $X_1 \rightarrow X_1 + Z_1$ (equivalently, by the Gaussian ordering of the “degraded” channel), hence $I(X_1; X_1 + Z_2/a) \geq I(X_1; X_1 + Z_1)$. The second inequality is symmetric using $b^2 \geq 1$. ■

Theorem 5.17. [Capacity region of the Gaussian IC under strong interference] For the scalar Gaussian IC with strong interference ($a^2 \geq 1, b^2 \geq 1$), the capacity region is

$$\begin{aligned} R_1 &\leq \frac{1}{2} \log(1 + P_1), \\ R_2 &\leq \frac{1}{2} \log(1 + P_2), \\ R_1 + R_2 &\leq \min\left\{\frac{1}{2} \log(1 + P_1 + b^2 P_2), \frac{1}{2} \log(1 + P_2 + a^2 P_1)\right\}. \end{aligned} \tag{5.14}$$

○

Proof. Achievability. Use simultaneous decoding with $Q = \emptyset$ and Gaussian inputs $X_1 \sim \mathcal{N}(0, P_1)$, $X_2 \sim \mathcal{N}(0, P_2)$. Receiver j decodes both messages; the MAC region at each receiver gives exactly (5.14).

Converse. We prove each of the three bounds. Throughout, the single-letterisation is by sub-additivity of the relevant functional over the memoryless channel (§1.13); the Gaussian step is Lemma 1.59.

Individual bound. By the cut-set bound and sub-additivity of $F = P_{X_1, X_2} \mapsto I(X_1; Y_1 | X_2)$,

$$nR_1 \leq I(X_1^n; Y_1^n | X_2^n) + n\epsilon_n \leq \sum_{i=1}^n I(X_{1i}; Y_{1i} | X_{2i}) + n\epsilon_n.$$

For each i , $X_{1i} \perp\!\!\!\perp X_{2i}$ (the two encoders are independent) and $Z_{1i} \perp\!\!\!\perp (X_{1i}, X_{2i})$, so

$$\begin{aligned} I(X_{1i}; Y_{1i}|X_{2i}) &= I(X_{1i}; X_{1i} + Z_{1i}) = h(X_{1i} + Z_{1i}) - h(Z_{1i}) \\ &\leq \frac{1}{2} \log 2\pi e(1 + \text{Var}(X_{1i})) - \frac{1}{2} \log 2\pi e = \frac{1}{2} \log(1 + \text{Var}(X_{1i})). \end{aligned}$$

By Jensen's inequality on the concave function $t \mapsto \frac{1}{2} \log(1 + t)$ and the power constraint,

$$\frac{1}{n} \sum_{i=1}^n \frac{1}{2} \log(1 + \text{Var}(X_{1i})) \leq \frac{1}{2} \log \left(1 + \frac{1}{n} \sum_{i=1}^n \text{Var}(X_{1i}) \right) \leq \frac{1}{2} \log(1 + P_1),$$

so $R_1 \leq \frac{1}{2} \log(1 + P_1)$; symmetrically $R_2 \leq \frac{1}{2} \log(1 + P_2)$.

Sum bound. It suffices to show $R_1 + R_2 \leq \frac{1}{2} \log(1 + P_2 + a^2 P_1)$ (the other sum bound is symmetric). By the strong-interference converse of §5.3.2, whose single-letterisation is sub-additivity of $F = P_{X_1, X_2} \mapsto I(X_1, X_2; Y_2)$,

$$n(R_1 + R_2) \leq I(X_1^n, X_2^n; Y_2^n) + 2n\epsilon_n \leq \sum_{i=1}^n I(X_{1i}, X_{2i}; Y_{2i}) + 2n\epsilon_n.$$

For each i , $X_{1i} \perp\!\!\!\perp X_{2i}$, so the cross-covariance vanishes and

$$\begin{aligned} I(X_{1i}, X_{2i}; Y_{2i}) &= h(aX_{1i} + X_{2i} + Z_2) - h(Z_2) \\ &\leq \frac{1}{2} \log 2\pi e(1 + a^2 \text{Var}(X_{1i}) + \text{Var}(X_{2i})) - \frac{1}{2} \log 2\pi e \\ &= \frac{1}{2} \log(1 + a^2 \text{Var}(X_{1i}) + \text{Var}(X_{2i})), \end{aligned}$$

by Gaussian optimality (Lemma 1.59), using $\text{Var}(aX_{1i} + X_{2i} + Z_2) = a^2 \text{Var}(X_{1i}) + \text{Var}(X_{2i}) + 1$. Jensen's inequality and the power constraints give

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \frac{1}{2} \log(1 + a^2 \text{Var}(X_{1i}) + \text{Var}(X_{2i})) \\ \leq \frac{1}{2} \log \left(1 + a^2 \frac{1}{n} \sum_i \text{Var}(X_{1i}) + \frac{1}{n} \sum_i \text{Var}(X_{2i}) \right) \leq \frac{1}{2} \log(1 + a^2 P_1 + P_2), \end{aligned}$$

so $R_1 + R_2 \leq \frac{1}{2} \log(1 + P_2 + a^2 P_1)$. (This is the direct route: bound $I(X_1, X_2; Y_2) = h(Y_2) - h(Z_2)$ by replacing the *first* term $h(Y_2)$ by its Gaussian maximum, never by substituting a Gaussian maximum into the *subtracted* entropy $h(X_2 + Z_2)$.) ■

Remark 5.4 (Corner points of the pentagon). Write $C_1 = \frac{1}{2} \log(1 + P_1)$, $C_2 = \frac{1}{2} \log(1 + P_2)$, $S_1 = \frac{1}{2} \log(1 + P_1 + b^2 P_2)$, $S_2 = \frac{1}{2} \log(1 + P_2 + a^2 P_1)$, and $S = \min\{S_1, S_2\}$. Then (5.14) is the pentagonal region

$$R_1 \leq C_1, \quad R_2 \leq C_2, \quad R_1 + R_2 \leq S,$$

whose corner points are

$$(0, 0), \quad (C_1, 0), \quad (0, C_2), \quad (C_1, \min\{C_2, S - C_1\}), \quad (\min\{C_1, S - C_2\}, C_2).$$

The caps $\min\{C_2, \cdot\}$ and $\min\{C_1, \cdot\}$ are essential: for example if $b^2 > 1 + P_1$, the corner $(C_1, S_1 - C_1)$ would lie above the line $R_2 = C_2$ and thus outside the region, unless the cap C_2 is applied. ○

5.6 The Injective Deterministic Interference Channel

A deterministic IC is one in which the outputs are deterministic functions of the inputs, $Y_1 = f_1(X_1, X_2)$, $Y_2 = f_2(X_1, X_2)$. A particularly tractable subclass is the *injective deterministic IC*, introduced by El Gamal and Costa, in which the interference each user causes is recoverable from the other user's output. This class is exactly the bridge that shows Han–Kobayashi coding is optimal (not merely an inner bound) for a nontrivial family of channels.

Definition 5.18. (Injective deterministic IC) A deterministic IC

$$Y_1 = y_1(X_1, T_2), \quad Y_2 = y_2(X_2, T_1), \quad T_1 = t_1(X_1), \quad T_2 = t_2(X_2),$$

is *injective* if for every x_1 the map $t_2 \mapsto y_1(x_1, t_2)$ is one-to-one, and for every x_2 the map $t_1 \mapsto y_2(x_2, t_1)$ is one-to-one. Equivalently, t_2 is determined by (x_1, y_1) and t_1 is determined by (x_2, y_2) . In particular,

$$H(Y_1|X_1) = H(T_2), \quad H(Y_2|X_2) = H(T_1).$$

○

Theorem 5.19. (Capacity region of the injective deterministic IC) The capacity region of the injective deterministic IC is the set of rate pairs (R_1, R_2) satisfying

$$\begin{aligned} R_1 &\leq H(Y_1|T_2, Q), \\ R_2 &\leq H(Y_2|T_1, Q), \\ R_1 + R_2 &\leq H(Y_1|Q) + H(Y_2|T_1, T_2, Q), \\ R_1 + R_2 &\leq H(Y_2|Q) + H(Y_1|T_1, T_2, Q), \\ R_1 + R_2 &\leq H(Y_1|T_1, Q) + H(Y_2|T_2, Q), \\ 2R_1 + R_2 &\leq H(Y_1|Q) + H(Y_1|T_1, T_2, Q) + H(Y_2|T_2, Q), \\ R_1 + 2R_2 &\leq H(Y_2|Q) + H(Y_2|T_1, T_2, Q) + H(Y_1|T_1, Q), \end{aligned} \tag{5.15}$$

for some $p(q)p(x_1|q)p(x_2|q)$.

○

Proof. Achievability. This is the Han–Kobayashi bound with $U_1 = T_1$ and $U_2 = T_2$. Because the channel is injective, decoder 1 recovers T_2^n jointly with X_1^n , and decoder 2 recovers T_1^n jointly with X_2^n ; hence T_1, T_2 can serve as the common auxiliary variables. For instance $R_1 < I(X_1; Y_1|T_2, Q) = H(Y_1|T_2, Q) - H(Y_1|X_1, T_2, Q) = H(Y_1|T_2, Q)$, using that Y_1 is a function of (X_1, T_2) ; the remaining inequalities of (5.15) follow identically from (5.11).

Converse. By symmetry it suffices to prove the four inequalities

$$\begin{aligned} R_1 &\leq H(Y_1|T_2, Q), \\ R_1 + R_2 &\leq H(Y_1|Q) + H(Y_2|T_1, T_2, Q), \\ R_1 + R_2 &\leq H(Y_1|T_1, Q) + H(Y_2|T_2, Q), \\ 2R_1 + R_2 &\leq H(Y_1|Q) + H(Y_1|T_1, T_2, Q) + H(Y_2|T_2, Q). \end{aligned} \tag{5.16}$$

(The remaining three are the symmetric counterparts.) Throughout, all functionals below are sub-additive over the memoryless deterministic channel (§1.13). (Each of these sub-additivities is the standard chain-rule bound $H(A^n) \leq \sum_{i=1}^n H(A_i)$ applied coordinate-wise; memorylessness of the channel suffices and no determinism beyond that is used. The converse we give follows the original argument of El Gamal–Costa [El Gamal and Costa \(1982\)](#), who determined this capacity region.)

First inequality. By Fano,

$$nR_1 \leq I(X_1^n; Y_1^n | T_2^n) + n\varepsilon_n = I(X_1^n; Y_1^n, T_2^n) + n\varepsilon_n = H(Y_1^n | T_2^n) + n\varepsilon_n \leq \sum_{i=1}^n H(Y_{1i} | T_{2i}) + n\varepsilon_n,$$

where the second step uses $T_2^n \perp\!\!\!\perp X_1^n$ (so $I(X_1^n; T_2^n) = 0$), the third uses that Y_1^n is a function of (X_1^n, T_2^n) , and the last is sub-additivity of $F = P_{X_1} \mapsto H(Y_1 | T_2)$. After time-sharing, $R_1 \leq H(Y_1 | T_2, Q)$.

Second inequality. By Fano and the injectivity identities $H(Y_1^n | X_1^n) = H(T_2^n)$ and $H(Y_2^n | X_2^n) = H(T_1^n)$,

$$n(R_1 + R_2) \leq I(X_1^n; Y_1^n) + I(X_2^n; Y_2^n) + n\varepsilon_n = H(Y_1^n) - H(T_2^n) + H(Y_2^n) - H(T_1^n) + n\varepsilon_n.$$

Now

$$H(Y_2^n) - H(T_1^n) - H(T_2^n) \leq H(Y_2^n | T_1^n, T_2^n),$$

because $H(Y_2^n | T_1^n, T_2^n) = H(Y_2^n, T_1^n, T_2^n) - H(T_1^n) - H(T_2^n) \geq H(Y_2^n) - H(T_1^n) - H(T_2^n)$ (using $T_1^n \perp\!\!\!\perp T_2^n$). Hence

$$n(R_1 + R_2) \leq H(Y_1^n) + H(Y_2^n | T_1^n, T_2^n) + n\varepsilon_n \leq \sum_{i=1}^n [H(Y_{1i}) + H(Y_{2i} | T_{1i}, T_{2i})] + n\varepsilon_n,$$

by the sub-additivity of $F = P_{X_1, X_2} \mapsto H(Y_1) + H(Y_2 | T_1, T_2)$. Time-sharing gives $R_1 + R_2 \leq H(Y_1 | Q) + H(Y_2 | T_1, T_2, Q)$.

Third inequality (the coupled bound). By Fano,

$$n(R_1 + R_2) \leq I(X_1^n; Y_1^n) + I(X_2^n; Y_2^n) + n\varepsilon_n = H(Y_1^n) - H(T_2^n) + H(Y_2^n) - H(T_1^n) + n\varepsilon_n.$$

Now the key coupled identity

$$\begin{aligned} H(Y_1^n) - H(T_2^n) + H(Y_2^n) - H(T_1^n) &= I(Y_1^n; T_1^n) + I(Y_2^n; T_2^n) + H(Y_1^n | T_1^n) + H(Y_2^n | T_2^n) \\ &\quad - H(T_1^n) - H(T_2^n) \leq H(Y_1^n | T_1^n) + H(Y_2^n | T_2^n), \end{aligned} \quad (5.17)$$

where the last step uses $I(Y_1^n; T_1^n) \leq H(T_1^n)$ and $I(Y_2^n; T_2^n) \leq H(T_2^n)$. (The two split functionals $H(Y_1^n) - H(T_2^n)$ and $H(Y_2^n) - H(T_1^n)$ are *not* individually sub-additive and there is no independence $Y_1 \perp\!\!\!\perp T_1$; the correct sub-additive functional here is the *coupled* pair $F = P_{X_1, X_2} \mapsto I(Y_1; T_1) + I(Y_2; T_2)$, which is exactly what (5.17) uses.) Then by sub-additivity of $H(Y_1 | T_1) + H(Y_2 | T_2)$,

$$n(R_1 + R_2) \leq \sum_{i=1}^n [H(Y_{1i} | T_{1i}) + H(Y_{2i} | T_{2i})] + n\varepsilon_n,$$

so $R_1 + R_2 \leq H(Y_1|T_1, Q) + H(Y_2|T_2, Q)$ after time-sharing.

Fourth inequality. By Fano and the injectivity identities,

$$n(2R_1 + R_2) \leq 2I(X_1^n; Y_1^n) + I(X_2^n; Y_2^n) + n\epsilon_n = 2H(Y_1^n) - 2H(T_1^n) + H(Y_2^n) - H(T_1^n) + n\epsilon_n.$$

Regroup the right-hand side as

$$2H(Y_1^n) - 2H(T_1^n) + H(Y_2^n) - H(T_1^n) = H(Y_1^n) + [H(Y_1^n) - H(T_1^n) - H(T_2^n)] + [H(Y_2^n) - H(T_2^n)],$$

and bound each bracket by a conditional entropy: $H(Y_1^n) - H(T_1^n) - H(T_2^n) \leq H(Y_1^n|T_1^n, T_2^n)$ (as in the second inequality) and $H(Y_2^n) - H(T_2^n) \leq H(Y_2^n|T_2^n)$ (equivalently $I(Y_2^n; T_2^n) \leq H(T_2^n)$, the valid single-term analogue). Hence

$$\begin{aligned} n(2R_1 + R_2) &\leq H(Y_1^n) + H(Y_1^n|T_1^n, T_2^n) + H(Y_2^n|T_2^n) + n\epsilon_n \\ &\leq \sum_{i=1}^n [H(Y_{1i}) + H(Y_{1i}|T_{1i}, T_{2i}) + H(Y_{2i}|T_{2i})] + n\epsilon_n, \end{aligned}$$

by the sub-additivity of $F = P_{X_1, X_2} \mapsto H(Y_1) + H(Y_1|T_1, T_2) + H(Y_2|T_2)$. Time-sharing gives

$$2R_1 + R_2 \leq H(Y_1|Q) + H(Y_1|T_1, T_2, Q) + H(Y_2|T_2, Q).$$

■

5.7 Outer Bounds

We now develop two general outer bounds. The first is a “genie-aided” bound that underlies the noisy-interference sum-capacity result of the next section; the second is the Etkin–Ordentlich bound.

5.7.1 A genie-based outer bound

The idea is to upper bound the original channel by a *genie-aided* channel in which each receiver is given additional side information. Giving side information to a receiver can only increase the set of achievable rate pairs, so any outer bound for the genie-aided channel is an outer bound for the original one. We now make this precise.

Why the genie is legitimate. Suppose receiver 1 is additionally shown T_1^n (a deterministic or noisy function of its own input X_1^n , or, in the Gaussian specialisation, a noisy version of the *interference* X_2^n , which is part of receiver 2’s observation Y_2^n), and receiver 2 is shown T_2^n (analogously). Any code for the original channel is also a code for the genie-aided channel: the encoders are unchanged, and each decoder can simply ignore the genie. Therefore

$$\mathcal{C}(\text{original}) \subseteq \mathcal{C}(\text{genie-aided}),$$

and it suffices to bound the genie-aided capacity region. In the single-letterisation below, the genie is introduced through the elementary inequality $I(X_1^n; Y_1^n) \leq I(X_1^n; Y_1^n, T_1^n)$ (conditioning enlarges mutual information), which is precisely the statement that the genie only helps. The genie channels are as declared in the theorem: $p(y_j, t_j | x_1, x_2) = p(t_j | x_j) p(y_j | x_1, x_2)$, each genie is a (possibly noisy) function of its own user's input only, the two genies are conditionally independent given their inputs, and the genie symbols are independent across time. The template below – Fano, sub-additive single-letterisation, and the concave-envelope/min-max bridge – is the auxiliary-receiver machinery of Gohari–Nair 2022, of which this genie bound is a special case; the Gaussian “smart genie” applications of §5.8 below import theorems of Shang–Kramer–Chen 2009 and Liu–Nair–Xia 2015 rather than being proved from the template alone.

Theorem 5.20. (Genie-based outer bound) Let $\lambda \geq 1$. Let T_1 and T_2 be auxiliary random variables satisfying

$$p(y_1, t_1 | x_1, x_2) = p(t_1 | x_1) p(y_1 | x_1, x_2), \quad p(y_2, t_2 | x_1, x_2) = p(t_2 | x_2) p(y_2 | x_1, x_2),$$

with the marginals consistent with the channel, i.e. $p(y_1 | x_1, x_2) = W(y_1 | x_1, x_2)$ and $p(y_2 | x_1, x_2) = W(y_2 | x_1, x_2)$. Then every achievable rate pair satisfies

$$R_1 + \lambda R_2 \leq \max_{p(x_1)p(x_2)} \left(I(X_1; Y_1, T_1) + \lambda I(X_2; Y_2, T_2) + \Gamma[f_1] + \Gamma[f_2] \right), \quad (5.18)$$

where

$$f_1 = I(X_1; T_1 | X_2, T_2) - \lambda I(X_1; Y_2 | X_2, T_2), \quad f_2 = I(X_2; T_2 | X_1, T_1) - I(X_2; Y_1 | X_1, T_1),$$

$\Gamma[f] = \mathfrak{C}[f] - f$, and $\mathfrak{C}[f]$ is the upper concave envelope of f over product distributions $p(x_1)p(x_2)$. $\circ$

Remark 5.5. The two functionals f_1, f_2 are asymmetric in λ : f_1 carries λ (the weighting of the Y_2 -term) while f_2 does not, because the weighted sum is $R_1 + \lambda R_2$. This asymmetry is essential and is fixed in the derivation below. $\circ$

Proof. (L1) Fano.

$$n(R_1 + \lambda R_2) \leq I(X_1^n; Y_1^n) + \lambda I(X_2^n; Y_2^n) + n\varepsilon_n.$$

(L1') Genie. Adding T_1^n to receiver 1 and T_2^n to receiver 2,

$$I(X_1^n; Y_1^n) \leq I(X_1^n; Y_1^n, T_1^n), \quad I(X_2^n; Y_2^n) \leq I(X_2^n; Y_2^n, T_2^n),$$

so, expanding by the chain rule,

$$\begin{aligned} n(R_1 + \lambda R_2) &\leq I(X_1^n; T_1^n) + I(X_1^n, X_2^n; Y_1^n | T_1^n) - I(X_2^n; Y_1^n | X_1^n, T_1^n) \\ &\quad + \lambda I(X_2^n; T_2^n) + \lambda I(X_1^n, X_2^n; Y_2^n | T_2^n) - \lambda I(X_1^n; Y_2^n | X_2^n, T_2^n) + n\varepsilon_n. \end{aligned} \quad (5.19)$$

(L2) Sub-additive single-letterisation. Since $T_1^n \perp\!\!\!\perp (X_2^n, T_2^n)$ and $T_2^n \perp\!\!\!\perp (X_1^n, T_1^n)$ (the genies are local), we have $I(X_1^n; T_1^n) \leq I(X_1^n; T_1^n | X_2^n, T_2^n)$ and $I(X_2^n; T_2^n) \leq I(X_2^n; T_2^n | X_1^n, T_1^n)$. Grouping,

$$\begin{aligned} n(R_1 + \lambda R_2) &\leq [I(X_1^n; T_1^n | X_2^n, T_2^n) - \lambda I(X_1^n; Y_2^n | X_2^n, T_2^n)] \\ &\quad + [I(X_2^n; T_2^n | X_1^n, T_1^n) - I(X_2^n; Y_1^n | X_1^n, T_1^n)] \\ &\quad + [(\lambda - 1)I(X_2^n; T_2^n) + I(X_1^n, X_2^n; Y_1^n | T_1^n) + \lambda I(X_1^n, X_2^n; Y_2^n | T_2^n)] + n\epsilon_n. \end{aligned} \quad (5.20)$$

The cyan (last) group is directly sub-additive (§1.13):

$$\begin{aligned} &(\lambda - 1)I(X_2^n; T_2^n) + I(X_1^n, X_2^n; Y_1^n | T_1^n) + \lambda I(X_1^n, X_2^n; Y_2^n | T_2^n) \\ &\leq \sum_{i=1}^n [(\lambda - 1)I(X_{2i}; T_{2i}) + I(X_{1i}, X_{2i}; Y_{1i} | T_{1i}) + \lambda I(X_{1i}, X_{2i}; Y_{2i} | T_{2i})]. \end{aligned}$$

For the red (first) group we use the chain rule and the Körner–Marton telescoping identity. Write

$$\begin{aligned} I(X_1^n; T_1^n | X_2^n, T_2^n) &= \sum_{i=1}^n [I(X_{1i}; T_{1i} | X_2^n, T_2^n, T_1^{i-1}, Y_{2,i+1}^n) + I(Y_{2,i+1}^n; T_{1i} | X_2^n, T_2^n, T_1^{i-1})], \\ I(X_1^n; Y_2^n | X_2^n, T_2^n) &= \sum_{i=1}^n [I(X_{1i}; Y_{2i} | X_2^n, T_2^n, T_1^{i-1}, Y_{2,i+1}^n) + I(T_1^{i-1}; Y_{2i} | X_2^n, T_2^n, Y_{2,i+1}^n)], \end{aligned}$$

and the Körner–Marton identity asserts

$$\sum_{i=1}^n I(Y_{2,i+1}^n; T_{1i} | X_2^n, T_2^n, T_1^{i-1}) = \sum_{i=1}^n I(T_1^{i-1}; Y_{2i} | X_2^n, T_2^n, Y_{2,i+1}^n).$$

Hence, since $\lambda \geq 1$,

$$I(X_1^n; T_1^n | X_2^n, T_2^n) - \lambda I(X_1^n; Y_2^n | X_2^n, T_2^n) \leq \sum_{i=1}^n [I(X_{1i}; T_{1i} | X_{2i}, T_{2i}, U_i) - \lambda I(X_{1i}; Y_{2i} | X_{2i}, T_{2i}, U_i)],$$

where $U_i = (X_2^{n \setminus i}, T_2^{n \setminus i}, T_1^{i-1}, Y_{2,i+1}^n)$ and the discarded $(\lambda - 1) \sum_i I(T_1^{i-1}; Y_{2i} | X_2^n, T_2^n, Y_{2,i+1}^n) \geq 0$ is dropped. By d -separation (independent inputs and a memoryless product channel), $X_{1i} - U_i - (X_{2i}, T_{2i})$ is a Markov chain. Symmetrically, for the blue (second) group,

$$I(X_2^n; T_2^n | X_1^n, T_1^n) - I(X_2^n; Y_1^n | X_1^n, T_1^n) \leq \sum_{i=1}^n [I(X_{2i}; T_{2i} | X_{1i}, T_{1i}, V_i) - I(X_{2i}; Y_{1i} | X_{1i}, T_{1i}, V_i)],$$

with $V_i = (X_1^{n \setminus i}, T_1^{n \setminus i}, T_2^{i-1}, Y_{1,i+1}^n)$ and $X_{2i} - V_i - (X_{1i}, T_{1i})$ a Markov chain.

(L3) Min–max / concave envelope. Assembling the three groups and letting $n \rightarrow \infty$, each single-letter functional is evaluated on the i -th letter of a product law, with the auxiliary conditioning U_i, V_i playing the role of time-sharing. The step replacing the averaged sliced functionals by the upper concave envelope over $p(x_1)p(x_2)$ is Jensen’s inequality applied

to the concave envelope (which majorises the functional and is concave by construction): $\mathbb{E}_U f_j(p^{(U)}) \leq \mathbb{E}_U \mathfrak{C}[f_j](p^{(U)}) \leq \mathfrak{C}[f_j](\mathbb{E}_U p^{(U)})$. The min-max theorem (Corollary 9.16 of the appendix) is the tool invoked in the applications of §5.8 below, where the “ $\sup_p \min_\lambda$ ” side of the bound is exchanged with the computable “ $\min_\lambda \sup_p$ ” side. This gives

$$R_1 + \lambda R_2 \leq \max_{p(x_1)p(x_2)} \left[(\lambda - 1)I(X_2; T_2) + \lambda I(X_1, X_2; Y_2|T_2) + I(X_1, X_2; Y_1|T_1) + \mathfrak{C}[f_1] + \mathfrak{C}[f_2] \right]. \quad (5.21)$$

It remains to rewrite (5.21) in the gap form (5.18). This is a term-by-term identity, verified as follows. Since $T_1 \perp\!\!\!\perp (X_2, T_2)$ and $T_2 \perp\!\!\!\perp (X_1, T_1)$,

$$I(X_1; T_1|X_2, T_2) = I(X_1; T_1), \quad I(X_2; T_2|X_1, T_1) = I(X_2; T_2),$$

and by the chain rule,

$$I(X_1, X_2; Y_1|T_1) = I(X_1; Y_1|T_1) + I(X_2; Y_1|X_1, T_1), \quad I(X_1, X_2; Y_2|T_2) = I(X_2; Y_2|T_2) + I(X_1; Y_2|X_2, T_2),$$

and $I(X_1; Y_1, T_1) = I(X_1; T_1) + I(X_1; Y_1|T_1)$, $I(X_2; Y_2, T_2) = I(X_2; T_2) + I(X_2; Y_2|T_2)$. Substituting these into

$$(\lambda - 1)I(X_2; T_2) + \lambda I(X_1, X_2; Y_2|T_2) + I(X_1, X_2; Y_1|T_1) + \mathfrak{C}[f_1] + \mathfrak{C}[f_2]$$

and subtracting $f_1 + f_2$ (which, with $\Gamma[f] = \mathfrak{C}[f] - f$, is the same as replacing $\mathfrak{C}[f_j]$ by $\Gamma[f_j]$), the T_1, T_2, Y_1, Y_2 terms cancel exactly, leaving

$$I(X_1; Y_1, T_1) + \lambda I(X_2; Y_2, T_2) + \Gamma[f_1] + \Gamma[f_2],$$

which is (5.18). ■

Remark 5.6 (Reduction to TIN). To reduce the genie bound (5.18) to the TIN bound

$$R_1 + \lambda R_2 \leq I(X_1; Y_1|Q) + \lambda I(X_2; Y_2|Q),$$

it suffices to choose the genies T_1, T_2 (the “smart genie”) so that

1. the gaps vanish: $\Gamma[f_1] = \Gamma[f_2] = 0$ (i.e. f_1, f_2 are already concave at the maximiser), and
2. the Markov chains $X_1 - Y_1 - T_1$ and $X_2 - Y_2 - T_2$ hold,

for then $I(X_1; Y_1, T_1) = I(X_1; Y_1)$ and $I(X_2; Y_2, T_2) = I(X_2; Y_2)$ and the max collapses to the TIN functional. This is exactly the route taken in the noisy-interference sum-capacity result below.

○

5.7.2 The auxiliary-receiver principle

The genie bounds above are instances of a more general and flexible technique for constructing outer bounds, developed in [Gohari and Nair \(2022\)](#). The idea is to introduce one or more *auxiliary receivers*: fictitious receivers whose outputs are generated from the system’s variables through auxiliary channels, and whose *messages need not be available to the legitimate receivers*. This last point is exactly where the technique generalises genie-aided bounds: a genie *gives* side information to a legitimate receiver, whereas an auxiliary receiver is a new node that only the analyst (not the code) can “observe”.

Why one needs auxiliary receivers. Consider erasing the output of every receiver with probability ϵ (the erasure event being known to the receivers). The standard single-letter outer bounds — the cut-set, UV , and genie bounds — all scale by a factor $(1 - \epsilon)$ under this transformation. The true capacity region, however, does *not* scale this way, as can be seen already from the achievability arguments. This discrepancy is the fingerprint of a missing ingredient in the classical bounds; auxiliary receivers supply it, because the outputs of the auxiliary channels need not undergo any erasure.

The template. Each auxiliary-receiver (“ J ”) bound is assembled from the same moves as the genie bounds of §5.7.1 — Fano, sub-additive single-letterisation with Gallager-type identification of the auxiliaries, and (where applicable) the min–max bridge of the appendix — together with one new step: *adding and subtracting the flow of information toward the auxiliary receiver*. Concretely,

- one bounds the flow of information *to the auxiliary receiver J* (an existing single-letter bound applied to J , or a cut-set-style bound), obtaining a term of the form $I(X; J, \cdot | \cdot)$;
- one bounds the *difference* of the flows to the legitimate receiver Y and to J , i.e. a term of the form $I(X; \cdot | W) - I(X; J | W)$ with a Gallager-type auxiliary W ;
- the sum of the two bounds the flow to the legitimate receiver.

A second benefit is structural: because the auxiliary receiver is fed through a *restricted* auxiliary channel (e.g. one that is forbidden from depending on certain inputs), the identification of the auxiliary random variables automatically inherits useful Markov chains, which is how the bounds sharpen the classical ones.

Example: the J -cutset bound for the relay channel. For the relay channel $p(y, y_r | x, x_r)$ of Chapter 8, the auxiliary-receiver version of the cut-set bound reads [Gohari and Nair \(2022\)](#): every achievable rate satisfies

$$R \leq \max_{p(x, x_r)} \min_{T_{J|X}} \max_{p(w|x, x_r)} \min \left\{ I(X, X_r; Y), I(X; Y, Y_r | X_r), \right. \\ \left. I(X; J, Y_r | X_r) + I(X, X_r; Y | W) - I(X, X_r; J | W) \right\},$$

where the auxiliary W satisfies the Markov chain $W \rightarrow (X, X_r) \rightarrow (Y, Y_r, J)$, and the auxiliary receiver J is generated from the transmitter’s input alone through an arbitrary auxiliary channel $T_{J|X}$ (so that $p(j, y, y_r | x, x_r) = T_{J|X}(j|x) p(y, y_r | x, x_r)$: the auxiliary receiver does not observe the relay’s input, and it does not affect the true channel). The first two terms are the two cut-set

terms of Chapter 8; the third term is the J -term, in which $I(X; J, Y_r | X_r)$ bounds the flow toward the auxiliary receiver J and $I(X, X_r; Y | W) - I(X, X_r; J | W)$ bounds the flow toward Y minus the flow toward J . This bound is strictly tighter than the cut-set bound for the Gaussian relay channel with non-zero channel gains [Gohari and Nair \(2022\)](#). A related strengthened-cutset upper bound for relay channels was developed by El Gamal, Gohari, and Nair (2022), using traditional weak-converse information inequalities and explicit auxiliary identification.

Scope of the technique. The auxiliary-receiver approach strictly improves classical outer bounds in all three basic settings: the relay channel (above), the Gaussian Z -interference channel (Chapter 5's subject, [Gohari and Nair \(2022\)](#)), and the two-receiver broadcast channel, where single- and two-auxiliary-receiver bounds strictly improve the UV outer bound of Chapter 4 [Gohari and Nair \(2022\)](#); [Gohari et al. \(2026\)](#). The enhanced genie-based outer bound for the sum-rate of interference channels of [Liu et al. \(2015\)](#) — which recovers all known capacity results for the scalar Gaussian interference channel from a single genie bound — is a close relative: there the “genie” is an auxiliary receiver whose message is handed to the legitimate receivers. In the language of §1.13 of Chapter 1, the auxiliary-receiver bounds are the most recent instalment of the same programme: declare the functional, control it by sub-additivity and data-processing inequalities, and let the min–max theorem of the appendix organise the boundary.

5.7.3 The Etkin–Ordentlich outer bound

Theorem 5.21. (Etkin–Ordentlich outer bound) The capacity region of any DM-IC is contained in

$$\begin{aligned} R_1 &\leq \min \{I(X_1; Y_1 | X_2, Q), I(U_2, X_1; Y_1 | Q), H(X_1 | U_1, Q)\}, \\ R_2 &\leq \min \{I(X_2; Y_2 | X_1, Q), I(U_1, X_2; Y_2 | Q), H(X_2 | U_2, Q)\}, \end{aligned} \quad (5.22)$$

for some $p(q)p(u_1, x_1 | q)p(u_2, x_2 | q)$. ○

Proof. We prove the three bounds on R_1 ; the bounds on R_2 are symmetric.

Cut-set bound. $R_1 \leq I(X_1; Y_1 | X_2, Q)$ is the cut-set bound (sub-additivity of $F = P_{X_1, X_2} \mapsto I(X_1; Y_1 | X_2)$, §1.13).

Second bound. Let $U_{2,i} = X_2^{i-1}$. By Fano and sub-additivity,

$$\begin{aligned} nR_1 &\leq I(X_1^n; Y_1^n) + n\epsilon_n \\ &= H(Y_1^n) - H(Y_1^n | X_1^n) \\ &= H(Y_1^n) - \sum_{i=1}^n H(Y_{1,i} | Y_1^{i-1}, X_1^n) \\ &\leq H(Y_1^n) - \sum_{i=1}^n H(Y_{1,i} | Y_1^{i-1}, X_2^{i-1}, X_1^n) \end{aligned}$$

$$\begin{aligned}
&\stackrel{(a)}{=} H(Y_1^n) - \sum_{i=1}^n H(Y_{1,i}|X_2^{i-1}, X_{1,i}) \\
&\leq \sum_{i=1}^n [H(Y_{1,i}) - H(Y_{1,i}|X_2^{i-1}, X_{1,i})] = \sum_{i=1}^n I(X_2^{i-1}, X_{1,i}; Y_{1,i}),
\end{aligned}$$

where (a) uses the Markov chain $(Y_1^{i-1}, X_1^{n \setminus i}) \rightarrow (X_{1,i}, X_2^{i-1}) \rightarrow Y_{1,i}$. With $U_2 = X_2^{Q-1}$ and time-sharing Q , this gives $R_1 \leq I(U_2, X_1; Y_1|Q)$. The functional here is $F = P_{X_1, X_2} \mapsto I(X_2^{i-1}, X_{1,i}; Y_{1,i})$; its sub-additivity is the entropy identity $H(Y_1^n) = \sum_i H(Y_{1,i}) - (\text{non-negative terms})$ used above.

