

Soft Covering Through the Lens of Hypothesis Testing

Neri Merhav

September 18, 2026

The Viterbi Faculty of Electrical and Computer Engineering
Technion - Israel Institute of Technology
Technion City, Haifa 3200003, ISRAEL
E-mail: merhav@technion.ac.il

Abstract

The soft covering lemma asserts that a random codebook whose rate R exceeds the input-output mutual information, $I(X; Y)$, of a given discrete memoryless channel (DMC), causes the output mixture distribution to be statistically indistinguishable from the i.i.d. output distribution. We study this phenomenon through the lens of Neyman–Pearson hypothesis testing: given a channel output sequence y^n , can one decide whether it was produced when the channel was driven by a random codeword, or generated independently from the output marginal? We establish tight liminf/limsup exponent envelopes for the jointly averaged false-alarm (FA) probability $\alpha_n(\tau, R_n)$ and missed-detection (MD) probability $\beta_n(\tau, R_n)$, as functions of the decision threshold τ and the codebook rate R , identifying the single-letter exponent functions $E_{\text{FA}}(\tau, R)$ and $E_{\text{MD}}(\tau, R)$. We identify the precise regularity condition on the threshold τ — that τ lies strictly inside a phase region — under which the limits exist and equal the single-letter formulas; no condition on the channel is needed. The analysis reveals a rich phase structure. For $R < I(X; Y)$, there is a genuine exponential tradeoff between the two error types over the interval $\tau \in (0, I(X; Y) - R)$. At $R = I(X; Y)$, this tradeoff interval collapses to the single point $\tau = 0$, where both error exponents simultaneously vanish, a fact which manifests the soft covering phenomenon in the Neyman–Pearson sense. For $R > I(X; Y)$, the same instantaneous collapse persists at $\tau = 0$; moreover, for every τ at least one exponent is zero: the FA exponent is zero for $\tau \leq 0$ (FA probability does not decay exponentially), and the MD exponent is zero for $\tau \geq 0$ (and finite, channel-specific for $\tau < 0$; see Remark 4). There is no interval of τ where both exponents are simultaneously positive. A sharp phase transition in the MD exponent occurs at $\tau^* = [I(X; Y) - R]_+$ for all rates.

Index Terms: soft covering, hypothesis testing, error exponents, phase transitions.

1 Introduction

The soft covering lemma, introduced by Han and Verdú [1], is a cornerstone of information theory. It asserts that a random codebook $\mathcal{C}$ of rate $R > I(X; Y)$ causes the channel output mixture distribution $P_{Y^n|\mathcal{C}}$ to become statistically indistinguishable from the i.i.d. output distribution $P_Y^{\otimes n}$, in the sense that their total variation distance vanishes exponentially. This result underpins key results in channel coding, wiretap secrecy, common randomness generation, and random-binning arguments. Despite its importance, the classical soft covering results address only the regime $R > I(X; Y)$. The regime $R < I(X; Y)$ — where the codebook rate is below the mutual information — has received comparatively little attention, and it turns out that it is considerably the more challenging and structurally richer regime.

In this paper, we study the soft covering phenomenon across *all* rates $R \geq 0$ by formulating it as a Neyman–Pearson hypothesis testing problem: Given an observed output sequence y^n , we ask: was it produced by the channel W driven by a randomly chosen codeword from $\mathcal{C}$ (hypothesis $\mathcal{H}_1$), or was it drawn independently from the marginal $P_Y^{\otimes n}$ (hypothesis $\mathcal{H}_0$)? The Neyman–Pearson test, based on the log-likelihood ratio (LLR) $\Lambda(y^n)$, decides in favor of $\mathcal{H}_1$ when $\Lambda(y^n) \geq \tau$, where τ is a threshold tuned to obtain a prescribed tolerable FA probability. The two error events are then: false alarm (FA): deciding in favor of $\mathcal{H}_1$ when $\mathcal{H}_0$ is true, and missed detection (MD): deciding in favor of $\mathcal{H}_0$ when $\mathcal{H}_1$ is true. We derive tight liminf/limsup exponent envelopes for the jointly averaged FA and MD error probabilities $\alpha_n(\tau, R_n)$ and $\beta_n(\tau, R_n)$, governed by the single-letter formulas $E_{\text{FA}}(\tau, R)$ and $E_{\text{MD}}(\tau, R)$ to be defined formally in Section 2.

The analysis reveals a considerably surprising phase structure in the (τ, R) plane, governed by two critical thresholds: $\tau = 0$ and $\tau^* = \max\{0, I(X; Y) - R\}$. For $\tau < 0$, the FA exponent vanishes (unless the channel has some singularity like the Z-channel) and the MD exponent is positive and channel-specific (see Remark 4). At $\tau = 0$, the picture depends critically on R . The FA exponent is always zero at $\tau = 0$ (with the above digression for singular channels), for every $R > 0$. The MD exponent at $\tau = 0$ undergoes a phase transition at $R = I(X; Y)$: it is strictly positive for $R < I(X; Y)$, and vanishes for $R \geq I(X; Y)$. The case $R = I(X; Y)$, $\tau = 0$, where *both* exponents simultaneously vanish, is the Neyman–Pearson exponent formulation of the soft covering phenomenon: at exactly the soft covering rate, neither error decays exponentially at threshold $\tau = 0$. For $\tau > 0$, both error exponents can be positive, and the picture depends on R . For $R < I(X; Y)$, there is a genuine tradeoff interval $\tau \in (0, I(X; Y) - R)$ where both error exponents are simultaneously positive. Beyond this interval (for $\tau \geq I(X; Y) - R$), the MD exponent drops to zero. The tradeoff interval has width $I(X; Y) - R$, which shrinks to zero as $R \nearrow I(X; Y)$. For $R \geq I(X; Y)$ at least one of the two exponents is zero no matter what the value of τ may be, in other words, there is no interval where both exponents are simultaneously positive. Specifically: for $\tau \leq 0$, the FA exponent is zero (FA probability does not decay exponentially); for $\tau \geq 0$, the MD exponent is zero (the codeword output is statistically easy to detect). The FA exponent is strictly positive for $\tau > 0$ and grows with τ , but over that same range the MD exponent is identically zero. The two zero-regions cover the entire real line, overlapping only at $\tau = 0$ where both exponents vanish simultaneously.

A structural asymmetry between the two exponents emerges from the formulas (stated precisely in Theorem 1). The FA exponent penalizes both the deviation of the channel output empirical distribution from the i.i.d. marginal, and the rate surplus of a codeword’s mutual information over R , reflecting the rarity of the event that a noise sequence looks like a codeword output. The MD exponent penalizes only the deviation of the empirical channel from the true channel W , with no explicit rate term, reflecting the rarity of the event that a codeword output goes undetected. The proofs are based on large deviations properties concerning type-class enumerators [2, Chapter 4]. These enumerators are

binomial random variables with exponentially many trials and exponentially decaying probabilities of success.

A few words on earlier related work are in order.

Han and Verdú [1] introduced channel resolvability and established that the minimum rate R for which the mixture distribution $P_{Y^n|\mathcal{C}}$ can approximate $P_Y^{\otimes n}$ (in total variation or normalized KL divergence) equals $I(X;Y)$; this is the first-order result that forms the foundation of all subsequent work. The soft-covering lemma first appeared as Theorem 6.3 of Wyner [3], where it serves as a technical tool for the achievability proof of the common information theorem; it was later recognized as a central technique in wiretap secrecy, identification coding, and channel synthesis, and made the subject of systematic study by Han and Verdú [1] under the name of channel resolvability. A substantial body of later work derives the exact *exponential* rate of convergence of $P_{Y^n|\mathcal{C}}$ to $P_Y^{\otimes n}$ for $R > I(X;Y)$, under various distance measures. Hayashi [4] obtained a lower bound to the exponent under KL divergence. Parizi *et al.* [5] derived the exact exponent of the KL divergence with application to the wiretap channel. Yu and Tan [6] characterized the exact exponent under Rényi divergence of order $\alpha \in [0, 2]$ (which includes the Kullback-Leibler (KL) divergence as the limiting case $\alpha \rightarrow 1$) for i.i.d. random codes. Yagli and Cuff [7] established the exact exponent under total variation distance. Recently, Li *et al.* [8] derived a strong-converse exponent under the KL divergence. On a somewhat different research route, Cuff [9, 10] moved beyond the expected-value analysis to show that soft covering holds with probability doubly exponentially close to unity over the random codebook, enabling applications via the union bound. Cuff [11] also developed the theory of distributed channel synthesis, which relies on and strengthens the soft-covering lemma.

As mentioned above, Yu and Tan [6] characterized the exact exponential decay of the Rényi divergence of order $\alpha \in (0, 2)$ and $R > I(X;Y)$. There is a precise connection to hypothesis testing: the Rényi divergence of order α evaluated at the optimal $\alpha \in (0, 1)$ corresponds to the *Chernoff exponent* (symmetric Bayesian error exponent) of testing $P_Y^{\otimes n}$ against $P_{Y^n|\mathcal{C}}$. In Neyman–Pearson terms, this is the exponent achieved at the *specific threshold* τ that equalizes $E_{\text{FA}}(\tau, R)$ and $E_{\text{MD}}(\tau, R)$ (the Chernoff point); it does not give the full Neyman–Pearson operating characteristics.

The most closely related prior work is [13], which considers a model where a transmitter either sends a codeword from a random fixed-composition codebook of rate R , or is silent, outputting the all-zero vector $\mathbf{0}$. The receiver must jointly detect whether transmission occurred and, if so, decode the message. The figures of merit are the FA probability (deciding transmission when silent), the MD probability (deciding silent when transmitting), and a decoding error probability. For a fixed composition random codebook, [13] derives the optimal detector/decoder in an extended Neyman–Pearson sense and characterizes the exact random coding exponents of all three error probabilities as functions of the rate R and two threshold parameters $\alpha, \beta \in \mathbb{R}$. The analysis in [13] is based on the type-class enumeration method, the same fundamental tool used here. The similarities with the present

work are as follows. Both papers formulate the problem model as a Neyman–Pearson binary hypothesis testing ($\mathcal{H}_0$: pure noise, $\mathcal{H}_1$: codeword output), both use random fixed-composition codebooks, and both derive exact random-coding exponents via type-class enumeration. However, at the same time, there are several differences. The first is that in [13], under $\mathcal{H}_0$ the transmitter output is a repetitive symbol 0 that designates silence (no transmission), so under the null hypothesis, the channel output y^n is a response to the all-zero input vector, unlike the present work where y^n is the channel response to an i.i.d. input source, P_X . The second difference is that the analysis in [13] is valid for all $R \geq 0$, but is most natural for $R > 0$ with a non-trivial decoding task. The regime $R < I(X; Y)$ is not highlighted, and the soft-covering threshold $R = I(X; Y)$ does not play a special role in [13] since the hypotheses are always different regardless of R . In the present work, $R = I(X; Y)$ is the central threshold: it is where both error exponents simultaneously vanish, characterizing the soft covering phenomenon in Neyman–Pearson terms. Third, in [13] there is also a characterization of the exponent of the decoding error probability, which has no analogue here. Finally, in [13] there is no phase transition analogous to the one at $\tau = 0$ in the present work. Here, the behavior of $E_{\text{MD}}(\tau, R)$ for $\tau \leq 0$ is a distinctive feature of the problem at hand.

The present paper differs from all prior work in the following respects.

