A Large Deviations approach to Shannon Random Coding

Candidate: Mattia Fedrigo

Supervisor: Professor Franco Flandoli

Scuola Normale Superiore, Pisa

1 February 2005

Acknowledgements

The present thesis describes the results of my Perfezionamento project at the Scuola Normale Superiore in Pisa, under the supervision of Professor Franco Flandoli of the University of Pisa.

Since september 2004, it was supported by the Institut für Biomathematik und Biometrie (IBB) of the Forschungszentrum für Umwelt und Gesundheit (GSF) in Neuherberg, München, within the framework of the Harmonic Analysis and Statistics for Signal and Image Processing program (HASSIP) - a Research Training Network funded by the European Union.

I thank my advisor, Professor Flandoli, for his guide throughout this project. I also thank Professor Hans-Otto Georgii for his patience and Professor Gerhard Winkler for his support.

Un abbraccio alla mia famiglia.

1 February 2005

Mattia Fedrigo

Contents

1 Introduction **5**

2 Coding theory **8**
2.1 Introduction 8
2.2 The general problem and its implementations 8
2.2.1 Modularity in information systems 9
2.2.2 Rate distortion theory 10
2.2.3 Source coding theory 10
2.2.4 Cryptography 12
2.3 Channel coding theory 13
2.4 Channel decoding in detail: the decision setup 16
2.4.1 wordwise decoding - "ML decoding" 17
2.4.2 bitwise decoding - "MAP decoding" 17
2.5 Random Coding 21
2.5.1 A geometric notion of a code's "goodness" 21
2.5.2 Random coding's fundamental idea 23
2.5.3 Algebraic structure versus randomness 24
2.5.4 "Annealed" cost functions 25
2.6 The binary noise - binary random coding problem 27
2.6.1 Some notations 27
2.6.2 The source space 28
2.6.3 The binary random code 28
2.6.4 The memoryless binary symmetric noise 29
2.6.5 The annealed bitwise error probability 29

3 The ansatz of statistical-mechanics **31**
3.1 Introduction 31
3.2 Definitions 31
3.2.1 Relative entropy or Kullbach-Leibler distance 31
3.2.2 Binary relative entropy and Gilbert-Varshamov distance . 32
3.2.3 Discrete entropy 33
3.3 Statistical mechanics at equilibrium: the finite canonical ensemble 34
3.3.1 The limiting behaviour of the measure μ_β 36
3.3.2 The partition function $Z(\beta)$ 37
3.3.3 The free energy $F(\beta)$ 38
3.3.4 The thermodynamic limit: the free energy density $f_n(\beta)$, phase transitions 40
3.3.5 Random systems 42
3.4 The statistical mechanic ansatz to random coding 44
3.4.1 Random coding summary 44
3.4.2 The fundamental identification 46
3.4.3 The "spin magnetisation" 47
3.4.4 The error exponent 48
3.5 The Shannon problem as a statistical mechanic system 51

3.5.1 The Shannon problem . . . 52
3.5.2 The statistical mechanic identification . . . 52
3.5.3 A code rotation . . . 52
3.5.4 The associated statistical-mechanic model . . . 53

4 Large Deviations **55**
4.1 Introduction . . . 55
4.2 A large deviation principle for random log-Laplace integrals . . . 55
4.2.1 LDP for positive fluctuations: a Varadhan-like lemma . . . 57
4.2.2 On the exponential tightness condition . . . 62
4.2.3 A simplified version on the real line . . . 64
4.2.4 LDP for negative fluctuations: studying joint events . . . 66
4.2.5 How to find $\widetilde{I}^{-}(x)$ . . . 69
4.2.6 Full LDP . . . 72
4.3 LDP for an "extended REM" model . . . 72
4.3.1 Preliminaries . . . 74
4.3.2 Rescaling the problem . . . 75
4.3.3 LDP for $\frac{1}{n}\log \widetilde{Z}^{\omega}_{\beta,n}$, positive fluctuations . . . 76
4.3.4 LDP for $\frac{1}{n}\log \widetilde{Z}^{\omega}_{\beta,n}$, negative fluctuations . . . 81
4.4 LDP for the Shannon problem . . . 84
4.4.1 Free energy density bounds of the error exponent . . . 85
4.4.2 The LDP of F^{*}_{n} . . . 88
4.4.3 The LDP of $F^{1}_{\beta,R,n}$ and $\hat{F}^{0}_{\beta,R,n}$ . . . 88
4.4.4 The analysis of $I_{\beta,R}(x)$ . . . 90
4.4.5 Computation of $I_{\beta,R}(0)$ . . . 92
4.4.6 The four cases breakdown . . . 94
4.4.7 The final result . . . 95
4.5 Appendix . . . 96
4.5.1 Appendix A: binomial computations . . . 96
4.5.2 Appendix B: the LDP-sum lemma . . . 100
4.5.3 Appendix C: the separation bound . . . 101