Third bound. Let $U_{1,i} = X_1^{i-1}$. By Fano and the data-processing inequality,

$$nR_1 \leq I(M_1; Y_1^n) + n\varepsilon_n \leq I(X_1^n; Y_1^n) + n\varepsilon_n \leq H(X_1^n) + n\varepsilon_n,$$

where the last step is $I(X_1^n; Y_1^n) \leq H(X_1^n)$. The sub-additivity step here is the *chain-rule (entropy) identity*

$$H(X_1^n) = \sum_{i=1}^n H(X_{1,i}|X_1^{i-1}) = \sum_{i=1}^n H(X_{1,i}|U_{1,i}),$$

which after time-sharing gives $R_1 \leq H(X_1|U_1, Q)$. This is the “private-message” style bound. ■

5.8 Very Weak Interference and the Noisy-Interference Sum Capacity

We now ask: for which DM-ICs is treating interference as noise (TIN) sum-rate optimal? This motivates a third notion of interference strength.

Definition 5.22. (Very weak interference) A DM-IC has *very weak interference* if

$$\begin{aligned}
I(U_1; Y_1) &\geq I(U_1; Y_2|X_2), & \forall p(u_1, x_1)p(x_2); \\
I(U_2; Y_2) &\geq I(U_2; Y_1|X_1), & \forall p(u_2, x_2)p(x_1).
\end{aligned} \tag{5.23}$$

○

Proposition 5.23. (HK sum rate under very weak interference) For a DM-IC with very weak interference, the Han–Kobayashi sum rate is the TIN sum rate:

$$SR_{\text{HK}} = \max_Q [I(X_1; Y_1|Q) + I(X_2; Y_2|Q)].$$

○

Proof sketch. In the Han–Kobayashi region (5.11), the sum rate is maximised by allocating all rate to the private messages, i.e. by setting $U_1 = U_2 = \emptyset$ (equivalently $R_{10} = R_{20} = 0$). The very-weak condition $I(U_1; Y_1) \geq I(U_1; Y_2|X_2)$ (and its symmetric counterpart) states precisely

that the “common” message of the other user is easier to treat as noise than to decode, so the common-message terms $I(U_2, X_1; Y_1|Q)$ in the sum constraints are dominated by the private TIN terms $I(X_1; Y_1|Q)$, and the HK sum rate collapses to $\max_Q [I(X_1; Y_1|Q) + I(X_2; Y_2|Q)]$. This proposition is due to Liu–Nair–Xia 2014, where the Fourier–Motzkin/inequality verification that the common-message terms are dominated by the private TIN terms is carried out in full; we state it as a cited HK-region fact. It identifies the *Han–Kobayashi* sum rate only: very-weak interference does not by itself establish the actual capacity of the DM-IC (for that one needs an outer bound, e.g. the genie bound of §5.7.1). ■

Remark 5.7 (Conjecture). For a very weak DM-IC it is conjectured that TIN achieves the actual sum capacity, i.e. $\max(R_1 + R_2) = \max_{p(x_1)p(x_2)} [I(X_1; Y_1) + I(X_2; Y_2)]$ (time-sharing Q may be needed in general). This is the natural companion to the strong-interference result and is known to be true for the Gaussian channel under the condition below. The class of very weak channels was introduced and studied by Liu et al. (2014), who give a genie-based outer bound on the sum-rate of a general DM-IC and use it to compute the sum-capacity of the binary skew-symmetric Z-channel in a non-trivial regime of parameters; the skew-symmetric Z-channel is worked out as an example in §5.10 below. ○

We now give the flagship Gaussian result, due to Shang, Kramer and Chen (2008), which identifies the exact sum capacity in the *noisy-interference* regime.

Theorem 5.24. (Noisy-interference sum capacity, Shang–Kramer–Chen) Consider the scalar Gaussian IC. If the *noisy-interference condition*

$$a(1 + b^2 P_2) + b(1 + a^2 P_1) \leq 1 \quad (5.24)$$

holds, then the sum capacity is

$$C_{\text{sum}} = \frac{1}{2} \log \left(1 + \frac{P_1}{1+b^2 P_2} \right) + \frac{1}{2} \log \left(1 + \frac{P_2}{1+a^2 P_1} \right). \quad (5.25)$$

○

Proof. Achievability (TIN). Use TIN with Gaussian inputs $X_1 \sim \mathcal{N}(0, P_1)$, $X_2 \sim \mathcal{N}(0, P_2)$. Then

$$I(X_1; Y_1) = \frac{1}{2} \log \left(1 + \frac{P_1}{1+b^2 P_2} \right), \quad I(X_2; Y_2) = \frac{1}{2} \log \left(1 + \frac{P_2}{1+a^2 P_1} \right),$$

which sum to (5.25).

Converse. Apply the genie bound of Theorem 5.20 with $\lambda = 1$ and a *smart* choice of genies (noisy versions of the inputs, $T_1 = X_1 + \tilde{Z}_1$, $T_2 = X_2 + \tilde{Z}_2$ with suitably scaled noise). Under the noisy-interference condition, the gaps $\Gamma[f_1] = \Gamma[f_2]$ vanish at the Gaussian maximiser and the Markov conditions $X_1 - Y_1 - T_1$, $X_2 - Y_2 - T_2$ hold, so the bound collapses to the TIN functional

$$R_1 + R_2 \leq \max_{p(x_1)p(x_2)} [I(X_1; Y_1) + I(X_2; Y_2)].$$

The maximisation over product inputs is achieved by Gaussians and evaluates to (5.25): for Gaussian inputs the two terms are as above, and the noisy-interference condition is exactly what makes the concave envelope of the TIN sum functional equal its value at independent Gaussian inputs (this is the “TIN sum-rate additivity” of Shang–Kramer–Chen). The Gaussian optimality here is part of the cited theorem’s verification, not an immediate consequence of Lemma 1.59: the TIN sum functional involves interference terms, so plain single-user Gaussian maximality does not apply to it directly. The detailed verification that the chosen genies make the gap vanish term-by-term is a substantial computation; we state the result and refer to Shang–Kramer–Chen 2009 for the complete derivation. ■

Remark 5.8 (Two consequences). 1. Taking $Q = \emptyset$ and $X_1 \sim \mathcal{N}(0, P_1)$, $X_2 \sim \mathcal{N}(0, P_2)$ achieves $\max_{p(q)p(x_1|q)p(x_2|q)} [I(X_1; Y_1|Q) + I(X_2; Y_2|Q)]$, so no time-sharing is needed in the Gaussian case.

2. The multi-letter TIN functional is additive in the noisy-interference regime:

$$\max_{p(x_1^n)p(x_2^n)} [I(X_1^n; Y_1^n) + I(X_2^n; Y_2^n)] = n \max_{p(x_1)p(x_2)} [I(X_1; Y_1) + I(X_2; Y_2)],$$

i.e. the outer bound obtained by sub-additivity is tight, with no loss.

○

Remark 5.9 (An enhanced genie bound for the sum-rate). The genie bound of §5.7.1 admits a refinement, due to Liu et al. (2015), that produces an outer bound on the *sum-rate* $R_1 + R_2$ of an arbitrary memoryless interference channel; specialised to the scalar Gaussian IC, this single bound recovers *all* the known Gaussian capacity results — strong interference (§5.3.2), the noisy-interference sum capacity of Theorem 5.24, the Z-interference results below — in one stroke. The key step over the classical genie argument is a careful choice of the genie’s correlation with the past, which makes the discarded sub-additivity terms vanish exactly in the regimes where the corresponding achievability is optimal. This is the natural bridge from the genie bounds treated in this chapter to the auxiliary-receiver bounds of §5.7.2. ○

5.9 Mixed Interference

Recall the scalar Gaussian IC of §5.5. The channel has *mixed interference* if one cross-gain is above 1 and the other below, say $a \geq 1$ and $b \leq 1$ (the case $a \leq 1, b \geq 1$ is the mirror image). Then receiver 2 experiences strong interference, $I(X_1; Y_2|X_2) \geq I(X_1; Y_1|X_2)$, while receiver 1 experiences weak interference. The natural scheme is: receiver 1 (weak) treats interference as noise, while receiver 2 (strong) decodes both messages.

The following useful sum bound (a genie outer bound) is stated without a full proof; it follows from the genie bound of §5.7.1 by taking the genies to be the auxiliary U_1, U_2 .

Lemma 5.25. For any DM-IC, the sum rate is upper bounded as

$$\begin{aligned} R_1 + R_2 &\leq I(U_2, X_1; Y_1|Q) + I(X_2; Y_2|U_2, X_1, Q), \\ R_1 + R_2 &\leq I(U_1, X_2; Y_2|Q) + I(X_1; Y_1|U_1, X_2, Q), \end{aligned} \quad (5.26)$$

where the first bound holds for $p(q)p(u_2, x_2|q)p(x_1|q)$ (i.e. $(U_2, X_2) \perp\!\!\!\perp X_1 | Q$) and the second for $p(q)p(u_1, x_1|q)p(x_2|q)$. $\circ$

Theorem 5.26. (Achievable region under mixed interference) For the Gaussian IC with mixed interference $a \geq 1, b \leq 1$, every rate pair (R_1, R_2) satisfying

$$\begin{aligned} R_1 &\leq \frac{1}{2} \log \left(1 + \frac{P_1}{1+b^2P_2} \right), \\ R_2 &\leq \frac{1}{2} \log(1 + P_2), \\ R_1 + R_2 &\leq \frac{1}{2} \log(1 + P_2 + a^2P_1), \end{aligned} \quad (5.27)$$

is achievable. $\circ$

Proof. Transmitter 1 uses $X_1 \sim \mathcal{N}(0, P_1)$; transmitter 2 uses $X_2 \sim \mathcal{N}(0, P_2)$. Receiver 1 (weak) treats X_2 as noise, giving

$$R_1 \leq I(X_1; Y_1) = \frac{1}{2} \log \left(1 + \frac{P_1}{1+b^2P_2} \right),$$

the unconditional mutual information under noise-plus-interference; the conditional quantity $I(X_1; Y_1|X_2)$ would instead be $\frac{1}{2} \log(1 + P_1)$, which is not the TIN rate. Receiver 2 (strong) decodes both messages by simultaneous decoding, giving the MAC constraints

$$R_2 \leq I(X_2; Y_2|X_1) = \frac{1}{2} \log(1 + P_2), \quad R_1 + R_2 \leq I(X_1, X_2; Y_2) = \frac{1}{2} \log(1 + P_2 + a^2P_1).$$

This yields (5.27). (Note that the first bound is the correct TIN rate $\frac{1}{2} \log(1 + \frac{P_1}{1+b^2P_2})$; the value $\frac{1}{2} \log(1 + P_1)$ would only be attainable if receiver 1 could remove the interference bX_2 , which it cannot do while user 2 transmits at power P_2 .) $\blacksquare$

5.10 Z-Interference and Binary Deterministic Examples

5.10.1 Clean Z-interference

An interference channel has *clean Z-interference* (CZI) if one of the sub-channels is clean, i.e. one receiver observes its own input exactly while the other receiver sees both inputs. We take $Y_2 = X_2$ (clean) and $Y_1 \sim W(y_1|x_1, x_2)$.

Example 5.27. (CZI has very weak interference) For a CZI with $Y_2 = X_2$, the very-weak condition holds. First,

$$I(U_1; Y_2|X_2) = I(U_1; X_2|X_2) = 0 \leq I(U_1; Y_1)$$

is immediate. Second, because $X_1 \perp\!\!\!\perp (U_2, X_2)$ and $U_2 \rightarrow X_2 \rightarrow Y_1$ is a Markov chain given X_1 , the data-processing inequality gives

$$I(U_2; Y_2) = I(U_2; X_2) = I(U_2; X_2|X_1) \geq I(U_2; Y_1|X_1),$$

where the middle equality uses $X_1 \perp\!\!\!\perp (U_2, X_2)$. Hence CZI channels fall under the very-weak class, and TIN is their HK sum-rate-optimal scheme.

Remark 5.10 (HK region for CZI). For the clean Z-interference channel, the Han–Kobayashi region reduces (setting $U_1 = \emptyset$) to

$$\begin{aligned} R_1 &< I(X_1; Y_1|U_2, Q), \\ R_2 &< H(X_2|Q), \\ R_1 + R_2 &< I(X_1, U_2; Y_1|Q) + H(X_2|U_2, Q). \end{aligned} \tag{5.28}$$

Indeed, with $Y_2 = X_2$ one takes $U_2 = X_2$ or a split of it; the terms $I(X_2; Y_2|\cdot)$ become $H(X_2|\cdot)$ and the user-1-common auxiliary U_1 is dropped because user 1 does not interfere with receiver 2. $\circ$

Lemma 5.28. (Han–Kobayashi weighted sum for CZI) For a CZI channel, for every $\lambda \geq 1$, the maximum of the weighted sum $R_1 + \lambda R_2$ over the Han–Kobayashi region (5.11) equals

$$\max_{p(x_1)p(x_2)} [I(X_1; Y_1) + \lambda I(X_2; Y_2)].$$

$\circ$

Proof sketch. Since $Y_2 = X_2$, receiver 2’s link is interference-free and $I(X_2; Y_2) = H(X_2)$. Under the very-weak property, the same Fourier–Motzkin verification that collapses the HK *sum* rate in §5.8 (due to Liu–Nair–Xia 2014) collapses the weighted sum to the private TIN terms, and the achievability is TIN. We state this as an *HK-region* fact only: promoting it to the actual weighted-sum capacity $\max(R_1 + \lambda R_2)$ of the channel requires an outer bound at the corresponding strength (e.g. the genie bound of §5.7.1), which is not established here. $\blacksquare$

5.10.2 The binary skew-symmetric Z-IC

Example 5.29. (Binary skew-symmetric Z-IC) Consider the binary deterministic one-sided channel

$$Y_1 = X_1 \oplus X_2, \quad Y_2 = X_2,$$

with $X_1, X_2 \in \{0, 1\}$ and $\oplus$ mod-2 addition. This is a clean Z-IC in which user 1’s signal is corrupted by user 2’s signal at receiver 1, while user 2’s signal reaches receiver 2 cleanly. It is an injective deterministic IC with $T_1 = \emptyset$ (user 1 does not reach receiver 2) and $T_2 = X_2$ (the map $x_2 \mapsto x_1 \oplus x_2$ is injective for every x_1). Its capacity region is

$$\mathcal{C} = \{(R_1, R_2) : R_1 \geq 0, R_2 \geq 0, R_1 + R_2 \leq 1\}.$$

Achievability. Take $X_1 \sim \text{Bern}(\frac{1}{2})$ uniform and X_2 with $H(X_2) \geq R_2$. Receiver 2 decodes M_2 cleanly at rate R_2 . Receiver 1 decodes *both* messages jointly: since $Y_1 = X_1 \oplus X_2$ has $H(Y_1) = 1$ bit when X_1 is uniform, the MAC constraint at receiver 1 is $R_1 + R_2 \leq H(Y_1) = 1$, so any split with $R_1 + R_2 \leq 1$ is achievable. (User 2’s message acts as a “common” message, decoded by both receivers.)

Converse. $R_1 + R_2 \leq I(X_1, X_2; Y_1|Q) = H(X_1 \oplus X_2|Q) \leq 1$ by the cut-set bound, and $R_2 \leq H(X_2|Q) \leq 1$, $R_1 \leq H(Y_1|X_2, Q) = H(X_1|Q) \leq 1$. The region $\{R_1 + R_2 \leq 1\}$ is the convex hull of the achievable union, so it is tight. The sum capacity is 1 bit, and it is achieved even though receiver 1 sees X_1 only through the XOR with the interference X_2 : the point is that receiver 1 decodes the *pair* (M_1, M_2) from the single bit Y_1 *subject to the sum constraint* $R_1 + R_2 \leq 1$. This should not be read as recovering two arbitrary independent one-bit streams from a one-bit observation; rather, Y_1 identifies one element of a message-pair set of size $2^{R_1+R_2} \leq 2$, which is exactly the MAC sum constraint.

5.10.3 The binary deterministic IC (El Gamal–Costa model)

A deterministic IC is *binary* if $\mathcal{X}_1 = \mathcal{X}_2 = \mathcal{Y}_1 = \mathcal{Y}_2 = \{0, 1\}$ and $Y_1 = f_1(X_1, X_2)$, $Y_2 = f_2(X_1, X_2)$. The El Gamal–Costa model is the class of deterministic ICs (typically one-sided, $Y_2 = f_2(X_2)$, i.e. a Z-IC); the binary skew-symmetric example above is its canonical instance. The general capacity region of an arbitrary binary deterministic IC is not known in closed form, but the Etkin–Ordentlich outer bound (Theorem 5.21) applies verbatim and is often tight.

For the one-sided deterministic IC $Y_1 = f_1(X_1, X_2)$, $Y_2 = f_2(X_2)$, the El Gamal–Costa analysis reduces the problem to a “Z” form: receiver 2’s clean link imposes $R_2 \leq H(X_2|Q)$, while receiver 1’s constraints are the MAC-like bounds from decoding the pair (X_1, X_2) . For the XOR example the three bounds collapse to the single constraint $R_1 + R_2 \leq 1$, giving the exact region above. For general f_1, f_2 the capacity region is obtained by specialising the injective-deterministic theorem of §5.6 (whenever f_1 is injective in X_2 for each X_1 and f_2 is injective in its argument), which is exactly the Han–Kobayashi region with $U_1 = t_1(X_1)$, $U_2 = t_2(X_2)$.

Remark 5.11 (Further reading: the Gaussian Z-interference channel). The Gaussian Z-interference channel ($b = 0$ in the model of §5.5, so that receiver 1 is interference-free while receiver 2 still sees aX_1) has been studied intensely because its capacity region is open while being considerably more tractable than the general IC. A non-exhaustive sampler: the optimality of TIN “around the corner” points and the structure of the Han–Kobayashi boundary [Costa and Nair \(2016\)](#); the noiseberg achievable region, which dominates both time-sharing and TIN, introduced in [Costa \(2011\)](#), and the proof of the noiseberg conjecture [Costa et al. \(2023\)](#); an information inequality motivated by the Z-channel [Gohari et al. \(2021\)](#); and the local optimality of Gaussian inputs for the Han–Kobayashi inner bound [Nair and Zhao \(2026\)](#). The auxiliary-receiver bounds of §5.7.2 currently give the tightest known outer bounds for this channel. $\circ$

6 Channels with state

The point-to-point channel of Chapter 2 is fixed: the same conditional law $W(y|x)$ governs every use of the channel. In many realistic links, the channel itself changes over time in a way that is partly outside the communication system's control. We model this by a *state* sequence S^n : at time i the channel is $W(y|x, S_i)$, and the state may be revealed — causally, non-causally, or not at all — to the transmitter and to the receiver.

This chapter has two distinct mathematical flavours, and we flag them honestly throughout. For the settings in which the state is revealed at the receiver, or in which it is merely averaged over, the converses are instances of the *sub-additivity* programme of Chapter 1 (§1.13): a Fano argument produces a multi-letter functional, and sub-additivity of that functional (plus, in the Gaussian cases, Gaussian optimality via Lemma 1.59) collapses it to a single-letter expression. By contrast, the two transmitter-only-state results — Shannon's causal-CSIT theorem and the Gel'fand–Pinsker theorem — use genuinely different techniques (a *strategy* auxiliary variable, and the *Csiszár sum identity* together with *binning*); they are *not* sub-additivity arguments, and we say so explicitly where they appear.

6.1 Setting and code definitions

A *discrete memoryless channel (DMC) with state* consists of finite alphabets $\mathcal{X}, \mathcal{Y}, \mathcal{S}$ and a stochastic matrix

$$W : \mathcal{X} \times \mathcal{S} \rightarrow \mathcal{Y}, \quad W(y|x, s) \triangleq \Pr(Y = y \mid X = x, S = s).$$

The state sequence $S^n = (S_1, \dots, S_n)$ is drawn i.i.d. according to a state distribution $p(s)$ that is *known to all parties*; the transmitter and receiver may additionally be given side information about the realisation S^n . The channel is memoryless in the sense that

$$\Pr(y^n \mid x^n, s^n) = \prod_{i=1}^n W(y_i \mid x_i, s_i).$$

Two degenerate extremes bracket the theory. If $S_1 = \dots = S_n = S$ is constant over the block (drawn once from $p(s)$), we obtain the *compound channel* (§6.2); if $S^n \sim p(s)^{\otimes n}$ is i.i.d., we obtain the *time-varying* channels of §6.3. Unless stated otherwise, “state” means the i.i.d. model.

Definition 6.1. (Code with state information) An $(n, 2^{nR})$ code for the state-dependent channel, with a prescribed state-information pattern, consists of:

- a message set $[1 : 2^{nR}]$, with M uniform on it and $M \perp\!\!\!\perp S^n$;
- an encoder that produces X^n from (M, S^n) through functions $x_i : [1 : 2^{nR}] \times S^n \rightarrow \mathcal{X}$ if the transmitter has *non-causal* state information, through $x_i : [1 : 2^{nR}] \times S^i \rightarrow \mathcal{X}$ if it has *causal* state information (X_i may depend on S^i but not on future S_{i+1}^n), or through $x_i : [1 : 2^{nR}] \rightarrow \mathcal{X}$ if it has no state information;
- a decoder $g : \mathcal{Y}^n \times S^n \rightarrow [1 : 2^{nR}]$ if the receiver has state information (causal or non-causal – since decoding is done only after the whole block Y^n is observed, the two are equivalent at the receiver), or $g : \mathcal{Y}^n \rightarrow [1 : 2^{nR}]$ if it does not.

The average probability of error is

$$P_e^{(n)} \triangleq \Pr(g(Y^n, \cdot) \neq M),$$

where the probability is over the uniform message, the state sequence, and the channel. A rate R is *achievable* if there is a sequence of $(n, 2^{nR})$ codes with $P_e^{(n)} \rightarrow 0$; the *capacity* is the supremum of achievable rates. $\circ$

For the compound channel the same definition applies with the understanding that $S_1 = \dots = S_n = S$, that the code is designed without knowing the value of S , and that the error probability is required to be small for *every* $s \in \mathcal{S}$ (see §6.2).

Remark 6.1 (The sub-additivity template). In every converse below the first step is Fano's inequality (Proposition 1.57). Written at the decoder's full observation it takes the form

$$nR \leq I(M; Y^n) + n\epsilon_n, \quad \epsilon_n \rightarrow 0,$$

when the receiver knows only Y^n (no CSIR), and

$$nR \leq I(M; Y^n, S^n) + n\epsilon_n$$

when the receiver knows the state (CSIR); the two templates are not interchangeable, and $nR \leq I(M; Y^n) + o(n)$ is not asserted for the receiver-known-state cases. The second step is the object of study: we must bound a multi-letter mutual information by a sum of single-letter functionals. For the receiver-known and averaged settings this second step is *sub-additivity* in the sense of §1.13; for the transmitter-only settings it is a different identity (a Markov/strategy argument, or the Csiszár sum identity), and the labelling below marks each step as such. $\circ$

6.2 The compound channel

Definition 6.2. In the *compound channel*, the state is fixed over the block, $S_1 = \dots = S_n = S$, with S drawn from $p(s)$. The transmitter does not know s ; the receiver *does* know s . A rate R is achievable if there is a single sequence of codes (independent of s) with $\sup_{s \in \mathcal{S}} P_e^{(n)}(s) \rightarrow 0$, where $P_e^{(n)}(s)$ is the error probability when the state happens to be s . $\circ$

Theorem 6.3. (Compound channel capacity)

$$C = \max_{p(x)} \min_{s \in \mathcal{S}} I(X; Y | S = s),$$

where $I(X; Y | S = s)$ is computed on the single-letter channel $W(\cdot|x, s)$ with input $X \sim p(x)$.

○

Proof. Converse. Fix a code with $\sup_s P_e^{(n)}(s) \leq \epsilon$. Since the transmitter does not see s , its input law $p(x^n)$ does not depend on s , and the decoder, knowing s , decodes over the channel $W_s(\cdot|x) = W(\cdot|x, s)$. For each fixed s , Fano's inequality (Proposition 1.57) gives

$$H(M | Y^n, S = s) \leq n\epsilon_n, \quad \epsilon_n \rightarrow 0,$$

so that

$$\begin{aligned} nR = H(M) &\leq I(M; Y^n | S = s) + n\epsilon_n \\ &\leq I(X^n; Y^n | S = s) + n\epsilon_n \quad (\text{data processing: } M \rightarrow X^n \rightarrow Y^n \text{ given } S = s) \\ &\leq \sum_{i=1}^n I(X_i; Y_i | S = s) + n\epsilon_n. \quad (\text{sub-additivity of } F_s = I(X; Y | S = s)) \end{aligned} \quad (6.1)$$

The labelled step is sub-additivity of the functional $F_s(p_{XY}) \triangleq I(X; Y | S = s)$ over the memoryless channel: $I(X^n; Y^n | S = s) \leq \sum_i I(X_i; Y_i | S = s)$, exactly the computation of §1.13 specialised to the fixed channel W_s (use the Markov chain $(X^n \setminus i, Y^{i-1}) \rightarrow X_i \rightarrow Y_i$ under $S = s$).

Since (6.1) holds for every s , taking the minimum over s gives

$$nR \leq \min_{s \in \mathcal{S}} \sum_{i=1}^n I(X_i; Y_i | S = s) + n\epsilon_n.$$

It remains to single-letterise. For each s , the functional $p(x) \mapsto I(X; Y | S = s)$ is concave in the input distribution (mutual information is concave in P_X ; the perturbation method of §1.13 proves it). Writing $\bar{p}(x) = \frac{1}{n} \sum_i p_{X_i}(x)$ for the averaged input law, concavity gives, for every s ,

$$\frac{1}{n} \sum_{i=1}^n I(X_i; Y_i | S = s) \leq I(X; Y | S = s) \Big|_{\bar{p}}.$$

The left and right sides are functions of s , the inequality holds pointwise in s with the *same* averaged input law $\bar{p}$ on the right, and pointwise domination passes through the minimum: minimising both sides over $s \in \mathcal{S}$ yields

$$\frac{1}{n} \min_s \sum_{i=1}^n I(X_i; Y_i | S = s) \leq \min_s I(X; Y | S = s) \Big|_{\bar{p}} \leq \max_{p(x)} \min_s I(X; Y | S = s),$$

hence

$$nR \leq n \max_{p(x)} \min_s I(X; Y | S = s) + n\epsilon_n.$$

Letting $n \rightarrow \infty$ yields $R \leq \max_{p(x)} \min_s I(X; Y | S = s)$. (This is the finite-family compound-channel formula; the minimisation over s and the time averaging commute through the pointwise bound above, and the *classical* version with an uninformed decoder is the compound-channel theorem of Blackwell–Breiman–Thomasian and Wolfowitz, cited below.)

Achievability. Let $p^*(x)$ attain the maximum (the simplex is compact and $I(X; Y | S = s)$ is continuous in $p(x)$ for finite alphabets, so a maximiser exists). Choose $R < \min_s I(X; Y | S = s)|_{p^*}$. Generate one random codebook C of 2^{nR} codewords $x^n(m)$ i.i.d. $\sim p^{*\otimes n}$, and reveal it to both parties; the decoder, knowing s , uses the joint-typicality decoder for W_s . For each fixed s , standard random coding over the channel W_s of rate $R < I(X; Y | S = s)|_{p^*}$ gives $E_C P_e^{(n)}(s, C) \rightarrow 0$. Because the state set is finite,

$$E_C \sum_{s \in \mathcal{S}} P_e^{(n)}(s, C) = \sum_{s \in \mathcal{S}} E_C P_e^{(n)}(s, C) \rightarrow 0,$$

so there exists a *single* codebook C whose errors under all states tend to zero (a statement that each state has a good random code alone does not select one common codebook). Thus any $R < \max_{p(x)} \min_s I(X; Y | S = s)$ is achievable. ■

Remark 6.2 (Broadcast interpretation). One may think of the compound channel as a broadcast channel with $|\mathcal{S}|$ “virtual” receivers, receiver s observing Y^n through W_s , all of them demanding the *common* message M . Capacity is then the common-message capacity of this broadcast channel, which is exactly $\max_{p(x)} \min_s I(X; Y | S = s)$. ○

Remark 6.3 (The classical compound channel). If the receiver also does *not* know s , the capacity is the same expression. The converse is identical (Fano per state plus the sub-additivity step (6.1)); the achievability is more delicate and is the classical compound-channel theorem of Blackwell–Breiman–Thomasian and Wolfowitz, which we do not need here. ○

6.3 Channels that change over time: the averaging argument

We now take $S^n \sim p(s)^{\otimes n}$ i.i.d.. In this section the state is either revealed at the receiver or merely averaged over; in all cases the capacity is obtained by an *averaging argument*, and the converse is a sub-additivity step (over the time index, over the state slices, or both). The two remaining cases, in which the state is known only at the transmitter, are the subject of §6.4 and §6.5.

6.3.1 State known to neither party

If neither the transmitter nor the receiver sees S^n , the system sees the *averaged channel*

$$\hat{W}(y|x) \triangleq \sum_{s \in \mathcal{S}} p(s) W(y|x, s).$$

The capacity is the ordinary point-to-point capacity of $\hat{W}$:

$$C = \max_{p(x)} I(X; Y), \quad \text{with } (X, Y) \sim p(x)\hat{W}(y|x).$$

This is immediate: the marginal output law induced by any input x^n is $\prod_i \hat{W}(y_i|x_i)$, so the state may be marginalised out. The converse is the sub-additivity of $I(X; Y)$ over the memoryless channel $\hat{W}$: $I(X^n; Y^n) \leq \sum_i I(X_i; Y_i)$ (§1.13), and the achievability is the ordinary channel coding theorem.

6.3.2 State known at the receiver only (CSIR)

The decoder observes S^n (causally or non-causally — the two coincide because decoding is block-wise), while the transmitter is ignorant of S^n .

Proposition 6.4. (CSIR capacity)

$$C = \max_{p(x)} I(X; Y | S) = \max_{p(x)} \sum_{s \in \mathcal{S}} p(s) I(X; Y | S = s),$$

where $X \perp\!\!\!\perp S$ (the transmitter cannot adapt X to S). ○

Proof. Converse. By Fano and the fact that the receiver's observation is the pair (Y^n, S^n) ,

$$nR \leq I(M; Y^n, S^n) + n\epsilon_n = I(M; Y^n | S^n) + n\epsilon_n$$

(because $M \perp\!\!\!\perp S^n$, so $I(M; S^n) = 0$). Data processing then gives

$$nR \leq I(X^n; Y^n | S^n) + n\epsilon_n \leq \sum_{i=1}^n I(X_i; Y_i | S_i) + n\epsilon_n,$$

where the last inequality is (*sub-additivity of $F = I(X; Y | S)$*): since the channel is memoryless, $I(X^n; Y^n | S^n) \leq \sum_i I(X_i; Y_i | S_i)$. Each term is evaluated on a law $p(x_i)p(s)W(y|x, s)$ with $X_i \perp\!\!\!\perp S_i$ (the encoder has no state), so each term is at most $\max_{p(x)} I(X; Y | S)$. Hence $R \leq \max_{p(x)} I(X; Y | S)$.

Achievability. Regard $X \mapsto (Y, S)$ as an ordinary DMC with output alphabet $\mathcal{Y} \times \mathcal{S}$ and law $p(y, s|x) = p(s)W(y|x, s)$. Its capacity is $\max_{p(x)} I(X; Y, S)$, and since $X \perp\!\!\!\perp S$ we have $I(X; Y, S) = I(X; S) + I(X; Y | S) = I(X; Y | S)$. The ordinary channel coding theorem applies (the decoder uses both Y_i and S_i). ■

Remark 6.4 (Averaging over state slices). The identity $I(X; Y | S) = \sum_s p(s)I(X; Y | S = s)$ is the “averaging over state slices” step: the CSIR capacity is the average of the conditional mutual informations, optimised over a single input law (which cannot depend on s). The converse's single-letter collapse is sub-additivity of the functional $F = I(X; Y | S)$; there is no binning anywhere in this argument. ○

6.3.3 State known at the transmitter and receiver (CSIT + CSIR)

Now the transmitter also knows the state. There are two sub-cases:

Non-causal CSIT: the encoder knows the whole S^n before transmitting. *Causal CSIT*: at time i the encoder knows only S^i .

Proposition 6.5. (CSIT+CSIR capacity) Both with causal and with non-causal CSIT, and with CSIR,

$$C = E[C_S] = \sum_{s \in S} p(s) C_s, \quad C_s \triangleq \max_{p(x)} I(X; Y | S = s)$$

(the point-to-point capacity of the channel W_s). ○

Sketch. Converse. Exactly as in Proposition 6.4, Fano and sub-additivity give

$$nR \leq \sum_{i=1}^n I(X_i; Y_i | S_i) + n\epsilon_n.$$

Now X_i may depend on S_i , so for each i ,

$$I(X_i; Y_i | S_i) = \sum_s p(s) I(X_i; Y_i | S_i = s) \leq \sum_s p(s) C_s = E[C_S],$$

using that $S_i \sim p(s)$ has the correct marginal and that each conditional term is at most the per-state capacity C_s . Hence $R \leq E[C_S] + o(1)$.

Achievability (Shannon multiplexing). For each state s with $p(s) > 0$ choose a small slack $0 < \delta_s < p(s)$, set

$$n_s = \lfloor n(p(s) - \delta_s) \rfloor,$$

and build a point-to-point code of length n_s and rate strictly below C_s for the channel W_s , with codewords drawn from the capacity-achieving $p_s^*(x)$; for states with $C_s = 0$ (or $p(s) = 0$) use a singleton message set rather than a negative-rate code. Split the message into submessages M_s of size equal to that code's message set, so submessage s carries $\approx n_s C_s$ bits. At time i , both parties know $S_i = s$, so the transmitter sends the next unused symbol of the codeword for M_s over channel W_s , and the receiver appends Y_i to the list of outputs for state s ; once the n_s codeword symbols of a state are exhausted, both parties send/ignore a fixed symbol for that state. By the law of large numbers, the number N_s of uses of state s satisfies $\Pr(N_s < n_s) \rightarrow 0$ (since S^n is i.i.d.), so with probability tending to one the receiver collects the full n_s outputs of each channel W_s and decodes every submessage. The total rate is $\sum_s \frac{n_s}{n} (C_s - \eta_s)$ for slacks $\eta_s \downarrow 0$, and letting $\delta_s, \eta_s \downarrow 0$ after $n \rightarrow \infty$ gives $\sum_s p(s) C_s = E[C_S]$. ■

The causal case is identical: to multiplex at time i it suffices to know the *current* state S_i (only the current state is needed to select the state-specific symbol stream), so causal CSIT is as good as non-causal CSIT here. ■

Remark 6.5 (Sub-additivity labels). The collapse $I(X^n; Y^n | S^n) \leq \sum_i I(X_i; Y_i | S_i)$ is sub-additivity of $F = I(X; Y | S)$, and the subsequent bound $I(X_i; Y_i | S_i) \leq E[C_S]$ is the “averaging over state slices” bound. No binning is used; these are cut-set flavoured steps. ○

6.4 State known causally at the transmitter: Shannon strategies

We now remove the receiver's knowledge of the state, and first consider *causal* CSIT: at time i the encoder knows (M, S^i) , the decoder knows only Y^n . This is Shannon's model (1958).