1. *Full Neyman–Pearson tradeoff for all τ .* All prior exponent results for resolvability correspond to a fixed scalar distance measure (KL, total variation, Rényi), which is the expected distance averaged over both $\mathcal{C}$ and y^n . We instead characterize the complete Neyman–Pearson operating characteristic: the pair $(E_{\text{FA}}(\tau, R), E_{\text{MD}}(\tau, R))$ for every threshold τ , giving a full tradeoff curve rather than a single operating point.
2. *All rates $R \geq 0$.* Every previous exponent result for resolvability requires $R > I(X; Y)$. We handle all $R \geq 0$, including the unexplored regime $R < I(X; Y)$ where the codebook output is more concentrated than i.i.d. and where there is a genuine Neyman–Pearson tradeoff between FA and MD.
3. *Proof method and tight envelopes.* Prior analyses apply Chernoff-type or Gallager-style bounding to the channel output distribution, which is a mixture of exponentially many conditional output distributions given the various input codewords. The Chernoff bound applied to a mixture of channel outputs is not automatically tight: prior papers establish tightness by a separate converse argument specific to each distance measure. In the present work, we avoid Chernoff bounds altogether. Instead, we apply the type-class enumeration method [2], which yields tight liminf/limsup exponent envelopes. We identify the precise regularity condition on τ — right-continuity of the exponent function at τ , equivalently τ strictly inside a phase region — under which the limits exist and equal the single-letter formulas. No condition on the channel W is needed.

A statistical-physics perspective on soft covering across all rates, including connections to phase transitions, appears in [14].

The outline of the remaining part of this work is as follows. Section 2 establishes notation conventions and provides some background. Section 3 states the main theorem, whose proof appears in the Appendix. Section 4 develops the phase structure and the connection to soft covering. Section 5 illustrates the results on the Z-channel. Finally, in Section 6, we summarize the main findings and speculate on possible future research directions.

2 Notation Conventions and Basic Background

Let $\mathcal{X}$ and $\mathcal{Y}$ denote finite input and output alphabets. Sequences of length n are written in boldface-free lowercase, i.e., $x^n = (x_1, \dots, x_n) \in \mathcal{X}^n$ and $y^n = (y_1, \dots, y_n) \in \mathcal{Y}^n$, where $\mathcal{X}^n$ and $\mathcal{Y}^n$ are the n -th Cartesian powers of $\mathcal{X}$ and $\mathcal{Y}$, respectively. The *type* of a sequence $x^n \in \mathcal{X}^n$ is the empirical probability distribution $\hat{P}_{x^n}(a) = \frac{1}{n} \# \{i : x_i = a\}$, $a \in \mathcal{X}$. The equivalence set of all sequences in $\mathcal{X}^n$ of type P is called the *type class* $\mathcal{T}(P) \subseteq \mathcal{X}^n$. For a sequence pair $(x^n, y^n) \in \mathcal{X}^n \times \mathcal{Y}^n$, the *joint type* Q_{XY} is the empirical distribution $Q_{XY}(a, b) = \frac{1}{n} \# \{i : x_i = a, y_i = b\}$, $a \in \mathcal{X}$, $b \in \mathcal{Y}$. Its marginals are denoted Q_X and Q_Y , and the conditional type $Q_{Y|X}$ is defined according to $Q_{XY}(a, b) = Q_X(a)Q_{Y|X}(b|a)$. We always restrict to joint types with $Q_X = P_X$, where P_X is some fixed input distribution. Throughout, $\hat{P}_{y^n}$ denotes the empirical distribution associated with y^n .

Information measures induced by a given probability distribution will be subscripted by the notation of this distribution. When this is an empirical distribution Q_{XY} , the subscript will be abbreviated by Q , in order to avoid cumbersome notation. Thus, $H_Q(Y) = -\sum_y Q_Y(y) \log Q_Y(y)$ is the marginal empirical entropy of an auxiliary random vector Y governed by Q_Y , $H_Q(Y|X) = -\sum_{x,y} Q_{XY}(x, y) \log Q_{Y|X}(y|x)$ is the conditional empirical entropy of Y given X , where (X, Y) are jointly governed by Q_{XY} , and $I_Q(X; Y) = H_Q(Y) - H_Q(Y|X)$ is the mutual information under Q_{XY} .

The discrete memoryless channel (DMC) with a finite input alphabet $\mathcal{X}$ and finite output alphabet $\mathcal{Y}$ will be denoted by $W : \mathcal{X} \rightarrow \mathcal{Y}$. The alphabets $\mathcal{X}$ and $\mathcal{Y}$ may have different sizes; no assumption $|\mathcal{X}| = |\mathcal{Y}|$ is made. When W is fed by an n -vector $x^n \in \mathcal{X}^n$, the corresponding channel output vector y^n is distributed according to

$$W^n(y^n|x^n) = \prod_{i=1}^n W(y_i|x_i). \quad (1)$$

The input distribution P_X is fixed throughout, with output marginal

$$P_Y(y) = \sum_x P_X(x)W(y|x) \quad (2)$$

and mutual information

$$I(X; Y) = \sum_{x,y} P_X(x)W(y|x) \log \frac{W(y|x)}{P_Y(y)}. \quad (3)$$

Remark 1 (Relation to channel capacity). *The mutual information $I(X;Y)$ is evaluated at the fixed input distribution P_X , not maximized over P_X . It is related to channel capacity $C = \max_{P_X} I(X;Y)$ by $I(X;Y) \leq C$, with equality iff P_X is a capacity-achieving distribution. All results in this paper hold for any fixed P_X , including when $R > C$ (in which case $R > I(X;Y)$ since $I(X;Y) \leq C$, and the soft covering regime applies regardless). The threshold $R = I(X;Y)$ is the soft-covering threshold for the chosen P_X ; with a capacity-achieving P_X it equals C .*

We also adopt the following shorthand notations.

$$D_m(Q_Y) = D(Q_Y \| P_Y) = \sum_y Q_Y(y) \log \frac{Q_Y(y)}{P_Y(y)} \quad (4)$$

is the KL divergence between Q_Y and P_Y defined in (2). Also,

$$D_c(Q_{XY}) = D(Q_{Y|X} \| W | P_X) = \sum_{x,y} P_X(x) Q_{Y|X}(y|x) \log \frac{Q_{Y|X}(y|x)}{W(y|x)} \quad (5)$$

is the conditional KL divergence between $Q_{Y|X}$ and the channel W with weighting P_X , and

$$\ell(Q_{XY}) = H_Q(Y|X) + D_c(Q_{XY}) \quad (6)$$

is a notation used throughout in the proofs.

The data processing inequality (DPI) of the KL divergence implies that

$$D_m(Q_Y) \leq D_c(Q_{XY}). \quad (7)$$

To see why this is true, observe that since $Q_X = P_X$,

$$\begin{aligned} D_c(Q_{XY}) &= D(Q_{XY} \| P_X \otimes W) - D(Q_X \| P_X) \\ &= D(Q_{XY} \| P_X \otimes W) - D(P_X \| P_X) \\ &= D(Q_{XY} \| P_X \otimes W) \\ &\geq D(Q_Y \| P_Y) \\ &= D_m(Q_Y), \end{aligned} \quad (8)$$

where the inequality is the DPI of the KL divergence applied to the marginalization map $(x, y) \mapsto y$.

For two positive sequences, $\{a_n\}_{n \geq 1}$ and $\{b_n\}_{n \geq 1}$, we write $a_n \doteq b_n$ to mean $\lim_{n \rightarrow \infty} \frac{1}{n} \log \frac{a_n}{b_n} = 0$. Accordingly, $a_n \doteq 0$ means that a_n decays faster than exponentially (e.g., doubly exponentially) and $a_n \doteq 1$ tells that a_n varies (grows or decays) at a sub-exponential rate, or even tends to a strictly positive constant. Throughout this paper, all logarithms are natural unless specified otherwise. We also denote the positive clipping operator by $[\cdot]_+$, which is defined by $[a]_+ = \max\{a, 0\}$ for any real a .

Throughout, we fix a sequence of codebook sizes $M_n \in \mathbb{N}$ satisfying $M_n \doteq e^{nR}$ for some rate $R \geq 0$, and set $R_n := \frac{1}{n} \ln M_n \rightarrow R$. A random codebook $\mathcal{C} = \{x^n(m)\}_{m=1}^{M_n}$ has codewords drawn independently at random under the uniform distribution across the type class $\mathcal{T}(P_X)$. (We assume

throughout that P_X is rational, so that $\mathcal{T}(P_X)$ is non-empty for all sufficiently large n ; equivalently, we work with a sequence of n -types $P_{X,n} \rightarrow P_X$, where $P_{X,n}$ is the n -type closest to P_X .) Once the codebook $\mathcal{C}$ has been randomly selected, the induced output mixture is given by

$$P_{Y^n|\mathcal{C}}(y^n) = \frac{1}{M_n} \sum_{m=1}^{M_n} W^n(y^n|x^n(m)). \quad (9)$$

We denote $P_Y^{\otimes n}(y^n) = \prod_{i=1}^n P_Y(y_i)$, where the single-letter output marginal P_Y is induced by P_X and W as in (2).

In this paper, we focus on the following Neyman–Pearson hypothesis testing problem: Under hypothesis $\mathcal{H}_0$, y^n is governed by $P_Y^{\otimes n}$, and under hypothesis $\mathcal{H}_1$, y^n is drawn by $P_{Y^n|\mathcal{C}}$. The log-likelihood ratio (LLR) statistic is defined as

$$\Lambda(y^n) = \frac{1}{n} \ln \frac{P_{Y^n|\mathcal{C}}(y^n)}{P_Y^{\otimes n}(y^n)}. \quad (10)$$

The likelihood ratio test (LRT) with threshold τ decides in favor of $\mathcal{H}_1$ if $\Lambda(y^n) \geq \tau$; otherwise, it accepts $\mathcal{H}_0$.

The two kinds of error probabilities associated with the LRT are as follows. The false-alarm (FA) probability is

$$\alpha_n(\tau, R_n) = \mathbb{E}_{\mathcal{C}} \left\{ \sum_{\{y^n: \Lambda(y^n) \geq \tau\}} P_Y^{\otimes n}(y^n) \right\}, \quad (11)$$

and the missed detection (MD) probability is

$$\beta_n(\tau, R_n) = \mathbb{E}_{\mathcal{C}} \left\{ \sum_{\{y^n: \Lambda(y^n) < \tau\}} P_{Y^n|\mathcal{C}}(y^n) \right\}, \quad (12)$$

where in both (11) and (12), $\mathbb{E}_{\mathcal{C}}\{\cdot\}$ denotes expectation with respect to (w.r.t.) the randomness of the codebook $\mathcal{C}$. Note that α_n and β_n depend on n through M_n (or equivalently R_n); the argument R_n makes this dependence explicit. We define the support-compatible set of joint distributions:

$$\mathcal{Q} := \mathcal{Q}(P_X, W) = \{Q \in \mathcal{P}(\mathcal{X} \times \mathcal{Y}) : Q_X = P_X, Q \ll P_X \otimes W\}. \quad (13)$$

This is the set of distributions that can be approximated by joint types in a codebook drawn from P_X over channel W ; only $Q \in \mathcal{Q}$ can contribute to $\alpha_n(\tau, R_n)$ and $\beta_n(\tau, R_n)$ at the exponential scale. For $Q \in \mathcal{Q}$, the divergences $D_m(Q_Y)$ and $D_c(Q_{XY})$ are finite and continuous in Q .

Here $P_{X,n}$ denotes the n -type closest to P_X (the unique n -type with $\|P_{X,n} - P_X\|_1 \leq |\mathcal{X}|/n$), and $\mathcal{Q}_n := \{Q \in \mathcal{P}_n(\mathcal{X} \times \mathcal{Y}) : Q_X = P_{X,n}, Q \ll P_X \otimes W\}$ is the set of admissible n -types. By the method of types, every $Q \in \mathcal{Q}$ can be approximated by a sequence of n -types $Q_n \in \mathcal{Q}_n$ with $Q_n \rightarrow Q$; no additional condition on W or P_X is needed.