A *strategy* is a function $t : \mathcal{S} \rightarrow \mathcal{X}$; let $\mathcal{T} = \mathcal{X}^{\mathcal{S}} = \{t : s \mapsto x\}$ be the set of all strategies. The central idea is that the encoder, rather than deciding X_i symbol by symbol, chooses at the start a whole strategy and then applies it causally: if the strategy t is chosen and the state at time i is S_i , transmit $X_i = t(S_i)$.

Theorem 6.6. (Shannon 1958; causal CSIT, no CSIR) The capacity of the state-dependent channel $W(y|x, s)$ with state $p(s)$ and state information known causally at the encoder only is

$$C = \max_{p(u), x(u,s)} I(U; Y),$$

where U is an auxiliary random variable (a *strategy* random variable), $x : \mathcal{U} \times \mathcal{S} \rightarrow \mathcal{X}$ is a deterministic function, and the induced DMC from U to Y is

$$p(y|u) = \sum_{s \in \mathcal{S}} p(s) W(y | x(u, s), s).$$

○

In other words, the original channel is replaced by a memoryless channel whose input is a strategy u (a full rule $s \mapsto x$) and whose output is Y ; the capacity is the ordinary capacity of this “strategy channel”.

Proof. Achievability. Fix $p(u)$ and $x(u, s)$. Regard $U \rightarrow Y$ with law $p(u)p(y|u)$ as a point-to-point DMC. By the channel coding theorem, any $R < I(U; Y)$ is achievable with a codebook $u^n(m)$ i.i.d. $\sim p(u)^{\otimes n}$ and joint-typicality decoding of $(u^n(m), y^n)$. To send m , at time i the encoder (knowing S_i) transmits $X_i = x(u_i(m), S_i)$. Because $S_i \sim p(s)$ is i.i.d. and independent of $u_i(m)$, the output Y_i has exactly the law $p(y|u_i(m))$; the encoder uses only S_i (causality), and the decoder decodes over the induced channel. Hence $R < I(U; Y)$ is achievable.

Converse. Consider any causal-CSIT code. Let

$$U_i \triangleq (M, S^{i-1}).$$

Then X_i is a function of $(M, S^i) = (U_i, S_i)$. Fano gives

$$nR \leq I(M; Y^n) + n\epsilon_n.$$

Now

$$I(M; Y^n) = \sum_{i=1}^n I(M; Y_i | Y^{i-1})$$

$$\begin{aligned}
&= \sum_{i=1}^n \left(I(M, Y^{i-1}; Y_i) - I(Y^{i-1}; Y_i) \right) \\
&\leq \sum_{i=1}^n I(M, Y^{i-1}; Y_i) && (\text{sub-additivity: drop } I(Y^{i-1}; Y_i) \geq 0) \\
&\leq \sum_{i=1}^n I(M, Y^{i-1}, S^{i-1}; Y_i) && (\text{monotonicity}) \\
&= \sum_{i=1}^n I(M, S^{i-1}; Y_i) && (I(Y^{i-1}; Y_i \mid M, S^{i-1}) = 0) \\
&= \sum_{i=1}^n I(U_i; Y_i). \tag{6.2}
\end{aligned}$$

For the step (6.2): by d-separation (§9.4 of the Appendix), the output Y_j touches the rest of the system only through (M, S_j) and its own channel noise, so conditioning on (M, S^{i-1}) cuts all paths from Y^{i-1} to Y_i ; hence $Y^{i-1} \perp\!\!\!\perp Y_i \mid (M, S^{i-1})$.

It remains to single-letterise the functional $F = I(U; Y)$. Introduce $Q \sim \text{Unif}[1 : n]$ independent of everything, and set $U = (U_Q, Q)$, $Y = Y_Q$, $X = X_Q$. Then

$$\sum_{i=1}^n I(U_i; Y_i) = n I(U_Q; Y_Q \mid Q) \leq n I(U_Q, Q; Y_Q) = n I(U; Y).$$

The induced channel from $U = (u, q)$ to Y is $p(y|u, q) = \sum_s p(s)W(y \mid x_q(u, s), s)$, i.e. a single DMC from the enlarged strategy alphabet with strategy $s \mapsto x_q(u, s)$. Therefore

$$nR \leq n I(U; Y) + n\epsilon_n \leq n \max_{p(u), x(u, s)} I(U; Y) + n\epsilon_n,$$

and $R \leq \max_{p(u), x(u, s)} I(U; Y)$ follows on letting $n \rightarrow \infty$. ■

Remark 6.6 (Strategy alphabet and cardinality). U ranges over the finite set of strategies $\mathcal{T} = \mathcal{X}^{\mathcal{S}}$, so $|\mathcal{U}| \leq |\mathcal{X}|^{|\mathcal{S}|}$ suffices. The strategy interpretation explains the achievable scheme: the encoder picks $u^n(m)$ (a sequence of strategies) ahead of time, then at time i simply evaluates $x(u_i(m), S_i)$; causality is automatic. The sub-additivity flavour of the converse is exactly the collapse $I(M; Y^n) \leq \sum_i I(U_i; Y_i)$ of the functional $F = I(U; Y)$. ○

Remark 6.7 (Why the strategy is needed). If the encoder tried to transmit a plain symbol X_i independent of the state, it would be limited to the averaged channel $\hat{W}$ of §6.3.1. The strategy gains because the transmitter can *adapt* X_i to the observed S_i ; the auxiliary U is the “name” of the adaptation rule. This is the causal analogue of the auxiliary U in the non-causal theorem next. ○

6.5 State known non-causally at the transmitter: Gelfand–Pinsker

The final and deepest case: the transmitter knows the entire S^n *before* transmitting (non-causal CSIT), the receiver knows nothing about the state. The transmitter can therefore pre-compensate for the interference S^n .

Theorem 6.7. (Gel’fand–Pinsker 1980) The capacity of the state-dependent channel $W(y|x, s)$ with i.i.d. state $S^n \sim p(s)^{\otimes n}$ and state information known non-causally at the encoder only is

$$C = \max_{p(u|s), x(u,s)} (I(U; Y) - I(U; S)),$$

where the auxiliary U ranges over an alphabet with $|\mathcal{U}| \leq |\mathcal{S}||\mathcal{X}|$, the joint law is

$$p(s) p(u|s) \mathbf{1}\{x = x(u, s)\} W(y|x, s),$$

and $x : \mathcal{U} \times \mathcal{S} \rightarrow \mathcal{X}$ is a deterministic function (equivalently a stochastic mapping $p(x|u, s)$; see Remark 6.8). $\circ$

Before the proof we record the one identity that drives the converse, and flag that it is *not* a sub-additivity statement.

Lemma 6.8. (Csiszár sum identity) For any random variables M, S^n, Y^n (no independence assumptions needed),

$$\sum_{i=1}^n I(S_{i+1}^n; Y_i | Y^{i-1}, M) = \sum_{i=1}^n I(Y^{i-1}; S_i | S_{i+1}^n, M),$$

where $S_{i+1}^n \triangleq (S_{i+1}, \dots, S_n)$ and $S_{n+1}^n \triangleq \emptyset$. $\circ$

Proof. Expand each term and telescope. Write the left-hand side as

$$\text{LHS} = \sum_{i=1}^n \left(H(S_{i+1}^n | Y^{i-1}, M) - H(S_{i+1}^n | Y^i, M) \right).$$

Using the chain rule twice,

$$H(S_{i+1}^n | Y^{i-1}, M) = H(S_i^n | Y^{i-1}, M) - H(S_i | S_{i+1}^n, Y^{i-1}, M),$$

so

$$\text{LHS} = \sum_{i=1}^n H(S_i^n | Y^{i-1}, M) - \sum_{i=1}^n H(S_{i+1}^n | Y^i, M) - \sum_{i=1}^n H(S_i | S_{i+1}^n, Y^{i-1}, M).$$

The first two sums telescope: writing $S_i^n = (S_i, \dots, S_n)$,

$$\sum_{i=1}^n H(S_i^n | Y^{i-1}, M) - \sum_{i=1}^n H(S_{i+1}^n | Y^i, M) = H(S_1^n | M) - H(S_{n+1}^n | Y^n, M)$$

$$= H(S^n | M).$$

The third sum, by the reverse chain rule, is $\sum_i H(S_i | S_{i+1}^n, M) = H(S^n | M)$; but the middle identity shows we should compare against $\sum_i H(S_i | S_{i+1}^n, Y^{i-1}, M)$. Assembling,

$$\begin{aligned} \text{LHS} &= H(S^n | M) - \sum_{i=1}^n H(S_i | S_{i+1}^n, Y^{i-1}, M) \\ &= \sum_{i=1}^n \left(H(S_i | S_{i+1}^n, M) - H(S_i | S_{i+1}^n, Y^{i-1}, M) \right) \quad (\text{reverse chain rule on } H(S^n | M)) \\ &= \sum_{i=1}^n I(Y^{i-1}; S_i | S_{i+1}^n, M) = \text{RHS}. \end{aligned}$$

■

Proof of Theorem 6.7 – converse. Consider any non-causal-CSIT code. By Fano,

$$nR \leq I(M; Y^n) + n\epsilon_n = \sum_{i=1}^n I(M; Y_i | Y^{i-1}) + n\epsilon_n.$$

Expanding each term via the chain rule and dropping the non-negative $I(Y^{i-1}; Y_i)$,

$$\begin{aligned} nR &\leq \sum_{i=1}^n \left(I(M, Y^{i-1}; Y_i) - I(Y^{i-1}; Y_i) \right) + n\epsilon_n \\ &\leq \sum_{i=1}^n I(M, Y^{i-1}; Y_i) + n\epsilon_n. \quad (\text{this is the only inequality of this type; cf. §6.4}) \end{aligned} \tag{6.3}$$

Define the auxiliary

$$U_i \triangleq (M, S_{i+1}^n, Y^{i-1}).$$

By the chain rule,

$$I(M, Y^{i-1}; Y_i) = I(M, Y^{i-1}, S_{i+1}^n; Y_i) - I(S_{i+1}^n; Y_i | M, Y^{i-1}) = I(U_i; Y_i) - I(S_{i+1}^n; Y_i | M, Y^{i-1}).$$

On the other hand, using $S_i \perp\!\!\!\perp S_{i+1}^n$ and $M \perp\!\!\!\perp S^n$,

$$\begin{aligned} I(U_i; S_i) &= I(M, Y^{i-1}, S_{i+1}^n; S_i) \\ &= I(S_{i+1}^n; S_i) + I(M, Y^{i-1}; S_i | S_{i+1}^n) \\ &= 0 + I(Y^{i-1}; S_i | M, S_{i+1}^n) + I(M; S_i | S_{i+1}^n) \\ &= I(Y^{i-1}; S_i | M, S_{i+1}^n). \end{aligned}$$

Subtracting,

$$I(U_i; Y_i) - I(U_i; S_i) = I(M, Y^{i-1}; Y_i) + I(S_{i+1}^n; Y_i | M, Y^{i-1}) - I(Y^{i-1}; S_i | M, S_{i+1}^n).$$

Summing over i and applying the Csiszár sum identity (Lemma 6.8), the last two sums cancel:

$$\sum_{i=1}^n (I(U_i; Y_i) - I(U_i; S_i)) = \sum_{i=1}^n I(M, Y^{i-1}; Y_i).$$

Substituting into (6.3),

$$nR \leq \sum_{i=1}^n (I(U_i; Y_i) - I(U_i; S_i)) + n\epsilon_n. \quad (6.4)$$

It remains to single-letterise. Introduce $Q \sim \text{Unif}[1 : n]$ independent of everything, and set $U = (U_Q, Q)$, $X = X_Q$, $Y = Y_Q$, $S = S_Q$. Then

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n (I(U_i; Y_i) - I(U_i; S_i)) &= I(U_Q; Y_Q | Q) - I(U_Q; S_Q | Q) \\ &= I(U_Q, Q; Y_Q) - I(Q; Y_Q) - I(U_Q, Q; S_Q) + I(Q; S_Q) \\ &\leq I(U; Y) - I(U; S) + I(Q; S_Q) \\ &= I(U; Y) - I(U; S), \end{aligned}$$

where we used $-I(Q; Y_Q) \leq 0$ and $I(Q; S_Q) = 0$ (the state is i.i.d. and Q is independent of S^n). Thus

$$nR \leq n(I(U; Y) - I(U; S)) + n\epsilon_n.$$

Finally we check feasibility of the auxiliary. The marginal of $S = S_Q$ is $p(s)$ (the state is i.i.d.). The joint law factors as $p(s) p(u|s) p(x|u, s) W(y|x, s)$, where $p(u|s)$ and $p(x|u, s)$ are the conditional laws induced by the code (X_i is a function of (M, S^n) , but only the conditional law $p(x|u, s)$ matters here). Moreover $U_i \rightarrow (X_i, S_i) \rightarrow Y_i$ is a Markov chain: Y_i is generated memorylessly from (X_i, S_i) through the channel W , and by d-separation (§9.4 of the Appendix) the only path from the auxiliary $U_i = (M, S_{i+1}^n, Y^{i-1})$ to Y_i runs through (X_i, S_i) . Hence (S, U, X, Y) lies in the feasible class of Theorem 6.7 (with a stochastic $p(x|u, s)$), and

$$I(U; Y) - I(U; S) \leq \max_{p(u|s), x(u, s)} (I(U; Y) - I(U; S)).$$

Therefore $R \leq \max_{p(u|s), x(u, s)} (I(U; Y) - I(U; S)) + \epsilon_n$. ■

Remark 6.8 (Stochastic versus deterministic auxiliary). The converse produced a stochastic mapping $p(x|u, s)$. The deterministic form $x = x(u, s)$ in the theorem is without loss: given any stochastic $p(x|u, s)$, realise it with an independent random seed R as $x = x'(u, s, R)$, and replace U by $U' = (U, R)$. Then $I(U'; S) = I(U; S)$ (since $R \perp\!\!\!\perp (U, S)$) and $I(U'; Y) \geq I(U; Y)$, so the deterministic form achieves at least the stochastic value. ○

Proof of Theorem 6.7 – achievability (binning). Fix $p(u|s)$, a deterministic function $x(u, s)$, and a target rate $R < I(U; Y) - I(U; S)$. We use the following two standard typicality lemmas (both are proved by the joint-typicality estimates of Chapter 2; short proofs are included for completeness).

Lemma 6.9. (Covering lemma) Let $(U, S) \sim p(u, s)$, and let $u^n(1), \dots, u^n(2^{nR'})$ be i.i.d. $\sim p(u)^{\otimes n}$. If $R' > I(U; S)$, then for $s^n \sim p(s)^{\otimes n}$,

$$\Pr(\exists j : (u^n(j), s^n) \in \mathcal{T}_\epsilon^{(n)}(p_{U,S})) \rightarrow 1.$$

○

Proof. The covering trials share the random state sequence S^n , so the joint-typicality events are not independent across trials unconditionally; condition on the state sequence. For a *fixed* typical s^n and any j , the codeword $u^n(j)$ is independent of s^n and has law $p(u)^{\otimes n}$, so

$$\Pr((u^n(j), s^n) \in \mathcal{T}_\epsilon^{(n)}(p_{U,S})) \geq 2^{-n(I(U;S)+\epsilon)}$$

for n large, uniformly over typical s^n (the joint-typicality probability estimate). The $2^{nR'}$ codewords are independent, so, for every typical s^n ,

$$\Pr(\forall j : (u^n(j), s^n) \notin \mathcal{T}_\epsilon^{(n)}) \leq (1 - 2^{-n(I(U;S)+\epsilon)})^{2^{nR'}} \leq e^{-2^{n(R'-I(U;S)-\epsilon)}} \rightarrow 0,$$

and the same bound controls the unconditional failure probability up to the vanishing atypicality of S^n :

$$P_{\text{fail}} \leq \Pr(S^n \text{ atypical}) + \sup_{s^n \text{ typical}} (1 - q_n(s^n))^{2^{nR'}} \rightarrow 0,$$

where $q_n(s^n)$ is the per-trial typicality probability above, whenever $R' > I(U; S) + \epsilon$. ■

Lemma 6.10. (Markov / conditional-typicality lemma) If $(u^n, s^n) \in \mathcal{T}_\epsilon^{(n)}(p_{U,S})$ and Y^n is generated from (u^n, s^n) by first forming $X_i = x(u_i, s_i)$ and then drawing $Y_i \sim W(\cdot | x_i, s_i)$ independently, then

$$\Pr((u^n, s^n, x^n, y^n) \in \mathcal{T}_\epsilon^{(n)'}(p_{U,S,X,Y})) \rightarrow 1$$

for a suitable $\epsilon' > \epsilon$; in particular $(u^n, y^n) \in \mathcal{T}_\epsilon^{(n)'}(p_{U,Y})$ with probability $\rightarrow 1$. ○

Proof. Because $U \rightarrow (X, S) \rightarrow Y$ is a Markov chain, the joint law is $p(u, s)p(x|u, s)W(y|x, s)$. The conditionally-typical lemma (Chapter 2) states that conditioning on (u^n, s^n) being typical and generating (x^n, y^n) by the conditional law makes the whole quadruple jointly typical with high probability. ■

Codebook. Let $R' = I(U; S) + \delta$ and $\tilde{R} = R + R'$, with $\delta > 0$ small. Generate $2^{n\tilde{R}}$ codewords $u^n(m, j)$, $m \in [1 : 2^{nR}]$, $j \in [1 : 2^{nR'}]$, each i.i.d. $\sim p(u)^{\otimes n}$. Reveal the codebook to both parties.

Encoding. Given the message m and the state sequence s^n , find a j such that $(u^n(m, j), s^n) \in \mathcal{T}_\epsilon^{(n)}(p_{U,S})$. If none exists, declare an encoding error; otherwise transmit $x_i = x(u_i(m, j), s_i)$, $i \in [1 : n]$.

Decoding. Upon observing y^n , declare $\hat{m}$ the unique message for which there exists j such that $(u^n(\hat{m}, j), y^n) \in \mathcal{T}_\epsilon^{(n)}(p_{U,Y})$.

Error analysis.

- *Encoding error.* By Lemma 6.9, since $R' = I(U; S) + \delta > I(U; S)$, with probability $\rightarrow 1$ the bin m contains a $u^n(m, j)$ jointly typical with s^n . Hence $P(\text{encoding error}) \rightarrow 0$.
- *Decoding error.* By Lemma 6.10, the transmitted $u^n(m, j)$ is jointly typical with y^n with probability $\rightarrow 1$. The decoder declares a *message* (a bin), not a codeword, so other typical codewords in the transmitted message's own bin are not decoding errors: only wrong bins matter. For a wrong bin $m' \neq m$, every codeword $u^n(m', j')$ is independent of s^n and hence of y^n , so

$$\Pr((u^n(m', j'), y^n) \in \mathcal{T}_\epsilon^{(n)}(p_{U,Y})) \leq 2^{-n(I(U;Y)-\epsilon)}.$$

A union bound over the $2^{n\tilde{R}}$ wrong-bin codewords gives

$$P(\text{decoding error}) \leq 2^{n\tilde{R}} 2^{-n(I(U;Y)-\epsilon)} + o(1) = 2^{-n(I(U;Y)-\tilde{R}-\epsilon)} + o(1) \rightarrow 0$$

provided $\tilde{R} < I(U; Y)$, i.e. $R < I(U; Y) - I(U; S) - (\delta + \epsilon)$.

Hence every rate $R < I(U; Y) - I(U; S)$ is achievable. $\blacksquare$

Remark 6.9 (Honest labelling: this is *not* sub-additivity). Per the programme of §1.13, label the flavours accurately.

- The *converse* is not a sub-additivity single-letterisation. Its engine is the *Csiszár sum identity* (Lemma 6.8), a side-information/state telescoping that cancels $\sum_i I(S_{i+1}^n; Y_i | Y^{i-1}, M)$ against $\sum_i I(Y^{i-1}; S_i | S_{i+1}^n, M)$. There is no single functional F whose sub-additivity is being invoked. The only step of the sub-additivity “shape” is (6.3) (dropping $I(Y^{i-1}; Y_i)$), and it is an inequality, not an equality chain.
- The *achievability* is *binning*: a covering step (encoding, Lemma 6.9) and a packing step (decoding). It is not a sub-additivity argument either.

Contrast with §6.2–§6.4, where the converses *are* sub-additivity arguments. $\circ$

Remark 6.10 (Comparison with CSIT+CSIR). Revealing the state to the receiver cannot hurt, so GP capacity is at most the CSIT+CSIR capacity. Indeed,

$$\begin{aligned} \max_{p(u|s), x(u,s)} (I(U; Y) - I(U; S)) &\leq \max_{p(u,x|s)} (I(U; Y, S) - I(U; S)) \\ &= \max_{p(u,x|s)} I(U; Y | S) \leq \max_{p(x|s)} I(X; Y | S), \end{aligned}$$

using $I(U; Y) \leq I(U; Y, S)$ and data processing ($U \rightarrow (X, S) \rightarrow Y$). The interesting phenomenon, to which we now turn, is that for the Gaussian additive-interference channel these two quantities are *equal*. $\circ$

Example 6.11. (Binary writing on dirty paper) Let $\mathcal{X} = \mathcal{S} = \mathcal{Y} = \{0, 1\}$ and $Y = X \oplus S$ (mod-2, noiseless), with $S \sim \text{Bern}(1/2)$ known non-causally at the transmitter and not at the receiver. The CSIT+CSIR capacity is $E[C_S] = 1$ bit/use (once S is known at both ends, each slot

is a clean binary channel). Gelfand–Pinsker also gives $C = 1$: take $U \sim \text{Bern}(1/2)$ independent of S and transmit $X = U \oplus S$. Then $Y = X \oplus S = U$, so

$$I(U; Y) - I(U; S) = H(U) - 0 = 1,$$

and $C \leq 1$ since $|\mathcal{Y}| = 2$. Here the auxiliary is independent of the state, $I(U; S) = 0$, so the bin (of size $\approx 2^{nI(U;S)} = 1$) degenerates and the scheme reduces to ordinary coding: the transmitter simply flips its bit when $S = 1$, pre-cancelling the state, and the receiver reads U without ever learning S . In the Gaussian example of §6.6 the optimal auxiliary is instead *correlated* with S ($I(U; S) = \frac{1}{2} \log(1 + \frac{\alpha^2 Q}{P}) > 0$), so the binning is genuinely needed there.

6.6 Costa’s dirty-paper coding

Consider the scalar Gaussian channel

$$Y = X + S + Z,$$

where $Z \sim \mathcal{N}(0, 1)$ is the additive white Gaussian noise, $S \sim \mathcal{N}(0, Q)$ is the *state* (the “dirt”), $S \perp\!\!\!\perp Z$, the transmitter knows S^n non-causally, and the input satisfies the average power constraint $\mathbb{E}[X^2] \leq P$. The receiver does *not* know S . Recall that all logs are base 2, so the Gaussian expressions below are in bits.

Theorem 6.12. (Costa 1983) For the Gaussian channel $Y = X + S + Z$ with $S \sim \mathcal{N}(0, Q)$ known non-causally at the transmitter and $\mathbb{E}[X^2] \leq P$,

$$C = \frac{1}{2} \log(1 + P),$$

independently of Q . The state, no matter how large its power, does not reduce the capacity below the clean-channel value. $\circ$

Proof. Upper bound (worst-case state). Give the receiver the state S^n as genie side information. This can only increase capacity, and the resulting setting is the CSIT+CSIR Gaussian channel of §6.3.3. For it,

$$I(X; Y | S) = h(X + S + Z | S) - h(Z) = h(X + Z | S) - h(Z) \leq h(X + Z) - h(Z),$$

where the last step is “conditioning reduces entropy.” Since $\mathbb{E}[(X + Z)^2] = \mathbb{E}[X^2] + 1 + 2\mathbb{E}[XZ] = \mathbb{E}[X^2] + 1 \leq P + 1$ (the cross term vanishes because the channel noise Z is independent of (X, S) , so $\mathbb{E}[XZ] = \mathbb{E}[X] \mathbb{E}[Z] = 0$), Gaussian maximality (Lemma 1.59) gives $h(X + Z) \leq \frac{1}{2} \log 2\pi e(P + 1)$, and $h(Z) = \frac{1}{2} \log 2\pi e$. Hence

$$I(X; Y | S) \leq \frac{1}{2} \log(1 + P),$$

so $C \leq \frac{1}{2} \log(1 + P)$. This upper bound is the *cut-set/Gaussian-optimality* route: reveal the state (the “worst case” for the receiver), then apply Gaussian maximality (the rotation-trick proof of Lemma 1.59, i.e. sub-additivity + double Markovity + Bernstein) to bound $h(X + Z)$.

Achievability (Gelfand–Pinsker with a Gaussian auxiliary). The finite-alphabet, unconstrained Theorem 6.7 does not literally cover the Gaussian cost-constrained model, so we invoke its standard Gaussian, cost-constrained extension (a continuous-alphabet/cost-constrained GP theorem; Costa’s dirty-paper proof supplies exactly this Gaussian version, (Costa, 1983)). The power constraint is imposed at code level under the *averaged-power* convention

$$\frac{1}{n} \mathbb{E} \left[\sum_{i=1}^n X_i^2 \right] \leq P,$$

where the expectation is over the uniform message and the state; stronger per-message or per-realisation constraints would require the corresponding version of the cited theorem. Within that extension, take the auxiliary

$$U = X + \alpha S, \quad X \sim \mathcal{N}(0, P), \quad X \perp\!\!\!\perp S,$$

and transmit $x = x(u, s) = u - \alpha s$ (so that $X = U - \alpha S$ is exactly the intended Gaussian input). The parameter α is chosen below. Compute

$$\begin{aligned} I(U; Y) - I(U; S) &= I(U; Y, S) - I(U; S | Y) - I(U; S) \\ &= I(U; Y | S) - I(U; S | Y). \end{aligned}$$

For the first term, given $S = s$ the pair $(U, Y) = (X + \alpha s, X + s + Z)$ is a translation of $(X, X + Z)$, so

$$I(U; Y | S) = I(X; X + Z | S) = I(X; X + Z) = \frac{1}{2} \log(1 + P),$$

using $X \perp\!\!\!\perp S$. For the second term,

$$I(U; S | Y) = I(X + \alpha S; S | X + S + Z).$$

Write $X + \alpha S = \alpha(X + S + Z) + ((1 - \alpha)X - \alpha Z)$. Given $X + S + Z$, this differs from $(1 - \alpha)X - \alpha Z$ by the constant $\alpha(X + S + Z)$, so

$$I(X + \alpha S; S | X + S + Z) = I((1 - \alpha)X - \alpha Z; S | X + S + Z).$$

Choose α to make $(1 - \alpha)X - \alpha Z$ independent of $X + S + Z$. All variables are jointly Gaussian, so independence is equivalent to orthogonality:

$$\mathbb{E} [((1 - \alpha)X - \alpha Z)(X + Z)] = (1 - \alpha) \mathbb{E}[X^2] - \alpha \mathbb{E}[Z^2] = (1 - \alpha)P - \alpha = 0 \iff \alpha = \frac{P}{P + 1}.$$

With this choice, $(1 - \alpha)X - \alpha Z \perp\!\!\!\perp (X + Z)$, and since $S \perp\!\!\!\perp (X, Z)$, we get $(1 - \alpha)X - \alpha Z \perp\!\!\!\perp (X + S + Z, S)$; therefore

$$I((1 - \alpha)X - \alpha Z; S | X + S + Z) = 0.$$

Thus

$$I(U; Y) - I(U; S) = \frac{1}{2} \log(1 + P) - 0 = \frac{1}{2} \log(1 + P),$$

and the Gaussian cost-constrained GP extension gives $C \geq \frac{1}{2} \log(1 + P)$.

Combining the two bounds, $C = \frac{1}{2} \log(1 + P)$. ■

Remark 6.11 (Writing on dirty paper). The result is colloquially “writing on dirty paper”: a writer who knows the pattern of dirt on a page can arrange the ink so that a reader, who sees only the sum of ink and dirt, still recovers the message as if the page were clean. Quantitatively, the transmitter pre-subtracts αS from its signal ($X = U - \alpha S$), and the choice $\alpha = P/(P+1)$ is the (MMSE) coefficient that makes the residual interference $(1-\alpha)X - \alpha Z$ orthogonal – hence independent – of the received sum. The capacity equals the clean capacity $\frac{1}{2} \log(1+P)$ for every Q , even as $Q \rightarrow \infty$: non-causal knowledge of the state completely neutralises it. $\circ$

Remark 6.12 (Optimality of $U = X + \alpha S$ and Gaussian framing). The achievability computation shows the Gaussian auxiliary $U = X + \alpha S$ with $\alpha = P/(P+1)$ saturates the universal upper bound, so it is optimal; no other auxiliary can do better. The Gaussian-optimality machinery of Chapter 1 (Lemma 1.59 and its rotation-trick proof) enters the *upper* bound (via $h(X+Z) \leq \frac{1}{2} \log 2\pi e(P+1)$), and the same equality conditions (Gaussian X , and orthogonality of the residual) are exactly what the achievability realises. The binning computation itself ($I(U; Y) - I(U; S)$) is *not* a sub-additivity step – it is the GP formula of §6.5. $\circ$

6.7 Extensions and remarks

We collect several short remarks and pointers.

Remark 6.13 (States known to neither party). This is §6.3.1: with no state information anywhere, the effective channel is the averaged channel $\hat{W}(y|x) = \sum_s p(s)W(y|x, s)$, and $C = \max_{p(x)} I(X; Y)$ over $\hat{W}$. It is the natural baseline against which every side-information gain below is measured. $\circ$

Remark 6.14 (The six cases in one table). For i.i.d. state $S^n \sim p(s)^{\otimes n}$ the taxonomy is:

Transmitter	Receiver	Capacity
no state	no state	$\max_{p(x)} I(X; Y)$ over $\hat{W}$
no state	state (CSIR)	$\max_{p(x)} I(X; Y S)$
state (causal or non-causal)	state	$E[C_S] = \sum_s p(s)C_s$
state, causal	no state	$\max_{p(u), x(u, s)} I(U; Y)$ (Shannon)
state, non-causal	no state	$\max_{p(u s), x(u, s)} (I(U; Y) - I(U; S))$ (GP)

The first three rows are sub-additivity/cut-set flavoured; the last two use strategies and the Csiszár sum identity / binning respectively. $\circ$

Remark 6.15 (Compound Gelfand–Pinsker). If the transmitter knows S^n non-causally but the state law $p(s)$ itself is unknown and drawn from a family $\{p_\theta\}$, the natural compound-GP problem is open in most regimes: the expression

$$\max_{p(u|s), x(u, s)} \min_{\theta} (I(U; Y_\theta) - I(U; S_\theta))$$

is a *candidate* single-letter form for discussion, not an established capacity theorem – the max and the min do not commute, the auxiliaries are state-law dependent, and the compound-GP

capacity remains open in general. The min–max theorem of the appendix (Corollary 9.16) is the natural tool for the convex cases, exactly as in the broadcast and interference chapters. ○

Remark 6.16 (Applications). Two canonical applications of this chapter’s results:

- *Precoding in MIMO broadcast (interference known at the transmitter).* The dirty-paper result is the single-user seed of Costa’s precoding: when a base station transmits to several users, it knows in advance the interference each symbol will cause to the others, and can pre-subtract it (“dirty-paper coding”), achieving the same rates as if the interference were absent.
- *Digital watermarking and information embedding.* Writing a message into a host signal S^n (an image, audio, or video frame) so that the message is recoverable while the host looks essentially unchanged is exactly the GP setting: the host is the state, known to the “writer” (encoder) but not necessarily to the decoder, and $I(U; Y) - I(U; S)$ quantifies the embedding rate above the cost of matching the host.

○

Remark 6.17 (Lossy and causal variants). The causal-vs-non-causal distinction is genuinely significant without CSIR: causal CSIT gives $\max I(U; Y)$ (Theorem 6.6) while non-causal CSIT gives $\max (I(U; Y) - I(U; S))$ (Theorem 6.7), and the latter is strictly larger in general. Shannon’s original treatment also studied the case where the decoder must reconstruct the *state* in addition to the message, a rate–distortion flavoured extension we do not develop here. ○

Remark 6.18 (Bibliography). The compound channel is due to Blackwell–Breiman–Thomasian and Wolfowitz; the causal-CSIT result to Shannon (1958); the non-causal result to Gel’fand–Pinsker (1980); the Gaussian writing-on-dirty-paper result to Costa (1983); and the Csiszár sum identity to Csiszár. See the bibliography. ○

7 Distributed source coding

Source coding is the *dual* problem to channel coding: instead of communicating a message reliably over a noisy channel, we compress a source so that it can be reconstructed reliably at a (possibly remote) decoder. In this chapter we study the *distributed* setting, in which several terminals observe correlated sources and encode them *separately* (the encoders do not cooperate), while one or more decoders reconstruct the sources — or a prescribed function of them. The central result is the Slepian–Wolf theorem: asymptotically, separate encoding can be as good as joint encoding, provided the decoder jointly exploits the correlation.

Every outer bound in this chapter is a *single-letterisation*, and every single-letterisation is an application of *sub-additivity of an information functional* (see §1.13 of Chapter 1). The two canonical sub-additivity facts — entropy sub-additivity $H(X_1, X_2) \leq H(X_1) + H(X_2)$ (equivalently, conditioning reduces entropy: $H(A|B, C) \leq H(A|B)$), and the KL/mutual-information contraction $D(v_Y \| \mu_Y) \leq D(v_X \| \mu_X)$ for $Y = WX$ (equivalently $I(U; Y) \leq I(U; X)$ for $U - X - Y$) — do all the work. The converses split into an honest taxonomy:

- *Saturated (additivity) bounds.* Slepian–Wolf uses i.i.d. entropy additivity (the sub-additivity inequalities hold with *equality*). Gray–Wyner uses i.i.d. additivity for its common-rate chain together with Fano and *genuine conditioning inequalities* for its private-rate bounds: conditioning on the common message can couple otherwise independent coordinates.
- *Genuinely sub-additive bounds.* The one-helper R_1 bound rests on the chain rule followed by a conditioning step with gap $I(X_i; Y^{i-1} | M_2, X^{i-1}) \geq 0$,

$$H(X^n | M_2) = \sum_i H(X_i | M_2, X^{i-1}) \geq \sum_i H(X_i | M_2, X^{i-1}, Y^{i-1}) = \sum_i H(X_i | U_i),$$

where the last equality is the *exact* identification $U_i = (M_2, Y^{i-1})$ obtained by d-separation (no $o(n)$ loss).

- *Contraction + telescoping bounds.* The Körner–Marton outer bound combines the data-processing contraction $I(M_2; X^n | M_1 Y^n) = 0$ with the Csiszár sum identity — the telescoping identity that makes sub-additivity inequalities *cancel to equality*.
- *The sharp endpoint.* Hypercontractivity is the *sharp* form of the contraction: the constant $S^*(X; Y)$ that measures how far $D(v_Y \| \mu_Y) \leq D(v_X \| \mu_X)$ is from saturation, and which gives a necessary test for the supporting normals of the Gray–Wyner region at the all-private corner on the restricted subfamily $U - X - Y$ (the full set of supporting normals is the hypercontractivity ribbon).

7.1 The distributed source coding problem

We begin with the simplest non-trivial setting: *two encoders, one decoder*. Two correlated memoryless sources $(X^n, Y^n) \sim \prod_{i=1}^n p_{XY}(x_i, y_i)$ are observed separately: Encoder 1 sees only X^n and Encoder 2 sees only Y^n . Each encoder maps its observation to an index

$$M_1 = f_1(X^n) \in [1 : 2^{nR_1}], \quad M_2 = f_2(Y^n) \in [1 : 2^{nR_2}],$$

so R_j is the rate (bits per source symbol) of encoder j . A decoder observes the pair (M_1, M_2) and produces an estimate $(\hat{X}^n, \hat{Y}^n) = g(M_1, M_2)$ of the *whole* pair. The probability of error is

$$P_e^{(n)} \triangleq \Pr((\hat{X}^n, \hat{Y}^n) \neq (X^n, Y^n)).$$

A rate pair (R_1, R_2) is *achievable* if there exists a sequence of $(2^{nR_1}, 2^{nR_2}, n)$ codes with $P_e^{(n)} \rightarrow 0$ as $n \rightarrow \infty$. The *optimal rate region* (we avoid the term “capacity”, which is reserved for channels) is the closure of the set of achievable pairs.




Figure 7.1: Distributed source coding with two encoders and one decoder.

As a warm-up we recall the point-to-point benchmark. If a single encoder saw both X^n and Y^n , Shannon’s lossless source coding theorem would say that $H(X, Y)$ bits per symbol suffice and are necessary. The surprise of Slepian–Wolf is that, even when the two encoders act *separately*, the same *total* rate $H(X, Y)$ suffices — only the *individual* rates must be rebalanced to reflect the conditional structure.

Throughout this chapter we use the standard strong-typicality toolkit: $\mathcal{T}_\epsilon^{(n)}(Z)$ denotes the ϵ -typical set of a sequence from a source p_Z , with the familiar consequences of the AEP and Sanov/type counting (Lemma 9.7 and Theorem 9.11 of the appendix):

$$|\mathcal{T}_\epsilon^{(n)}(X, Y)| \leq 2^{n(H(X, Y) + \delta)}, \quad |\{x^n : (x^n, y^n) \in \mathcal{T}_\epsilon^{(n)}(X, Y)\}| \leq 2^{n(H(X|Y) + \delta)}, \quad (7.1)$$

$$\Pr((X^n, Y^n) \in \mathcal{T}_\epsilon^{(n)}(X, Y)) \rightarrow 1, \quad (7.2)$$

for any fixed $\delta > 0$ and sufficiently large n ; all constants below are stated up to such δ -factors.

Two consequences of the AEP are used repeatedly, and we isolate them here.

Lemma 7.1. (Covering lemma) Let $(U^n, Y^n) \sim \prod_i p_{U,Y}(u_i, y_i)$, and let $U^n(1), \dots, U^n(M)$ be independent random sequences, each $\sim \prod_i p_U(u_i)$, independent of Y^n . If $M \geq 2^{n(I(U;Y)+\delta)}$ for some $\delta > 0$, then

$$\Pr(\exists m \in [1 : M] \text{ with } (U^n(m), Y^n) \in \mathcal{T}_\epsilon^{(n)}(U, Y)) \rightarrow 1.$$

○

Lemma 7.2. (Packing lemma) Let $(U^n, Y^n) \sim \prod_i p_{U,Y}(u_i, y_i)$, and let $X^n(m)$, $m \in [1 : M]$, be independent random sequences, each conditionally generated according to $p(x^n(m)|u^n) = \prod_i p_{X|U}(x_i(m)|u_i)$. Suppose Y^n is generated through the memoryless channel $p(y^n|x^n(1), u^n) = \prod_i p_{Y|X,U}(y_i|x_i(1), u_i)$, so that the codewords $X^n(m)$, $m \neq 1$, are independent of Y^n given U^n . If $M \leq 2^{n(I(X;Y|U)-\delta)}$ for some $\delta > 0$, then

$$\Pr(\exists m \neq 1 \text{ with } (X^n(m), U^n, Y^n) \in \mathcal{T}_\epsilon^{(n)}(X, U, Y)) \rightarrow 0.$$

○

For the covering lemma, condition on the observed sequence $Y^n = y^n$ (typical with probability $\rightarrow 1$) and separate the rate slack from the typicality slack: for a fixed typical y^n , each covering trial is jointly typical with y^n with probability $q_n(y^n) \geq 2^{-n(I(U;Y)+\delta/2)}$ for sufficiently small typicality tolerance and large n (a uniform typical-sequence bound). The M covering trials are independent given y^n , so the product bound for independent trials gives, uniformly over typical y^n ,

$$(1 - q_n(y^n))^M \leq e^{-Mq_n(y^n)} \leq e^{-2^{n\delta/2}} \rightarrow 0$$

whenever $M \geq 2^{n(I(U;Y)+\delta)}$ with $\delta > 0$; adding the vanishing probability of an atypical Y^n gives the stated convergence. (The two slacks must be kept distinct: using δ in both places only gives $Mq_n \geq 1$, which does not force divergence.) The packing lemma follows by the union bound with the conditional-typicality estimate for a codeword independent of Y^n given U^n (the exponent is $I(X;Y|U)$, and the observation Y^n is essential: conditionally generated codewords are themselves jointly typical with U^n with high probability, so a test that omits Y^n cannot succeed). Throughout, “ $\rightarrow$ ” means convergence to 0 (or 1) as $n \rightarrow \infty$ with $\epsilon = \epsilon_n \rightarrow 0$ slowly and δ fixed; the sub-exponential ϵ -factors never affect the exponents.

7.2 Slepian–Wolf: lossless coding of two correlated sources

Theorem 7.3. (Slepian–Wolf (Slepian and Wolf, 1973)) Let (X^n, Y^n) be i.i.d. following p_{XY} . The optimal rate region for lossless reconstruction of (X^n, Y^n) is

$$R_1 \geq H(X|Y), \quad R_2 \geq H(Y|X), \quad R_1 + R_2 \geq H(X, Y). \quad (7.3)$$

○

Proof. Converse. We show the R_1 and sum-rate inequalities; R_2 is symmetric.

First inequality. Since M_1 is a function of X^n and, given Y^n , the estimate $(\hat{X}^n, \hat{Y}^n) = g(M_1, f_2(Y^n))$ is a function of M_1 alone, we have

$$\begin{aligned}
nR_1 &\geq H(M_1) \geq H(M_1|Y^n) \geq I(M_1; \hat{X}^n \hat{Y}^n | Y^n) \\
&= I(X^n M_1; \hat{X}^n \hat{Y}^n | Y^n) && ((\hat{X}^n, \hat{Y}^n) \text{ is a function of } (M_1, Y^n)) \\
&\geq I(X^n; \hat{X}^n \hat{Y}^n | Y^n) \\
&= H(X^n | Y^n) - H(X^n | \hat{X}^n \hat{Y}^n Y^n) \\
&\geq H(X^n | Y^n) - n\epsilon_n && (\text{Fano's inequality, Proposition 1.57}) \\
&= \sum_{i=1}^n H(X_i | Y_i) - n\epsilon_n = nH(X|Y) - n\epsilon_n.
\end{aligned} \tag{7.4}$$

The last equality is the *saturation* of sub-additivity: for i.i.d. sources $H(X^n | Y^n) = \sum_i H(X_i | Y_i)$ holds with equality.

Sum rate. This is the step that is often garbled. The clean argument uses only that (M_1, M_2) is a function of (X^n, Y^n) and that $(\hat{X}^n, \hat{Y}^n)$ is a function of (M_1, M_2) :

$$\begin{aligned}
n(R_1 + R_2) &\geq H(M_1 M_2) \\
&\geq I(M_1 M_2; \hat{X}^n \hat{Y}^n) \\
&= H(\hat{X}^n \hat{Y}^n) && ((\hat{X}^n, \hat{Y}^n) \text{ is a function of } (M_1, M_2)) \\
&\geq H(X^n Y^n) - n\epsilon_n && (\text{Fano: } H(X^n Y^n | \hat{X}^n \hat{Y}^n) \leq n\epsilon_n) \\
&= \sum_{i=1}^n H(X_i Y_i) - n\epsilon_n = nH(X, Y) - n\epsilon_n.
\end{aligned} \tag{7.5}$$

In the fourth line we used $H(X^n Y^n) = H(X^n Y^n | \hat{X}^n \hat{Y}^n) + I(X^n Y^n; \hat{X}^n \hat{Y}^n) \leq n\epsilon_n + H(\hat{X}^n \hat{Y}^n)$.

Remark 7.1. Both converses are *additivity* arguments in the taxonomy of §1.13: the chain rule plus Fano, with the sub-additivity inequalities saturated (equalities) because the source is memoryless. There is no genuine sub-additivity gap to expose. $\circ$

Achievability (random binning).

- *Codebook and encoding.* For each typical $x^n \in \mathcal{T}_\epsilon^{(n)}(X)$, assign it uniformly at random to a bin $b_1(x^n) \in [1 : 2^{nR_1}]$; declare an error if X^n is atypical. Encoder 1 sends $M_1 = b_1(X^n)$. Independently, bin each typical $y^n \in \mathcal{T}_\epsilon^{(n)}(Y)$ into $[1 : 2^{nR_2}]$ and send $M_2 = b_2(Y^n)$. Note that binning is done *irrespective* of the other source; joint typicality is used only at the decoder.
- *Decoding.* Upon receiving (M_1, M_2) , declare $(\hat{X}^n, \hat{Y}^n)$ to be the *unique* jointly typical pair in bin (M_1, M_2) .

- *Error analysis.* Let $\mathcal{E}_0 = \{(X^n, Y^n) \notin \mathcal{T}_\epsilon^{(n)}(X, Y)\}$; by the AEP, $\Pr(\mathcal{E}_0) \rightarrow 0$. Let (X^n, Y^n) denote the transmitted pair. The remaining error events involve the other sequences in the decoded bins (the bin contents, not an i.i.d. enumeration of codewords):

$$\mathcal{E}_1 = \{\exists \hat{x}^n \neq X^n \text{ in bin } M_1 : (\hat{x}^n, Y^n) \in \mathcal{T}_\epsilon^{(n)}(X, Y)\}, \quad (7.6)$$

$$\mathcal{E}_2 = \{\exists \hat{y}^n \neq Y^n \text{ in bin } M_2 : (X^n, \hat{y}^n) \in \mathcal{T}_\epsilon^{(n)}(X, Y)\}, \quad (7.7)$$

$$\mathcal{E}_{12} = \{\exists \hat{x}^n \neq X^n \text{ in bin } M_1, \hat{y}^n \neq Y^n \text{ in bin } M_2 : (\hat{x}^n, \hat{y}^n) \in \mathcal{T}_\epsilon^{(n)}(X, Y)\}. \quad (7.8)$$

For $\mathcal{E}_1$: for a fixed typical y^n there are at most $2^{nH(X|Y)}$ sequences x^n with (x^n, y^n) jointly typical, and each lands in bin M_1 with probability 2^{-nR_1} ; hence

$$\Pr(\mathcal{E}_1) \leq 2^{-nR_1} \cdot 2^{nH(X|Y)} = 2^{n(H(X|Y) - R_1)} \rightarrow 0$$

if $R_1 > H(X|Y)$. Symmetrically $\Pr(\mathcal{E}_2) \rightarrow 0$ if $R_2 > H(Y|X)$. Finally, for $\mathcal{E}_{12}$, there are at most $2^{nH(X,Y)}$ jointly typical pairs, each landing in bin (M_1, M_2) with probability $2^{-n(R_1+R_2)}$, so

$$\Pr(\mathcal{E}_{12}) \leq 2^{-n(R_1+R_2)} \cdot 2^{nH(X,Y)} = 2^{n(H(X,Y) - (R_1+R_2))} \rightarrow 0$$

if $R_1 + R_2 > H(X, Y)$. All probabilities are up to sub-exponential ϵ -factors, which do not affect the exponents. ■

The result extends to an arbitrary number of sources. Let $S_1, \dots, S_k$ be k correlated memoryless sources observed separately, to be reconstructed at a common decoder at rates $R_1, \dots, R_k$.

Theorem 7.4. (Multi-source Slepian–Wolf) For every non-empty $I \subseteq \{1, \dots, k\}$,

$$\sum_{i \in I} R_i \geq H(S_I | S_{I^c}),$$

where $S_I = (S_i)_{i \in I}$ and $S_{I^c} = (S_i)_{i \notin I}$. ○

Proof sketch. The converse is obtained by grouping the sources in I into a single super-source and those in I^c into a single super-side-information, then applying (7.4). Achievability follows by the same random-binning scheme applied to all k sources, with the natural generalisation of the error events $\mathcal{E}_J$ ($J \subseteq [k]$ a subset of “wrong” encoders), each of which is small whenever $\sum_{j \in J} R_j > H(S_J | S_{J^c})$. ■

Example 7.5. (A correlated binary source) Let $X \sim \text{Bern}(\frac{1}{2})$ and $Y = X \oplus N$ with $N \sim \text{Bern}(p)$ independent of X (a doubly symmetric binary source with correlation $\rho = 1 - 2p$). Then $H(X) = H(Y) = 1$ and $H(X|Y) = H(Y|X) = H_b(p)$, so the Slepian–Wolf region is

$$R_1 \geq H_b(p), \quad R_2 \geq H_b(p), \quad R_1 + R_2 \geq 1 + H_b(p).$$

In particular, for $p = 0.1$ (with $H_b(0.1) \approx 0.469$), the pair $R_1 = R_2 = 0.75$ is achievable even though each source alone has entropy 1: neither encoder alone can compress below 1 bit, but jointly they use only 1.5 bits per pair (the sum-rate floor is $1 + H_b(0.1) \approx 1.469$, so $R_1 = R_2 = 0.7$ would *not* be achievable). The corner $(R_1, R_2) = (1, 0.469)$ corresponds to Encoder 1 sending X^n uncompressed while Encoder 2 only resolves the residual $Y|X$; by symmetry $(0.469, 1)$ is also achievable, and the whole region is the convex hull of these corners together with the sum-rate boundary. This is the Slepian–Wolf “surprise”: the total rate $H(X, Y) = 1 + H_b(p)$ equals the joint-entropy benchmark, but the individual rates trade off through the conditional entropies.

7.3 The one-helper problem: source coding with side information at the decoder

In the one-helper setting the decoder needs only X^n ; Encoder 1 observes X^n (rate R_1), while Encoder 2 is a *helper* that observes the correlated source Y^n (rate R_2) and supplies side information. This is the lossless source coding problem with (coded) side information at the decoder, solved by Wyner (Wyner, 1975b) and, in the helper form, by Ahlswede and Körner (Ahlswede and Körner, 1975) (see also the textbook treatment (Csiszar and Körner, 2011)).




Figure 7.2: The one-helper problem: the decoder reconstructs X^n only.

Theorem 7.6. (One-helper (Wyner–Ahlswede–Körner)) The optimal rate region is the set of pairs (R_1, R_2) satisfying

$$R_1 \geq H(X|U), \quad R_2 \geq I(U; Y), \quad (7.9)$$

for some conditional distribution $p(u|y)$, i.e. with $U - Y - X$ forming a Markov chain, and with $|\mathcal{U}| \leq |\mathcal{Y}| + 1$. $\circ$

Proof. Converse. The auxiliary must be generated from the *helper’s* observation Y ; the correct identification is $U_i = (M_2, Y^{i-1})$, with the past block X^{i-1} eliminated exactly by d-separation.

The R_2 bound (saturated). Since M_2 is a function of Y^n ,

$$\begin{aligned} nR_2 &\geq H(M_2) \geq I(M_2; Y^n) \\ &= \sum_{i=1}^n I(M_2; Y_i | Y^{i-1}) \end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^n [H(Y_i|Y^{i-1}) - H(Y_i|M_2Y^{i-1})] \\
&= \sum_{i=1}^n [H(Y_i) - H(Y_i|M_2Y^{i-1})] \quad (\text{memoryless: } Y_i \perp Y^{i-1}) \\
&= \sum_{i=1}^n I(M_2Y^{i-1}; Y_i) \\
&= \sum_{i=1}^n I(U_i; Y_i), \tag{7.10}
\end{aligned}$$

with the auxiliary $U_i = (M_2, Y^{i-1})$. This step is saturated (an exact chain-rule identity).

The R_1 bound. By the same Fano/chain-rule steps as in the Slepian–Wolf converse,

$$\begin{aligned}
nR_1 &\geq H(X^n|M_2) - n\epsilon_n = \sum_{i=1}^n H(X_i|M_2, X^{i-1}) - n\epsilon_n \\
&\geq \sum_{i=1}^n H(X_i|M_2, X^{i-1}, Y^{i-1}) - n\epsilon_n \quad (\text{conditioning reduces entropy}) \\
&= \sum_{i=1}^n H(X_i|U_i) - n\epsilon_n, \tag{7.11}
\end{aligned}$$

with $U_i = (M_2, Y^{i-1})$. The last equality is the identification step, and it is *exact* (no $o(n)$ loss). Expand

$$H(X_i|M_2, X^{i-1}, Y^{i-1}) = H(X_i|U_i) - I(X_i; X^{i-1}|U_i).$$

By the data-processing inequality applied to the Markov chain $X_i - Y_i - (U_i, X^{i-1})$, which is the coordinate-star independence $X_i \perp\!\!\!\perp (M_2, Y^{i-1}, X^{i-1}) \mid Y_i$ of the Appendix (Corollary 9.19: in the graph $Y_j \rightarrow X_j, Y_j \rightarrow M_2$ every X_j is a leaf, so conditioning on Y_i separates X_i from every other node),

$$I(X_i; X^{i-1}|U_i) \leq I(Y_i; X^{i-1}|U_i).$$

The right-hand side vanishes by the disjoint-form coordinate-star independence

$$X^{i-1} \perp\!\!\!\perp (M_2, Y_i^n, X_i) \mid Y^{i-1},$$

again from Corollary 9.19 (the path from any $X_j, j < i$, leaves through its unique neighbour $Y_j \in Y^{i-1}$, a conditioned non-collider, so every path is blocked; the conditioning set Y^{i-1} and the right-hand block Y_i^n are disjoint). Hence $I(Y_i; X^{i-1}|M_2, Y^{i-1}) = 0$, $I(X_i; X^{i-1}|U_i) = 0$, and $H(X_i|M_2, X^{i-1}, Y^{i-1}) = H(X_i|U_i)$ exactly.

Markovity and single-letterisation. The auxiliary satisfies $U_i - Y_i - X_i$: given Y_i , the pair (M_2, Y^{i-1}) is a function of the other coordinates $Y^{\setminus i} = (Y_j : j \neq i)$ ($M_2 = f_2(Y^n)$ with Y_i fixed), which are independent of X_i . Introduce $Q \sim \text{Unif}[1 : n]$ independent of everything, and set $X = X_Q$, $Y = Y_Q$, $U = (U_Q, Q)$. Then

$$\frac{1}{n} \sum_{i=1}^n H(X_i|U_i) = H(X_Q|U_Q, Q) = H(X|U), \quad \frac{1}{n} \sum_{i=1}^n I(U_i; Y_i) = I(U_Q; Y_Q|Q) = I(U; Y),$$

the latter because $I(U_Q Q; Y_Q) = I(U_Q; Y_Q|Q) + I(Q; Y_Q)$ and $I(Q; Y_Q) = 0$. Dividing (7.11)–(7.10) by n and letting $n \rightarrow \infty$ gives (7.9) with $U - Y - X$.

Cardinality. $|\mathcal{U}| \leq |\mathcal{Y}| + 1$ is the standard support-lemma bound of the Wyner–Ahlsvede–Körner converse: preserve the marginal $p(y)$ ($|\mathcal{Y}| - 1$ degrees of freedom), the conditional entropy $H(X|U)$, and the mutual information $I(U; Y) = H(Y) - H(Y|U)$ – in total $|\mathcal{Y}| + 1$ scalar quantities, whence the support lemma (Fenchel–Eggleston–Carathéodory) bounds the auxiliary alphabet by $|\mathcal{Y}| + 1$. The bound is governed by the helper’s alphabet size $|\mathcal{Y}|$, not $|\mathcal{X}|$. ■

Achievability. Fix $p(u|y)$ and its marginal $p(u)$; choose $R_1 > H(X|U)$, $R_2 > I(U; Y)$.

- *Codebook and encoding.* Encoder 1 bins each typical $x^n \in \mathcal{T}_\epsilon^{(n)}(X)$ into $[1 : 2^{nR_1}]$ and sends the bin index M_1 . Encoder 2 generates 2^{nR_2} sequences $u^n(m_2) \sim \prod_i p(u_i)$ and sends the index M_2 of a sequence that is jointly typical with Y^n (declaring an error if none exists).
- *Decoding.* Given (M_1, M_2) , recover $u^n(M_2)$, then declare $\hat{X}^n$ to be the unique sequence in bin M_1 that is jointly typical with $u^n(M_2)$.
- *Error analysis.* The source atypicality and the helper’s failure are controlled by the AEP and the covering lemma: $\Pr(\mathcal{E}_0 \cup \mathcal{E}_2) \rightarrow 0$ if $R_2 > I(U; Y)$, where $\mathcal{E}_2 = \{(Y^n, u^n(m_2)) \notin \mathcal{T}_\epsilon^{(n)}(Y, U) \forall m_2\}$. For the decoder, the *Markov lemma* applies: since $U - Y - X$ and $(u^n(M_2), Y^n)$ and (X^n, Y^n) are jointly typical, the triple $(X^n, u^n(M_2), Y^n)$ is jointly typical with high probability, so the true X^n is a candidate. Finally, the number of $\hat{x}^n \neq X^n$ with $(\hat{x}^n, u^n) \in \mathcal{T}_\epsilon^{(n)}(X, U)$ is at most $2^{nH(X|U)}$ and each is in bin M_1 with probability 2^{-nR_1} , so

$$\Pr(\mathcal{E}_1) \leq 2^{n(H(X|U) - R_1)} \rightarrow 0$$

if $R_1 > H(X|U)$. ■

Remark 7.2 (Lossy counterpart). A related lossy problem is Wyner–Ziv coding, in which the decoder directly observes uncoded side information Y^n and reconstructs X^n within distortion D ; the rate–distortion function is

$$R(D) = \min_{p(u|x), \hat{x}(u,y): \mathbb{E} d(X, \hat{X}) \leq D} [I(X; U) - I(Y; U)],$$

with $U - X - Y$; see (Wyner and Ziv, 1976). Note the scope of this formula: it is the standard Wyner–Ziv setting in which the decoder observes the side information Y^n *uncoded*. The general rate-limited-helper lossy-source-coding problem, in which Y^n is itself conveyed to the decoder through a rate-limited link (a compression of Y^n rather than the raw observation), is a different and harder problem, and is not the region displayed above. We only need the lossless version in this chapter. ○

7.4 One source, two helpers

It is natural to extend the one-helper problem to two helpers: a decoder wishes to reconstruct a source, while two terminals observe correlated side information and encode it separately. The *general* problem — in which the decoder wants a source S that is merely *correlated* with the helpers' observations, rather than a deterministic function of them — is famously *open*. Körner and Marton (Körner and Marton, 1979) isolated the cleanest solvable instance: the helpers observe binary X and Y , and the decoder wants the modulo-two sum $Z = X \oplus Y$ (with the image-of-a-set method of (Körner and Marton, 1977b) behind the outer bound). We develop the weak (Fano-based) converses here and defer the full region analysis to §7.6.

Setting. Encoder 1 observes X^n , Encoder 2 observes Y^n (binary), and the decoder estimates $\hat{Z}^n$ of $Z^n = X^n \oplus Y^n$ at rates (R_1, R_2) , with $\Pr(\hat{Z}^n \neq Z^n) \rightarrow 0$.

Theorem 7.7. (Weak converses for the sum problem) Any achievable pair must satisfy

$$R_1 \geq H(Z|Y) = H(X|Y), \quad R_2 \geq H(Z|X) = H(Y|X), \quad R_1 + R_2 \geq H(Z).$$

○

Proof. First inequality. Give the decoder Y^n as a genie (this can only help, so it yields an outer bound). Then

$$\begin{aligned} nR_1 &\geq H(M_1) \geq I(M_1; Z^n | M_2, Y^n) \\ &= I(M_1 M_2; Z^n | Y^n) && (M_2 \text{ is a function of } Y^n) \\ &= H(Z^n | Y^n) - H(Z^n | M_1 M_2 Y^n) \\ &\geq H(Z^n | Y^n) - n\epsilon_n && (\text{Fano}) \\ &= nH(Z|Y) - n\epsilon_n. \end{aligned} \tag{7.12}$$

Here $H(Z|Y) = H(X \oplus Y|Y) = H(X|Y)$ because, given Y , the map $x \mapsto x \oplus Y$ is a bijection. *Second inequality.* Symmetric, with Y and X interchanged (note: it is $R_2 \geq H(Z|X)$, not R_1). *Sum rate.*

$$n(R_1 + R_2) \geq H(M_1 M_2) \geq I(M_1 M_2; Z^n) = H(Z^n) - H(Z^n | M_1 M_2) \geq nH(Z) - n\epsilon_n.$$

■

Remark 7.3 (Doubly symmetric binary source). For the doubly symmetric binary source (DSBS) with

$$\Pr(X=0, Y=0) = \Pr(X=1, Y=1) = \frac{1+\rho}{4}, \quad \Pr(X=0, Y=1) = \Pr(X=1, Y=0) = \frac{1-\rho}{4},$$

($-1 < \rho < 1$), both X and Y are uniform and $Z = X \oplus Y$ is independent of X and of Y ; hence $H(Z|X) = H(Z|Y) = H(Z)$, and the weak converses collapse to

$$R_1 \geq H(Z), \quad R_2 \geq H(Z)$$

(the sum-rate bound is then redundant). This corner point is *achievable* by a random linear code over $\text{GF}(2)$, which we develop in §7.6: take a random $k \times n$ binary matrix A , send $M_1 = AX^n$ and $M_2 = AY^n$; the decoder observes $M_1 \oplus M_2 = A(X^n \oplus Y^n) = AZ^n$ and recovers the unique typical Z^n whenever $k/n > H(Z)$. $\circ$

7.5 The Gray–Wyner system

The Gray–Wyner system (Wyner, 1975a) is the distributed coding problem in which a *common* description of (X^n, Y^n) is available to *both* decoders, while each decoder additionally receives a private description of its own source. A single encoder sees (X^n, Y^n) and sends a common index M_0 at rate R_0 to both decoders; Encoder 1 sends M_1 (rate R_1) to Decoder 1; Encoder 2 sends M_2 (rate R_2) to Decoder 2. Decoder 1 reconstructs $\hat{X}^n$ from (M_0, M_1) and Decoder 2 reconstructs $\hat{Y}^n$ from (M_0, M_2) .




Figure 7.3: The Gray–Wyner system.

Theorem 7.8. (Gray–Wyner) The optimal rate region is the set of triples (R_0, R_1, R_2) satisfying

$$R_0 \geq I(U; XY), \quad R_1 \geq H(X|U), \quad R_2 \geq H(Y|U), \quad (7.13)$$

for some $p(u|x, y)$ (equivalently, an auxiliary U jointly distributed with (X, Y)), with $|\mathcal{U}| \leq |\mathcal{X}||\mathcal{Y}| + 2$. $\circ$

Proof. Converse. Identify $U_i = (M_0, X^{i-1}, Y^{i-1})$. First,

$$\begin{aligned}
 nR_0 &\geq H(M_0) \geq I(M_0; X^n Y^n) \\
 &= \sum_{i=1}^n I(M_0; X_i Y_i | X^{i-1} Y^{i-1}) \\
 &= \sum_{i=1}^n I(M_0 X^{i-1} Y^{i-1}; X_i Y_i) && ((X_i, Y_i) \perp\!\!\!\perp (X^{i-1}, Y^{i-1})) \\
 &= \sum_{i=1}^n I(U_i; X_i Y_i).
 \end{aligned} \quad (7.14)$$

For the private rate,

$$\begin{aligned}
nR_1 &\geq H(M_1) \geq H(M_1|M_0) \\
&= I(M_1; X^n|M_0) + H(M_1|X^n M_0) \\
&= H(X^n|M_0) - H(X^n|M_1 M_0) + H(M_1|X^n M_0) \\
&= H(X^n|M_0) - H(X^n|\hat{X}^n M_1 M_0) + H(M_1|X^n M_0) \quad (\hat{X}^n \text{ is a function of } (M_0, M_1)) \\
&\geq H(X^n|M_0) - n\epsilon_n \quad (\text{Fano; } H(M_1|X^n M_0) \geq 0) \\
&= \sum_{i=1}^n H(X_i|M_0 X^{i-1}) - n\epsilon_n \\
&\geq \sum_{i=1}^n H(X_i|M_0 X^{i-1} Y^{i-1}) - n\epsilon_n \quad (\text{conditioning reduces entropy}) \\
&= \sum_{i=1}^n H(X_i|U_i) - n\epsilon_n, \tag{7.15}
\end{aligned}$$

and symmetrically $nR_2 \geq \sum_i H(Y_i|U_i) - n\epsilon_n$. The converse is a chain-rule-plus-Fano argument. Note the two different steps: the chain rule $H(X^n|M_0) = \sum_i H(X_i|M_0, X^{i-1})$ is an exact identity, but the conditioning step in the penultimate line is a genuine inequality, with gap $I(X_i; Y^{i-1}|M_0, X^{i-1}) \geq 0$; this gap need *not* vanish, because conditioning on M_0 can couple the coordinates (the i.i.d. structure of the unconditioned source does not survive the common description). The single-letterisation $U_i = (M_0, X^{i-1}, Y^{i-1})$ is an exact identification by the chain rule: the R_0 bound uses $(X_i, Y_i) \perp\!\!\!\perp (X^{i-1}, Y^{i-1})$ to write $I(M_0; X_i Y_i | X^{i-1} Y^{i-1}) = I(U_i; X_i Y_i)$ exactly, while the conditioning gap above is absorbed as a nonnegative term in the R_1 and R_2 bounds.

Introduce $Q \sim \text{Unif}[1 : n]$, $X = X_Q$, $Y = Y_Q$, $U = (U_Q, Q)$; as in §7.3, $\frac{1}{n} \sum_i I(U_i; X_i Y_i) = I(U; XY)$, $\frac{1}{n} \sum_i H(X_i|U_i) = H(X|U)$, $\frac{1}{n} \sum_i H(Y_i|U_i) = H(Y|U)$, giving (7.13). The cardinality bound follows from the support lemma with the $|X||Y|$ constraints fixing $p(x, y)$, plus $H(X|U), H(Y|U), I(U; XY)$. ■

Achievability. Fix $p(u|x, y)$ and $R_0 > I(U; XY)$, $R_1 > H(X|U)$, $R_2 > H(Y|U)$.

- *Codebook.* Generate 2^{nR_0} sequences $u^n(m_0) \sim \prod_i p(u_i)$. Bin each typical x^n into $[1 : 2^{nR_1}]$ and each typical y^n into $[1 : 2^{nR_2}]$ (independently and uniformly).
- *Encoding.* Encoder 0 finds m_0 with $(X^n, Y^n, u^n(m_0)) \in \mathcal{T}_\epsilon^{(n)}(X, Y, U)$ and sends it; Encoders 1, 2 send the bin indices of X^n, Y^n .
- *Decoding.* Decoder 1 declares the unique $\hat{X}^n$ in bin M_1 with $(\hat{X}^n, u^n(M_0)) \in \mathcal{T}_\epsilon^{(n)}(X, U)$; Decoder 2 symmetrically for Y .
- *Error analysis.* With

$$\mathcal{E}_0 = \{(X^n, Y^n, u^n(m_0)) \notin \mathcal{T}_\epsilon^{(n)}(X, Y, U) \forall m_0\}, \tag{7.16}$$

$$\mathcal{E}_1 = \{(X^n(k), u^n(1)) \in \mathcal{T}_\epsilon^{(n)}(X, U) \text{ for some } k \neq 1\}, \tag{7.17}$$

$$\mathcal{E}_2 = \{(Y^n(k), u^n(1)) \in \mathcal{T}_\epsilon^{(n)}(Y, U) \text{ for some } k \neq 1\}, \quad (7.18)$$

the *covering lemma* gives $\Pr(\mathcal{E}_0) \rightarrow 0$ if $R_0 > I(U; XY)$: each of the 2^{nR_0} i.i.d. u^n sequences is jointly typical with (X^n, Y^n) with probability $\approx 2^{-nI(U; XY)}$, so at least one succeeds with high probability. For $\mathcal{E}_1$, direct type counting applies (no packing lemma is needed): the number of $\hat{x}^n$ with $(\hat{x}^n, u^n) \in \mathcal{T}_\epsilon^{(n)}(X, U)$ is at most $2^{nH(X|U)}$ and each lands in bin M_1 with probability 2^{-nR_1} , so $\Pr(\mathcal{E}_1) \leq 2^{n(H(X|U)-R_1)} \rightarrow 0$ if $R_1 > H(X|U)$, and symmetrically $\Pr(\mathcal{E}_2) \leq 2^{n(H(Y|U)-R_2)} \rightarrow 0$ if $R_2 > H(Y|U)$. ■

7.5.1 The supporting hyperplane and the corner points

A supporting hyperplane of the Gray–Wyner region with normal $(\lambda_0, \lambda_1, \lambda_2) \geq 0$, $\lambda_0 > 0$, characterises a boundary point through

$$\lambda_0 R_0 + \lambda_1 R_1 + \lambda_2 R_2 = \min_{p(u|x,y)} [\lambda_0 I(U; XY) + \lambda_1 H(X|U) + \lambda_2 H(Y|U)],$$

and, normalising $\lambda_0 = 1$ (valid only when $\lambda_0 > 0$), write

$$F(\lambda_1, \lambda_2) \triangleq \min_{p(u|x,y)} [I(U; XY) + \lambda_1 H(X|U) + \lambda_2 H(Y|U)].$$

The elementary weight regimes of F are as follows.

Regime	Optimizer	Value
$\lambda_1 + \lambda_2 \leq 1$	$U = \emptyset$	$\lambda_1 H(X) + \lambda_2 H(Y)$
$\lambda_1 \geq 1, 0 \leq \lambda_2 \leq 1$	$U = X$	$H(X) + \lambda_2 H(Y X)$
$\lambda_2 \geq 1, 0 \leq \lambda_1 \leq 1$	$U = Y$	$H(Y) + \lambda_1 H(X Y)$
$\lambda_1, \lambda_2 \geq 1$	$U = (X, Y)$	$H(X, Y)$

The first regime: using $I(U; XY) = H(XY) - H(XY|U)$, $H(X|U) = H(XY|U) - H(Y|XU)$, $H(Y|U) = H(XY|U) - H(X|YU)$,

$$I(U; XY) + \lambda_1 H(X|U) + \lambda_2 H(Y|U) = H(XY) + (\lambda_1 + \lambda_2 - 1)H(XY|U) - \lambda_1 H(Y|XU) - \lambda_2 H(X|YU),$$

and since $H(Y|XU) \leq H(Y|X)$, $H(X|YU) \leq H(X|Y)$, $H(XY|U) \leq H(XY)$, all three remainder terms have favourable signs when $\lambda_1 + \lambda_2 \leq 1$, with simultaneous equality at $U = \emptyset$. The second regime follows from the expansion

$$\begin{aligned} I(U; XY) + \lambda_1 H(X|U) + \lambda_2 H(Y|U) \\ = H(X) + \lambda_2 H(Y | X) + (\lambda_1 - 1)H(X|U) + (1 - \lambda_2)I(U; Y | X) + \lambda_2 I(X; Y | U), \end{aligned}$$

whose remainder is nonnegative for $\lambda_1 \geq 1, 0 \leq \lambda_2 \leq 1$, with equality at $U = X$; the third regime is symmetric. For $\lambda_1, \lambda_2 \geq 1$,

$$I(U; XY) + \lambda_1 H(X|U) + \lambda_2 H(Y|U) = H(X, Y) + I(X; Y | U) + (\lambda_1 - 1)H(X|U) + (\lambda_2 - 1)H(Y|U),$$

so $U = (X, Y)$ attains the lower bound $H(X, Y)$.