The corresponding exponent functions are defined as the liminf and limsup of the normalized

log-probabilities (see Theorem 1 for conditions under which limits exist):

$$E_{\text{FA}}(\tau, R) := \inf_{Q \in \mathcal{Q}: \lambda(Q, R) \geq \tau} \{D_{\text{m}}(Q_Y) + [I_Q(X; Y) - R]_+\}, \quad (14)$$

$$E_{\text{MD}}(\tau, R) := \inf_{\substack{Q \in \mathcal{Q}: \lambda(Q, R) < \tau \\ \Delta(Q_Y, R) < \tau}} D_{\text{c}}(Q_{XY}), \quad (15)$$

with the conventions $\inf \emptyset = +\infty$ and $\sup \emptyset = -\infty$. These are the natural single-letter exponent candidates, and our main theorem shows they govern the liminf/limsup behavior of the error sequences.

The following quantity will be useful in the proofs. Given the randomly selected codebook, $\mathcal{C}$, and a channel output vector, y^n , the type-class enumerator (TCE) associated with type Q_{XY} and y^n is defined as

$$N(Q_{XY}|y^n) = \#\{m : (x^n(m), y^n) \in \mathcal{T}(Q_{XY})\}. \quad (16)$$

Clearly, the randomness of $\mathcal{C}$ induces randomness of $N(Q_{XY}|y^n)$. In particular, due to the independent random selection, $N(Q_{XY}|y^n)$ is a binomial random variable with M_n trials and success rate given by the probability that a single randomly chosen codeword from $\mathcal{T}(P_X)$ happens to have, together with y^n , the given joint type Q_{XY} . This probability is of the exponential order of $e^{-nI_Q(X; Y)}$. The unnormalized mixture under $\mathcal{H}_1$ can be easily expressed in terms of the TCE's as follows:

$$S(y^n) = \sum_{m=1}^{M_n} W^n(y^n|x^n(m)) = \sum_{Q_{XY}} N(Q_{XY}|y^n) e^{-n\ell(Q_{XY})}, \quad (17)$$

and so, $P_{Y^n|\mathcal{C}}(y^n) = S(y^n)/M_n$. For y^n of type $\hat{P}_{y^n} = Q_Y$, using $P_Y^{\otimes n}(y^n) = e^{-n[H_Q(Y) + D_{\text{m}}(Q_Y)]}$:

$$\Lambda(y^n) = \frac{1}{n} \log S(y^n) - R_n + H_Q(Y) + D_{\text{m}}(Q_Y). \quad (18)$$

To analyze the two kinds of probability of error, we shall invoke results concerning the large deviations behavior of $\{N(Q_{XY})\}$. Since these are binomial random variables, the following theorems from [2] (with a slight change in notation) will be useful.

(T1) **Theorem 4.1 of [2].** For $N \sim \text{Binomial}(e^{nA}, e^{-nB})$ ($A > 0$ and $B > 0$) and $C \in \mathbb{R}$:

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \ln \Pr\{N > e^{nC}\} = \begin{cases} [B - A]_+ & [A - B]_+ \geq C, \\ \infty & \text{elsewhere,} \end{cases} \quad (19)$$

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \ln \Pr\{N < e^{nC}\} = \begin{cases} 0 & A - B < C, \\ \infty & A - B > C. \end{cases} \quad (20)$$

When $A > B$, $\Pr\{N > e^{nC}\} \doteq 1$ for $C \leq A - B$, and $\Pr\{N < e^{nC}\}$ is doubly exponentially small for $C < A - B$. When $A < B$ and $C < 0$, $\Pr\{N > e^{nC}\} \doteq e^{-n(B-A)}$. We note in passing that the case $A = B$ is the case of asymptotic Poissonianity. In particular, if $N \sim \text{Binomial}(e^{nA}, \mu e^{-nA})$, then as $n \rightarrow \infty$, the distribution of N tends to a Poissonian random variable with parameter μ , i.e., $\Pr\{N = k\} \rightarrow \frac{\mu^k e^{-\mu}}{k!}$, $k = 0, 1, 2, \dots$

(T2) **Theorem 4.3 of [2].** Let $N_j \sim \text{Binomial}(e^{nA_j}, e^{-nB_j})$, $j = 1, \dots, k_n$, with $k_n \doteq 1$ (i.e., k_n grows sub-exponentially in n). Then for any thresholds $C_j \in \mathbb{R}$:

$$\Pr\left\{\bigcap_j \{N_j \leq e^{nC_j}\}\right\} \doteq \mathbf{1}\left\{\min_j (B_j - A_j + [C_j]_+) > 0\right\}. \quad (21)$$

When $\min_j (B_j - A_j + [C_j]_+) > 0$, every $\Pr\{N_j > e^{nC_j}\}$ is either exponentially or doubly exponentially small (by Theorem 4.1 of [2]), and so, the union bound on the complementary events yields an intersection probability tending to 1. On the other hand, when $\min_j (B_j - A_j + [C_j]_+) \leq 0$, then at least for some j , $B_j - A_j + [C_j]_+ \leq 0$, and for that j , $\Pr\{N_j \leq e^{nC_j}\} \doteq 0$, which implies that the intersection probability is a-fortiori $\doteq 0$. Note that there are no assumptions on the statistical dependence or independence among $\{N_j\}$.

It should be noted that in [2, Theorem 4.3], the assertion is formulated with a common threshold parameter C , rather than various thresholds C_j , as stated here. Nevertheless, the proof in [2] extends straightforwardly to allow different thresholds.

In view of those results, it is apparent that given y^n , a type Q_{XY} with $I_Q(X; Y) < R$ is a type that is typically populated with exponentially many codewords, in particular, $N(Q_{XY}|y^n)$ concentrates at the exponential order of $e^{n[R - I_Q(X; Y)]}$. Such a type will be henceforth referred to as a *bulk type*. By contrast, a type for which $I_Q(X; Y) > R$, is rarely populated, as the probability that $N(Q_{XY}|y^n) \geq 1$ is of the exponential order of $e^{-n[I_Q(X; Y) - R]} \rightarrow 0$. Such a type will be referred to as a *sparse type*.

3 Main Result

For a joint type $Q = Q_{XY}$ with $Q_X = P_X$ and rate $R \geq 0$, let us define

$$\lambda(Q_{XY}, R) = D_m(Q_Y) - D_c(Q_{XY}) + [I_Q(X; Y) - R]_+. \quad (22)$$

By the DPI (7), $D_m(Q_Y) \leq D_c(Q_{XY})$. Hence for bulk types ($I_Q(X; Y) \leq R$): $\lambda(Q_{XY}, R) = D_m(Q_Y) - D_c(Q_{XY}) \leq 0$. For the true channel $Q_{Y|X} = W$, $D_c(P_X \otimes W) = D(W\|W|P_X) = 0$ and $D_m(Q_Y) = D(P_Y\|P_Y) = 0$, and so, $\lambda(P_X \otimes W, R) = [I(X; Y) - R]_+$, which is positive for $R < I(X; Y)$, and zero for $R \geq I(X; Y)$.

Our main theorem, whose proof appears in the appendix, is the following.

Theorem 1. Let $M_n \in \mathbb{N}$ with $M_n \doteq e^{nR}$, and set $R_n := \frac{1}{n} \ln M_n \rightarrow R$.

(i) **(FA lower envelope, always).** For all $R \geq 0$, $\tau \in \mathbb{R}$, and any sequence $M_n \doteq e^{nR}$:

$$\liminf_{n \rightarrow \infty} \left(-\frac{1}{n} \ln \alpha_n(\tau, R_n) \right) \geq E_{FA}(\tau, R) := \inf_{Q \in \mathcal{Q}: \lambda(Q, R) \geq \tau} \{D_m(Q_Y) + [I_Q(X; Y) - R]_+\}. \quad (23)$$

(ii) **(FA upper envelope).** For every fixed $\eta > 0$:

$$\limsup_{n \rightarrow \infty} \left(-\frac{1}{n} \ln \alpha_n(\tau, R_n) \right) \leq E_{FA}(\tau + \eta, R) := \inf_{Q \in \mathcal{Q}: \lambda(Q, R) \geq \tau + \eta} \{D_m(Q_Y) + [I_Q(X; Y) - R]_+\}. \quad (24)$$

(iii) (**MD lower envelope**). For every fixed $\eta > 0$ and $R > 0$:

$$\liminf_{n \rightarrow \infty} \left(-\frac{1}{n} \ln \beta_n(\tau, R_n) \right) \geq E_{\text{MD}}^\eta(\tau, R) := \inf_{\substack{Q \in \mathcal{Q}: \lambda(Q, R) < \tau + \eta \\ \Delta_<(Q_Y, R) < \tau + \eta}} D_c(Q_{XY}). \quad (25)$$

(iv) (**MD upper envelope**). For $R > 0$:

$$\limsup_{n \rightarrow \infty} \left(-\frac{1}{n} \ln \beta_n(\tau, R_n) \right) \leq E_{\text{MD}}(\tau, R) := \inf_{\substack{Q \in \mathcal{Q}: \lambda(Q, R) < \tau \\ \Delta(Q_Y, R) < \tau}} D_c(Q_{XY}), \quad (26)$$

where

$$\Delta(Q_Y, R) = \sup_{\substack{Q' \in \mathcal{Q}: Q'_Y = Q_Y \\ I_{Q'}(X; Y) \leq R}} [D_m(Q_Y) - D_c(Q'_{XY})], \quad (27)$$

and its strict-inequality counterpart

$$\Delta_<(Q_Y, R) = \sup_{\substack{Q' \in \mathcal{Q}: Q'_Y = Q_Y \\ I_{Q'}(X; Y) < R}} [D_m(Q_Y) - D_c(Q'_{XY})]. \quad (28)$$

Note $\Delta_<(Q_Y, R) \leq \Delta(Q_Y, R)$, with equality whenever the supremum in (27) is not achieved at $I_{Q'} = R$. For $\tau > 0$, both constraints $\Delta_<(Q_Y, R) < \tau$ and $\Delta(Q_Y, R) < \tau$ are redundant (by the DPI (7)) and may be dropped.

The FA and MD exponents exhibit a rich phase structure (fully characterized in Section 4) governed by the following critical threshold values. Let $Q_{XY}^* \in \mathcal{Q}$ denote the unconstrained minimizer of $D_m(Q_Y) + [I_Q(X; Y) - R]_+$ over $\mathcal{Q}$. Next, define:

$$\lambda_{\max}(R) := \sup_{Q \in \mathcal{Q}} \lambda(Q, R), \quad (29)$$

$$\lambda_{\min}(R) := \inf_{Q \in \mathcal{Q}} \lambda(Q, R), \quad (30)$$

$$\tau^*(R) := [I(X; Y) - R]_+, \quad (31)$$

$$\tau_{\text{flat}}(R) := \lambda(Q_{XY}^*, R), \quad (32)$$

$$\tau_{\text{kink}}(R) := \text{the unique } \tau \in (\lambda_{\min}(R), 0) \text{ at which } I_{Q^*(\tau)}(X; Y) = R, \quad (33)$$

where $Q^*(\tau)$ is the minimizer of (15) at threshold τ . Here: $\lambda_{\max}(R)$ and $\lambda_{\min}(R)$ are the largest and smallest LLR values achievable by types in $\mathcal{Q}$; $\tau^*(R)$ is the MD threshold above which $E_{\text{MD}}(\tau, R) = 0$; $\tau_{\text{flat}}(R)$ is the FA threshold below which $E_{\text{FA}}(\tau, R)$ is constant (the unconstrained minimum is also feasible); and $\tau_{\text{kink}}(R) < 0$ is the point in the MD active region where the minimizing type transitions from bulk ($I_{Q^*} \leq R$) to sparse ($I_{Q^*} > R$).

Remark 2 (Regularity condition for existence of limits). *All four parts of Theorem 1 hold for every DMC W and every $\tau \in \mathbb{R}$, $R \geq 0$, with no condition on W . The only condition needed for the limits to exist (rather than mere liminf/limsup envelopes) is a regularity condition on τ : for FA, the function $g_{\text{FA}}(t) := \inf_{Q \in \mathcal{Q}: \lambda(Q, R) \geq t} \{D_m(Q_Y) + [I_Q(X; Y) - R]_+\}$ must be right-continuous at $t = \tau$; for MD,*

the analogous infimum $g_{\text{MD}}(t) := \inf_{Q \in \mathcal{Q}: \lambda(Q, R) < t, \Delta(Q_Y, R) < t} D_c(Q_{XY})$ must be left-continuous at $t = \tau$. Both conditions hold if and only if τ is strictly inside a phase region (see (29)–(33)). All known counterexamples to limit existence are at these phase boundaries, confirming that this is both necessary and sufficient.

Corollary 1 (Existence of limits). *The limit $\lim_{n \rightarrow \infty} (-\frac{1}{n} \ln \alpha_n(\tau, R_n))$ exists and equals $E_{\text{FA}}(\tau, R)$ if and only if the function $g(t) := \inf_{Q \in \mathcal{Q}: \lambda(Q, R) \geq t} \{D_m(Q_Y) + [I_Q(X; Y) - R]_+\}$ is right-continuous at $t = \tau$ (or $g(\tau) = +\infty$). Similarly, the limit $-\frac{1}{n} \ln \beta_n(\tau, R_n) \rightarrow E_{\text{MD}}(\tau, R)$ exists whenever $E_{\text{MD}}(\tau, R)$ equals the same infimum with $\Delta_{<}$ in place of Δ in (26), i.e., whenever the infimum in (15) is not attained on the boundary $\lambda(Q, R) = \tau$ or $\Delta(Q_Y, R) = \tau$.*

In particular, the limits exist for all τ strictly inside a phase region (see (29)–(33)), independently of any property of the channel W .

Remark 3 (Sensitivity at phase boundaries and M_n subexponential factors). *The M_n sensitivity and the phase-boundary condition are two faces of the same phenomenon.*

At τ strictly inside a phase region, the optimal type Q^ satisfies $\lambda(Q^*, R) > \tau + \eta$ for some $\eta > 0$ (FA case). The threshold comparison $\{\Lambda(y^n) \geq \tau\}$ is decided with an exponential gap of size η . Replacing $M_n = \lfloor k_1 e^{nR} \rfloor$ by $\lfloor k_2 e^{nR} \rfloor$ shifts the LLR by $\frac{1}{n} \ln(k_2/k_1) = o(1)$, which is subexponential and is dominated by the gap η . Hence the limit is independent of the prefactor k .*

At a phase boundary (i.e., $\tau = \lambda_{\max}(R)$, $\tau = \tau^(R)$, or $\tau = \lambda_{\min}(R)$; see (29)–(33)), the relevant exponent function is discontinuous at τ , and the optimal type has $\lambda(Q^*, R) = \tau$ exactly. The threshold comparison is at the knife's edge: a subexponential shift of $\frac{1}{n} \ln k$ is decisive, and the limit depends on k . For the Z-channel at $\tau = \lambda_{\max}(R)$ with $M_n = \lfloor k e^{nR} \rfloor$, a direct calculation gives FA exponent $D_m(0) + k(\ln 2 - R)$, varying with k (see Section 5).*

Thus: the M_n sensitivity at phase boundaries is not a separate issue requiring a separate assumption on M_n . The single condition — τ strictly inside a phase region — is both necessary and sufficient for the limits to exist and be independent of the subexponential factors in M_n .

Remark 4 (MD for $\tau < 0$). *For $\tau < 0$, missed detection requires the transmitted codeword to have a rare joint type, which is an unlikely event whose probability decays exponentially according to $D_c(Q_{XY})$. Given such an atypical transmitted pair, the interfering codewords are too weak to push the likelihood ratio above the (negative) threshold, and so, missed detection occurs with probability $\doteq 1$ conditional on this rare type. Hence $E_{\text{MD}}(\tau, R)$ is finite and positive for $\tau < 0$ (when $R < I(X; Y)$), governed by the same formula (26) as for $\tau > 0$. As $\tau \rightarrow -\infty$: the feasible set eventually empties and $E_{\text{MD}}(\tau, R) \rightarrow +\infty$.*

Remark 5 ($R = 0$: one codeword). *For $R = 0$ ($M_n = 1$), the problem reduces to simple hypothesis testing of $W^n(\cdot | x^n(1))$ vs. $P_Y^{\otimes n}$ averaged over $x^n(1) \sim \text{Uniform}(\mathcal{T}(P_X))$. With a single codeword, $S(y^n) = e^{-n\ell(Q_{XY})}$ is deterministic, and both formulas (23)–(26) hold with $R = 0$ (the Δ constraint is*

vacuous since there are no interfering codewords). In particular, $E_{MD}(\tau, 0)$ is finite for all $\tau \in \mathbb{R}$ with no discontinuity, in contrast to $R > 0$.

Remark 6 (Self-averaging behavior and mismatched detection). *The exponents $E_{FA}(\tau, R)$ and $E_{MD}(\tau, R)$ characterize not only the annealed (average-over-codebook) error probabilities, but also the almost-sure (quenched) behavior at every τ strictly inside a phase region: the companion paper [12] establishes that $-\frac{1}{n} \ln \alpha_n(\tau, R_n) \rightarrow E_{FA}(\tau, R)$ and $-\frac{1}{n} \ln \beta_n(\tau, R_n) \rightarrow E_{MD}(\tau, R)$ almost surely over the random codebook, and develops the statistical-physics perspective of these results. The companion paper [12] also treats the mismatched setting in which the detector employs a surrogate channel $\tilde{W} \neq W$, deriving the self-averaging exponents for that case.*

4 Properties and Phase Structure

The single-letter functions $E_{FA}(\tau, R)$ and $E_{MD}(\tau, R)$ are the exponent candidates whose structure we now describe. By Corollary 1, at every τ strictly inside a phase region (i.e., $\tau \neq \lambda_{\max}(R)$ for FA, and $\tau \neq \tau^*(R)$, $\tau \neq \lambda_{\min}(R)$ for MD), the limits exist and the formulas below give the true error exponents. No condition on the channel W is needed. We begin with an elementary observation that follows directly from Theorem 1.

Corollary 2 (Soft covering). *For all $R \geq I(X; Y)$: $E_{FA}(0, R) = E_{MD}(0, R) = 0$.*

To see why this is true, observe that the type $Q_{XY} = P_X \otimes W$ has $D_m(Q_Y) = 0$, $I_Q(X; Y) = I(X; Y) \leq R$, $D_c(Q_{XY}) = 0$, and $\lambda(Q_{XY}, R) = 0$, so it is feasible for $\tau = 0$ with zero cost, hence $E_{FA}(0, R) = 0$. Likewise, $E_{MD}(0, R) = 0$ since types arbitrarily close to $P_X \otimes W$ with $\lambda(Q_{XY}, R) < 0$ and $D_c(Q_{XY}) \rightarrow 0$ exist by continuity.

This is the Neyman–Pearson implication of the soft covering lemma [1]: at $\tau = 0$, both error exponents simultaneously vanish when $R \geq I(X; Y)$, meaning the two hypotheses are exponentially indistinguishable. In other words, even if the FA and MD probabilities decay as $n \rightarrow \infty$, in this case, the rates of their decay are definitely slower than exponential.

4.1 Phase transitions

Both exponents have several regions of behavior as functions of τ . The following two propositions describe them.

Proposition 1 (FA: flat and active regions). *Let Q_{XY}^* , $\tau_{\text{flat}}(R)$, and $\lambda_{\max}(R)$ be as defined in (32) and (29). $E_{FA}(\tau, R)$ has three regions:*

- (i) Flat: $\tau \leq \tau_{\text{flat}}(R)$: $E_{FA}(\tau, R) = D_m(Q_{XY}^*) + [I_Q^*(X; Y) - R]_+$ (constant in τ).
- (ii) Active: $\tau \in (\tau_{\text{flat}}(R), \lambda_{\max}(R))$: $E_{FA}(\tau, R)$ strictly increases in τ .
- (iii) Infinite: $\tau > \lambda_{\max}(R)$: $E_{FA}(\tau, R) = +\infty$.

Proof. (i) For $\tau \leq \tau_{\text{flat}}(R) = \lambda(Q_{XY}^*, R)$, the unconstrained minimizer Q_{XY}^* satisfies $\lambda(Q_{XY}^*, R) \geq \tau$ and hence is feasible. Since it achieves the global minimum, the constrained minimum equals the unconstrained one; $E_{\text{FA}}(\tau, R)$ is flat. (ii) For $\tau > \tau_{\text{flat}}(R)$, Q_{XY}^* is infeasible. The constrained minimum strictly exceeds the unconstrained one and strictly increases with τ as the feasible set shrinks. (iii) For $\tau > \lambda_{\text{max}}(R)$, every type has $\lambda(Q_{XY}, R) \leq \lambda_{\text{max}}(R) < \tau$, so the feasible set is empty and $E_{\text{FA}}(\tau, R) = +\infty$ by convention. $\square$

Proposition 2 (MD: zero, active, and divergent regions). *Let $\tau^*(R)$ and $\lambda_{\text{min}}(R)$ be as in (31) and (30). For $R > 0$, $E_{\text{MD}}(\tau, R)$ has three regions:*

- (i) Zero: $\tau \geq \tau^*(R)$: $E_{\text{MD}}(\tau, R) = 0$.
- (ii) Active: $\tau \in (\lambda_{\text{min}}(R), \tau^*(R))$: $E_{\text{MD}}(\tau, R)$ is finite, positive, and strictly decreasing in τ .
- (iii) Infinite: $\tau \leq \lambda_{\text{min}}(R)$: $E_{\text{MD}}(\tau, R) = +\infty$ (the feasible set $\{\lambda(Q_{XY}, R) < \tau\}$ is empty).

Within the active region, there is a kink at $\tau_{\text{kink}}(R)$ (defined in (33)) where the minimizing type transitions from bulk ($I_{Q^}(X; Y) \leq R$) to sparse ($I_{Q^*}(X; Y) > R$).*

Proof. (i) For $\tau \geq \tau^*(R)$, types near $P_X \otimes W$ (with $D_c(Q_{XY}) \rightarrow 0$ and $\lambda(Q_{XY}, R) \rightarrow \tau^*(R)^-$) are feasible, giving $\inf D_c(Q_{XY}) = 0$ and $E_{\text{MD}}(\tau, R) = 0$. (ii) Monotonicity in τ is immediate: the feasible set $\{\lambda(Q_{XY}, R) < \tau\}$ grows as τ increases, so $E_{\text{MD}}(\tau, R)$ is non-increasing. Strict decrease and positivity follow from the formula (26). (iii) For $\tau \leq \lambda_{\text{min}}(R)$, every type has $\lambda(Q_{XY}, R) \geq \lambda_{\text{min}}(R) \geq \tau$, so the feasible set is empty and $E_{\text{MD}}(\tau, R) = +\infty$. $\square$

4.2 The Neyman–Pearson tradeoff zone

The results above combine into a clean picture of the Neyman–Pearson tradeoff (Corollary 3 and Proposition 2).