The remaining regime $0 < \lambda_1, \lambda_2 < 1$ with $\lambda_1 + \lambda_2 > 1$ is nontrivial. The hypercontractivity ribbon identifies which of these weight vectors keep the constant auxiliary optimal; evaluating the other supporting values requires the full minimization over $p(u | x, y)$. For a *fixed* U jointly distributed with (X, Y) ,

$$[I(U; XY) + \lambda_1 H(X|U) + \lambda_2 H(Y|U)] - [\lambda_1 H(X) + \lambda_2 H(Y)] = [I(U; XY) - \lambda_1 I(U; X) - \lambda_2 I(U; Y)],$$

so the constant choice $U = \emptyset$ is optimal among unrestricted $U - (X, Y)$ only if $I(U; XY) \geq \lambda_1 I(U; X) + \lambda_2 I(U; Y)$ for *every* such U . Restricting this test to the Markov class $U - X - Y$ (where $I(U; XY) = I(U; X)$) yields a *necessary* condition for unrestricted optimality of the constant auxiliary, not a sufficient one:

$$I(U; Y) \leq \frac{1 - \lambda_1}{\lambda_2} I(U; X) \quad \text{for all } U - X - Y,$$

which, by the sharpened data-processing inequality of §7.7, holds exactly when $(1 - \lambda_1)/\lambda_2 \geq S^*(X; Y)$, i.e.

$$\lambda_1 + S^*(X; Y) \lambda_2 \leq 1.$$

Remark 7.4 (Scope of the contraction connection). The hypercontractivity ribbon characterises the supporting normals at the all-private Gray–Wyner point $(0, H(X), H(Y))$: the full condition is

$$I(U; X, Y) \geq \lambda_1 I(U; X) + \lambda_2 I(U; Y) \quad \text{for every } U \text{ jointly distributed with } (X, Y),$$

and the ribbon is the set of weights satisfying it. The scalar SDPI constants give *necessary tests* on the restricted Markov subfamilies: restricting to $U - X - Y$ gives $\lambda_1 + S^*(X; Y) \lambda_2 \leq 1$, and restricting to $U - Y - X$ gives the symmetric test $S^*(Y; X) \lambda_1 + \lambda_2 \leq 1$. These restricted tests are not a general substitute for the full ribbon, and neither the scalar constants alone nor the supporting cone at this single point evaluates the entire nonlinear Gray–Wyner boundary (Ahlsweede and Han, 1983). $\circ$

Remark 7.5 (Common information). The rate R_0 in the Gray–Wyner system is the rate at which one extracts a *common* description U of (X, Y) . Wyner’s *common information*

$$C_W(X; Y) \triangleq \min_{U: X \perp\!\!\!\perp Y | U} I(U; XY),$$

equivalently over Markov chains $X - U - Y$, is read off the Gray–Wyner region as the minimum common rate on the *minimum-total-rate face* $R_0 + R_1 + R_2 = H(X, Y)$: for the auxiliary description,

$$I(U; X, Y) + H(X | U) + H(Y | U) = H(X, Y) + I(X; Y | U),$$

so minimum total rate forces $X \perp\!\!\!\perp Y | U$, and minimising the common rate over such U gives $C_W(X; Y)$. (The minimisation is over U jointly distributed with (X, Y) subject to $X - U - Y$; it is not characterised by the two chains $X - Y - U$ and $Y - X - U$.) The all-common point of the region is $(H(X, Y), 0, 0)$ – encoder 0 transmits the full joint description and the private

rates vanish — while the all-private point is $(0, H(X), H(Y))$: with no common description the two decoders receive only their own encoder's bits, so each needs the full entropy of its source. In particular the minimum common rate at private rates $(H(X), H(Y))$ is 0, not $C_W(X; Y)$, the point $(C_W(X; Y), H(X), H(Y))$ is dominated by $(0, H(X), H(Y))$, and the Slepian–Wolf point $(0, H(X | Y), H(Y | X))$ is not generally inside the Gray–Wyner region at all (decoding jointly from both descriptions is not available to the individual Gray–Wyner decoders). See (Wyner, 1975a). $\circ$

7.6 The Körner–Marton modulo-two sum problem

We return to the setting of §7.4: Encoders 1, 2 observe binary X^n, Y^n and the decoder wants $Z^n = X^n \oplus Y^n$. The linear-code achievability is due to Körner–Marton (Körner and Marton, 1979), the outer bound below is due to Ahlswede–Han (Ahlswede and Han, 1983) (building on the image-of-a-set method of (Körner and Marton, 1977b)), and the sharpened analysis to Sefidgaran–Gohari–Aref (Sefidgaran et al., 2015).

7.6.1 Linear-code achievability

Theorem 7.9. The rate pair $R_1 = R_2 = H(Z)$ (indeed any $R_1, R_2 > H(Z)$) is achievable. $\circ$

Proof. • *Code generation.* Let $A \in \mathbb{F}_2^{k \times n}$ be a random matrix whose kn entries are i.i.d. $\text{Bern}(\frac{1}{2})$. Each encoder's message is a k -bit string, so the rate is $R_1 = R_2 = k/n$. (For arbitrary $R_1, R_2 > H(Z)$, first choose k with $k/n > H(Z)$ and use the *same* $k \times n$ matrix A at both encoders — taking different row subsets at the two encoders would prevent the decoder from forming the single common syndrome $A(X^n \oplus Y^n)$ — and then append independent dummy bits to one or both messages until the declared rates R_1, R_2 are reached; the decoder ignores the dummy bits. This uses the upward closure of the rate region.)

- *Encoding.* Encoder 1 sends $M_1 = AX^n \in \mathbb{F}_2^k$; Encoder 2 sends $M_2 = AY^n$.
- *Decoding.* The decoder forms $M_1 \oplus M_2 = A(X^n \oplus Y^n) = AZ^n$ and declares $\hat{Z}^n$ to be the unique typical sequence with $A\hat{Z}^n = AZ^n$.
- *Error analysis.* With high probability $Z^n \in \mathcal{T}_\epsilon^{(n)}(Z)$. For any $\hat{z}^n \neq Z^n$ in $\mathcal{T}_\epsilon^{(n)}(Z)$, writing $w^n = \hat{z}^n \oplus Z^n \neq 0$,

$$\Pr(A\hat{z}^n = AZ^n) = \Pr(Aw^n = 0) = \frac{1}{2^k},$$

because the rows of A are independent and each $a_i \cdot w^n$ is an unbiased Bernoulli bit whenever $w^n \neq 0$ (some coordinate of w^n is 1, and flipping the corresponding column flips $a_i \cdot w^n$). There are at most $2^{nH(Z)}$ typical sequences, so by the union bound

$$\Pr(\mathcal{E}_1) \leq 2^{nH(Z)} \cdot 2^{-k} = 2^{n(H(Z) - k/n)} = 2^{n(H(Z) - R)} \rightarrow 0$$

if $R = k/n > H(Z)$. ■

Note the crucial scaling: the rate is $R = k/n$ (transmitted bits per source symbol), and the condition is $R = k/n > H(Z)$, not n/k .

7.6.2 The Ahlswede–Han inner bound

Theorem 7.10. (Ahlswede–Han (Ahlswede and Han, 1983)) A rate pair (R_1, R_2) is achievable if

$$R_1 \geq I(U; X|V) + H(Z|UV), \quad (7.19)$$

$$R_2 \geq I(V; Y|U) + H(Z|UV), \quad (7.20)$$

$$R_1 + R_2 \geq I(UV; XY) + 2H(Z|UV), \quad (7.21)$$

for some $U - X - Y - V$ with $p(u, v|x, y) = p(u|x)p(v|y)$. ○

Proof sketch. The scheme interpolates between two extremes. Fix $p(u|x)$, $p(v|y)$. Encoder 1 first quantises X^n to a typical U^n and sends its index (at rate $I(U; X)$); Encoder 2 quantises Y^n to a typical V^n (at rate $I(V; Y)$). Conditioned on the pair (U^n, V^n) , each encoder then transmits the residual information about the modulo-two sum through a linear code, at rate $H(Z|U, V)$ each (as in §7.6 but conditioned on (U, V)); the decoder combines the two linear syndromes. The three inequalities reflect, respectively: the cost of describing U given V plus the linear residual; the cost of describing V given U plus the linear residual; and the joint cost of (U, V) plus the two residuals.

At $U = V = \emptyset$ the region reduces to the linear-code bound ($R_1, R_2 \geq H(Z)$, $R_1 + R_2 \geq 2H(Z)$); at $U = X$, $V = Y$ the conditional entropy $H(Z|XY) = 0$ vanishes (since $Z = X \oplus Y$), and the region reduces to the Slepian–Wolf region ($R_1 \geq H(X|Y)$, $R_2 \geq H(Y|X)$, $R_1 + R_2 \geq H(XY)$). The general region is obtained by binning and covering exactly as in §7.5, with the two auxiliary descriptions played off against the two linear syndromes. ■

7.6.3 An outer bound via the Csiszár sum identity

Theorem 7.11. (Outer bound) Any achievable pair (R_1, R_2) must satisfy, for every $\lambda \geq 1$,

$$R_1 + \lambda R_2 \geq H(XY) + \min_{U-X-Y} [\lambda H(Z|U) - H(Y|U)], \quad (7.22)$$

$$\lambda R_1 + R_2 \geq H(XY) + \min_{V-Y-X} [\lambda H(Z|V) - H(X|V)]. \quad (7.23)$$

○

Proof sketch (following Ahlswede–Han and Körner–Marton). We prove the first inequality; the second is symmetric. The derivation follows (Ahlswede and Han, 1983; Körner and Marton, 1977b); the long entropy identities below are reproduced for completeness and should be read as a cited proof sketch, with the Fano decoding conditions spelled out where they are used. The only non-elementary input is the following telescoping identity.

Lemma 7.12. (Csiszár sum identity (Csiszar and Körner, 2011)) For any Y^n, Z^n and side information M ,

$$\sum_{i=1}^n I(Z_i; Y^{i-1} | M, Z_{i+1}^n) = \sum_{i=1}^n I(Y_i; Z_{i+1}^n | M, Y^{i-1}).$$

○

Proof. By the chain rule for mutual information,

$$\begin{aligned} \sum_i I(Z_i; Y^{i-1} | M, Z_{i+1}^n) &= \sum_i \sum_{j < i} I(Z_i; Y_j | M, Y^{j-1}, Z_{i+1}^n) \\ &= \sum_j \sum_{i > j} I(Z_i; Y_j | M, Y^{j-1}, Z_{i+1}^n) = \sum_j I(Z_{j+1}^n; Y_j | M, Y^{j-1}). \end{aligned}$$

■

Now, for $\lambda \geq 1$,

$$\begin{aligned} n(R_1 + \lambda R_2) &\geq H(M_1) + \lambda H(M_2) \\ &= H(M_1 M_2) + (\lambda - 1)H(M_2) + I(M_1; M_2) \\ &\geq H(M_1 M_2) + (\lambda - 1)H(M_2 | M_1) \\ &\geq I(M_1 M_2; X^n Y^n) + (\lambda - 1)H(M_2 | M_1) \\ &= H(X^n Y^n) - H(X^n Y^n | M_1 M_2) + (\lambda - 1)H(M_2 | M_1). \end{aligned} \tag{7.24}$$

We now expand the last two terms. Using $H(X^n Y^n) = H(X^n Z^n)$ (the map $(x^n, y^n) \mapsto (x^n, z^n)$ is a bijection with $z^n = x^n \oplus y^n$), and $Z^n = X^n \oplus Y^n$:

$$\begin{aligned} &-H(X^n Y^n | M_1 M_2) + (\lambda - 1)H(M_2 | M_1) \\ &= -H(Y^n | M_1) + I(Y^n; M_2 | M_1) - H(X^n | M_1 M_2 Y^n) \\ &\quad + \lambda H(Z^n M_2 | M_1) - \lambda H(Z^n | M_1 M_2) - H(M_2 | M_1). \end{aligned} \tag{7.25}$$

To verify (7.25), expand $H(Z^n M_2 | M_1) = H(Z^n | M_1) + H(M_2 | M_1 Z^n)$ and $I(Y^n; M_2 | M_1) = H(M_2 | M_1) - H(M_2 | M_1 Y^n)$; then the $H(M_2 | M_1)$ terms cancel and, using $H(X^n | M_1 M_2 Y^n) = H(Z^n | M_1 M_2 Y^n)$ and the bijection $(X^n, Y^n) \leftrightarrow (Z^n, Y^n)$, the identity reduces to the chain rule for $H(X^n Y^n | M_1 M_2)$.

Substituting into (7.24) and applying Fano's inequality to both $H(X^n | M_1 M_2 Y^n)$ and $H(Z^n | M_1 M_2)$ (both are decodable: Z^n from (M_1, M_2) , and $X^n = Z^n \oplus Y^n$ from (M_1, M_2, Y^n)):

$$n(R_1 + \lambda R_2) \geq H(X^n Y^n) - H(Y^n | M_1) + I(Y^n; M_2 | M_1)$$

$$\begin{aligned}
& + \lambda H(Z^n|M_1) + \lambda H(M_2|M_1 Z^n) - H(M_2|M_1) - (1 + \lambda)n\epsilon_n \\
& \stackrel{(*)}{\geq} H(X^n Y^n) - H(Y^n|M_1) + \lambda H(Z^n|M_1) - (1 + \lambda)n\epsilon_n.
\end{aligned} \tag{7.26}$$

To justify (*), it suffices to show $I(Y^n; M_2|M_1) + \lambda H(M_2|M_1 Z^n) - H(M_2|M_1) \geq 0$. Indeed

$$\begin{aligned}
& I(Y^n; M_2|M_1) + \lambda H(M_2|M_1 Z^n) - H(M_2|M_1) \\
& = \lambda H(M_2|M_1 Z^n) - H(M_2|M_1 Y^n) \geq H(M_2|M_1 Z^n) - H(M_2|M_1 Y^n),
\end{aligned}$$

and $H(M_2|M_1 Z^n) \geq H(M_2|M_1 Y^n)$ follows from

$$H(M_2|M_1 Y^n) = H(M_2|M_1 Y^n X^n) = H(M_2|M_1 Y^n Z^n) \leq H(M_2|M_1 Z^n),$$

where the first equality uses $I(M_2; X^n|M_1 Y^n) = 0$ – by d-separation (§9.4 of the Appendix), the functional node $M_2 = f_2(Y^n)$ has all its dependence through its parent Y^n , so $M_2 - Y^n - X^n$ given M_1 – and the second uses the bijection $(X^n, Y^n) \leftrightarrow (Z^n, Y^n)$. This step is the *data-processing contraction* of the converse.

It remains to single-letterise

$$\Phi \triangleq H(X^n Y^n) - H(Y^n|M_1) + \lambda H(Z^n|M_1).$$

Using $H(X^n Y^n) = \sum_i H(X_i Y_i)$ (memoryless), the forward chain rule $H(Y^n|M_1) = \sum_i H(Y_i|M_1 Y^{i-1})$, and the backward chain rule $H(Z^n|M_1) = \sum_i H(Z_i|M_1 Z_{i+1}^n)$:

$$\Phi = \sum_i \left[H(X_i Y_i) - H(Y_i|M_1 Y^{i-1}) + \lambda H(Z_i|M_1 Z_{i+1}^n) \right].$$

With $U_i = (M_1, Y^{i-1}, Z_{i+1}^n)$,

$$H(Y_i|M_1 Y^{i-1}) = H(Y_i|U_i) + I(Y_i; Z_{i+1}^n|M_1 Y^{i-1}), \quad H(Z_i|M_1 Z_{i+1}^n) = H(Z_i|U_i) + I(Z_i; Y^{i-1}|M_1 Z_{i+1}^n),$$

so

$$\begin{aligned}
\Phi &= \sum_i \left[H(X_i Y_i) - H(Y_i|U_i) + \lambda H(Z_i|U_i) \right] \\
&\quad - \sum_i I(Y_i; Z_{i+1}^n|M_1 Y^{i-1}) + \lambda \sum_i I(Z_i; Y^{i-1}|M_1 Z_{i+1}^n) \\
&= \sum_i \left[H(X_i Y_i) - H(Y_i|U_i) + \lambda H(Z_i|U_i) \right] + (\lambda - 1) \sum_i I(Y_i; Z_{i+1}^n|M_1 Y^{i-1}) \\
&\geq \sum_i \left[H(X_i Y_i) - H(Y_i|U_i) + \lambda H(Z_i|U_i) \right],
\end{aligned} \tag{7.27}$$

where the second line used the Csiszár sum identity (with $M = M_1$) and the last inequality used $\lambda \geq 1$. This is the *Csiszár-sum telescoping*: the identity makes the sub-additivity corrections cancel, leaving the per-letter bound with a non-negative remainder.

It remains to check the Markov structure $U_i - X_i - Y_i$. Since $M_1 = f_1(X^n)$ and the source is memoryless, given X_i the variable Y_i is independent of $(M_1, X^{i-1}, X_{i+1}^n, Y^{i-1}, Y_{i+1}^n)$ and hence of $U_i = (M_1, Y^{i-1}, X_{i+1}^n \oplus Y_{i+1}^n)$; thus $U_i - X_i - Y_i$.

Finally introduce $Q \sim \text{Unif}[1 : n]$, $X = X_Q$, $Y = Y_Q$, $Z = Z_Q$, $U = (U_Q, Q)$. Then (7.26)–(7.27) give

$$R_1 + \lambda R_2 \geq H(XY) - H(Y|U) + \lambda H(Z|U) = H(XY) + [\lambda H(Z|U) - H(Y|U)],$$

with $U - X - Y$, and the minimisation over U completes the proof. $\blacksquare$

Remark 7.6 (Sharpness at $\lambda = 1$). At $\lambda = 1$ the first inequality reads $R_1 + R_2 \geq H(XY) + \min_{U-X-Y} [H(Z|U) - H(Y|U)]$. For the DSBS ($Y = X \oplus Z$ with $Z \sim \text{Bern}(p)$ independent of X and the transition translation-invariant), the minimiser is $U = \emptyset$, and the verification is direct: for any $U - X - Y$,

$$p(z | x, u) = p_{Y|X}(x \oplus z | x) = p(z)$$

(the DSBS transition is translation-invariant), so $Z \perp\!\!\!\perp (X, U)$ and $H(Z|U) = H(Z)$; since Y is binary, $H(Y|U) \leq 1$. Hence

$$H(Y|U) - H(Z|U) \leq 1 - H(Z) = H(Y) - H(Z),$$

with equality at $U = \emptyset$, which proves the minimisation and recovers the correct DSBS sum-rate $R_1 + R_2 \geq H(XY) + H(Z) - 1 = 2H(Z)$. $\circ$

Conjecture 7.13. (Gohari) For the doubly symmetric binary source,

$$\min_{U-X-Y-V} [I(UV; XY) + 2H(Z|UV)] = \min\{H(XY), 2H(Z)\},$$

i.e. the optimal sum-rate of the Ahlswede–Han inner bound equals $\min\{H(XY), 2H(Z)\}$; see (Sefidgaran et al., 2015). $\circ$

Remark 7.7 (Further reading: the weighted-sum rates of the modulo-sum problem). The supporting-hyperplane computation of §7.5.1 can be carried out for the Körner–Marton problem as well: the optimal weighted-sum rates — the *minimal* values $\min(R_1 + \lambda R_2)$, since the rate region is upward closed and the weighted sum is a cost to be minimised — for the modulo-two sum problem are determined in (Wang and Nair, 2020), again by a combination of the Csiszár sum identity (the contraction telescoping of this chapter) and the min–max bridge of the appendix. It is a good exercise to reproduce the two extreme regimes ($\lambda = 1$ recovering the sum-rate, and the linear-code regime) from that paper. $\circ$

7.7 Hypercontractivity: the sharp contraction calculus

The data-processing inequality (DPI) for KL divergence, Lemma 1.14, says that for a channel $W : \mathcal{X} \rightarrow \mathcal{Y}$ and two input distributions $\nu_X \ll \mu_X$,

$$D(\nu_Y \| \mu_Y) \leq D(\nu_X \| \mu_X), \quad \nu_Y = W\nu_X, \quad \mu_Y = W\mu_X.$$

Equivalently, for $U - X - Y$, $I(U; Y) \leq I(U; X)$. Hypercontractivity is the *sharp* form of this contraction: the constant that measures how far the contraction is from saturation. Let $(X, Y) \sim \mu_X W_{Y|X}$.

Definition 7.14. (Contraction coefficient) The sharp (strong data-processing) contraction constant of $(X; Y)$ is

$$S^*(X; Y) \triangleq \sup_{v_X: v_X \ll \mu_X, D(v_X \| \mu_X) > 0} \frac{D(W v_X \| W \mu_X)}{D(v_X \| \mu_X)} \in [0, 1]. \quad (7.28)$$

○

By the DPI, $S^* \leq 1$; non-negativity is clear. The two extremes are: $S^*(X; Y) = 1$ when $X = Y$ (identity channel, contraction is exact), and $S^*(X; Y) = 0$ when $X \perp\!\!\!\perp Y$ (constant channel, all inputs collapse to μ_Y).

Theorem 7.15. (Sharpened DPI (Anantharam et al., 2013, 2014)) For every $U - X - Y$,

$$I(U; Y) \leq S^*(X; Y) I(U; X),$$

and the constant is sharp:

$$S^*(X; Y) = \sup_{U-X-Y, I(U; X) > 0} \frac{I(U; Y)}{I(U; X)}.$$

○

Proof. The first assertion is immediate from (7.28): writing $v_{X|u} = p(x|u)$,

$$I(U; Y) = \sum_u p(u) D(W v_{X|u} \| W \mu_X) \leq S^* \sum_u p(u) D(v_{X|u} \| \mu_X) = S^* I(U; X).$$

For sharpness we prove the following equivalence. ■

Lemma 7.16. For $\lambda \geq 0$,

$$\left\{ \lambda : I(U; Y) \leq \lambda I(U; X) \ \forall U - X - Y \right\} = \left\{ \lambda : D(W v_X \| W \mu_X) \leq \lambda D(v_X \| \mu_X) \ \forall v_X \ll \mu_X \right\}.$$

○

Proof. (RHS $\Rightarrow$ LHS). Done in the proof above, using $I(U; Y) = \sum_u p(u) D(p_{Y|u} \| p_Y)$ and $I(U; X) = \sum_u p(u) D(p_{X|u} \| p_X)$.

(LHS $\Rightarrow$ RHS). Fix $v_X \ll \mu_X$. Define a binary U_ϵ with $\Pr(U_\epsilon=1) = \epsilon$, $\Pr(U_\epsilon=0) = 1 - \epsilon$, and

$$p(x|U_\epsilon=1) = v_X(x), \quad p(x|U_\epsilon=0) = \frac{\mu_X(x) - \epsilon v_X(x)}{1 - \epsilon},$$

which is a valid distribution for all sufficiently small ϵ (since $v_X \ll \mu_X$). Then $p_X = \mu_X$, so that $U_\epsilon - X - Y$ holds. By hypothesis,

$$f(\epsilon) \triangleq I(U_\epsilon; Y) - \lambda I(U_\epsilon; X) \leq 0,$$

while $f(0) = 0$ (at $\epsilon = 0$, U_ϵ is degenerate). Hence $f'(0) \leq 0$. Now, since $p_X = \mu_X$,

$$I(U_\epsilon; X) = (1 - \epsilon)D(q_\epsilon \| \mu_X) + \epsilon D(v_X \| \mu_X),$$

where $q_\epsilon(x) = \frac{\mu_X(x) - \epsilon v_X(x)}{1 - \epsilon} = \mu_X(x) + O(\epsilon)$. Because $q \mapsto D(q \| \mu_X)$ is minimised (to first order) at $q = \mu_X$, we have $D(q_\epsilon \| \mu_X) = O(\epsilon^2)$, so

$$I(U_\epsilon; X) = \epsilon D(v_X \| \mu_X) + O(\epsilon^2), \quad I(U_\epsilon; Y) = \epsilon D(W v_X \| W \mu_X) + O(\epsilon^2).$$

Therefore $f'(0) = D(W v_X \| W \mu_X) - \lambda D(v_X \| \mu_X) \leq 0$, as claimed. $\blacksquare$

Remark 7.8 (Erkip–Cover, corrected). Erkip and Cover ([Erkip and Cover, 1998](#)) claimed that $\sup_{U-X-Y} I(U; Y)/I(U; X)$ equals the *squared* maximal correlation $\rho_m^2(X; Y)$, where

$$\rho_m(X; Y) \triangleq \sup_{f, g} \mathbb{E}[f(X)g(Y)] \quad \text{subject to} \quad \mathbb{E} f(X) = \mathbb{E} g(Y) = 0, \mathbb{E} f^2(X) = \mathbb{E} g^2(Y) = 1.$$

This is *incorrect*. Anantharam–Gohari–Kamath–Nair ([Anantharam et al., 2013, 2014](#)) showed that the correct sharp constant is $S^*(X; Y)$ (the contraction coefficient), not ρ_m^2 . $\circ$

7.7.1 Reverse hypercontractivity and the Bonami–Beckner / log-Sobolev statements

The contraction coefficient S^* was introduced by Ahlswede and Gács ([Ahlswede and Gács, 1976](#)) through the *hypercontraction* of the Markov (conditional-expectation) operator. Let

$$(Tg)(x) \triangleq \mathbb{E}[g(Y) | X = x],$$

the adjoint of the channel W , and write $\|g\|_q = (\mathbb{E} |g(Y)|^q)^{1/q}$ (allowing $q \leq 1$ for non-negative g).

Definition 7.17. (Reverse hypercontractivity) For $g \geq 0$, (X, Y) is (p, q) -reverse-hypercontractive if

$$\|Tg\|_p \geq \|g\|_q.$$

$\circ$

Caution on the exponent convention. The naive version of this definition is trivial in a large exponent range: for $p \geq 1$, $0 < q \leq 1$, and $g \geq 0$,

$$\|Tg\|_p \geq \mathbb{E}[Tg] = \mathbb{E}[g] \geq \|g\|_q,$$

so those exponents work for *every* channel, and an infimum over them cannot recover a nontrivial contraction coefficient. We therefore do *not* record an endpoint formula here: a meaningful statement requires the precise exponent domain and normalisation of the reverse ribbon, for which the reader is referred to Ahlswede–Gács (Ahlswede and Gács, 1976) and the information-measure derivations of (Anantharam et al., 2013, 2014; Kamath, 2015). The classical *forward* hypercontractivity is the safer optional connection: for the DSBS noise operator T_ρ ($Y = X \oplus N$ with $\rho = \mathbb{E}[(-1)^X(-1)^Y] = 1 - 2\epsilon$), Bonami–Beckner gives

$$\|T_\rho f\|_q \leq \|f\|_p \quad \text{for } 1 < p \leq q < \infty \text{ with } \rho^2(q-1) \leq p-1,$$

equivalently $\rho^2 \leq (p-1)/(q-1)$ (endpoint cases $p = 1, q = \infty$ are treated separately by continuity/duality). This inequality gets the *direction* right (T_ρ improves integrability), and its squared constant is $\rho^2 = \rho_m^2$, the maximal correlation of the DSBS.

Remark 7.9 (Log-Sobolev). The infinitesimal form of Bonami–Beckner hypercontractivity is the *log-Sobolev inequality* for the noise semigroup, and stating it requires fixing the semigroup and the normalisation of the Dirichlet form. For the discrete heat semigroup $P_t = e^{-t(I-T_\rho)}$ on the two-point space, with the Dirichlet form $\mathcal{E}(f) = \mathbb{E}[f(X)(f(X) - T_\rho f(X))]$ and entropy $\text{Ent}(f^2) = \mathbb{E}[f^2 \log f^2] - \mathbb{E}[f^2] \log \mathbb{E}[f^2]$,

$$\text{Ent}(f^2) \leq c(\rho) \mathcal{E}(f),$$

with sharp constant $c(\rho)$ obtained by differentiating the Bonami–Beckner ribbon at $p = q = 2$ (Gross); the intrinsic log-Sobolev constant is not a universal number but depends on the correlation parameter ρ and the chosen normalisation. We only need the contraction-coefficient form in this course; the operator/norm statements are recorded here for the connection and are treated fully in the literature (Ahlswede and Gács, 1976; Anantharam et al., 2013, 2014). $\circ$

Remark 7.10 (The sub-additivity endpoint). Hypercontractivity is the sub-additivity theme at its endpoint. The DPI $D(v_Y \| \mu_Y) \leq D(v_X \| \mu_X)$ is a KL contraction; $S^*(X; Y)$ is the *best* constant in $D(v_Y \| \mu_Y) \leq S^* D(v_X \| \mu_X)$, i.e. it measures how far the contraction is from the saturated (equality) case. Its equivalent form $I(U; Y) \leq S^* I(U; X)$ is exactly the sharpened version of the mutual-information sub-additivity $I(U; Y) \leq I(U; X)$ of §1.13, and it is this sharp constant — not the maximal correlation — that supplies the necessary restricted-family test for the supporting normals of the Gray–Wyner region at the all-private corner (§7.5.1, Remark on the scope of the contraction connection) and that drives the Körner–Marton outer bound (§7.6). $\circ$

Remark 7.11 (Further reading). The contraction calculus of this section is developed at length in the literature: equivalent information-measure formulations of hypercontractivity (including the reverse direction needed for Ahlswede–Gács) are given in (Nair, 2014); the evaluation of hypercontractivity parameters for specific channels (e.g. the binary erasure channel and its reverse-hypercontractivity region) is carried out in (Nair and Wang, 2016, 2017); and the general sharpened data-processing theory that underlies these evaluations is developed in (Anantharam et al., 2014). The auxiliary-receiver bounds of the interference channel chapter and the relay chapter are current applications of exactly these sharpened contractions. $\circ$

7.8 Closing: the sub-additivity / contraction calculus

The chapter can be read as a single story about a small number of moves. Each converse first lands on an n -letter functional of the induced law, then a *single-letterisation* collapses it. The moves are:

1. *The chain rule* $H(X^n|M) = \sum_i H(X_i|M, X^{i-1})$: the saturated form of entropy sub-additivity. Slepian–Wolf uses i.i.d. entropy additivity throughout; Gray–Wyner uses i.i.d. additivity for its common-rate chain together with Fano and genuine conditioning inequalities for its private-rate bounds (conditioning on the common message can couple otherwise independent coordinates).
2. *Conditioning reduces entropy* $H(X_i|M, X^{i-1}) \geq H(X_i|M, X^{i-1}, Y^{i-1})$: the genuine sub-additivity step of the one-helper R_1 bound, with gap $I(X_i; Y^{i-1}|M_2, X^{i-1})$; the auxiliary $U_i = (M_2, Y^{i-1})$ is then obtained by the exact d-separation identification.
3. *Data-processing contraction* $I(M_2; X^n|M_1 Y^n) = 0$ (equivalently $D(v_Y||\mu_Y) \leq D(v_X||\mu_X)$): the engine of the Körner–Marton outer bound.
4. *The Csiszár sum identity* $\sum_i I(Z_i; Y^{i-1}|M, Z_{i+1}^n) = \sum_i I(Y_i; Z_{i+1}^n|M, Y^{i-1})$: the telescoping identity that makes the sub-additivity corrections of the Körner–Marton bound cancel, turning an n -letter expression into a clean per-letter sum.
5. *The sharp endpoint* $S^*(X; Y)$: the best constant in the contraction, i.e. the sharp form of move (3); it supplies the necessary restricted-family test for the Gray–Wyner supporting normals at the all-private corner, whose full set is the hypercontractivity ribbon.

The taxonomy is honest: Slepian–Wolf is a *saturated* (additivity) bound; Gray–Wyner combines an additivity chain for the common rate with genuine conditioning inequalities for the private rates; the one-helper R_1 bound is *genuinely sub-additive*; the Körner–Marton outer bound is a *contraction + telescoping* bound; and hypercontractivity upgrades the contraction to its *sharp* constant $S^*(X; Y)$. This is the same calculus announced in §1.13 of Chapter 1 and instantiated in the min–max and perturbation machinery of the appendix (§9.3): *sub-additivity of an information functional is the single idea behind every outer bound of this chapter*.

7.9 Exercises

Exercise 2 (Slepian–Wolf corner achievability): For the doubly symmetric binary source of §7.2, describe explicitly the binning scheme that achieves the corner point $(R_1, R_2) = (1, H_b(p))$, and verify that the three error events $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_{12}$ all have vanishing probability.

Exercise 3 (One-helper saturation): Show that in the one-helper converse of §7.3, the R_2 bound $\sum_i I(M_2 Y^{i-1}; Y_i) = \sum_i I(U_i; Y_i)$ is an *equality* (the saturated step), whereas the R_1 bound $\sum_i H(X_i|M_2 X^{i-1}) \geq \sum_i H(X_i|U_i)$ has gap $\sum_i I(X_i; Y^{i-1}|M_2 X^{i-1})$, followed by the *exact* identification $H(X_i|M_2 X^{i-1} Y^{i-1}) = H(X_i|M_2 Y^{i-1})$ of d-separation (§9.4 of the Appendix). Con-

clude that the one-helper region is, in general, strictly larger than the Slepian–Wolf corner $\{(R_1, R_2) : R_1 \geq H(X|Y), R_2 \geq H(Y)\}$.

Exercise 4 (Cardinality bound): Use the support lemma (Fenchel–Eggleston–Carathéodory) to prove the cardinality bound $|\mathcal{U}| \leq |\mathcal{Y}| + 1$ in the one-helper theorem, preserving $p(y)$, $H(X|U)$ and $I(U; Y)$.

Exercise 5 (Gray–Wyner supporting hyperplane): Verify the four elementary weight regimes of §7.5.1 in detail: (i) $\lambda_1 + \lambda_2 \leq 1 \Rightarrow U = \emptyset$ with value $\lambda_1 H(X) + \lambda_2 H(Y)$; (ii) $\lambda_1 \geq 1, 0 \leq \lambda_2 \leq 1 \Rightarrow U = X$ with value $H(X) + \lambda_2 H(Y|X) = H(XY) - (1 - \lambda_2)H(Y|X)$; (iii) $\lambda_2 \geq 1, 0 \leq \lambda_1 \leq 1 \Rightarrow U = Y$ with the symmetric value; (iv) $\lambda_1, \lambda_2 \geq 1 \Rightarrow U = (X, Y)$ with value $H(X, Y)$. Check that the only remaining interior weight regime is $0 < \lambda_1, \lambda_2 < 1$ with $\lambda_1 + \lambda_2 > 1$.

Exercise 6 (Körner–Marton linear code): Prove that for any nonzero $w^n \in \mathbb{F}_2^n$ and a random $A \in \mathbb{F}_2^{k \times n}$ with i.i.d. $\text{Bern}(\frac{1}{2})$ entries, Aw^n is uniform on $\mathbb{F}_2^k$, so $\Pr(Aw^n = 0) = 2^{-k}$. Conclude the error exponent $2^{n(H(Z)-R)}$ in the linear-code achievability of §7.6.

Exercise 7 (Contraction-coefficient extremes): Verify directly from (7.28) that $S^*(X; Y) = 1$ when $X = Y$ a.s. and $S^*(X; Y) = 0$ when $X \perp\!\!\!\perp Y$. For the DSBS with correlation ρ (i.e. $Y = X \oplus N$, $N \sim \text{Bern}(p)$, $\rho = 1 - 2p$), show that $S^* = \rho^2$. *Hint:* exhibit a *sequence* of binary symmetric auxiliaries $U_\delta = X - Y$ whose ratios approach ρ^2 : for the binary symmetric auxiliary with crossover probability $1/2 - \delta$ ($0 < \delta < 1/2$), verify

$$\frac{I(U; Y)}{I(U; X)} = \frac{1 - H_b(1/2 - \rho\delta)}{1 - H_b(1/2 - \delta)} \longrightarrow \rho^2 \quad (\delta \downarrow 0),$$

which need not be attained by any nondegenerate U ; bound S^* from above using that the KL contraction coefficient of a BSC(p) is $(1 - 2p)^2$ (a classical fact; it follows from the Bonami–Beckner inequality of §7.7).