Corollary 3. *Let $\tau^*(R)$ be as in (31).*

- (i) *For $R < I(X; Y)$: $E_{\text{MD}}(\tau, R) > 0$ for all $\tau < \tau^*(R)$, and $E_{\text{MD}}(\tau, R) = 0$ for all $\tau \geq \tau^*(R)$. Both $E_{\text{FA}}(\tau, R) > 0$ and $E_{\text{MD}}(\tau, R) > 0$ simultaneously for all $\tau \in (0, \tau^*(R))$.*
- (ii) *For $R \geq I(X; Y)$: $\tau^*(R) = 0$, so $E_{\text{MD}}(\tau, R) = 0$ for all $\tau \geq 0$ and $E_{\text{FA}}(\tau, R) = 0$ for all $\tau \leq 0$. Both exponents simultaneously vanish at $\tau = 0$ (Corollary 2), and the Neyman–Pearson tradeoff zone is empty.*

Proof. $E_{\text{MD}}(\tau, R) = 0$ iff $\tau \geq \tau^*(R)$. The only type with $D_c(Q_{XY}) = 0$ is $Q_{XY} = P_X \otimes W$, which has $\lambda(P_X \otimes W, R) = \tau^*(R)$. For $\tau \geq \tau^*(R)$, types approaching $P_X \otimes W$ have $D_c(Q_{XY}) \rightarrow 0$ and $\lambda \rightarrow \tau^*(R)^- < \tau$, so they are feasible with $D_c(Q_{XY}) \rightarrow 0$, giving $E_{\text{MD}}(\tau, R) = 0$. For $\tau < \tau^*(R)$, every type with small $D_c(Q_{XY})$ has $Q_{Y|X}$ close to W and hence $\lambda(Q_{XY}, R)$ close to $\tau^*(R) > \tau$, so it is not feasible; thus $\inf_{\lambda(Q_{XY}, R) < \tau} D_c(Q_{XY}) > 0$ and $E_{\text{MD}}(\tau, R) > 0$.

$E_{\text{FA}}(\tau, R) > 0$ for all $\tau > 0$. A zero-cost type for $E_{\text{FA}}(\tau, R)$ requires $Q_Y = P_Y$ and $I_Q(X; Y) \leq R$. Any such type has $\lambda(Q_{XY}, R) = -D_c(Q_{XY}) \leq 0$, so it is feasible for $E_{\text{FA}}(\tau, R)$ (i.e. $\lambda(Q_{XY}, R) \geq \tau$) only if $\tau \leq 0$. Hence for $\tau > 0$ no zero-cost type is feasible and $E_{\text{FA}}(\tau, R) > 0$.

Conclusion. For $R < I(X; Y)$: both $E_{\text{FA}}(\tau, R) > 0$ (since $\tau > 0$) and $E_{\text{MD}}(\tau, R) > 0$ (since $\tau < \tau^*(R)$) hold simultaneously iff $\tau \in (0, \tau^*(R))$. For $R \geq I(X; Y)$: $\tau^*(R) = 0$, so $E_{\text{MD}}(\tau, R) = 0$ for all $\tau \geq 0$ and $E_{\text{FA}}(\tau, R) = 0$ for all $\tau \leq 0$ (since $P_X \otimes W$ has $D_c(Q_{XY}) = 0$, $Q_Y = P_Y$, $I_Q(X; Y) = I(X; Y) \leq R$, giving a feasible zero-cost type for $E_{\text{FA}}(\tau, R)$ whenever $\tau \leq 0$). Thus the two zero regions cover all of $\mathbb{R}$ and the tradeoff zone is empty. $\square$

5 Numerical Illustration: Z-Channel

In this section, we demonstrate the behavior of both exponents for a numerical example of a Z-channel. We verify that our numerical results are at parameter values where the limits in Theorem 1 exist (i.e., τ strictly inside a phase region; see Corollary 1).

Let $\mathcal{X} = \mathcal{Y} = \{0, 1\}$ and consider the Z-channel with the following input-output transition probabilities:

$$W(0|0) = 1, \quad (34)$$

$$W(1|0) = 0, \quad (35)$$

$$W(0|1) = w, \quad \text{and} \quad (36)$$

$$W(1|1) = 1 - w, \quad (37)$$

with $w = 0.45$ and $P_X(0) = P_X(1) = 0.5$. The corresponding mutual information is $I(X; Y) = 0.2441$ nats per channel use. Joint types are parameterized by $q = Q(0|1) \in (0, 1)$, with $Q(y|0) = W(y|0)$ forced. Defining the binary entropy function and the binary KL divergence as

$$h_b(u) = -u \ln u - (1 - u) \ln(1 - u) \quad (38)$$

$$D_b(u\|v) = u \ln \frac{u}{v} + (1 - u) \ln \frac{1 - u}{1 - v}, \quad (39)$$

where $(u, v) \in [0, 1]^2$, this gives

$$I_Q(X; Y) = h_b\left(\frac{1+q}{2}\right) - \frac{1}{2}h_b(q), \quad (40)$$

$$D_c(Q_{XY}) = \frac{1}{2}D_b(q\|w), \quad \text{and} \quad (41)$$

$$D_m(Q_Y) = D_b\left(\frac{1+q}{2}\left\|\frac{1+w}{2}\right.\right). \quad (42)$$

Validity of limits for the numerical results

By Corollary 1, the FA exponent limit exists and equals $E_{\text{FA}}(\tau, R)$ at all $\tau \neq \lambda_{\text{max}}(R)$. For the Z-channel, a direct computation shows that $\lambda_{\text{max}}(R) = \lambda(0, R) = \ln 2 - \frac{1}{2} \ln(1 + w) - R$ is achieved uniquely at $q = 0$, and all FA curves (Figs. 1, 2, and 6) are plotted at τ values strictly below $\lambda_{\text{max}}(R)$

(i.e., within the flat or active regions, not at the singular boundary). Similarly, the MD exponent limit exists at all $\tau \neq \tau^*(R) = [I(X; Y) - R]_+$ and $\tau \neq \lambda_{\min}(R)$, and all MD curves (Figs. 3, 4, 5, and 7) are plotted at τ values strictly inside the active region. Hence all the numerical results displayed below are valid exact exponents.

Seven figures follow, one per page, and each one contains a brief description and discussion, in addition to the figure caption.




Figure 1: **FA exponent** $E_{\text{FA}}(\tau, R)$ vs. τ for four rates.

In Fig. 1, for each rate $R \in \{0.05, 0.10, 0.15, 0.20\}$ (all below $I(X; Y) = 0.2441$), $E_{\text{FA}}(\tau, R)$ is flat at $E_{\text{FA,flat}}(R) = D_m(Q^*) + [I_{Q^*}(X; Y) - R]_+ > 0$ for $\tau \leq \tau_{\text{flat}}(R)$ (filled dots, dotted verticals), then strictly increases for $\tau > \tau_{\text{flat}}(R)$, diverging to $+\infty$ at $\lambda_{\text{max}}(R)$ (dashed verticals). $E_{\text{FA}}(\tau, R)$ is continuous throughout. The flat value $E_{\text{FA,flat}}(R) > 0$ (rather than 0) is unique to singular channels like the Z-channel with $W(1|0) = 0$: a typical output y^n (which has many 1's) can only be explained by codewords that have '1' at every position where y^n has '1', and with only M_n codewords drawn from $\mathcal{T}(P_X)$, the probability that even a single such compatible codeword exists is exponentially small. Hence $P_{Y^n|C}(y^n) = 0$ with high probability over the ensemble of codes, $\Lambda(y^n) = -\infty$, and FA is exponentially rare for *any* finite τ , even $\tau \rightarrow -\infty$. See Figure 2 for a zoom on the active region.




Figure 2: **Zoom into the active region of $E_{\text{FA}}(\tau, R)$.** Normalized axes: $x = 0$ at $\tau_{\text{flat}}(R)$, $x = 1$ at $\lambda_{\text{max}}(R)$. Only $R = 0.05$ (width ≈ 0.032 nats) is visible.

Fig. 2 provides a zoom into the active region $\tau \in (\tau_{\text{flat}}(R), \lambda_{\text{max}}(R))$ for $R = 0.05$. At $\tau_{\text{flat}}(R) \approx 0.033$ (filled dot): $E_{\text{FA}}(\tau, R)$ lifts off its flat value ≈ 0.111 and begins to increase. At $\lambda_{\text{max}}(R) \approx 0.457$ (dashed vertical): $E_{\text{FA}}(\tau, R) \rightarrow +\infty$. $E_{\text{FA}}(\tau, R)$ is strictly convex and monotone increasing throughout the active region.




Figure 3: **MD exponent** $E_{\text{MD}}(\tau, R)$ vs. τ for four rates $R \in \{0.05, 0.10, 0.15, 0.20\}$. Filled circles: $\tau^*(R) = [I(X; Y) - R]_+$ (onset of $E_{\text{MD}}(\tau, R) = 0$).

In Fig. 3 we see the MD analogue of Fig. 1. $E_{\text{MD}}(\tau, R)$ is monotone non-increasing in τ for each one of four rates as before. The curves start from $\tau \approx \lambda_{\min} \approx -0.078$ nats (minimum achievable $\lambda(Q_{XY}, R)$; $E_{\text{MD}}(\tau, R) = +\infty$ for smaller τ).

Common branch. For τ sufficiently negative (below the leftmost kink at $\tau_{\text{kink}}(R = 0.05) \approx -0.047$), all four curves are *exactly identical*: the minimizing type $Q^*(\tau)$ has $I_{Q^*}(X; Y) < R$ for all four rates, so the rate constraint is inactive and $\lambda(Q_{XY}^*, R) = D_{\text{m}}(Q_Y^*) - D_{\text{c}}(Q_{XY}^*)$ is the same for all R .

Sequential splitting. As τ increases, the curves split one by one at the kink points $\tau_{\text{kink}}(R)$ (filled squares), where $I_{Q^*}(X; Y) = R$ exactly and the minimizing type transitions from bulk to sparse. This is a first-order phase transition: the slope changes discontinuously, with the sparse branch (higher rate of decrease) taking over above the kink.

For $0 < \tau < \tau^*(R) = [I(X; Y) - R]_+$: $E_{\text{MD}}(\tau, R) > 0$. At $\tau = \tau^*(R)$ (filled circles, dotted verticals): $E_{\text{MD}}(\tau, R) = 0$.




Figure 4: **Zoom into the active region of $E_{\text{MD}}(\tau, R)$, $R = 0.05$.** The kink at $\tau_{\text{kink}}(R) \approx -0.047$ (filled square, dotted vertical) marks the transition from the common bulk branch (left) to the sparse branch (right, rate-dependent). Filled circle: $\tau^*(R) = I(X; Y) - R \approx 0.194$ (onset of $E_{\text{MD}}(\tau, R) = 0$), dashed vertical.

Fig. 4 provides a zoom-in picture on $E_{\text{MD}}(\tau, R)$ for $R = 0.05$, revealing the full structure of the curve from $\lambda_{\min}(R)$ to $\tau^*(R)$. Left of the kink $\tau_{\text{kink}} \approx -0.047$: minimizing type is bulk, common to all rates. Right of the kink: minimizer becomes sparse, curves split by rate. $E_{\text{MD}}(\tau, R)$ is monotone non-increasing throughout, reaching zero at $\tau^*(R)$.




Figure 5: **Neyman–Pearson tradeoff curve:** $E_{\text{MD}}(\tau, R)$ vs. $E_{\text{FA}}(\tau, R)$, parametrized by τ (τ increasing: $E_{\text{FA}}(\tau, R) \nearrow$, $E_{\text{MD}}(\tau, R) \searrow$). **Left:** raw parametric curve; vertical segments arise because $E_{\text{FA}}(\tau, R)$ is flat while $E_{\text{MD}}(\tau, R)$ decreases. Triangles: top of each vertical. **Right:** upper-envelope curve (each flat segment collapsed to its highest point).

Fig. 5 gives the parametric curve $(E_{\text{FA}}(\tau, R), E_{\text{MD}}(\tau, R))$ as τ sweeps the tradeoff range. It is shown for four rates. Each curve has a vertical segment where $E_{\text{FA}}(\tau, R)$ is flat (Proposition 1) while $E_{\text{MD}}(\tau, R)$ decreases: during this portion there is no genuine tradeoff, only a loss in MD performance at fixed $E_{\text{FA}}(\tau, R)$. The right panel collapses each vertical segment to its highest point, revealing the genuine Neyman–Pearson tradeoff: to achieve larger $E_{\text{FA}}(\tau, R)$ one must sacrifice $E_{\text{MD}}(\tau, R)$.