Exercise 8 (Csiszár sum identity): Prove the Csiszár sum identity of §7.6 from the chain rule for mutual information, and verify that it is exactly the telescoping needed in (7.27) to cancel the forward/backward conditioning corrections.

8 Relay Channels

The relay channel is the simplest communication network that forces us to use, simultaneously, several of the coding ideas developed in earlier chapters: *superposition coding*, *block Markov coding*, and *binning*. It is also the setting in which the unifying theme of these notes — *sub-additivity of an information functional* as the engine behind every converse (§1.13) — appears in its cleanest form: the *cut-set outer bound* is nothing but two applications of the time sub-additivity of mutual information. By contrast, the achievable schemes of this chapter (decode–forward and compress–forward) are *typicality-based* random coding arguments, not sub-additivity arguments; we flag this distinction explicitly at the end of the chapter.

The setting was introduced by van der Meulen and put on a firm footing by Cover and El Gamal (1979), who established the cut-set bound, the decode–forward and compress–forward lower bounds, and the capacity of the degraded relay channel. The capacity of the general (even Gaussian) relay channel remains open; see the remarks in §8.5. Our exposition follows Cover and El Gamal (1979); El Gamal and Kim (2012).

8.1 The relay channel: setting and code definition

A *relay channel* is a memoryless channel with one sender, one receiver, and one relay that both *listens to* the channel and *helps* the sender by transmitting its own signal. Formally the channel is specified by four finite alphabets $\mathcal{X}$ (sender input), $\mathcal{X}_r$ (relay input), $\mathcal{Y}_r$ (relay output), $\mathcal{Y}$ (receiver output), and a transition kernel

$$p(y, y_r | x, x_r),$$

where $x \in \mathcal{X}$, $x_r \in \mathcal{X}_r$, $y_r \in \mathcal{Y}_r$, $y \in \mathcal{Y}$. The pair (x, x_r) is sent into the channel and the pair (y, y_r) comes out: y_r is observed only by the relay, y only by the destination. The channel is memoryless: in each use i the output pair depends on the current inputs through the kernel $p(y_i, y_{r,i} | x_i, x_{r,i})$ only. Under a *relay code*, however, the induced joint law of the block is *not* the product of single-use kernels: the relay chooses its inputs adaptively, so future relay inputs can depend on past outputs. The correct causal factorisation of the joint law of a code is

$$p(m, x^n, x_r^n, y^n, y_r^n) = p(m) \mathbf{1}\{x^n = x^n(m)\} \cdot \prod_{i=1}^n \mathbf{1}\{x_{r,i} = f_i(y_r^{i-1})\} p(y_i, y_{r,i} | x_i, x_{r,i}).$$

Here $x^n(m)$ is the sender’s codeword for message m and f_i is the relay’s causal encoding function at time i (a constant when $i = 1$).

The distinctive feature of the relay is its *strict causality*: at time i the relay has observed only Y_r^{i-1} , and on that basis produces the input symbol $X_{r,i}$. The sender, by contrast, knows the message M ahead of time and can encode the entire block.




Figure 8.1: The relay channel. The relay observes Y_r^n causally and, at time i , emits $X_{r,i}$ as a function of Y_r^{i-1} only.

Definition 8.1. (Relay channel code) A $(2^{nR}, n)$ code for the relay channel consists of

- a message set $[1 : 2^{nR}]$ with M uniformly distributed over it;
- an *encoder* that maps each message to a codeword,

$$x^n : [1 : 2^{nR}] \rightarrow \mathcal{X}^n, \quad m \mapsto x^n(m);$$

- a set of *relay encoding functions* (strictly causal),

$$f_i : \mathcal{Y}_r^{i-1} \rightarrow \mathcal{X}_r, \quad i \in [1 : n],$$

so that $x_{r,i} = f_i(y_r^{i-1})$; here f_1 is a constant, since the relay has observed nothing at time 1;

- a *decoder* that maps the received sequence to an estimate of the message,

$$g : \mathcal{Y}^n \rightarrow [1 : 2^{nR}], \quad \hat{m} = g(y^n).$$

The *average probability of error* is

$$P_e^{(n)} \triangleq \Pr\{\hat{M} \neq M\},$$

where the channel is used in the obvious way: $X^n = x^n(M)$, $X_{r,i} = f_i(Y_r^{i-1})$, and (Y^n, Y_r^n) is the channel output. $\bigcirc$

Definition 8.2. (Capacity) A rate R is *achievable* if there exists a sequence of $(2^{nR}, n)$ codes with $P_e^{(n)} \rightarrow 0$ as $n \rightarrow \infty$. The *capacity* C of the relay channel is the supremum of achievable rates. $\bigcirc$

Notational convention. We abbreviate $I((X, X_r); Y)$ as $I(X, X_r; Y)$, and conditional quantities similarly. The joint law of (X, X_r, Y, Y_r) always factors as $p(x, x_r) p(y, y_r | x, x_r)$, where $p(x, x_r)$ is the input distribution and $p(y, y_r | x, x_r)$ is the channel.

8.2 The cut-set outer bound

The fundamental outer bound for the relay channel is the *cut-set bound*. It is the canonical instance of the sub-additivity principle of §1.13: each *cut* of the network defines an information-theoretic functional (a mutual information), the time sub-additivity of that functional collapses the n -letter quantity to a single letter, and the outer bound is the minimum over the cuts followed by the maximum over the input distribution.

A cut of a network is a partition of its nodes into two sets S and S^c ; the *cut value* is the total rate at which information can flow across the cut. For the relay channel there are two nontrivial cuts (the source is node 1, the relay node 2, the destination node 3):

- the *destination cut* $S = \{1, 2\}$, $S^c = \{3\}$, which separates the pair (source, relay) from the destination; and
- the *source cut* $S = \{1\}$, $S^c = \{2, 3\}$, which separates the source from the pair (relay, destination).

Theorem 8.3. (Cut-set bound for the relay channel) The capacity of the discrete memoryless relay channel $p(y, y_r|x, x_r)$ is upper bounded by

$$C \leq \max_{p(x, x_r)} \min \{ I(X, X_r; Y), I(X; Y, Y_r|X_r) \}.$$

○

Proof. Fix a $(2^{nR}, n)$ code with $P_e^{(n)} = \Pr\{\hat{M} \neq M\} \rightarrow 0$. By Fano's inequality (Proposition 1.57), there is a sequence $\epsilon_n \rightarrow 0$ such that

$$H(M|Y^n) \leq H(M|\hat{M}) \leq n\epsilon_n.$$

Both cuts are proved with *causal histories*: each per-letter bound conditions only on the past (Y^{i-1}, Y_r^{i-1}) and on the current inputs, which are determined by (M, Y_r^{i-1}) through the causal encoding functions $X_i = x_i(M)$ and $X_{r,i} = f_i(Y_r^{i-1})$. No future relay input is conditioned on; future relay inputs can depend on past outputs, so conditioning on all of X_r^n as if it were exogenous is not available.

The destination (broadcast) cut. By Fano,

$$\begin{aligned} nR &\leq I(M; Y^n) + n\epsilon_n = \sum_{i=1}^n [H(Y_i | Y^{i-1}) - H(Y_i | M, Y^{i-1})] + n\epsilon_n \\ &\leq \sum_{i=1}^n [H(Y_i) - H(Y_i | M, Y^{i-1}, Y_r^{i-1})] + n\epsilon_n \\ &= \sum_{i=1}^n I(X_i, X_{r,i}; Y_i) + n\epsilon_n. \end{aligned} \tag{8.1}$$

The second inequality uses $H(Y_i|Y^{i-1}) \leq H(Y_i)$ and conditioning reduces entropy. For the equality: given (M, Y^{i-1}, Y_r^{i-1}) the current inputs $(X_i, X_{r,i}) = (x_i(M), f_i(Y_r^{i-1}))$ are determined, and by the memoryless kernel Y_i is conditionally independent of (M, Y^{i-1}, Y_r^{i-1}) given $(X_i, X_{r,i})$, so $I(M, Y^{i-1}, Y_r^{i-1}; Y_i) = I(X_i, X_{r,i}; Y_i)$. Introducing $Q \sim \text{Unif}[1 : n]$ independent of everything, and writing $X = X_Q, X_r = X_{r,Q}, Y = Y_Q$,

$$\frac{1}{n} \sum_{i=1}^n I(X_i, X_{r,i}; Y_i) = I(X_Q, X_{r,Q}; Y_Q | Q) \leq I(X_Q, X_{r,Q}; Y_Q),$$

where the inequality uses

$$I(X_Q, X_{r,Q}; Y_Q) = I(X_Q, X_{r,Q}; Y_Q | Q) + I(Q; Y_Q),$$

obtained from $I(Q; Y_Q | X_Q, X_{r,Q}) = 0$, which follows from the memoryless chain $Q - (X_Q, X_{r,Q}) - Y_Q$; the term $I(Q; Y_Q) \geq 0$ need not vanish (the code may induce different input and output laws at different times), and its nonnegativity is all that is used. Hence

$$R \leq \max_{p(x, x_r)} I(X, X_r; Y) + \epsilon_n. \quad (8.2)$$

The source (multiple-access) cut. Write $O_i = (Y_i, Y_{r,i})$ for the pair of outputs. By Fano, $H(M|O^n) \leq H(M|Y^n) \leq n\epsilon_n$, so

$$\begin{aligned} nR &\leq I(M; O^n) + n\epsilon_n = \sum_{i=1}^n [H(O_i | O^{i-1}) - H(O_i | M, O^{i-1})] + n\epsilon_n \\ &\leq \sum_{i=1}^n [H(O_i | X_{r,i}) - H(O_i | X_i, X_{r,i})] + n\epsilon_n \\ &= \sum_{i=1}^n I(X_i; Y_i, Y_{r,i} | X_{r,i}) + n\epsilon_n. \end{aligned} \quad (8.3)$$

For the two steps: first, $X_{r,i} = f_i(Y_r^{i-1})$ is a function of O^{i-1} , so $H(O_i | O^{i-1}) = H(O_i | O^{i-1}, X_{r,i}) \leq H(O_i | X_{r,i})$. Second, the subtracted entropy satisfies the *equality*

$$H(O_i | M, O^{i-1}) = H(O_i | X_i, X_{r,i}) :$$

the history (M, O^{i-1}) determines the current inputs $(X_i, X_{r,i}) = (x_i(M), f_i(Y_r^{i-1}))$, and by the memoryless kernel the current output pair O_i has law $p(\cdot | X_i, X_{r,i})$ independent of the message and of the past history given the current inputs. (A mere inequality here would point the wrong way, since this entropy is subtracted.) Time-sharing as before gives

$$\frac{1}{n} \sum_{i=1}^n I(X_i; O_i | X_{r,i}) = I(X_Q; O_Q | X_{r,Q}, Q) \leq I(X_Q; O_Q | X_{r,Q}),$$

using $I(Q; O_Q | X_Q, X_{r,Q}) = 0$, which follows from $Q - (X_Q, X_{r,Q}) - O_Q$, and $I(X_Q; O_Q | X_{r,Q}) = I(X_Q; O_Q | X_{r,Q}, Q) + I(Q; O_Q | X_{r,Q}, Q)$, hence

$$R \leq \max_{p(x, x_r)} I(X; Y, Y_r | X_r) + \epsilon_n. \quad (8.4)$$

Assembly. The *pre-maximisation* versions of the two per-letter bounds hold simultaneously for the *same* input law, namely the time-averaged distribution $p(x_Q, x_{r,Q}) = \frac{1}{n} \sum_i p(x_i, x_{r,i})$ induced by the code: from the chains above, $R \leq I(X_Q, X_{r,Q}; Y_Q) + \epsilon_n$ and $R \leq I(X_Q; O_Q | X_{r,Q}) + \epsilon_n$. (Separately maximised cut inequalities would not imply the max-of-min bound; it is the common law that does.) Therefore

$$\begin{aligned} R &\leq \min \left\{ I(X_Q, X_{r,Q}; Y_Q), I(X_Q; O_Q | X_{r,Q}) \right\} + \epsilon_n \\ &\leq \max_{p(x, x_r)} \min \left\{ I(X, X_r; Y), I(X; Y, Y_r | X_r) \right\} + \epsilon_n. \end{aligned}$$

Letting $n \rightarrow \infty$ (so $\epsilon_n \rightarrow 0$) completes the proof. $\blacksquare$

Remark on the two cuts. Each cut is proved by a causal-history chain rule: the destination cut keeps the receiver output Y_i and conditions on the past history (Y^{i-1}, Y_r^{i-1}) together with the message, while the source cut keeps the output pair $O_i = (Y_i, Y_{r,i})$ and conditions on the relay's causal input $X_{r,i}$ and the source input X_i . Both chains use only past observations and current inputs, which is essential because future relay inputs can depend on past outputs (a data-processing chain through the entire X_r^n is not valid). The outer bound keeps the tightest of the two per-letter bounds, which is why the min appears.

Remark 8.1 (Min-max evaluation of the cut-set bound). Write $A(p) = I(X, X_r; Y)$ and $B(p) = I(X; Y, Y_r | X_r)$, both as functions of the input distribution $p = p(x, x_r)$. For a fixed channel each of A and B is *concave* in p (mutual information is concave in the input distribution), and so is their pointwise minimum: the hypograph of $\min\{A, B\}$ is the intersection of the hypographs of the two concave functions, hence convex. Consequently the cut-set computation is a concave program as written:

$$\max_{p(x, x_r)} \min\{A(p), B(p)\} = \max \left\{ t : t \leq A(p), t \leq B(p) \text{ for some } p \right\},$$

and the outer-bound region $\mathcal{R} = \{R : R \leq \min\{A(p), B(p)\} \text{ for some } p\}$ is convex (it is the lower set of the concave function $\max_p \min\{A, B\}$). The finite-vector min-max corollary of the Appendix (Theorem 9.15 and Corollary 9.16) supplies an equivalent dual-weight representation: since $\min_{\lambda \in [0,1]} [\lambda A + (1-\lambda)B] = \min\{A, B\}$ and the lower set $\{(a, b) : a \leq A(p), b \leq B(p) \text{ for some } p\}$ is convex by simultaneous concavity (mixtures of witnessing inputs witness the midpoint componentwise), the exchange

$$\max_p \min_{\lambda \in [0,1]} [\lambda A(p) + (1-\lambda)B(p)] = \min_{\lambda \in [0,1]} \max_p [\lambda A(p) + (1-\lambda)B(p)]$$

is justified, the right-hand side being a family of concave programs in p (one per λ). The pointwise minimum of *non-concave* cut functionals would not be concave; in such settings the lower-set convexity must be supplied by an admissible time-sharing or concave-envelope construction (the role §1.13 assigns to that machinery), and sub-additivity alone does not establish it. $\circ$

8.3 Decode-and-forward

The relay can help in two essentially different ways. In *decode-forward* (DF) the relay *fully decodes* the message from its own observation Y_r^n and then re-encodes and forwards it. For this to help, the relay must be able to decode at the rate being sent, which requires $R < I(X; Y_r | X_r)$; and the destination, seeing both X^n and the relay's re-transmission X_r^n , can then decode at the cooperative rate $R < I(X, X_r; Y)$. The scheme is a *block Markov* code: the relay decodes the message of block $b - 1$ and cooperates in block b .

Theorem 8.4. (Decode-forward lower bound) The capacity of the discrete memoryless relay channel is lower bounded by

$$C \geq \max_{p(x, x_r)} \min \{ I(X; Y_r | X_r), I(X, X_r; Y) \}.$$

○

Proof. Fix $p(x, x_r)$, i.e. $p(x_r)p(x|x_r)$. We send $B - 1$ messages over B blocks, each block of length n . The rate of the scheme, once we let $B \rightarrow \infty$, tends to R .

Codebook generation. Generate 2^{nR} sequences $x_r^n(m)$, $m \in [1 : 2^{nR}]$, each i.i.d. according to $\prod_i p(x_{r,i})$. For each m , generate 2^{nR} sequences $x^n(m'|m)$, $m' \in [1 : 2^{nR}]$, each i.i.d. according to $\prod_i p(x_i | x_{r,i} = x_{r,i}(m))$. Reveal the codebook to the sender, the relay, and the destination.

Encoding. Let $m_0 = m_B = 1$ be dummy messages and let $m_1, \dots, m_{B-1}$ be the messages to send. In block $b \in [1 : B]$,

- the source transmits $x^n(m_b | m_{b-1})$;
- the relay transmits $x_r^n(m_{b-1})$.

Thus in block b the relay forwards the *previous* message m_{b-1} (which it has already decoded from block $b - 1$), while the source sends the *fresh* message m_b “on top of” the relay's codeword.

Relay decoding. At the end of block $b \in [1 : B - 1]$ the relay has received $y_r^n(b)$ and knows m_{b-1} (hence its own codeword $x_r^n(m_{b-1})$). It declares $\hat{m}_b$ to be the unique index such that

$$(x^n(\hat{m}_b | m_{b-1}), x_r^n(m_{b-1}), y_r^n(b)) \in \mathcal{T}_\epsilon^{(n)}(X, X_r, Y_r).$$

Destination decoding (backward). After all B blocks are received, the destination decodes backwards from block B to block 1. At block b , knowing m_b (already decoded from block $b + 1$; and $m_B = 1$), it declares $\hat{m}_{b-1}$ to be the unique index such that

$$(x^n(m_b | \hat{m}_{b-1}), x_r^n(\hat{m}_{b-1}), y^n(b)) \in \mathcal{T}_\epsilon^{(n)}(X, X_r, Y).$$

Analysis of the probability of error. *Relay.* By the law of large numbers (Proposition 2.6), the true triple $(x^n(m_b|m_{b-1}), x_r^n(m_{b-1}), y_r^n(b))$ is jointly ϵ -typical with probability tending to 1. For a wrong index $\hat{m}_b \neq m_b$, condition on the known relay cloud $x_r^n(m_{b-1})$: the wrong satellite $x^n(\hat{m}_b|m_{b-1})$ was generated conditionally on that cloud and is therefore *conditionally* independent of $y_r^n(b)$ given $x_r^n(m_{b-1})$ (the output $y_r^n(b)$ was produced through the channel from the true satellite and the same cloud). By the conditional joint-typicality bound (Proposition 2.7), the probability that such a wrong satellite is jointly typical with $(x_r^n(m_{b-1}), y_r^n(b))$ is at most

$$2^{-n(I(X;Y_r|X_r)-\delta(\epsilon))},$$

where $\delta(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$. (The packing statement is conditional on the cloud: a wrong satellite is independent of Y_r^n only given X_r^n , and the exponent is $I(X; Y_r|X_r)$, not $I(X, X_r; Y_r)$.) A union bound over the $2^{nR} - 1$ wrong indices gives

$$\Pr\{\text{relay error at block } b\} \leq 2^{nR} 2^{-n(I(X;Y_r|X_r)-\delta(\epsilon))} \rightarrow 0$$

provided $R < I(X; Y_r|X_r) - \delta(\epsilon)$.

Destination. Similarly, the true triple $(x^n(m_b|m_{b-1}), x_r^n(m_{b-1}), y^n(b))$ is jointly typical with probability tending to 1 (Proposition 2.6). For a wrong index $\hat{m}_{b-1} \neq m_{b-1}$, the corresponding pair of codewords is independent of $y^n(b)$, distributed as $\prod_i p(x_i, x_{r,i})$, while $y^n(b)$ has marginal $\prod_i p(y_i)$. By Proposition 2.7, a wrong pair is jointly typical with $y^n(b)$ with probability at most

$$2^{-n(I(X, X_r; Y)-\delta(\epsilon))},$$

so a union bound over the $2^{nR} - 1$ wrong indices gives

$$\Pr\{\text{destination error at block } b\} \leq 2^{nR} 2^{-n(I(X, X_r; Y)-\delta(\epsilon))} \rightarrow 0$$

provided $R < I(X, X_r; Y) - \delta(\epsilon)$.

By the union bound over the B blocks (with B fixed) and over the two decoding stages, the overall probability of error tends to 0 as $n \rightarrow \infty$ whenever

$$R < \min \{ I(X; Y_r|X_r), I(X, X_r; Y) \} - \delta(\epsilon).$$

Letting $\epsilon \rightarrow 0$ and then $B \rightarrow \infty$ (so that $\frac{B-1}{B}R \rightarrow R$), the rate

$$R^* = \min \{ I(X; Y_r|X_r), I(X, X_r; Y) \}$$

is achievable. Maximizing over $p(x, x_r)$ completes the proof. ■

Remark 8.2 (Why backward decoding?). Backward decoding is used only for convenience. In block b the source's codeword $x^n(m_b|m_{b-1})$ carries *two* messages, the fresh m_b and the relay's message m_{b-1} ; decoding m_{b-1} at the destination would, with forward decoding, require knowledge of m_b . Decoding backwards avoids this by exploiting that m_b is already known when block b is processed. Equivalently, one may use sliding-window or simultaneous (joint typicality) decoding; all variants achieve the same rate $\min\{I(X; Y_r|X_r), I(X, X_r; Y)\}$. ○

Optimality of decode–forward for the degraded relay. There is a class of channels for which DF is exactly optimal and meets the cut-set bound.

Definition 8.5. (Degraded relay channel) The relay channel $p(y, y_r|x, x_r)$ is *degraded* if $p(y, y_r|x, x_r) = p(y_r|x, x_r) p(y|x_r, y_r)$, i.e. if

$$X \rightarrow (X_r, Y_r) \rightarrow Y$$

forms a Markov chain. In words: the destination’s output is a physically degraded version of the relay’s output. $\circ$

Corollary 8.6. (Capacity of the degraded relay channel) If the relay channel is degraded, then

$$C = \max_{p(x, x_r)} \min \{ I(X, X_r; Y), I(X; Y_r|X_r) \},$$

achieved by decode–forward. $\circ$

Proof. Under the Markov chain $X \rightarrow (X_r, Y_r) \rightarrow Y$ we have $I(X; Y|X_r, Y_r) = 0$, hence

$$I(X; Y, Y_r|X_r) = I(X; Y_r|X_r) + I(X; Y|X_r, Y_r) = I(X; Y_r|X_r).$$

Therefore the source-cut term of the cut-set bound (Theorem 8.3) coincides with the relay-decoding term of the DF bound (Theorem 8.4), while the destination-cut term $I(X, X_r; Y)$ is common to both. The two bounds agree, giving the stated capacity. $\blacksquare$

8.4 Compress-and-forward

When the relay’s observation is too noisy for it to decode the message ($I(X; Y_r|X_r)$ is too small), the relay can still help by *forwarding a compressed description* of what it heard, rather than the message itself. In *compress–forward* (CF) the relay quantizes its received block Y_r^n into a reconstruction $\hat{Y}_r^n$ and sends the *bin index* of the quantization to the destination; the destination uses its own observation Y^n as side information to recover $\hat{Y}_r^n$, and then decodes the message from the pair $(Y^n, \hat{Y}_r^n)$. The compression step is exactly a *Wyner–Ziv* problem with side information Y^n at the decoder.

Theorem 8.7. (Cover–El Gamal compress–forward lower bound) The capacity of the discrete memoryless relay channel is lower bounded by

$$C \geq \max_{p(x)p(x_r)p(\hat{y}_r|x_r, y_r)} \min \left\{ I(X, X_r; Y) - I(\hat{Y}_r; Y_r|X, X_r, Y), I(X; Y, \hat{Y}_r|X_r) \right\},$$

where the maximum is over all distributions of the form $p(x)p(x_r)p(\hat{y}_r|x_r, y_r)$ and $\hat{Y}_r$ takes values in some finite alphabet $\hat{\mathcal{Y}}_r$. This min-form is the Cover–El Gamal compress–forward bound (1979); we present the *sequential* CF scheme below and cite the equivalence between the sequential form and the displayed optimized form (Lim–Kim–El Gamal–Chung [Lim et al. \(2011\)](#)). $\circ$

Proof sketch: sequential CF. Fix $p(x)$, $p(x_r)$, and a test channel $p(\hat{y}_r|x_r, y_r)$ (so that $\hat{Y}_r - (Y_r, X_r) - (X, Y)$). Let R_0 be the rate of the relay's description (the bin index) and $\hat{R}$ the rate of the compression codebook.

Codebook generation.

- Generate 2^{nR} sequences $x^n(m)$, $m \in [1 : 2^{nR}]$, i.i.d. according to $\prod_i p(x_i)$.
- Generate 2^{nR_0} sequences $x_r^n(l)$, $l \in [1 : 2^{nR_0}]$, i.i.d. according to $\prod_i p(x_{r,i})$.
- For each l , generate $2^{n\hat{R}}$ sequences $\hat{y}_r^n(k|l)$, $k \in [1 : 2^{n\hat{R}}]$, i.i.d. according to $\prod_i p(\hat{y}_{r,i}|x_{r,i} = x_{r,i}(l))$. Partition the index set $[1 : 2^{n\hat{R}}]$ into 2^{nR_0} bins of equal size $\mathcal{B}(l')$, $l' \in [1 : 2^{nR_0}]$.

Encoding (block Markov, explicit indices). The scheme runs over B blocks plus a final flushing block. Use independent block-specific codebooks, with the following index bookkeeping: l_{b-1} is the relay-input index (a bin index) used in block b ; k_b is the compression index selected *after* block b , and l_b is the bin containing k_b , transmitted by the relay in block $b + 1$. In block $b \in [1 : B]$ the source transmits $x_b^n(m_b)$ and the relay transmits $x_{r,b}^n(l_{b-1})$, with l_0 fixed and dummy messages at the boundaries. After receiving $y_r^n(b)$, the relay selects a k_b with $(\hat{y}_{r,b}^n(k_b|l_{b-1}), x_{r,b}^n(l_{b-1}), y_r^n(b))$ jointly typical (covering step; note the compression sequence for block b is generated conditionally on the cloud $x_{r,b}^n(l_{b-1})$ actually used in that block, not on the next block's input index).

Decoding (delayed sequential). The bin l_b is transmitted in block $b + 1$, so the destination decodes it only at the end of block $b + 1$; the decoder therefore runs one block behind the encoder.

1. *Bin index (one block late).* At the end of block $b + 1$ the destination declares $\hat{l}_b$ the unique index with $(x_{r,b+1}^n(\hat{l}_b), y^n(b + 1))$ jointly typical (the codeword indexed by l_b produced the observation $y^n(b + 1)$, since the relay transmitted $x_{r,b+1}^n(l_b)$ in that block).
2. *Compression (retrieval).* It retrieves the saved observation $y^n(b)$ and the known previous cloud $x_{r,b}^n(l_{b-1})$, and finds the unique $k_b \in \mathcal{B}_b(l_b)$ such that $(\hat{y}_{r,b}^n(k_b|l_{b-1}), x_{r,b}^n(l_{b-1}), y^n(b))$ is jointly typical, using the block- b compression codebook generated for that cloud.
3. *Message.* It declares $\hat{m}_b$ the unique index with

$$(x_b^n(\hat{m}_b), x_{r,b}^n(l_{b-1}), \hat{y}_{r,b}^n(k_b|l_{b-1}), y^n(b)) \in \mathcal{T}_\epsilon^{(n)}(X, X_r, \hat{Y}_r, Y).$$

A final flushing block conveys the last compression bin l_B without a new message; the rate loss from the flushing block vanishes as $B \rightarrow \infty$.

Rate constraints. The standard joint-typicality analysis (covering, conditional packing, and the Markov lemma $\hat{Y}_r - (Y_r, X_r) - (X, Y)$) gives the following four constraints.

1. *Covering.* A jointly typical $\hat{y}_r^n$ exists provided $\hat{R} > I(\hat{Y}_r; Y_r | X_r)$.
2. *Compression (Wyner–Ziv binning).* Conveying the quantisation $\hat{Y}_r^n$ to a decoder that has side information (Y^n, X_r^n) requires a bin rate

$$R_0 > I(\hat{Y}_r; Y_r | X_r, Y) = I(\hat{Y}_r; Y_r | X_r) - I(\hat{Y}_r; Y | X_r),$$

where the equality is the correct chain-rule identity under $\hat{Y}_r - (Y_r, X_r) - (X, Y)$ (not $I(\hat{Y}_r; Y_r | X, X_r, Y)$, which drops the term $I(\hat{Y}_r; X | X_r, Y)$). This is the Wyner–Ziv step, invoked rather than re-proved.

3. *Relay–destination link.* The bin index l_b is decoded at rate

$$R_0 < I(X_r; Y).$$

(This is an achievability constraint from the packing bound, not a MAC converse; it states that the decoder can resolve the relay’s description.)

4. *Message decoding given the compression.* Once $\hat{Y}_r^n$ is recovered, the message is decoded from the effective channel $X \rightarrow (Y, \hat{Y}_r)$ given the known cloud X_r , so

$$R < I(X; Y, \hat{Y}_r | X_r).$$

The sequential CF region is therefore

$$R < I(X; Y, \hat{Y}_r | X_r), \quad I(\hat{Y}_r; Y_r | X_r, Y) < I(X_r; Y),$$

the second inequality being the feasibility condition obtained by eliminating R_0 between constraints 2 and 3. The min-form displayed in the theorem is the *optimized* version of this region: it is not a fixed-distribution equality, but the cited equivalence (Lim–Kim–El Gamal–Chung [Lim et al. \(2011\)](#)) that optimising the sequential region over $p(x)p(x_r)p(\hat{y}_r|x_r, y_r)$ produces the displayed minimum. We cite that equivalence rather than re-derive the optimization. ■

Remark 8.3 (Honest status of the CF proof). The sequential scheme above is the classical Cover–El Gamal compress–forward construction: the new ingredient is Wyner–Ziv binning (constraint 2), quoted rather than derived, and the block-Markov timing with the index bookkeeping l_{b-1}, k_b, l_b is essential to strict causality. The passage from the sequential rate region to the theorem’s min-form is a genuine optimization step and is cited to [Lim et al. \(2011\)](#); it is not a chain-rule rearrangement of the fixed-distribution rates. Covering, packing, and the Markov lemma are standard joint-typicality estimates. ○

8.5 Remarks

Remark 8.4 (The relay channel without a relay). If the relay is ignored (the relay input is constant, or the relay is removed), the channel reduces to the point-to-point channel $p(y|x)$, whose capacity is $\max_{p(x)} I(X; Y)$. This *direct-transmission* rate is always achievable separately (it is the point-to-point coding theorem). Note the following careful reading of the bounds: with X_r constant, decode–forward (Theorem 8.4) gives $\min\{I(X; Y_r), I(X; Y)\}$, which can be *smaller* than $I(X; Y)$ when the relay observes a worse channel than the destination, so DF with a constant relay does not by itself recover direct transmission. Direct transmission is nevertheless achievable, and it is recovered by compress–forward with a constant $\hat{Y}_r$ (the compression carries no information and the second term reduces to $I(X; Y|X_r) = I(X; Y)$). The relay is useful only when it can strictly improve on $\max_{p(x)} I(X; Y)$. $\circ$

Remark 8.5 (Capacity is open in general). Unlike the point-to-point, multiple-access, and degraded-broadcast channels, the capacity of the general relay channel is *not known*, and this remains one of the central open problems in network information theory. Even for the *Gaussian* relay channel $Y_r = X + Z_r$, $Y = X + X_r + Z$ (with independent Gaussian noises) the capacity is unknown; the cut-set bound and the DF/CF bounds do not coincide in general, and the question is known as Cover’s open problem for the Gaussian relay channel Cover (1987). Recent progress — including the fact that the cut-set bound is loose for Gaussian relay networks, and reverse-hypercontractivity based converses — is surveyed in El Gamal and Kim (2012), Ch. 16. Two generalisations of the cut-set bound deserve mention as optional reading. El Gamal, Gohari, and Nair (2022) develop a *strengthened cut-set* upper bound using mutual-information inequalities and Gallager-type identification of auxiliaries. It strictly improves previously known bounds for the Gaussian relay channel with nonzero channel gains; for relay channels with orthogonal receiver components, it resolves Kim’s conjecture for a class of deterministic relays. See also the auxiliary-receiver approach of Gohari and Nair (2022), discussed in §5.7.2 of the interference channel chapter. Both are instalments of the sub-additivity programme of §1.13. $\circ$

Remark 8.6 (The Cover–El Gamal converse and its successors). The cut-set bound (Theorem 8.3) treats the relay as if it could choose X_r^n with full (non-causal) knowledge of Y_r^n . The relay is in fact *strictly causal*: $X_{r,i}$ may depend only on Y_r^{i-1} . Three converses should be distinguished in the literature (the structured classes below are developed at book length in El Gamal and Kim (2012), Ch. 16). First, the original Cover–El Gamal converses (1979) establish the cut-set bound as a general upper bound and settle two structured classes exactly, with distinct formulas: for the *degraded* relay channel (the destination observes a degraded version of the relay’s output, Definition 8.5) the capacity is $\max_{p(x, x_r)} \min\{I(X, X_r; Y), I(X; Y_r|X_r)\}$, achieved by decode–forward; for the *reversely degraded* relay channel (the relay observes a degraded version of the destination’s output, $X \rightarrow (X_r, Y) \rightarrow Y_r$) the capacity is $\max_{x_r} \max_{p(x)} I(X; Y | X_r = x_r)$: the relay holds its input at the symbol giving the best direct channel and is otherwise useless. The two formulas must not be interchanged (with $Y = X$ noiseless and Y_r constant the channel is reversely degraded with capacity 1, while the DF formula evaluates to 0). Second, the *strengthened cut-set bound* of El Gamal, Gohari, and Nair (2022) exploits the causal structure more finely through mutual-information inequalities with explicit auxiliary identification, and

is strictly tighter than the cut-set bound for every Gaussian relay channel with non-zero channel gains. Third, the auxiliary-receiver (J -)cutset bound of [Gohari and Nair \(2022\)](#) (stated in §5.7.2 of the interference channel chapter) is again strictly tighter. We omit the statements; the point for this course is that the cut-set bound is *not* the last word on the relay channel. $\bigcirc$

Remark 8.7 (Where sub-additivity does and does not apply). This chapter is a study in contrast for the sub-additivity theme of §1.13. The cut-set bound is a chain-rule argument with the same per-letter shape as time sub-additivity: each cut keeps a mutual-information functional of the current inputs, and the two causal-history chains (8.1) and (8.3) bound nR by the sum of those per-letter functionals, exactly as in the point-to-point converse of Chapter 2 — with the essential difference that under a causal relay only past observations and current inputs may be conditioned on (conditioning on the entire X_r^n as if exogenous is not valid). The min-max structure $\max_p \min\{A(p), B(p)\}$ is a concave program, handled with the dual-weight representation of Corollary 9.16. By contrast, the *achievability* results of this chapter — decode-forward and compress-forward — are *not* chain-rule arguments: they are block Markov coding, superposition, and binning, all analyzed by joint typicality. Chain rules are the language of *converses*; typicality is the language of *achievability*. The relay channel is the cleanest place to see both languages side by side. $\bigcirc$

8.6 Exercises

Exercise 9: Prove directly that for a memoryless relay channel the destination cut alone gives $C \leq \max_{p(x, x_r)} I(X, X_r; Y)$, following the causal-history chain of (8.1) (condition only on past observations and current inputs, never on future relay inputs).

Exercise 10: Let the relay channel be *reversely degraded*, i.e. $X \rightarrow (X_r, Y) \rightarrow Y_r$ forms a Markov chain (the relay observes a degraded version of the destination's output). Show, following Cover–El Gamal (1979), that the capacity is

$$C = \max_{x_r} \max_{p(x)} I(X; Y \mid X_r = x_r),$$

i.e. the relay holds its input at the symbol giving the best direct channel and is otherwise useless. (Note: this is *not* the decode-forward formula of Definition 8.5 — that formula belongs to the degraded class in which the destination observes a degraded version of the relay's output. Decisive check: binary X with $Y = X$ noiseless and Y_r constant is reversely degraded, has capacity 1, while the DF formula evaluates to 0. For the reversely degraded Gaussian relay the formula reduces to the direct-transmission capacity.)

Exercise 11: Show that for the degraded relay channel of Definition 8.5, the cut-set bound and the DF bound coincide, and conclude Corollary 8.6 directly without appealing to the single-letter comparison in its proof.

Exercise 12: Consider the Gaussian relay channel $Y_r = X + Z_r$, $Y = X + X_r + Z$, with independent noises $Z_r \sim \mathcal{N}(0, N_r)$, $Z \sim \mathcal{N}(0, N)$ and average power constraints P on X and P_r on X_r . Write down the cut-set bound, the DF rate, and the CF rate as optimizations over the input powers and the quantization noise. (The Gaussian relay capacity itself is open.)

9 Appendix

9.1 Diffusion Equation, Heat Flow and Fisher Information

Given a box of gas with initial density $g(\mathbf{x})$, one might wonder what is the law of the density evolution along time $t \geq 0$. Let's denote the density of gas at location $\mathbf{x} \in \Omega$ and time $t \geq 0$ as $f(\mathbf{x}, t)$, where $\Omega \subseteq \mathbb{R}^d$ is the domain containing the gas (below we standardise to the whole space $\Omega = \mathbb{R}^d$), with initial condition $f(\mathbf{x}, 0) = g(\mathbf{x})$.

Consider a homogeneous isotropic medium with constant diffusivity $c > 0$. Fick's law of diffusion is a *flux law*: the flux is $j = -c\nabla f$, i.e. mass flows from regions of high density to regions of low density at a rate proportional to the density gradient. Combining this flux law with mass conservation gives the diffusion equation $\partial_t f(\mathbf{x}, t) = c\Delta f(\mathbf{x}, t)$. Here $\Delta = \sum_{i=1}^d \partial_{x_i}^2$ is the Laplace operator, and one can interpret it as a measure of difference in density between $\mathbf{x}$ and its neighborhood.

$$\partial_t f(\mathbf{x}, t) = c\Delta f(\mathbf{x}, t) \quad (9.1)$$

Here the constant c measures the diffusivity of the medium. Moreover, it turns out that the evolution of the temperature of a homogeneous isotropic medium also follows the same PDE. Hence it's also called the heat equation. Formally, a heat equation is composed of three components.

1. (initial condition) the initial density $g(\mathbf{x})$ for $\mathbf{x} \in \Omega$;
2. (boundary condition) the rule on the boundary of the domain, e.g., fixed boundary values, $f(\mathbf{x}, t)|_{\partial\Omega} = u(\mathbf{x}), \forall t \geq 0$;
3. (differential equation) the PDE $\partial_t f(\mathbf{x}, t) = c\Delta f(\mathbf{x}, t)$ for the interior.

We are especially interested in the heat equation over the whole space $\mathbb{R}^d$ and therefore we don't need to specify boundary conditions. In this probabilistic heat flow we consider nonnegative solutions of mass one, with the initial condition understood as the appropriate limit: $f(\cdot, t)$ converges weakly to the initial law as $t \downarrow 0$. Merely checking the PDE does not verify the initial condition or unrestricted uniqueness, and the initial Dirac mass below is a distribution (measure), not an ordinary density. Then we have the following PDE.

$$\begin{cases} \partial_t f(\mathbf{x}, t) = c\Delta f(\mathbf{x}, t), & \forall \mathbf{x} \in \Omega, t > 0; \\ f_t(\mathbf{x}) d\mathbf{x} \Rightarrow P_G, & t \downarrow 0, \end{cases} \quad (9.2)$$

where P_G is the initial probability law and $\Rightarrow$ denotes weak convergence of measures. For a sufficiently regular initial density one may separately state the stronger classical initial condition $f(\mathbf{x}, 0) = g(\mathbf{x})$; the measure-valued formulation above also covers singular initial laws such as the Dirac mass below.

This PDE can be solved by first considering the special case where the initial condition is the Dirac measure $g = \delta_0$ (an impulse at $\mathbf{x} = 0$; it is a distribution, not an ordinary density), and denote its solution as $\Phi(\mathbf{x}, t)$. To solve the special case, we will need tools from Fourier analysis, hence the proof is omitted. It turns out that $\Phi(\mathbf{x}, t)$ follows the Gaussian distribution.

$$\Phi(\mathbf{x}, t) = \frac{1}{(4\pi ct)^{d/2}} \exp\left(-\frac{\mathbf{x}^T \mathbf{x}}{4ct}\right)$$

However, to get a general solution with arbitrary initial condition $g(\mathbf{x})$, it suffices to notice that due to the property of convolution, we have

$$(\partial_t - c\Delta)(\Phi(\mathbf{x}, t) * g(\mathbf{x})) = [(\partial_t - c\Delta)\Phi(\mathbf{x}, t)] * g(\mathbf{x}) = 0.$$

Hence $g(\mathbf{x}) * \Phi(\mathbf{x}, t)$ is the solution to Equation (9.2). In other words, if we have $\mathbf{G} \sim g(\mathbf{x})$ and $\mathbf{Z} \sim \mathcal{N}(0, I)$, with $\mathbf{G}$ and $\mathbf{Z}$ independent, then we have $\sqrt{2ct}\mathbf{Z} \sim \Phi(\mathbf{x}, t)$ and $\mathbf{G} + \sqrt{2ct}\mathbf{Z} \sim g(\mathbf{x}) * \Phi(\mathbf{x}, t)$ (with the initial condition understood as the weak limit as $t \downarrow 0$). Hence the heat flow is just the original distribution with additive Gaussian noise that increases over time.

From now on, we consider $c = \frac{1}{2}$ and the heat flow $\mathbf{G}_t \triangleq \mathbf{G} + \sqrt{t}\mathbf{Z}$. By the heat equation, we know that $\partial_t f(\mathbf{x}, t) = \frac{1}{2}\Delta f(\mathbf{x}, t)$. We would like to show the following connection between the Fisher information $J(\mathbf{G} + \sqrt{t}\mathbf{Z})$ and the change rate of entropy $\partial_t h(\mathbf{G} + \sqrt{t}\mathbf{Z})$.

Lemma 9.1. (de Bruijn's identity) Let $\mathbf{G}$ be a random vector in $\mathbb{R}^d$ with finite second moment, and let $\mathbf{Z} \sim \mathcal{N}(0, I_d)$ be standard normal, independent of $\mathbf{G}$. Write $\mathbf{G}_t = \mathbf{G} + \sqrt{t}\mathbf{Z}$. With entropy in nats and the standard natural-log Fisher information, for every $t > 0$,

$$\frac{d}{dt}h(\mathbf{G}_t) = \frac{1}{2}J(\mathbf{G}_t). \quad (9.3)$$

For a classical right derivative at $t = 0$, assume additionally that $\mathbf{G}$ has a density, finite entropy and Fisher information, and regularity ensuring the endpoint de Bruijn identity. Under those additional hypotheses,

$$\left.\frac{d}{dt}h(\mathbf{G}_t)\right|_{0+} = \frac{1}{2}J(\mathbf{G}).$$

○

Remark 9.1 (Convention). Throughout this section, entropy is measured in nats and the score is defined using $\ln$, which is what makes the constant $\frac{1}{2}$ in de Bruijn's identity exact; this matches the convention of the Fisher-information section of Chapter 1, where J is base-independent. The vector form of the Fisher information is $J(\mathbf{X}) = \mathbb{E}[\|\nabla \ln f_{\mathbf{X}}(\mathbf{X})\|^2]$; the scalar computation below is the special case of a one-dimensional $\mathbf{X}$. ○

Before proving it, we remind the readers that the Fisher information of a family of distributions $\mathcal{F} \triangleq \{f(\mathbf{x}; \theta), \theta \in \Omega\}$, parametrized by a vector $\theta \in \Omega \subseteq \mathbb{R}^k$, is the Fisher information *matrix*

$$\mathcal{I}(\theta) = \mathbb{E}_\theta [s_\theta s_\theta^T], \quad s_\theta = \nabla_\theta \ln f(\mathbf{X}; \theta).$$

Under the usual Cramér–Rao regularity assumptions and nonsingular $\mathcal{I}(\theta)$, any unbiased estimator $\widehat{\theta}$ satisfies the matrix bound $\text{Cov}_\theta(\widehat{\theta}) \succeq \mathcal{I}(\theta)^{-1}$. The scalar quantity $\mathbb{E} \|s_\theta\|^2$ is the *trace* of $\mathcal{I}(\theta)$, not the matrix itself, and the reciprocal of the trace is *not* the general vector covariance bound. For a scalar parameter this reduces to $\text{Var}(\widehat{\theta}) \geq 1/\mathcal{I}(\theta)$ (an *at least* inequality, not a strict one). Given any random scalar X , we will consider the translation group associated, i.e., $\mathcal{F} \triangleq \{f_X(x - \theta), \theta \in \mathbb{R}\}$. Then the Fisher information is just

$$\begin{aligned} J(\theta) &= \int_{\mathbb{R}} f_X(x - \theta) (\partial_\theta \ln f_X(x - \theta))^2 dx \\ &= \int_{\mathbb{R}} f_X(x - \theta) (\partial_x \ln f_X(x - \theta))^2 dx \\ &= \int_{\mathbb{R}} f_X(x) (\partial_x \ln f_X(x))^2 dx \end{aligned}$$

Note this quantity is independent of the translation parameter θ . We will abuse the notation a little to denote $J(X) \triangleq \mathbb{E}_X (\partial_x \ln f_X(x))^2$. And by Fisher information of a r.v. X we mean the fisher information of the translation group associated with X .

Proof of Lemma 9.1. We prove the scalar case of Equation (9.3). The vector case can be proved using similar idea. We can calculate $\partial_t h(G_t)$ as the following.

$$\begin{aligned} \partial_t h(G_t) &= -\partial_t \int_{x \in \mathbb{R}} f_{G_t}(x) \ln f_{G_t}(x) dx \\ &= -\int_{x \in \mathbb{R}} \partial_t f_{G_t}(x) \ln f_{G_t}(x) dx - \int_{x \in \mathbb{R}} f_{G_t}(x) \partial_t (\ln f_{G_t}(x)) dx \\ &= -\frac{1}{2} \int_{x \in \mathbb{R}} \partial_x^2 f_{G_t}(x) \ln f_{G_t}(x) dx - \int_{x \in \mathbb{R}} \partial_t f_{G_t}(x) dx \\ &= -\frac{1}{2} \int_{x \in \mathbb{R}} \ln f_{G_t}(x) d(\partial_x f_{G_t}(x)) - \partial_t \left(\int_{x \in \mathbb{R}} f_{G_t}(x) dx \right) \\ &= -\left(\frac{1}{2} \partial_x f_{G_t}(x) \ln f_{G_t}(x) \right) \Big|_{-\infty}^{+\infty} + \frac{1}{2} \int_{x \in \mathbb{R}} \partial_x f_{G_t}(x) d(\ln f_{G_t}(x)) \\ &= \frac{1}{2} \int_{x \in \mathbb{R}} f_{G_t}(x) (\partial_x \ln f_{G_t}(x))^2 dx \\ &= \frac{1}{2} J(G_t) \end{aligned}$$

This proves the lemma for $t > 0$. The formal steps (differentiating under the integral and the boundary term $\partial_x f_{G_t} \ln f_{G_t}|_{-\infty}^{+\infty}$) are justified by a standard cutoff argument using moment/score estimates, since for $t > 0$ the density f_{G_t} is the convolution of the law of G with a positive Gaussian kernel and hence positive and smooth (its tails need not decay faster than polynomially). The $t = 0$ conclusion requires the stated separate endpoint justification. ■

Both heat flow and Fisher information will play a role in the proof of EPI.

9.2 Sanov's inequality in discrete spaces: Types and Typical Sequences

Let Σ be a finite space. For the purposes of these notes, let $\Sigma = \{1, 2, \dots, M\}$, or equivalently $\Sigma = [1 : M]$. All logarithms in this section are base 2, so entropies are in bits and divergences carry the factor $2^{-nD(\cdot\|\cdot)}$ in exponentiated bounds. The extended conventions used throughout are $\inf \emptyset = +\infty$ and $\log 0 = -\infty$ (so that $2^{-\infty} = 0$), and divergences take the value $+\infty$ when the first argument is not absolutely continuous with respect to the second. Denote $\mathcal{M}(\Sigma)$ as the set of probability mass functions on Σ . For a given probability mass function μ , let us denote $\Sigma_\mu = \{i : \mu(i) > 0\}$ to be the support of μ . Thus $\Sigma_\mu \subseteq \Sigma$.

9.2.1 Types

Given a sequence $y^n \in \Sigma^n$, we define the **type** of y^n , $T_{y^n} \in \mathcal{M}(\Sigma)$, as the probability mass function given by

$$T_{y^n}(i) := \frac{1}{n} \sum_{k=1}^n 1_{\{y_k=i\}}, \quad 1 \leq i \leq M.$$

It is, equivalently, the empirical measure induced by the sequence y^n . Let $T_n \subset \mathcal{M}(\Sigma)$ denote the collection of all types, i.e.

$$T_n := \{\mu : \mu = T_{y^n} \text{ for some } y^n \in \Sigma^n\}.$$

Lemma 9.2. The following statements hold:

- i $|T_n| = \binom{n+M-1}{M-1} \leq (n+1)^{M-1}$.
- ii For any $\mu \in \mathcal{M}(\Sigma)$, there exists $\nu \in T_n$ such that $|\mu(i) - \nu(i)| \leq \frac{1}{n}$. Consequently, $d_{TV}(\mu, T_n) \leq \frac{M}{2n}$, where $d_{TV}(\mu, T_n) := \min_{\nu \in T_n} d_{TV}(\mu, \nu)$ and $d_{TV}(\mu, \nu) = \frac{1}{2} \sum_{i=1}^M |\mu(i) - \nu(i)|$. Further $\Sigma_\nu \subseteq \Sigma_\mu$.

○

Proof. Note that every $\mu \in T_n$ is in one-to-one correspondence with non-negative integer sequences $\{a_1, \dots, a_M\}$ such that $\sum_{i=1}^M a_i = n$. The count of the latter is a problem in elementary combinatorics, and the count is essentially a bijection to choosing the identities of $M-1$ dividers from $n+M-1$ locations. Note that $\binom{n+M-1}{M-1} \leq (n+1)^{M-1}$ is an equality for $M=1$ and for $M>1$, we have

$$\binom{n+M-1}{M-1} = \prod_{k=1}^{M-1} \frac{n+k}{k} \leq \prod_{k=1}^{M-1} (n+1) = (n+1)^{M-1},$$

as $\frac{n+k}{k} \leq n+1, \forall k \geq 1$.

Given a $\mu \in \mathcal{M}(\Sigma)$, let us define two non-negative integer sequences according to

$$k_l(i) = \lfloor n\mu(i) \rfloor, \quad k_u(i) = \lceil n\mu(i) \rceil, \quad 1 \leq i \leq M.$$

The following estimates are clear:

$$n\mu(i) - 1 \leq k_l(i) \leq n\mu(i) \leq k_u(i) \leq n\mu(i) + 1.$$

Summing up over i , we obtain

$$n - M \leq \sum_{i=1}^M k_l(i) \leq n \leq \sum_{i=1}^M k_u(i) \leq n + M.$$

Therefore, we can find a sequence of non-negative integers, $k_{int}(i)$ such that $k_l(i) \leq k_{int}(i) \leq k_u(i)$ such that $\sum_{i=1}^M k_{int}(i) = n$. This is essentially like a discrete intermediate-value-theorem, which a greedy algorithm (starting from k_l and increasing value at each coordinate by one while obeying the bounds) can easily establish. Now define $\nu(i) = \frac{k_{int}(i)}{n}$. Note that $\nu \in \mathcal{T}_n$. We know that

$$n\mu(i) - 1 \leq k_l(i) \leq k_{int}(i) \leq k_u(i) \leq n\mu(i) + 1.$$

Therefore $|\nu(i) - \mu(i)| \leq \frac{1}{n}$ and $d_{TV}(\mu, \nu) \leq \frac{M}{2n}$. If $\mu(i) = 0$, then observe that $k_u(i) = 0$ implying $k_{int}(i) = 0$ and hence $\nu(i) = 0$. This establishes the relationship between the supports. ■

For $\nu \in \mathcal{T}_n$, we define the **type-class** by $\mathcal{Y}_n(\nu) = \{y^n \in \Sigma^n : \mathcal{T}_{y^n} = \nu\}$. Note that $\mathcal{Y}_n(\nu)$ is the collection of permutations of a generic string y^n whose empirical measure is ν , and the cardinality of $\mathcal{Y}_n(\nu)$ is the multinomial co-efficient associated with the empirical counts, i.e. $|\mathcal{Y}_n(\nu)| = \binom{n}{n\nu(1), n\nu(2), \dots, n\nu(M)}$.

Let $\mathbb{P}_\mu$ be the probability law associated with an infinite sequence of i.i.d. random variables $Y_1, Y_2, \dots$, distributed according to $\mu \in \mathcal{M}(\Sigma)$.

In the following, for $\nu, \mu \in \mathcal{M}(\Sigma)$ with $\Sigma = [1 : M]$:

$$H(\nu) = \sum_{i=1}^M -\nu(i) \log \nu(i),$$

$$D(\nu \parallel \mu) = \sum_{i=1}^M \nu(i) \log \frac{\nu(i)}{\mu(i)},$$

with the convention: $0 \log 0 = 0$ and if $\nu \not\ll \mu$, then $D(\nu \parallel \mu) = \infty$.

Lemma 9.3.

$$\mathbb{P}_\mu[(Y_1, Y_2, \dots, Y_n) = y^n] = 2^{-n(H(\mathcal{T}_{y^n}) + D(\mathcal{T}_{y^n} \parallel \mu))}.$$

○

Proof. Note that

$$\begin{aligned}\mathbb{P}_\mu[(Y_1, Y_2, \dots, Y_n) = y^n] &= \prod_{i=1}^M (\mu(i))^{n\tau_{y^n}(i)} \\ &= 2^{-n \left(-\sum_{i=1}^M \tau_{y^n}(i) \log_2 \tau_{y^n}(i) + \sum_{i=1}^M \tau_{y^n}(i) \log_2 \frac{\tau_{y^n}(i)}{\mu(i)} \right)} \\ &= 2^{-n(H(\tau_{y^n}) + D(\tau_{y^n} \parallel \mu))}.\end{aligned}$$

■

Lemma 9.4. Let $m \in \{0, 1, 2, \dots\}$ and $l \in \{1, 2, \dots\}$. Then $\frac{m!}{l!} \geq l^{m-l}$.

○

Proof. If $m > l$, then $\frac{m!}{l!} = \prod_{k=l+1}^m k \geq l^{m-l}$. If $m < l$, then $\frac{m!}{l!} = \prod_{k=m+1}^l \frac{1}{k} \geq \frac{1}{l}^{l-m} = l^{m-l}$. Finally, equality holds for $m = l$ (and for $m = 0$ the claim is $1/l! \geq l^{-l}$, which follows from the $m < l$ case).

■

Corollary 9.5. For $\gamma, \nu \in \mathbb{T}_n$,

$$\frac{|\mathcal{Y}_n(\nu)|}{|\mathcal{Y}_n(\gamma)|} \geq 2^{n(H(\nu) - D(\gamma \parallel \nu) - H(\gamma))}.$$

○

Proof. If $\gamma \not\ll \nu$, then $D(\gamma \parallel \nu) = +\infty$ and the right-hand side is $2^{-\infty} = 0$, so the inequality is immediate. Assume $\gamma \ll \nu$ and set $S = \text{supp } \nu$. Outside S both counts vanish and the factorial ratios are $0!/0! = 1$. Therefore

$$\frac{|\mathcal{Y}_n(\nu)|}{|\mathcal{Y}_n(\gamma)|} = \prod_{i \in S} \frac{(n\gamma(i))!}{(n\nu(i))!} \geq \prod_{i \in S} (n\nu(i))^{n(\gamma(i) - \nu(i))} = \prod_{i \in S} (\nu(i))^{n(\gamma(i) - \nu(i))}.$$

The cancellation of the powers of n uses $\sum_{i \in S} \gamma(i) = \sum_{i \in S} \nu(i) = 1$. It is immediate that

$$\prod_{i \in S} (\nu(i))^{n(\gamma(i) - \nu(i))} = 2^{n(H(\nu) - D(\gamma \parallel \nu) - H(\gamma))}.$$

■

Lemma 9.6. For every $\nu \in \mathbb{T}_n$,

$$\frac{1}{|\mathbb{T}_n|} 2^{nH(\nu)} \leq |\mathcal{Y}_n(\nu)| \leq 2^{nH(\nu)}$$

○

Proof. Let $\mathbb{P}_v$ be the probability law associated with an infinite sequence of i.i.d. random variables $Y_1, Y_2, \dots$, distributed according to v .

$$\sum_{y^n \in \mathcal{Y}_n(v)} \mathbb{P}_v[(Y_1, Y_2, \dots, Y_n) = y^n] \leq 1,$$

implying (from Lemma 9.3) that

$$|\mathcal{Y}_n(v)| 2^{-nH(v)} \leq 1.$$

For the lower bound, sum only over the types supported on Σ_v (the terms with $\gamma \not\ll v$ carry no probability under v): put $T_n(v) = \{\gamma \in T_n : \gamma \ll v\}$. Using the factorial inequality of Lemma 9.4 with $m = n\gamma(i) \geq 0$ and $l = nv(i) \geq 1$ on each coordinate of the support (zero counts of γ cause no 0^0 expression),

$$\begin{aligned} 1 &= \sum_{\gamma \in T_n(v)} \sum_{y^n \in \mathcal{Y}_n(\gamma)} \mathbb{P}_v[(Y_1, Y_2, \dots, Y_n) = y^n] \\ &= \sum_{\gamma \in T_n(v)} |\mathcal{Y}_n(\gamma)| 2^{-n(H(\gamma) + D(\gamma \| v))} \\ &\stackrel{(a)}{\leq} \sum_{\gamma \in T_n(v)} |\mathcal{Y}_n(v)| 2^{-n(H(v) - D(\gamma \| v) - H(\gamma))} 2^{-n(H(\gamma) + D(\gamma \| v))} \\ &= \sum_{\gamma \in T_n(v)} |\mathcal{Y}_n(v)| 2^{-nH(v)} \\ &\leq |T_n| |\mathcal{Y}_n(v)| 2^{-nH(v)}. \end{aligned}$$

Here, (a) follows from Corollary 9.5, and on $T_n(v)$ all divergences are finite so the cancellations are legitimate. Rearranging gives the lower bound $|T_n|^{-1} 2^{nH(v)} \leq |\mathcal{Y}_n(v)|$. ■

Lemma 9.7. For any $v \in T_n$ and any probability law μ on the finite alphabet Σ ,

$$\frac{1}{|T_n|} 2^{-nD(v \| \mu)} \leq \mathbb{P}_\mu(\widehat{P}_n = v) \leq 2^{-nD(v \| \mu)},$$

where $\widehat{P}_n = T_{Y^n}$ is the random empirical measure of an i.i.d. sequence drawn from μ . If $v \not\ll \mu$, the right-hand bound is $2^{-\infty} = 0$ and is immediate; if $v \ll \mu$ the displayed bounds hold as stated. ○

Proof. From Lemma 9.3, we see that

$$\mathbb{P}_\mu(\widehat{P}_n = v) = |\mathcal{Y}_n(v)| 2^{-n(H(v) + D(v \| \mu))}.$$

The proof is completed by applying Lemma 9.6. ■

Lemma 9.8. Let $\Sigma_{v_n}, \Sigma_v \subseteq \Sigma_\mu$ and $v_n \rightarrow v$. Then $D(v_n \| \mu) \rightarrow D(v \| \mu)$. ○

Proof. This follows as, for $i = 1, \dots, M$, $v_n(i) \rightarrow v(i)$ and hence $v_n(i) \log \frac{v_n(i)}{\mu(i)} \rightarrow v(i) \log \frac{v(i)}{\mu(i)}$. ■

Theorem 9.9. (Sanov) For every $\Gamma \subseteq \mathcal{M}(\Sigma_\mu)$,

$$-\inf_{v \in \Gamma^r} D(v \parallel \mu) \leq \liminf_n \frac{1}{n} \log \mathbb{P}_\mu(T_{Y^n} \in \Gamma) \leq \limsup_n \frac{1}{n} \log \mathbb{P}_\mu(T_{Y^n} \in \Gamma) \leq -\inf_{v \in \Gamma} D(v \parallel \mu).$$

Here, $\Gamma^r = \{v \in \Gamma : \exists v_n \in T_n \cap \Gamma \text{ for all sufficiently large } n, v_n \rightarrow v\}$ (finitely many initial n are irrelevant for the accessible set). $\circ$

Proof. From Lemma 9.7, we have

$$\begin{aligned} \mathbb{P}_\mu(T_{Y^n} \in \Gamma) &= \sum_{\gamma \in T_n \cap \Gamma} \mathbb{P}_\mu(T_{Y^n} = \gamma) \\ &\leq \sum_{\gamma \in T_n \cap \Gamma} 2^{-nD(\gamma \parallel \mu)} \\ &\leq \sum_{\gamma \in T_n \cap \Gamma} 2^{-n \inf_{v \in \Gamma} D(v \parallel \mu)} \\ &= |\Gamma \cap T_n| 2^{-n \inf_{v \in \Gamma} D(v \parallel \mu)} \\ &\leq (n+1)^{(M-1)} 2^{-n \inf_{v \in \Gamma} D(v \parallel \mu)}. \end{aligned}$$

Therefore,

$$\frac{1}{n} \log \mathbb{P}_\mu(T_{Y^n} \in \Gamma) \leq \frac{M-1}{n} \log(n+1) - \inf_{v \in \Gamma} D(v \parallel \mu).$$

Taking $\limsup_n$ on both sides yields the upper bound.

Given $v \in \Gamma^r$. Let $\hat{v}_n \in T_n \cap \Gamma$, such that $\hat{v}_n \rightarrow v$. Then

$$\begin{aligned} \mathbb{P}_\mu(T_{Y^n} \in \Gamma) &= \sum_{\gamma \in T_n \cap \Gamma} \mathbb{P}_\mu(T_{Y^n} = \gamma) \\ &\geq \mathbb{P}_\mu(T_{Y^n} = \hat{v}_n) \\ &\geq \frac{1}{|T_n|} 2^{-nD(\hat{v}_n \parallel \mu)}, \end{aligned}$$

where the last inequality follows from Lemma 9.7.

Therefore

$$\frac{1}{n} \log \mathbb{P}_\mu(T_{Y^n} \in \Gamma) \geq -\frac{1}{n} \log |T_n| - D(\hat{v}_n \parallel \mu) \geq -\frac{M-1}{n} \log(n+1) - D(\hat{v}_n \parallel \mu).$$

Taking $\liminf$ and using Lemma 9.8, we obtain

$$\liminf_n \frac{1}{n} \log \mathbb{P}_\mu(T_{Y^n} \in \Gamma) \geq -D(v \parallel \mu).$$

Since, this holds for all $v \in \Gamma^r$, we get

$$\liminf_n \frac{1}{n} \log \mathbb{P}_\mu(T_{Y^n} \in \Gamma) \geq -\inf_{v \in \Gamma^r} D(v \parallel \mu).$$

■

Given two sets A and B , the Hausdorff distance is defined as

$$d_H(A, B) = \max\{\sup_{x \in A} \inf_{y \in B} d(x, y), \sup_{y \in B} \inf_{x \in A} d(x, y)\}.$$

In the space of probability distributions, let us consider the underlying metric to be the total variation distance.

Lemma 9.10. Let $\{\Gamma_n\}_{n \geq 1}$ and Γ be nonempty subsets of $\mathcal{M}(\Sigma_\mu)$ (nonempty eventually in n is enough), and $d_H(\Gamma_n, \Gamma) \rightarrow 0$. Then

$$\lim_n \inf_{v \in \Gamma_n} D(v \| \mu) = \inf_{v \in \Gamma} D(v \| \mu).$$

○

Proof. On the compact support face $\mathcal{M}(\Sigma_\mu)$, the function $F(v) = D(v \| \mu)$ is uniformly continuous; let ω_F be its modulus of continuity. Nonemptiness guarantees that the infima and the approximate nearest points used below are defined. For each direction, approximate a nearest point within $d_H(\Gamma_n, \Gamma) + 1/n$:

$$\left| \inf_{\Gamma_n} F - \inf_{\Gamma} F \right| \leq \omega_F(d_H(\Gamma_n, \Gamma) + 1/n) \rightarrow 0.$$

This yields both one-sided estimates simultaneously and completes the proof. (The sets need not be closed.) ■

Theorem 9.11. (Finite-alphabet moving-set Sanov bounds) Let μ be a law on the finite alphabet Σ of size M . Let Γ and Γ_n be nonempty subsets of $\mathcal{M}(\Sigma_\mu)$ with $d_H(\Gamma_n, \Gamma) \rightarrow 0$. Define the *accessible set*

$$R = \{v \in \Gamma : \exists v_n \in \Gamma_n \text{ for all sufficiently large } n, v_n \rightarrow v\}.$$

With base-2 logarithms and $\inf \emptyset = +\infty$,

$$-\inf_{v \in R} D(v \| \mu) \leq \liminf_n \frac{1}{n} \log_2 \mathbb{P}_\mu(\widehat{P}_n \in \Gamma_n) \leq \limsup_n \frac{1}{n} \log_2 \mathbb{P}_\mu(\widehat{P}_n \in \Gamma_n) \leq -\inf_{v \in \Gamma} D(v \| \mu). \quad (9.4)$$

If $\Gamma_n \subseteq \Gamma$ for all sufficiently large n , then $R = \Gamma$ and the rate limit is $-\inf_{v \in \Gamma} D(v \| \mu)$. ○

Remark 9.2 (No interior inclusion). Hausdorff convergence of moving sets does not imply that R contains the interior Γ° . For a closed-set counterexample, take Bernoulli $\mu = (1/2, 1/2)$, parameterise its simplex by $p \in [0, 1]$, choose an irrational $\alpha \in (0, 1)$, and set $\Gamma = [0, 1]$, $\Gamma_n = \{(j + \alpha)/(n + 1) : j = 0, \dots, n\}$. Each Γ_n is finite (hence compact) and $d_H(\Gamma_n, \Gamma) \leq 1/(n + 1)$, yet every member of Γ_n is irrational whereas every n -type parameter is k/n , so $\Gamma_n \cap \Gamma = \emptyset$ and $\mathbb{P}_\mu(\widehat{P}_n \in \Gamma_n) = 0$ while Γ° is nonempty. The main bound below is therefore stated with the accessible set R . ○

Proof. Note that

$$\begin{aligned}
\mathbb{P}_\mu(\widehat{P}_n \in \Gamma_n) &= \sum_{\gamma \in \mathbb{T}_n \cap \Gamma_n} \mathbb{P}_\mu(\widehat{P}_n = \gamma) \\
&\stackrel{(a)}{\leq} \sum_{\gamma \in \mathbb{T}_n \cap \Gamma_n} 2^{-nD(\gamma \parallel \mu)} \\
&\leq \sum_{\gamma \in \mathbb{T}_n \cap \Gamma_n} 2^{-n \inf_{v \in \Gamma_n} D(v \parallel \mu)} \\
&= |\Gamma_n \cap \mathbb{T}_n| 2^{-n \inf_{v \in \Gamma_n} D(v \parallel \mu)} \\
&\leq (n+1)^{(M-1)} 2^{-n \inf_{v \in \Gamma_n} D(v \parallel \mu)}.
\end{aligned}$$

Here (a) follows from Lemma 9.7. Therefore,

$$\frac{1}{n} \log_2 \mathbb{P}_\mu(\widehat{P}_n \in \Gamma_n) \leq \frac{M-1}{n} \log_2(n+1) - \inf_{v \in \Gamma_n} D(v \parallel \mu).$$

Taking $\limsup_n$ on both sides and using Lemma 9.10 yields the upper bound.

For any $v \in R$, choose its sequence $v_n \in \mathbb{T}_n \cap \Gamma_n$ (valid for all sufficiently large n) with $v_n \rightarrow v$. Then

$$\begin{aligned}
\mathbb{P}_\mu(\widehat{P}_n \in \Gamma_n) &= \sum_{\gamma \in \mathbb{T}_n \cap \Gamma_n} \mathbb{P}_\mu(\widehat{P}_n = \gamma) \\
&\geq \mathbb{P}_\mu(\widehat{P}_n = v_n) \\
&\stackrel{(a)}{\geq} \frac{1}{|\mathbb{T}_n|} 2^{-nD(v_n \parallel \mu)}.
\end{aligned}$$

Here, again, (a) follows from Lemma 9.7. Therefore

$$\frac{1}{n} \log_2 \mathbb{P}_\mu(\widehat{P}_n \in \Gamma_n) \geq -\frac{1}{n} \log_2 |\mathbb{T}_n| - D(v_n \parallel \mu).$$

Taking $\liminf$ and using Lemma 9.8 (continuity of divergence on the support face) and the polynomial type count, we obtain

$$\liminf_n \frac{1}{n} \log_2 \mathbb{P}_\mu(\widehat{P}_n \in \Gamma_n) \geq -D(v \parallel \mu).$$

Since this holds for all $v \in R$, we get

$$\liminf_n \frac{1}{n} \log_2 \mathbb{P}_\mu(\widehat{P}_n \in \Gamma_n) \geq -\inf_{v \in R} D(v \parallel \mu).$$

This proves (9.4).

If $\Gamma_n \subseteq \mathbb{T}_n$ eventually, then Hausdorff convergence supplies approximating points in Γ_n for every $v \in \Gamma$ (for all sufficiently large n), so $R = \Gamma$, and the lower and upper bounds coincide. ■

9.2.2 Maximal couplings

Let p_X, p_Y be two distributions supported on a finite alphabet Σ . Let $\{\omega_n\}$ be a non-negative sequence such that $\omega_n \rightarrow 0$ as $n \rightarrow \infty$ and $\omega_n \sqrt{n} \rightarrow \infty$ as $n \rightarrow \infty$. Define $A_n \subseteq \Sigma^n$ as

$$A_n = \{g^n \in \Sigma^n : |T_{g^n}(a) - p_X(a)| \leq \omega_n p_X(a), \forall a \in \Sigma\}.$$

Remark 9.3. The set A_n is usually called the typical sequences (or strongly-typical sequences) corresponding to distribution p_X . These sets play an important role in network information theory, particularly in the proofs of the channel coding theorems. $\circ$

Similarly, let

$$B_n = \{g^n \in \Sigma^n : |T_{g^n}(a) - p_Y(a)| \leq \omega_n p_Y(a), \forall a \in \Sigma\}.$$

Define $\hat{\Gamma}_n = \{\nu \in \mathcal{M}(\Sigma \times \Sigma) : \nu = T_{(g_1^n, g_2^n)} \text{ for some } (g_1^n, g_2^n) \in A_n \times B_n\}$ and $\hat{\Gamma} = \Pi(p_X, p_Y)$, the set of all couplings with the given marginals.