Figure 6: **FA exponent** $E_{\text{FA}}(\tau, R)$: phase diagram in the (τ, R) plane (Proposition 1). Region **III** (left of blue curve, $\tau \leq \tau_{\text{flat}}(R)$): E_{FA} is flat in τ (for fixed R) but varies with R . Region **II** (between blue and cyan curves, $\tau_{\text{flat}}(R) < \tau < \lambda_{\text{max}}(R)$): E_{FA} strictly increasing. Region **I** (grey, right of cyan curve, $\tau > \lambda_{\text{max}}(R)$): $E_{\text{FA}} = +\infty$. Blue curve: $\tau_{\text{flat}}(R)$; cyan curve: $\lambda_{\text{max}}(R)$. Blue square: cusp in $\tau_{\text{flat}}(R)$ at $R \approx 0.106$ nats. White star: soft-covering point $(0, I(X; Y))$, at the boundary between Regions II and III.

Fig. 6 displays the three regions of Proposition 1 in the (τ, R) plane. Moving from left to right in τ : Region III (left of the blue curve $\tau_{\text{flat}}(R)$, where E_{FA} is flat in τ for fixed R but varies with R) transitions into Region II (strictly increasing E_{FA} , between the blue and cyan curves), and finally into Region I ($E_{\text{FA}} = +\infty$, grey, to the right of the cyan curve $\lambda_{\text{max}}(R)$). The white star marks the soft-covering point $(\tau, R) = (0, I(X; Y))$; it lies at the boundary between Regions II and III (since $\tau_{\text{flat}}(R) \rightarrow 0$ as $R \nearrow I(X; Y)$) and corresponds to $E_{\text{FA}} = E_{\text{FA,flat}}(R) > 0$, not zero. The blue square marks the cusp in $\tau_{\text{flat}}(R)$ at $R = R_{\text{cr}} \approx 0.106$ nats, where the unconstrained minimizer Q_{XY}^* transitions from a sparse type ($I_{Q^*} > R$, slope -1) to a bulk type, creating the visible kink in the blue curve.




Figure 7: **MD exponent** $E_{\text{MD}}(\tau, R)$: phase diagram. Shaded region ($\tau \leq 0$): $E_{\text{MD}}(\tau, R)$ finite and positive when the feasible set $\{\lambda(Q_{XY}, R) < \tau\}$ is non-empty (Remark 4); $E_{\text{MD}}(\tau, R) = +\infty$ otherwise. Colored region ($\tau > 0$): $E_{\text{MD}}(\tau, R)$ finite, with sharp transition to 0 at $\tau^*(R) = [I(X; Y) - R]_+$ (red line). White dashed line: $\tau = 0$. White star: soft-covering point $(0, I(X; Y))$.

In Fig. 7 the shaded region ($\tau \leq 0$) has finite $E_{\text{MD}}(\tau, R)$ given by $\inf_{\lambda(Q_{XY}, R) < \tau} D_c(Q_{XY})$ when the feasible set is non-empty, and $E_{\text{MD}}(\tau, R) = +\infty$ otherwise (Remark 4). The white dashed line at $\tau = 0$ marks $E_{\text{MD}}(\tau, R) = 0$. The colored region ($\tau > 0$) shows finite $E_{\text{MD}}(\tau, R)$, with a sharp transition to 0 at $\tau^*(R) = [I(X; Y) - R]_+$ (red line, slope -1). Together, the two phase diagrams show that the Neyman–Pearson tradeoff zone $0 < \tau < \tau^*(R)$ shrinks as $R \nearrow I(X; Y)$ and vanishes at $R = I(X; Y)$, the soft-covering threshold.

6 Conclusion and Outlook

We studied the soft covering lemma through the lens of Neyman–Pearson hypothesis testing, asking: can an observer distinguish the random codebook mixture $P_{Y^n|C}$ from the i.i.d. distribution $P_Y^{\otimes n}$? We derived tight single-letter liminf/limsup exponent envelopes for the FA and MD probabilities $\alpha_n(\tau, R_n)$ and $\beta_n(\tau, R_n)$, for all rates $R \geq 0$ and all thresholds $\tau \in \mathbb{R}$, for all τ strictly inside a phase region. The necessary and sufficient condition under which the limits exist and equal $E_{\text{FA}}(\tau, R)$ and $E_{\text{MD}}(\tau, R)$ is identified in Corollary 1: τ strictly inside a phase region, with no condition on W . The main findings are:

1. *Soft covering as a phase transition.* At $R = I(X; Y)$ and $\tau = 0$, both exponents simultaneously

vanish — the Neyman–Pearson exponent characterization of the soft covering phenomenon. For $R < I(X; Y)$, a genuine Neyman–Pearson tradeoff interval $(0, I(X; Y) - R)$ exists where both exponents are positive. For $R \geq I(X; Y)$, the two zero regions cover all of $\mathbb{R}$, leaving no tradeoff interval.

2. *Rich phase structure.* The FA exponent $E_{\text{FA}}(\tau, R)$ is flat (constant in τ) below $\tau_{\text{flat}}(R)$ and strictly increasing above it, diverging at $\lambda_{\text{max}}(R)$. The MD exponent $E_{\text{MD}}(\tau, R)$ is strictly decreasing in τ , with a kink at $\tau_{\text{kink}}(R) < 0$ (a first-order phase transition where the minimizing type changes from bulk to sparse), and reaches zero at $\tau^*(R) = [I(X; Y) - R]_+$.
3. *Structural asymmetry.* The FA exponent penalizes both the empirical output deviation from P_Y and the rate surplus $[I_Q(X; Y) - R]_+$. The MD exponent penalizes only the channel divergence $D_c(Q_{XY})$, with an additional constraint $\Delta(Q_Y, R) < \tau$ that identifies which transmitted types contribute — automatically satisfied for $\tau > 0$ by the DPI.

Several directions are open for future work.

1. *Compound channels.* Extending the present liminf/limsup envelopes to compound channels where W is unknown remains an open direction.
2. *Finite-block-length refinements.* The exponent results give the leading exponential term. Sub-exponential (polynomial) pre-factors, as in [2], would sharpen the approximation and may reveal further phase transitions at the second order.
3. *Other channels and source models.* The analysis here is for discrete memoryless channels. Extensions to continuous alphabet channels (most notably, Gaussian), channels with memory, or multi-terminal settings (broadcast, multiple access) would broaden the scope and may uncover new structural phenomena.

Appendix – Proof of Theorem 1

6.1 The FA exponent

We prove the liminf bound (23) and the limsup bound (24), for $R > 0$. (The case $R = 0$ follows by a direct type-counting argument; see Remark 5.)

Write

$$\alpha_n(\tau, R_n) = \sum_{y^n} P_Y^{\otimes n}(y^n) \cdot \Pr\{S(y^n) \geq e^{n\theta_n(y^n)}\}, \quad (\text{A.1})$$

where $\theta_n(y^n) = \tau + R_n - H_{Q_Y}(Y) - D_{\text{m}}(Q_Y)$ with $Q_Y = \hat{P}_{y^n}$ and $R_n = \frac{1}{n} \ln M_n$. Note that θ_n depends on y^n only through its type Q_Y .

For a fixed n -type $q \in \mathcal{P}_n(\mathcal{Y})$ and fixed $y^n \in \mathcal{T}(q)$, let $\mathcal{Q}_n(q) := \{Q \in \mathcal{P}_n(\mathcal{X} \times \mathcal{Y}) : Q_X = P_{X,n}, Q_Y = q, Q \ll P_X \otimes W\}$ be the set of admissible joint n -types with output marginal q . Using the decomposition $S(y^n) = \sum_{Q \in \mathcal{Q}_n(q)} N(Q|y^n) e^{-n\ell(Q)}$ and noting that $S(y^n) \geq e^{n\theta_n}$ implies

$N(Q|y^n)e^{-n\ell(Q)} \geq e^{n\theta_n}/|\mathcal{Q}_n(q)|$ for at least one Q (where $|\mathcal{Q}_n(q)| = \text{poly}(n)$), we have:

$$\Pr\{S(y^n) \geq e^{n\theta_n}\} \leq \sum_{Q \in \mathcal{Q}_n(q)} \Pr\{N(Q|y^n) \geq e^{n[\theta_n + \ell(Q) - o(1)]}\}, \quad (\text{A.2})$$

where the $o(1)$ absorbs the logarithm of $|\mathcal{Q}_n| = \text{poly}(n)$. Conversely, for any single $Q \in \mathcal{Q}_n(Q_Y)$:

$$\Pr\{S(y^n) \geq e^{n\theta_n}\} \geq \Pr\{N(Q|y^n)e^{-n\ell(Q)} \geq e^{n\theta_n}\} = \Pr\{N(Q|y^n) \geq e^{n[\theta_n + \ell(Q)]}\}. \quad (\text{A.3})$$

Using the identity

$$\begin{aligned} \theta_n + \ell(Q) &= \tau + R_n - H_Q(Y) - D_m(Q_Y) + H_Q(Y|X) + D_c(Q_{XY}) \\ &= \tau - (D_m(Q_Y) - D_c(Q_{XY})) + R_n - I_Q(X; Y), \end{aligned} \quad (\text{A.4})$$

we apply Theorem 4.1 of [2] (see (19)) to $N(Q|y^n) \sim \text{Binomial}(M_n, e^{-nI_Q(X; Y) + o(n)})$.

Liminf bound (23). Fix any $E < E_{\text{FA}}(\tau, R)$ and let $\eta > 0$ be such that $E < E_{\text{FA}}(\tau, R) - \eta$. By the compactness of $\mathcal{Q}$ and continuity of $\lambda(\cdot, R)$ and $F(\cdot, R) := D_m(Q_Y) + [I_Q(X; Y) - R]_+$ on $\mathcal{Q}$, there exists $\delta > 0$ such that: for all $Q \in \mathcal{Q}$ with $F(Q, R) \leq E$, we have $\lambda(Q, R) \leq \tau - \delta$. This implies that for all large n and all n -types $Q_n \rightarrow Q$: $\lambda_n(Q_n, R_n) \leq \tau - \delta/2$. By the identity (A.4), $\theta_n + \ell(Q_n) = \tau - (D_m(Q_n) - D_c(Q_n)) + R_n - I_{Q_n}(X; Y)$. We bound this from below in each case:

- *Bulk* ($I_{Q_n}(X; Y) \leq R_n$): $\lambda_n(Q_n, R_n) = D_m(Q_n) - D_c(Q_n) \leq \tau - \delta/2$, so $-(D_m - D_c) \geq -(\tau - \delta/2)$, and $R_n - I_{Q_n} \geq 0$. Hence $\theta_n + \ell(Q_n) \geq \tau - (\tau - \delta/2) + 0 = \delta/2$, and since $[R_n - I_{Q_n}]_+ \geq 0$, $\theta_n + \ell(Q_n) \geq \delta/2 \geq \delta/4 + [R_n - I_{Q_n}]_+$ for large n .
- *Sparse* ($I_{Q_n}(X; Y) > R_n$): $\lambda_n(Q_n, R_n) = D_m - D_c + I_{Q_n} - R_n \leq \tau - \delta/2$, so $-(D_m - D_c) + (R_n - I_{Q_n}) \geq -(\tau - \delta/2) - 0$. Hence $\theta_n + \ell(Q_n) = \tau - (D_m - D_c) + R_n - I_{Q_n} = \tau + [-(D_m - D_c) + (R_n - I_{Q_n})] \geq \tau - (\tau - \delta/2) = \delta/2$. Since $[R_n - I_{Q_n}]_+ = 0$ for sparse types, $\theta_n + \ell(Q_n) \geq \delta/2 = \delta/4 + \delta/4 \geq \delta/4 + [R_n - I_{Q_n}]_+$.

In both cases $\theta_n + \ell(Q_n) \geq \delta/4 + [R_n - I_{Q_n}]_+$ for large n . By (19), such types contribute with probability $\Pr\{N(Q_n|y^n) \geq e^{n[\theta_n + \ell(Q_n) - o(1)]}\}$ that is super-exponentially small. For the remaining types with $F(Q_n, R_n) > E$: $\Pr\{N(Q_n|y^n) \geq e^{n[\theta_n + \ell(Q_n)]}\} \leq e^{-n[I_Q(X; Y) - R_n]_+}$, and $P_Y^{\otimes n}(\mathcal{T}(Q_Y)) \leq e^{-nD_m(Q_Y)}$, so the contribution of each type to α_n is at most $e^{-nF(Q_n, R_n)} \leq e^{-nE}$. Combining, $\alpha_n(\tau, R_n) \leq \text{poly}(n) \cdot e^{-nE}$ for all large n , giving $\liminf_{n \rightarrow \infty} (-\frac{1}{n} \ln \alpha_n) \geq E$. Since $E < E_{\text{FA}}(\tau, R)$ is arbitrary, this proves (23).

Limsup bound (24). Fix $\eta > 0$ and let $Q^* \in \mathcal{Q}$ achieve $F(Q^*, R) = E_{\text{FA}}(\tau + \eta, R)$ (attained since $\mathcal{Q}$ is compact and F is continuous). Since $\lambda(Q^*, R) \geq \tau + \eta$, by the standard method-of-types approximation there exist n -types $Q_n^* \in \mathcal{Q}_n$ with $Q_n^* \rightarrow Q^*$ and $\lambda_n(Q_n^*, R_n) \geq \tau + \eta/2$ for all large n . By (A.4), $\theta_n + \ell(Q_n^*) \leq [R_n - I_{Q_n^*}]_+ - \eta/4$ for large n . By (A.3): $\Pr\{S(y^n) \geq e^{n\theta_n}\} \geq \Pr\{N(Q_n^*|y^n) \geq e^{n[\theta_n + \ell(Q_n^*)]}\} \geq e^{-n[I_{Q_n^*} - R_n]_+ + o(n)}$ by (19) (since the threshold exponent is $\leq -\eta/4 < 0$, giving upper-tail probability $\geq e^{-n[I_{Q_n^*} - R_n]_+ - o(n)}$). Averaging over $y^n \in \mathcal{T}((Q_n^*)_Y)$ via $P_Y^{\otimes n}$: $\alpha_n(\tau, R_n) \geq e^{-n[D_m((Q_n^*)_Y) + [I_{Q_n^*} - R_n]_+ + o(n)]} = e^{-nF(Q_n^*, R_n) + o(n)}$, and $F(Q_n^*, R_n) \rightarrow F(Q^*, R) = E_{\text{FA}}(\tau + \eta, R)$. Hence $\limsup_{n \rightarrow \infty} (-\frac{1}{n} \ln \alpha_n(\tau, R_n)) \leq E_{\text{FA}}(\tau + \eta, R)$. Since $\eta > 0$ is arbitrary, this proves (24).

6.2 The MD exponent

We prove the liminf bound (25) and the limsup bound (26) for $R > 0$. The case $R = 0$ follows by a direct argument (see Remark 5).

Owing to the symmetry of the random coding mechanism, we assume, without loss of generality that the transmitted codeword is $x^n(1)$. Under $\mathcal{H}_1$, the pair $(x^n(1), y^n)$ is drawn from $P_{X^n} \times W^n$. By the method of types:

$$\Pr\{(x^n(1), y^n) \in \mathcal{T}(Q_{XY})\} = e^{-n[D_c(Q_{XY}) + o(1)]} \quad (\text{A.5})$$

for any joint type Q_{XY} with $Q_X = P_X$. We condition on this type and treat the remaining $M_n - 1$ codewords $x^n(m)$, $m = 2, 3, \dots, M_n$ as random.

We next show that if $\lambda(Q_{XY}, R) \geq \tau$ then the MD event is virtually impossible. Given that $(x^n(1), y^n)$ has joint type Q_{XY} (so y^n has marginal type Q_Y), define $\theta := \tau + R_n - H_Q(Y) - D_m(Q_Y)$ (the threshold exponent for this type class, as in (A.4)). Decompose

$$S(y^n) = e^{-n\ell(Q_{XY})} + \tilde{S}(y^n), \quad (\text{A.6})$$

where $e^{-n\ell(Q_{XY})} = W^n(y^n|x^n(1))$ is the deterministic contribution of the transmitted codeword and $\tilde{S}(y^n) = \sum_{m=2}^{M_n} W^n(y^n|x^n(m))$ is the contribution of the remaining $M_n - 1$ interfering codewords. The MD event is equivalent to $S(y^n) < e^{n\theta}$. Once again, we treat bulk types and sparse types separately.

Considering sparse types first, observe that since $S(y^n) \geq e^{-n\ell(Q_{XY})}$, it suffices to show $-\ell(Q_{XY}) \geq \theta$. Expanding using the definitions of $\ell(Q_{XY})$ and θ :

$$\begin{aligned} -\ell(Q_{XY}) \geq \theta &\iff -H_Q(Y|X) - D_c(Q_{XY}) \geq \tau + R_n - H_Q(Y) - D_m(Q_Y) \\ &\iff D_m(Q_Y) - D_c(Q_{XY}) + I_Q(X; Y) - R_n \geq \tau \\ &\iff \lambda(Q_{XY}, R_n) \geq \tau, \end{aligned}$$

where the last equivalence uses $\lambda(Q_{XY}, R_n) = D_m(Q_Y) - D_c(Q_{XY}) + I_Q(X; Y) - R_n$ for sparse types ($I_Q(X; Y) > R_n$), and $\lambda(Q_{XY}, R_n) \geq \tau$ holds by assumption (since $\lambda(Q_{XY}, R_n) \rightarrow \lambda(Q_{XY}, R)$ and $R_n \rightarrow R$). So $S(y^n) \geq e^{n\theta}$ deterministically given $(x^n(1), y^n)$, and so, the MD event is impossible.

Passing on to bulk types, to show that the MD event is virtually impossible when $\lambda(Q_{XY}, R) \geq \tau$, we first treat the strict case $\lambda(Q_{XY}, R) > \tau$; the boundary $\lambda(Q_{XY}, R) = \tau$ is then handled by continuity of $\lambda(Q_{XY}, R)$ in Q_{XY} : types with λ at distance $\eta > 0$ above τ already show super-exponentially small MD probability, so taking $\eta \rightarrow 0$ covers the boundary. Consider the interfering codewords that share the *same* joint type Q_{XY} with y^n as the transmitted one, $x^n(1)$. Their contribution to $\tilde{S}(y^n)$ gives

$$S(y^n) \geq \tilde{S}(y^n) \geq \tilde{N}(Q_{XY}|y^n) \cdot e^{-n\ell(Q_{XY})}, \quad (\text{A.7})$$

where here $\tilde{N}(Q_{XY}|y^n) \sim \text{Binomial}(M_n - 1, e^{-nI_Q(X; Y)})$ (and $M_n - 1$ has the same exponential rate

e^{nR_n} as M_n , since $M_n - 1 \geq M_n/2$ for large n). Expanding $\theta + \ell(Q_{XY})$ from the definitions, we have

$$\begin{aligned}\theta + \ell(Q_{XY}) &= \tau + R_n - H_Q(Y) - D_m(Q_Y) + H_Q(Y|X) + D_c(Q_{XY}) \\ &= \tau - (D_m(Q_Y) - D_c(Q_{XY})) + R_n - I_Q(X; Y).\end{aligned}\tag{A.8}$$

Since $D_m(Q_Y) - D_c(Q_{XY}) = \lambda(Q_{XY}, R) \geq \tau$ and $R_n - I_Q(X; Y) \geq 0$, we get $\theta + \ell(Q_{XY}) \leq R_n - I_Q(X; Y)$, with strict inequality when $\lambda(Q_{XY}, R) > \tau$. By Theorem 4.1 of [2] (20), $\Pr\{\tilde{N}(Q_{XY}|y^n) < e^{n(\theta + \ell(Q_{XY}))}|x^n(1), y^n\}$ is doubly exponentially small whenever $\theta + \ell(Q_{XY}) < R_n - I_Q(X; Y)$, i.e. whenever $\lambda(Q_{XY}, R_n) > \tau$. Hence with doubly-exponentially-high probability (i.e., all but a doubly exponentially small event),

$$S(y^n) \geq \tilde{N}(Q_{XY}|y^n) \cdot e^{-n\ell(Q_{XY})} \geq e^{n[\theta + \ell(Q_{XY})]} \cdot e^{-n\ell(Q_{XY})} = e^{n\theta},\tag{A.9}$$

so the MD event probability decays doubly exponentially.

In both cases, $\lambda(Q_{XY}, R) \geq \tau$ implies $\Pr\{\text{MD}|x^n(1), y^n\}$ is doubly exponentially small. Hence only types Q_{XY} of $(x^n(1), y^n)$ with $\lambda(Q_{XY}, R) < \tau$ can contribute to the MD probability on the exponential scale. We next evaluate this contribution.

Now, fix a type Q_{XY} of $(x^n(1), y^n)$ with $\lambda(Q_{XY}, R) < \tau$. As mentioned earlier, the threshold $\theta = \tau + R_n - H_Q(Y) - D_m(Q_Y)$ depends on y^n only through its marginal Q_Y , and on τ and R_n , but not on any other property of Q_{XY} . For each possible joint type Q'_{XY} of an interfering codeword $x^n(m)$, $m = 2, 3, \dots, M_n$, with y^n , the enumerator $\tilde{N}(Q'_{XY}|y^n) \sim \text{Binomial}(M_n - 1, e^{-nI_{Q'}(X; Y)})$. The event $\{\tilde{S}(y^n) < e^{n\theta}\}$ is equivalent to the event that $\tilde{N}(Q'_{XY}|y^n) \leq e^{n[\theta + \ell(Q'_{XY})]}$ for all Q'_{XY} simultaneously. To assess the probability of this intersection probability, we invoke Theorem 4.3 of [2] (see (21)). To this end, we compute $\theta + \ell(Q'_{XY})$ for each type Q'_{XY} of an interfering codeword. Since all interfering types Q'_{XY} share the marginal Q_Y of y^n , we have $D_m(Q'_Y) = D_m(Q_Y)$, and so,

$$\begin{aligned}\theta + \ell(Q'_{XY}) &= \tau + R_n - H_Q(Y) - D_m(Q_Y) + H_{Q'}(Y|X) + D_c(Q'_{XY}) \\ &= \tau - [D_m(Q'_Y) - D_c(Q'_{XY})] + R_n - I_{Q'}(X; Y).\end{aligned}\tag{A.10}$$

Note that this depends on Q'_{XY} of the interfering codewords, and on τ , but *not* on the type Q_{XY} of $(x^n(1), y^n)$ beyond its marginal Q_Y , which is already absorbed into $D_m(Q'_Y) = D_m(Q_Y)$. Upon applying Theorem 4.3 of [2], we obtain

$$\Pr\left[\bigcap_{Q'_{XY}} \{\tilde{N}(Q'_{XY}|y^n) \leq e^{n[\theta + \ell(Q'_{XY})]}\middle| x^n(1), y^n\}\right] \doteq \mathbf{1}\left\{\min_{Q'_{XY}} [I_{Q'}(X; Y) - R + [\theta + \ell(Q'_{XY})]_+] > 0\right\}.\tag{A.11}$$

Let us define then

$$f(Q'_{XY}) = I_{Q'}(X; Y) - R + [\theta + \ell(Q'_{XY})]_+.\tag{A.12}$$

Using the simple identity, $a + [b - a]_+ \equiv \max\{a, b\}$ with $a = I_{Q'}(X; Y) - R$ and $b = \tau - [D_m(Q_Y) - D_c(Q'_{XY})]$, and (A.10), we may rewrite the function f as

$$f(Q'_{XY}) = \max\{I_{Q'}(X; Y) - R, \tau - [D_m(Q_Y) - D_c(Q'_{XY})]\}.\tag{A.13}$$

We next simplify $f(Q'_{XY})$. For sparse interfering types ($I_{Q'}(X;Y) > R$), $f(Q'_{XY}) \geq I_{Q'}(X;Y) - R > 0$ always. Hence sparse types never ‘threaten’ the occurrence of the event $\{\min_{Q'_{XY}} f(Q'_{XY}) > 0\}$ and can be excluded. For bulk interfering types, $I_{Q'}(X;Y) - R \leq 0$, and so, the occurrence of the event $\{\min_{Q'_{XY}} f(Q'_{XY}) > 0\}$ depends on the sign of $\tau - [D_m(Q_Y) - D_c(Q'_{XY})]$. Thus the condition $\min_{Q'_{XY}} f(Q'_{XY}) > 0$ over all Q'_{XY} reduces to

$$\max_{\substack{Q'_{XY}: Q'_X = P_X, Q'_Y = Q_Y \\ I_{Q'}(X;Y) \leq R}} [D_m(Q_Y) - D_c(Q'_{XY})] < \tau, \quad (\text{A.14})$$

which is exactly the condition $\Delta(Q_Y, R) < \tau$ as defined in the theorem.