Lemma 9.12. Let $\hat{\Gamma}_n$ and $\hat{\Gamma}$ be as described above. Then

$$d_H(\hat{\Gamma}_n, \hat{\Gamma}) \rightarrow 0.$$

$\circ$

Proof. Write $S_X = \text{supp } p_X$ and $S_Y = \text{supp } p_Y$ for the supports (zero-mass rows and columns stay zero throughout). Let $\nu \in \hat{\Gamma}$, i.e. $\nu \in \Pi(p_X, p_Y)$, supported on $S_X \times S_Y$. From Lemma 9.2 there exists a type $\hat{\nu}_n \in T_n$ such that $|\hat{\nu}_n(a, b) - \nu(a, b)| \leq \frac{1}{n}$ for all $(a, b) \in S_X \times S_Y$, with $\hat{\nu}_n(a, b) = 0$ whenever $\nu(a, b) = 0$ (the rounded support is contained in $\text{supp } \nu$). The joint coordinate error is at most $1/n$, so the marginal errors satisfy

$$\left| \sum_{b \in S_Y} (\hat{\nu}_n(a, b) - \nu(a, b)) \right| \leq \sum_{b \in S_Y} |\hat{\nu}_n(a, b) - \nu(a, b)| \leq \frac{|S_Y|}{n}, \quad \left| \sum_{a \in S_X} (\hat{\nu}_n(a, b) - \nu(a, b)) \right| \leq \frac{|S_X|}{n}.$$

Because $n\omega_n \rightarrow \infty$ (which follows from, and is weaker than, the stated $\sqrt{n}\omega_n \rightarrow \infty$), for all sufficiently large n depending only on the fixed marginals,

$$\frac{|S_Y|}{n} \leq \omega_n \min_{a \in S_X} p_X(a), \quad \frac{|S_X|}{n} \leq \omega_n \min_{b \in S_Y} p_Y(b),$$

uniformly over all couplings ν (the thresholds do not depend on ν ; the pointwise assertion in terms of $\nu(a, b)$ alone does not by itself explain the Hausdorff supremum). Hence $\hat{\nu}_n \in \hat{\Gamma}_n$, i.e. any sequence of the joint type $\hat{\nu}_n$ lies in $A_n \times B_n$. Moreover $d_{TV}(\hat{\nu}_n, \nu) \leq \frac{|S_X||S_Y|}{2n}$. Since this holds for every $\nu \in \hat{\Gamma}$,

$$\sup_{\nu \in \hat{\Gamma}} \inf_{\nu_n \in \hat{\Gamma}_n} d_{TV}(\nu, \nu_n) \leq \frac{|S_X||S_Y|}{2n} \rightarrow 0.$$

For the opposite direction, take $q_n \in \hat{\Gamma}_n$, an admissible joint type. Its marginals satisfy

$$d_{TV}((q_n)_X, p_X) \leq \omega_n/2, \quad d_{TV}((q_n)_Y, p_Y) \leq \omega_n/2.$$

Repair each marginal by a maximal coupling, applying the corresponding conditional kernels to the two coordinates: the repaired joint law q has exactly the marginals p_X, p_Y and

$$d_{TV}(q_n, q) \leq d_{TV}((q_n)_X, p_X) + d_{TV}((q_n)_Y, p_Y) \leq \omega_n.$$

Hence

$$\sup_{q_n \in \hat{\Gamma}_n} \inf_{q \in \hat{\Gamma}} d_{TV}(q_n, q) \leq \omega_n \rightarrow 0,$$

and altogether, for all large n ,

$$d_H(\hat{\Gamma}_n, \hat{\Gamma}) \leq \max \left\{ \frac{|S_X||S_Y|}{2n}, \omega_n \right\} \rightarrow 0,$$

as desired. ■

Lemma 9.13. (Data-Processing) Let $W_{Y|X}$ be a stochastic mapping (channel). Let p_X, q_X be two distributions on $\mathcal{X}$ and $p_Y = \sum_x W_{Y|X}(y|x)p_X$ and $q_Y = \sum_x W_{Y|X}(y|x)q_X$ be the two induced distributions on $\mathcal{Y}$. Then $d_{TV}(p_X, q_X) \geq d_{TV}(p_Y, q_Y)$. ○

Proof. Observe the following:

$$\begin{aligned} \sum_y |p_Y(y) - q_Y(y)| &= \sum_y \left| \sum_x W_{Y|X}(y|x)(p_X(x) - q_X(x)) \right| \\ &\leq \sum_{x,y} W_{Y|X}(y|x) |p_X(x) - q_X(x)| \\ &= \sum_x |p_X(x) - q_X(x)| \end{aligned}$$
■

Theorem 9.14. Let p_X and p_Y be two distributions having finite supports S_X, S_Y on an Abelian group $\mathbb{G}$. Let $\{\omega_n\}$ be a non-negative sequence such that $\omega_n \rightarrow 0$ as $n \rightarrow \infty$ and $\omega_n \sqrt{n} \rightarrow \infty$ as $n \rightarrow \infty$ (for the cardinality arguments below $n\omega_n \rightarrow \infty$ would suffice; the stronger assumption additionally supports the usual probability-one typicality interpretation). Define $A_n \subseteq \mathbb{G}^n$, as

$$A_n = \{g^n \in \mathbb{G}^n : |k_{g^n}(a) - np_X(a)| \leq np_X(a)\omega_n, \forall a \in \mathbb{G}\}.$$

Here $k_{g^n}(a) := |\{i : g_i = a, 1 \leq i \leq n\}|$. Similarly, let

$$B_n = \{g^n \in \mathbb{G}^n : |k_{g^n}(a) - np_Y(a)| \leq np_Y(a)\omega_n, \forall a \in \mathbb{G}\}.$$

$$\lim_n \frac{1}{n} \log_2 |A_n + B_n| = \max_{q \in \Pi(p_X, p_Y)} H_q(X + Y),$$

with base-2 logarithms, where $\Pi(p_X, p_Y)$ is the set of joint distributions (couplings) such that the marginals are p_X and p_Y , respectively. ○

Proof. Choose the finite ambient alphabet explicitly as $S \supseteq S_X \cup S_Y \cup (S_X + S_Y)$ (this covers the sum supports for *every* coupling, not a single unspecified p_{X+Y}), and let μ be the uniform distribution on S . Let $C_n = A_n + B_n$, and $\Gamma_n = \{\nu \in \mathcal{M}(S) : \nu = T_{g^n} \text{ for some } g^n \in C_n\}$. Let $\Gamma = \{\nu \in \mathcal{M}(S) : \nu = q_{X+Y}, q_{X,Y} \in \Pi(p_X, p_Y)\}$. Define $\hat{\Gamma}_n = \{\nu \in \mathcal{M}(S \times S) : \nu = T_{(g_1^n, g_2^n)} \text{ for some } (g_1^n, g_2^n) \in A_n \times B_n\}$ and $\hat{\Gamma} = \Pi(p_X, p_Y)$, the set of all couplings with the given marginals.

For $\nu \in \Gamma_n$, we had defined the type-class by $\mathcal{Y}_n(\nu) = \{y^n \in S^n : T_{y^n} = \nu\}$. Hence, $\nu \in \Gamma_n$ if and only if $\mathcal{Y}_n(\nu) \subseteq C_n$. Note that, by definition, the sets A_n , B_n , and C_n are permutation invariant. Therefore, by permutation invariance,

$$\mathbb{P}_\mu(\hat{P}_n \in \Gamma_n) = \mathbb{P}_\mu(Y^n \in C_n) = \frac{|C_n|}{|S|^n},$$

where $\hat{P}_n = T_{Y^n}$ is the random empirical measure of an i.i.d. sequence drawn from μ . From Lemma 9.12, $d_H(\hat{\Gamma}_n, \hat{\Gamma}) \rightarrow 0$. The pushforward $q \mapsto q_{X+Y}$ is a deterministic TV contraction, so by Lemma 9.13 it sends the uniform joint-type Hausdorff convergence to output-type Hausdorff convergence: $d_H(\Gamma_n, \Gamma) \rightarrow 0$. For each n , Γ_n consists of n -types. Therefore the type-valued case of the moving-set theorem (Theorem 9.11) applies, and its accessible set is the Hausdorff limit Γ (which is generally a continuum of distributions, not a set of types). The theorem yields

$$\lim_n \frac{1}{n} \log_2 \mathbb{P}_\mu(\hat{P}_n \in \Gamma_n) = - \inf_{\nu \in \Gamma} D(\nu \| \mu)$$

or equivalently

$$\lim_n \frac{1}{n} \log_2 \frac{|C_n|}{|S|^n} = - \inf_{\nu \in \Gamma} (\log_2 |S| - H_\nu(X + Y)) = \sup_{\nu \in \Gamma} H_\nu(X + Y) - \log_2 |S|.$$

The maximum $\max_{q \in \Pi(p_X, p_Y)} H_q(X + Y)$ exists by compactness of the finite coupling polytope and continuity of finite-alphabet entropy, so $\sup_{\nu \in \Gamma} H_\nu(X + Y) = \max_{q \in \Pi(p_X, p_Y)} H_q(X + Y)$, and we are done. ■

Remark 9.4 (The translation programme). The entropy-of-sum identity above is the seed of the *translation programme* between additive combinatorics and information theory: sumset cardinality estimates on one side translate into entropy (or entropy-power) inequalities on the other. A short survey of this correspondence is [Nair \(2026\)](#); the maximal-coupling proof technique used above underlies the information inequality for sums on finite subsets of Abelian groups in [Lau et al. \(2026\)](#), and further entropy inequalities for sumsets obtained by these ideas are in [Lau and Nair \(2023\)](#). These are natural optional reading after this appendix. ○

9.3 A Min-Max Theorem

The min-max corollary below is one of the key analytic tools of this course: it converts an outer-bound functional with a dual-weight minimisation into the computable single-letter

expression. The standard template, used repeatedly in later chapters, is the following. One identifies a functional of the form $f(\lambda, p) = \sum_i \lambda_i T_i(p)$; once convexity of the componentwise lower-set region has been established — for example through simultaneous concavity on a common convex domain or an admissible time-sharing construction — the corollary yields

$$\sup_p \min_{\lambda \in \Lambda_d} \sum_i \lambda_i T_i(p) = \min_{\lambda \in \Lambda_d} \sup_p \sum_i \lambda_i T_i(p),$$

where the left-hand side is the quantity bounded in the converse and the right-hand side is the computable single-letter expression. Compactness of the weight simplex and affine dependence on the weights supply the topological hypotheses automatically; no compactness of the distribution class is assumed in this corollary. Subadditivity is a separate property used to single-letterise the resulting bound; it does not by itself establish lower-set convexity. See §1.13 of Chapter 1 and the UV/UVW outer bounds in the broadcast channel chapter.

Theorem 9.15. (Theorem 3 of *Terkelsen (1972)*) Let X be a nonempty compact connected space, let Y be a nonempty set, and let $f : X \times Y \mapsto \mathbb{R}$ be a function satisfying:

- (i) For any $y_1, y_2 \in Y$ there exists $y_0 \in Y$ such that

$$f(x, y_0) \geq \frac{1}{2} (f(x, y_1) + f(x, y_2)), \forall x \in X.$$

- (ii) Every finite intersection of sets of the form $\{x \in X : f(x, y) \leq \alpha\}$ with $(y, \alpha) \in Y \times \mathbb{R}$ is closed and connected.

Then

$$\sup_{y \in Y} \min_{x \in X} f(x, y) = \min_{x \in X} \sup_{y \in Y} f(x, y).$$

○

We now present a corollary of the above theorem in the form in which it is used in information theory.

Corollary 9.16. (Finite-vector lower-set min-max) Let Λ_d be the probability simplex in $\mathbb{R}^d$ (of affine dimension $d - 1$), i.e. $\lambda_i \geq 0$ and $\sum_{i=1}^d \lambda_i = 1$. Let $\mathcal{P} \neq \emptyset$ be any set, and let $T : \mathcal{P} \rightarrow \mathbb{R}^d$ be finite-valued, i.e. $T_i(p)$ finite for $i = 1, \dots, d$ and all $p \in \mathcal{P}$. Suppose the componentwise lower-set region $\mathcal{A}$, defined by

$$\mathcal{A} = \{(a_1, a_2, \dots, a_d) \in \mathbb{R}^d : a_i \leq T_i(p) \text{ for all } i, \text{ for some } p \in \mathcal{P}\},$$

is a convex set (the same p witnesses all the inequalities in the definition of $\mathcal{A}$). Then

$$\sup_{p \in \mathcal{P}} \min_{\lambda \in \Lambda_d} \sum_{i=1}^d \lambda_i T_i(p) = \min_{\lambda \in \Lambda_d} \sup_{p \in \mathcal{P}} \sum_{i=1}^d \lambda_i T_i(p).$$

The right-hand minimum is attained; the common value may be $+\infty$.

○

Proof. Let $f(\lambda, p) = \sum_{i=1}^d \lambda_i T_i(p)$. It suffices to verify that $f(\lambda, p)$ satisfies the conditions of Theorem 9.15, with the compact connected space being the weight simplex Λ_d . For any $p_1, p_2 \in \mathcal{P}$, both $T(p_1)$ and $T(p_2)$ lie in $\mathcal{A}$ (their coordinates are dominated by themselves, so $T(p_i) \in \mathcal{A}$ componentwise). Since $\mathcal{A}$ is convex, their midpoint $[T(p_1) + T(p_2)]/2$ lies in $\mathcal{A}$, so by the definition of the lower-set region some $p_0 \in \mathcal{P}$ satisfies

$$T_i(p_0) \geq \frac{1}{2}(T_i(p_1) + T_i(p_2)), i = 1, \dots, d.$$

Hence (using linearity in λ and non-negativity of λ_i) we have

$$f(\lambda, p_0) \geq \frac{1}{2}(f(\lambda, p_1) + f(\lambda, p_2)), \forall \lambda \in \Lambda_d.$$

Since $f(\lambda, p)$ is an affine function of λ (linear plus constant zero), it is immediate that the set

$$\mathcal{B}(p, \alpha) = \{\lambda \in \Lambda_d : f(\lambda, p) \leq \alpha\}$$

is closed for every pair $(p, \alpha) \in \mathcal{P} \times \mathbb{R}$, and convex. Every finite intersection of closed convex sets is closed and convex, hence connected (convexity of the finite intersection follows from linearity of the individual sublevel sets). Therefore finite intersections of the sets of the form $\mathcal{B}(p, \alpha)$ are closed and connected, and condition (i) holds simultaneously for every $\lambda \in \Lambda_d$ by the componentwise domination above. Thus the Corollary 9.16 follows from Theorem 9.15. Moreover, for each fixed p , $\lambda \mapsto f(\lambda, p)$ is continuous and affine, so $\sup_{p \in \mathcal{P}} \lambda \cdot T(p)$ is lower semicontinuous as a supremum of continuous functions; compactness of Λ_d gives the indicated minimum, and nonemptiness with finite $T(p)$ provides a finite lower bound at each λ . ■

Remark 9.5. Corollary 9.16 is the form of the min–max theorem used in the outer bounds of this course. The convexity hypothesis on $\mathcal{A}$ must be verified in each application — for example through simultaneous concavity on a common convex domain, or through an admissible time-sharing/concave-envelope construction whose admissibility and objective behavior are checked. Sub-additivity of the functionals T_i (in the sense of §1.13 of Chapter 1) is a separate property and does *not* by itself supply convexity of $\mathcal{A}$: with two admissible parameter values satisfying $T(p_0) = (1, 0)$ and $T(p_1) = (0, 1)$, the lower-set region is the union of two downward rectangles and is not convex, even though each coordinate can be extended to an exactly additive multi-letter functional by summing its one-letter values. See the broadcast channel chapter for explicit instances. ○

9.4 d-separation: proof of the graphical independence rule

This section proves the d-separation rule used in Chapter 1 to read conditional independences off a dependency graph. All random variables are finite-valued; functional nodes (deterministic functions of their parents) are allowed as degenerate conditionals.

Let G be a directed acyclic graph on vertices $V = \{v_1, \dots, v_m\}$, and let $\text{pa}(v)$ denote the parents of v . A law $p(v)$ factorises over G if

$$p(v) = \prod_{v \in V} p(v \mid \text{pa}(v)).$$

A path between two sets is *blocked* by C if some node on it belongs to C without being a collider, or some collider on it has no conditioned descendant (neither itself nor any of its descendants lies in C); here a *collider* is a node at which two consecutive edges of the path meet head-to-head. The set C *d-separates* A from B if every path between A and B is blocked by C .

Theorem 9.17. (Directed global Markov property) Let A, B, C be pairwise disjoint subsets of V . If $p(v)$ factorises over the directed acyclic graph G and C d-separates A from B in G , then $A \perp\!\!\!\perp B \mid C$ (interpreted on conditioning values of positive probability). $\circ$

Proof. We first prove the undirected case and then reduce the directed case to it by ancestral moralisation.

Undirected case. Suppose $p(v) = \prod_{D \in \mathcal{D}} \psi_D(v_D)$ factors over the cliques $\mathcal{D}$ of an undirected graph H , and that S separates A from B in H (every path meets S). Let A' be the union of the connected components of $V \setminus S$ that meet A , and $B' = V \setminus (S \cup A')$, so that $A \subseteq A'$, $B \subseteq B'$ and no edge of H joins A' to B' . Each clique D lies entirely in $A' \cup S$ or entirely in $B' \cup S$: if D contained $u \in A'$ and $w \in B'$, the clique edge $u-w$ would contradict the definition of the components. Collecting the clique factors accordingly, $p(v) = F(v_{A' \cup S}) G(v_{B' \cup S}) K(v_S)$. Summing out $A' \setminus A$ and $B' \setminus B$ gives $p(a, b, s) = \tilde{F}(a, s) \tilde{G}(b, s) K(s)$ for suitable $\tilde{F}, \tilde{G}$, and normalising each side of s yields $p(a, b \mid s) = p(a \mid s) p(b \mid s)$, i.e. $A \perp\!\!\!\perp B \mid S$.

Directed case: reduction to the ancestral moral graph. Let $W = \text{An}(A \cup B \cup C)$ be the set of ancestors of the queried and conditioned vertices (the vertices themselves included). Marginalise out the vertices outside W one at a time in reverse topological order: at each step the eliminated vertex is a sink of the remaining graph, so its factor $p(v \mid \text{pa}(v))$ sums to 1 and the remaining law keeps a DAG factorisation over the induced graph $G[W]$. Moralise $G[W]$: connect every pair of co-parents and drop the arrowheads, obtaining the undirected graph H . Each factor $p(v \mid \text{pa}(v))$ of the factorisation of the marginal law over $G[W]$ is a function of the set $\{v\} \cup \text{pa}(v)$, which is a clique of H ; hence the marginal law $p(v_W)$ factors over the cliques of H .

The bridge between the directed and undirected blocking rules is the following standard equivalence, due to Verma and Pearl (Verma and Pearl, 1988) (see also Lauritzen (Lauritzen, 1996)).

Lemma 9.18. (Ancestral moralisation) With H as above, C d-separates A from B in G if and only if C separates A from B in H . $\circ$

Proof of one direction; the other is cited. Suppose some path in G between $a \in A$ and $b \in B$ is active given C . Every node of an active path is an ancestor of $A \cup B \cup C$: walking the path, each maximal directed run of consecutive edges is traversed in a consistent direction, and the run ends at one of the endpoints a, b or at an active collider; every node of the run is an ancestor

of that terminal node, and since each active collider has a descendant in C , every node of the run lies on a directed path to $A \cup B \cup C$. Hence the path lies in $G[W]$. Replacing each collider segment $u \rightarrow c \leftarrow w$ by the moral edge $u-w$ and keeping the other edges produces a path in H between a and b whose non-collider nodes lie outside C (activity) and whose colliders have been removed; it avoids C . Thus, if C fails to d-separate in G , then C fails to separate in H ; equivalently, separation in H implies d-separation in G . The reverse implication — d-separation in G implies separation in H — is the nontrivial half of the Verma–Pearl equivalence and is cited rather than reproduced. ■

By the d-separation hypothesis and the cited direction of Lemma 9.18, C separates A from B in H ; since $p(v_W)$ factors over the cliques of H , the undirected case gives $A \perp\!\!\!\perp B \mid C$ for the marginal law over W , and hence for the original law. ■

Remark 9.6 (Caution). D-separation guarantees conditional independence only after the actual joint law has been shown to factorise over the proposed DAG. Conditioning on the parents separates a node from its non-descendants, not generally from its descendants. Failure of d-separation does not itself prove dependence; deterministic relations may create additional independences. ○

Corollary 9.19. (Coordinate star) Let (X_i, Y_i) be independent pairs and let $M = f(Y^n)$ be a function of the observations. Then, for each i ,

$$X^{i-1} \perp\!\!\!\perp (M, Y_i^n, X_i) \mid Y^{i-1}, \quad X_i \perp\!\!\!\perp (M, Y^{i-1}, X^{i-1}) \mid Y_i,$$

equivalently in chain notation $X^{i-1} \rightarrow Y^{i-1} \rightarrow (M, Y_i^n, X_i)$ and $X_i - Y_i - (M, Y^{i-1}, X^{i-1})$. ○

Proof. Orient the graph as $Y_j \rightarrow X_j$ and $Y_j \rightarrow M$ for every j ; the joint law $\prod_j p(y_j) p(x_j \mid y_j) \mathbf{1}\{M = f(y^n)\}$ factorises over this DAG by construction, and every X_j is a leaf (its only neighbour is Y_j). Every path from a node of X^{i-1} — that is, from some X_j with $j < i$ — to any other node starts with the edge $X_j - Y_j$; the node Y_j has only outgoing edges ($Y_j \rightarrow X_j$ and $Y_j \rightarrow M$), so on any path through it Y_j is a non-collider, and $Y_j \in Y^{i-1}$ is conditioned, which blocks the path. Likewise every path from X_i starts with $X_i - Y_i$ and is blocked by the conditioned non-collider Y_i . By Theorem 9.17 the two displayed independences follow. ■

Bibliography

- Ahlsvede, R. and P. Gács (1976). Spreading of sets in product spaces and hypercontraction of the Markov operator. *Annals of Probability* 4, 925–939.
- Ahlsvede, R. and T. Han (1983). On source coding with side information via a multiple-access channel and related problems in multi-user information theory. *IEEE Transactions on Information Theory* 29(3), 396–412.
- Ahlsvede, R. F. and J. Körner (November, 1975). Source coding with side information and a converse for degraded broadcast channels. *IEEE Trans. Info. Theory* IT-21(6)(6), 629–637.
- Anantharam, V., A. Gohari, S. Kamath, and C. Nair (2014, June). On hypercontractivity and a data processing inequality. In *2014 IEEE International Symposium on Information Theory (ISIT'2014)*, Honolulu, USA, pp. 3022–3026.
- Anantharam, V., A. Gohari, and C. Nair (2019, March). On the evaluation of marton's inner bound for two-receiver broadcast channels. *IEEE Transactions on Information Theory* 65(3), 1361–1371.
- Anantharam, V., A. A. Gohari, S. Kamath, and C. Nair (2013). On maximal correlation, hypercontractivity, and the data processing inequality studied by erkip and cover. *CoRR abs/1304.6133*.
- Anantharam, V., A. A. Gohari, and C. Nair (2013, July). Improved cardinality bounds on the auxiliary random variables in marton's inner bound. In *2013 IEEE International Symposium on Information Theory (ISIT'2013)*, Istanbul, Turkey.
- Anantharam, V., V. S. Jog, and C. Nair (2022). Unifying the brascamp-lieb inequality and the entropy power inequality. *IEEE Transactions on Information Theory* 68(12), 7665–7684.
- Balister, P. and B. Bollobás (2012, Mar). Projections, entropy and sumsets. *Combinatorica* 32(2), 125–141.
- Bergmans, P. F. (March, 1973, 3). Coding theorem for broadcast channels with degraded components. *IEEE Trans. Info. Theory* IT-15, 197–207.
- Bergmans, P. P. (1974). A simple converse for broadcast channels with additive white gaussian noise. *IEEE Transactions on Information Theory* IT-20(2), 279–280.
- Bernstein, S. N. (1941). On a property which characterizes a gaussian distribution. *Proc. Leningrad. Polytech. Inst.* 217(3), 21–22. Also in: *Sobranie Sochinenii*, Vol. IV, Nauka, Moscow, 1964.
- Chong, H.-F., M. Motani, H. Garg, and H. El Gamal (2008, july). On the han-kobayashi region for the interference channel. *Information Theory, IEEE Transactions on* 54(7), 3188–3195.
- Costa, M. (1983, may). Writing on dirty paper (corresp.). *Information Theory, IEEE Transactions on* 29(3), 439–441.

- Costa, M. H. M. (2011). Noisebergs in Z Gaussian interference channels. *Information Theory and Applications Workshop (ITA)*.
- Costa, M. H. M. and A. El Gamal (1987). The capacity region of the discrete memoryless interference channel with strong interference (corresp.). *Information Theory, IEEE Transactions on* 33(5), 710–711.
- Costa, M. H. M., A. Gohari, C. Nair, and D. Ng (2023). A proof of the noiseberg conjecture for the Gaussian Z-interference channel. In *Proc. IEEE International Symposium on Information Theory (ISIT)*.
- Costa, M. H. M. and C. Nair (2016). Gaussian z-interference channel: around the corner. *Information Theory and Applications Workshop (ITA)*, 1–6.
- Cover, T. (1987). The capacity of the relay channel. In *Open Problems in Communication and Computation*, pp. 72–73. Springer.
- Cover, T. and A. El Gamal (1979, Sep.). Capacity theorems for the relay channel. *IEEE Transactions on Information Theory* 25(5), 572–584.
- Csiszar, I. and J. Körner (2011). *Information theory: coding theorems for discrete memoryless systems*. Cambridge University Press.
- El Gamal, A. (March, 1979, 3). The capacity of a class of broadcast channels. *IEEE Trans. Info. Theory* IT-25, 166–169.
- El Gamal, A. and M. H. M. Costa (1982). The capacity region of a class of deterministic interference channels. *IEEE Transactions on Information Theory* IT-28(2), 343–346.
- El Gamal, A., A. Gohari, and C. Nair (2022, Aug). A strengthened cutset upper bound on the capacity of the relay channel and applications. *IEEE Transactions on Information Theory* 68, 5012–5043.
- El Gamal, A. and Y.-H. Kim (2012). *Network Information Theory*. Cambridge University Press.
- Erkip, E. and T. M. Cover (1998). The efficiency of investment information. *IEEE Transactions on Information Theory* 44(3), 1026–1040.
- Fiedler, M. (1971). Bounds for the determinant of the sum of Hermitian matrices. *Proceedings of the American Mathematical Society* 30(1), 27–31.
- Gallager, R. G. (1974). Capacity and coding for degraded broadcast channels. *Probl. Peredac. Inform.* 10(3)(3), 3–14.
- Geng, Y. and C. Nair (2014, April). The capacity region of the two-receiver Gaussian vector broadcast channel with private and common messages. *IEEE Transactions on Information Theory* 60(4), 2087–2104.
- Gohari, A., Y. Liu, and C. Nair (2025). A conjecture regarding the optimizers of Marton’s inner bound for the two-receiver broadcast channel. A conference version was presented at ISIT 2025.
- Gohari, A., Y. Liu, and C. Nair (2026). A two auxiliary receiver outer bound to the capacity region of a two-receiver discrete memoryless broadcast channel. In *Proc. IEEE International Symposium on Information Theory (ISIT)*, Guangzhou, China.
- Gohari, A. and C. Nair (2022, Feb). Outer bounds for multiuser settings: The auxiliary receiver approach. *IEEE Transactions on Information Theory* 68(2), 701–736.

- Gohari, A., C. Nair, and D. Ng (2021). An information inequality motivated by the gaussian z -interference channel. In *Proc. of IEEE ISIT*.
- Kamath, S. (2015). Reverse hypercontractivity using information measures. In *Proceedings of the 53rd Annual Allerton Conference on Communications, Control and Computing*.
- Körner, J. and K. Marton (1977a). General broadcast channels with degraded message sets. *IEEE Transactions on Information Theory* IT-23(1), 60–64.
- Korner, J. and K. Marton (1979, Mar). How to encode the modulo-two sum of binary sources (corresp.). *IEEE Transactions on Information Theory* 25(2), 219–221.
- Körner, J. and K. Marton (Nov, 1977b, 1). Images of a set via two channels and their role in mult-user communication. *IEEE Trans. Info. Theory* IT-23, 751–761.
- Lau, K. and C. Nair (2023). Information inequalities via ideas from additive combinatorics. In *Proc. IEEE International Symposium on Information Theory (ISIT)*.
- Lau, K., C. Nair, and C. Yao (2022). Uniqueness of local maximizers for some non-convex log-determinant optimization problems using information theory. In *Proc. IEEE International Symposium on Information Theory (ISIT)*, pp. 432–437.
- Lau, K., C. Nair, and Z. Zhu (2026). A maximal-coupling information inequality of sums on finite subsets of Abelian groups. In *Proc. IEEE International Symposium on Information Theory (ISIT)*, Guangzhou, China.
- Lauritzen, S. L. (1996). *Graphical Models*. Oxford University Press.
- Ledoux, M., C. Nair, and Y. Wang (2021). *Log-convexity of Fisher information along heat flow*. Available at <http://chandra.ie.cuhk.edu.hk/pub/papers/NIT/Log-cvx.pdf>.
- Liang, Y., G. Kramer, and S. Shamai (2008). Capacity outer bounds for broadcast channels. *Proceedings of IEEE Inf. Theory Workshop, Porto, Portugal*, 2–4.
- Lim, S. H., Y. Kim, A. El Gamal, and S. Chung (2011, May). Noisy network coding. *IEEE Transactions on Information Theory* 57(5), 3132–3152.
- Liu, S., C. Nair, and L. Xia (2014, June). Interference channels with very weak interference. In *2014 IEEE International Symposium on Information Theory (ISIT'2014)*, Honolulu, USA, pp. 1031–1035.
- Liu, S., C. Nair, and L. Xia (2015). An enhanced genie-based outer bound for sum-rate of interference channels. In *Proc. Information Theory and Applications Workshop (ITA)*.
- Marton, K. (1979, May). A coding theorem for the discrete memoryless broadcast channel. *IEEE Transactions on Information Theory* 25(3), 306–311.
- Nair, C. (2011). A note on outer bounds for broadcast channel. *CoRR abs/1101.0640*.
- Nair, C. (2013, Mar). Upper concave envelopes and auxiliary random variables. *International Journal of Advances in Engineering Sciences and Applied Mathematics* 5(1), 12–20.
- Nair, C. (2014). Equivalent formulations of hypercontractivity using information measures. In *Proc. International Zurich Seminar on Communications (IZS)*.
- Nair, C. (2020). On marton’s achievable region: Local tensorization for product channels with a binary component. In *2020 Information Theory and Applications Workshop (ITA)*, pp. 1–7.

- Nair, C. (2026). Inequalities in additive combinatorics and information theory: A short survey of the translation programme. *IEEE BITS the Information Theory Magazine*. To appear.
- Nair, C. and A. El Gamal (2009, Oct.). The capacity region of a class of three-receiver broadcast channels with degraded message sets. *Information Theory, IEEE Transactions on* 55(10), 4479–4493.
- Nair, C. and A. El Gamal (January, 2007, 1). An outer bound to the capacity region of the broadcast channel. *IEEE Trans. Info. Theory* IT-53, 350–355.
- Nair, C. and Y. N. Wang (2016, July). Evaluating hypercontractivity parameters using information measures. In *2016 IEEE International Symposium on Information Theory (ISIT)*, pp. 570–574.
- Nair, C. and Y. N. Wang (2017, June). Reverse hypercontractivity region for the binary erasure channel. In *2017 IEEE International Symposium on Information Theory (ISIT) (ISIT'2017)*, Aachen, Germany, pp. 938–942.
- Nair, C. and Z. V. Wang (2011). The capacity region of the three receiver less noisy broadcast channel. *IEEE Transactions on Information Theory* 57(7), 4058–4062.
- Nair, C., Z. V. Wang, and Y. Geng (2010, June). An information inequality and evaluation of marton's inner bound for binary input broadcast channels. In *2010 IEEE International Symposium on Information Theory*, pp. 550–554.
- Nair, C. and L. Xia (2012). On three-receiver more capable channels. In *2012 IEEE International Symposium on Information Theory Proceedings*, pp. 378–382.
- Nair, C., L. Xia, and M. Yazdanpanah (2015, February). Sub-optimality of the Han–Kobayashi Achievable Region for Interference Channels. *ArXiv e-prints*.
- Nair, C. and M. Yazdanpanah (2017, June). Sub-optimality of superposition coding region for three receiver broadcast channel with two degraded message sets. In *2017 IEEE International Symposium on Information Theory (ISIT) (ISIT'2017)*, Aachen, Germany, pp. 1038–1042.
- Nair, C. and J. Zhao (2026). On the local optimality of Gaussian distributions for the Han–Kobayashi inner bound for the Gaussian z -interference channel. In *Proc. IEEE International Symposium on Information Theory (ISIT)*, Guangzhou, China.
- Sato, H. (1981, Nov). The capacity of the Gaussian interference channel under strong interference. *IEEE Transactions on Information Theory* 27, 786–788.
- Sefidgaran, M., A. Gohari, and M. R. Aref (2015, May). On korner-marton's sum modulo two problem. In *2015 Iran Workshop on Communication and Information Theory (IWCIT)*, pp. 1–6.
- Shang, X., G. Kramer, and B. Chen (2009, feb.). A new outer bound and the noisy-interference sum-rate capacity for gaussian interference channels. *Information Theory, IEEE Transactions on* 55(2), 689–699.
- Shannon, C. E. (July and October, 1948). A mathematical theory of communication. *Bell System Technical Journal* 27, 379–423 and 623–656.
- Slepian, D. and J. Wolf (1973, July). Noiseless coding of correlated information sources. *IEEE Transactions on Information Theory* 19(4)(4), 471–480.
- Stam, A. (1959). Some inequalities satisfied by the quantities of information of fisher and shannon. *Information and Control* 2(2), 101–112.

- Terkelsen, F. (1972). Some minimax theorems. *Mathematica Scandinavica* 31, 405–413.
- Verma, T. and J. Pearl (1988). Causal networks: Semantics and expressiveness. *Proceedings of the Fourth Conference on Uncertainty in Artificial Intelligence*, 352–359.
- Wang, Y. and C. Nair (2020). On optimal weighted-sum rates for the modulo sum problem. In *Proc. IEEE International Symposium on Information Theory (ISIT)*.
- Weingarten, H., Y. Steinberg, and S. Shamai (2006). The capacity region of the gaussian multiple-input multiple-output broadcast channel. *IEEE transactions on information theory* 52(9), 3936–3964.
- Willems, F. M. J. (1990). The maximal-error and average-error capacity region of the broadcast channel are identical. *Problems of Control and Information Theory* 19(4), 339–347.
- Wyner, A. (1975a, March). The common information of two dependent random variables. *IEEE Transactions On Information Theory* 21(2), 163–179.
- Wyner, A. (1975b, May). On source coding with side information at the decoder. *IEEE Transactions on Information Theory* 21, 294–300.
- Wyner, A. and J. Ziv (1976, January). The rate-distortion function for source coding with side information at the decoder. *IEEE Transactions on Information Theory* 22(1), 1–10.