For $\tau > 0$, the DPI (7) gives $D_c(Q'_{XY}) \geq D_m(Q'_Y) = D_m(Q_Y)$, and so, $D_m(Q_Y) - D_c(Q'_{XY}) \leq 0 < \tau$, which means that the condition $\Delta(Q_Y, R) < \tau$ holds automatically for every Q_{XY} , and so, the corresponding constraint is slack. Thus $\Pr\{\tilde{S}(y^n) < e^{n\theta}\} \rightarrow 1$ super-exponentially fast (doubly exponentially) for every feasible transmitted type. On the other hand, for $\tau \leq 0$, $D_m(Q_Y) - D_c(Q'_{XY})$ may be larger than or equal to τ , making $\Delta(Q_Y, R) \geq \tau$ possible, in which case the indicator in (A.11) equals 0 and then $\Pr\{\tilde{S}(y^n) < e^{n\theta} | x^n(1), y^n\}$ is doubly exponentially small.

The conclusion is that for $\tau > 0$, $\Pr\{\tilde{S}(y^n) < e^{n\theta} | x^n(1), y^n\} \rightarrow 1$ super-exponentially for every transmitted type Q_{XY} with $\lambda(Q_{XY}, R) < \tau$, while the MD event is doubly exponentially unlikely when $\lambda(Q_{XY}, R) \geq \tau$. This heuristically gives (at the exponential scale):

$$\beta_n(\tau, R_n) \asymp \sum_{\substack{Q \in \mathcal{Q}_n: \lambda(Q, R_n) < \tau \\ \Delta(Q_Y, R_n) < \tau}} e^{-n[D_c(Q_{XY}) + o(1)]}, \quad (\text{A.15})$$

suggesting the exponent $E_{\text{MD}}(\tau, R)$. The rigorous liminf/limsup argument follows. For the *limsup bound* (upper envelope), we show

$$\limsup_{n \rightarrow \infty} \left(-\frac{1}{n} \ln \beta_n(\tau, R_n) \right) \leq E_{\text{MD}}(\tau, R).$$

Fix any $E > E_{\text{MD}}(\tau, R)$; we exhibit a sequence of transmitted types contributing to β_n with cost $\leq E$. Since $E_{\text{MD}}(\tau, R) = \inf_{\lambda(Q, R) < \tau, \Delta(Q_Y, R) < \tau} D_c(Q_{XY})$, there exists $Q^* \in \mathcal{Q}$ with $\lambda(Q^*, R) < \tau$, $\Delta(Q_Y^*, R) < \tau$, and $D_c(Q_{XY}^*) < E$. (If the infimum is not attained, take a sequence approaching it; the argument applies to any such Q^* with $D_c(Q_{XY}^*) < E$.) Let $\eta := \min(\tau - \lambda(Q^*, R), \tau - \Delta(Q_Y^*, R))/2 > 0$. By the standard method of types, there exist n -types $Q_n^* \in \mathcal{Q}_n$ with $Q_n^* \rightarrow Q^*$, and for large n : $\lambda_n(Q_n^*, R_n) < \tau - \eta/2$ and $\Delta_n((Q_n^*)_Y, R_n) < \tau - \eta/2$ (using continuity of λ and Δ on the compact set $\mathcal{Q}$). With this strict gap, Theorem 4.3 of [2] applied to the interferer counts gives $\Pr\{\tilde{S}(y^n) < e^{n\theta} | Q_n^*\} \rightarrow 1$ super-exponentially fast, so the transmitted codeword of type Q_n^* contributes to β_n with cost $D_c((Q_n^*)_{XY}) \rightarrow D_c(Q_{XY}^*) < E$. The contribution satisfies $\beta_n(\tau, R_n) \geq e^{-nD_c((Q_n^*)_{XY}) + o(n)}$, giving $\limsup_n (-\frac{1}{n} \ln \beta_n(\tau, R_n)) \leq D_c(Q_{XY}^*) < E$. Since $E > E_{\text{MD}}(\tau, R)$ is arbitrary, $\limsup_n (-\frac{1}{n} \ln \beta_n(\tau, R_n)) \leq E_{\text{MD}}(\tau, R)$. For the *liminf bound* (25), we show $\liminf_n (-\frac{1}{n} \ln \beta_n) \geq E_{\text{MD}}^\eta(\tau, R)$ for every fixed $\eta > 0$. Fix any $E < E_{\text{MD}}^\eta(\tau, R)$. We claim $\beta_n(\tau, R_n) \leq \text{poly}(n) \cdot e^{-nE}$ for all large n , establishing $\liminf_n (-\frac{1}{n} \ln \beta_n(\tau, R_n)) \geq E$.

By the definition (25), every $Q \in \mathcal{Q}$ with $\lambda(Q, R) < \tau + \eta$ and $\Delta_{<}(Q_Y, R) < \tau + \eta$ satisfies $D_c(Q_{XY}) \geq E_{\text{MD}}^\eta(\tau, R) > E$. We split the types contributing to β_n into two groups:

- *Types with $\lambda(Q, R) \geq \tau$:* contribute zero to β_n at the exponential scale (by the argument above for sparse and bulk types).
- *Types with $\lambda(Q, R) < \tau$ but $D_c(Q_{XY}) > E$:* their contribution to β_n is at most $\text{poly}(n) \cdot e^{-nE}$ each, which is acceptable.
- *Types with $\lambda(Q, R) < \tau$ and $D_c(Q_{XY}) \leq E$:* these must satisfy $\Delta_{<}(Q_Y, R) \geq \tau + \eta$ (since otherwise they would be feasible for E_{MD}^η with cost $\leq E < E_{\text{MD}}^\eta$, a contradiction). For such types, $\Delta_{<}(Q_Y, R) \geq \tau + \eta$ means there exists a bulk interfering type Q' with $I_{Q'}(X; Y) < R$ (strict) and $D_m(Q_Y) - D_c(Q') \geq \tau + \eta$, i.e., $\lambda(Q', R) = D_m - D_c(Q') \geq \tau + \eta > \tau$. By compactness of $\mathcal{Q}$ and continuity of λ , $\Delta_{<}$, I_Q , and D_c on $\mathcal{Q}$ (all are continuous functions of Q on the compact set $\mathcal{Q}$), for all large n and n -types $Q_n \rightarrow Q$: there exists an n -type $(Q')_n$ with $(I_{(Q')_n}(X; Y) - R_n) \leq -\eta/4 < 0$ and $\lambda_n((Q')_n, R_n) \geq \tau + \eta/2$, so $f((Q')_n) = \max\{I_{(Q')_n} - R_n, \tau - [D_m - D_c((Q')_n)]\} \leq -\eta/4 < 0$. Hence $\min_{Q''} f(Q'') \leq f((Q')_n) < 0$, and Theorem 4.3 of [2] gives $\Pr\{\tilde{S}(y^n) < e^{n\theta_n} | Q_n\}$ super-exponentially small. Therefore types with $\Delta_{<}(Q_Y, R) \geq \tau + \eta$ contribute zero at the exponential scale.

Combining all three groups: $\beta_n(\tau, R_n) \leq \text{poly}(n) \cdot e^{-nE}$ for all large n , giving $\liminf_n (-\frac{1}{n} \ln \beta_n) \geq E$. Since $E < E_{\text{MD}}^\eta(\tau, R)$ is arbitrary and $\eta > 0$ is fixed, this proves (25). This establishes both bounds in Theorem 1 for the MD exponent.

References

- [1] T. Han and S. Verdú, “Approximation theory of output statistics,” *IEEE Trans. Inf. Theory*, vol. 39, no. 3, pp. 752–772, 1993.
- [2] N. Merhav and N. Weinberger, “A Toolbox for Refined Information-Theoretic Analyses with Applications,” *Foundations and Trends in Comm. and Inf. Theory*, vol. 22, no. 1, pp. 1–184, 2025.
- [3] A. D. Wyner, “The common information of two dependent random variables,” *IEEE Trans. Inf. Theory*, vol. 21, no. 2, pp. 163–179, 1975.
- [4] M. Hayashi, “General nonasymptotic and asymptotic formulas in channel resolvability and identification capacity and their application to the wiretap channel,” *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1562–1575, 2006.
- [5] M. Bastani Parizi, E. Telatar and N. Merhav, “Exact random coding secrecy exponents for the wiretap channel,” *IEEE Trans. Inform. Theory*, vol. 63, no. 1, pp. 509–531, January 2017.

- [6] L. Yu and V. Y. F. Tan, “Rényi resolvability and its applications to the wiretap channel,” *IEEE Trans. Inf. Theory*, vol. 65, no. 3, pp. 1862–1897, 2019.
- [7] S. Yagli and P. Cuff, “Exact exponent for soft covering,” *IEEE Trans. Inf. Theory*, vol. 65, no. 12, pp. 7635–7654, 2019.
- [8] S.-B. Li, K. Li, and L. Yu, “Two-parameter Rényi information quantities with applications to privacy amplification and soft covering,” submitted for publication, 2026. Available on-line at: <https://arxiv.org/abs/2511.02297>.
- [9] P. Cuff, “A stronger soft-covering lemma and applications,” *Proc. 2nd Workshop on Physical-Layer Methods for Wireless Security* (co-located with IEEE CNS), Florence, Italy, pp. 40–43, September 2015.
- [10] P. Cuff, “Soft covering with high probability,” in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, pp. 2963–2967, 2016.
- [11] P. Cuff, “Distributed channel synthesis,” *IEEE Trans. Inf. Theory*, vol. 59, no. 11, pp. 7071–7096, 2013.
- [12] N. Merhav, “Soft Covering via Hypothesis Testing: Typical-Code Exponents and Mismatched Detection,” *IEEE Trans. Inf. Theory*, submitted 2026. [Online]. Available: <https://arxiv.org/pdf/2606.08124>
- [13] N. Weinberger and N. Merhav, “Codeword or noise? Exact random coding exponents for joint detection and decoding,” *IEEE Trans. Inf. Theory*, vol. 60, no. 9, pp. 5077–5094, September 2014.
- [14] N. Merhav, “A statistical-physics refinement of soft covering,” submitted for publication, 2026. Available on-line at: <https://arxiv.org/pdf/2605.01839>