5 Conclusions **103**
5.1 Capacity . . . 103
5.2 Error exponent . . . 104
5.3 Towards a choice of β . . . 105
5.4 Open problems . . . 106

1 Introduction

After the second World War, the development of digital communication, the invention of computers and the evolution of automatic control prompted the birth of a new and wide research area spanning programming, data filtering and signal theory. Within this vast research effort, information theory studies the abstract bases of data processing and communication. As such, it is one of its most formal branches with a pervasive use of mathematics.

Since the birth of the discipline has been induced by a technological breakthrough, the foundational work has often been done on a case-by-case application-driven basis. Although this pioneering approach has been spectacularly prolific in results and effective in applications, it often resorts to use a self-contained dialect of the mathematical language that on one side is parsimonious, adapted and intuitively clear, but on the other side lacks the generalisation and problem-identification capabilities of a a full-blown mathematical theory. Furthermore quantities and concepts have been defined, that were already well known in apparently unrelated topics - like statistical mechanics. Reasons for this localisation phenomenon might include the fact that the authors of the mentioned fundamentals works are not always mathematicians or physicists, or even more simply that the relevant mathematical theories were not yet developed. The classic example of pioneering work in information theory is Claude E. Shannon's 1948 paper entitled "A Mathematical Theory of Communication" [18] which gives also the first formulation of the channel coding theorem. It is a very general and groundbreaking work, and at the same time it is built on simple analysis and probability theory knowledge. It defines critical notions like entropy and random coding, and it works its way from there, through successive assumptions and all the required computations, to its final theorems. Much work has followed, with a continuous increase in depth, generalisation and mathematical precision. The books by Gallager [11] and Csiszar and Korner [5] can be considered two landmarks of this process. Nevertheless the theory persists to be mostly self contained, both in results and techniques.

This observation is the primary reason behind the research effort contained in the present thesis. In this work, the fundamental random coding problem originally solved by Shannon is rewritten in the precise mathematical language of large deviation theory, and according to an approach aimed at recovering the "know-how" in random systems developed in statistical mechanics.

Random coding was originally devised by Shannon only as a technique to prove the channel coding theorem, not as a practical coding strategy. In the mentioned paper, he studied the problem of communication with a binary random code of block length n and rate R through a memoryless binary symmetric channel (BSC) with error probability p. He was in this way able to prove its homonimous theorem on channel capacity: error-free communication is asymptotically possible if and only if the rate R is smaller than a capacity C, function only of the channel parameter p. Despite some early results by Robert Gallager on the feasibility and performance of a class of random codes named "low density parity check" (LDPC) codes [12], the topic has been soon discarded by

developers in favor of algebraic codes, because of the decoding problem: the structure of algebraic codes allows fast decoding, while the absence of structure in the originally proposed random codes requires a decoding complexity which grows exponentially with the block length n. Although incapable to achieve performances close to the theoretical bound defined by Shannon's channel theorem, algebraic codes have been the main focus of applied research until 1993. In that year, the article "Near Shannon Limit Error-Correcting Coding and Decoding: Turbo Codes" [2] by Berrou, Glavieux, and Thitimajshima put random codes back on the highlights with the introduction of the so called turbo codes, which - as the name implies - allow for a fast iterative decoding strategy. The paper sparked a burst of research, both theoretical and applied, which also rediscovered and fully appreciated Gallagers LDPC codes.

At the same time of this revival, there was an increasing perception of the general, multidisciplinar nature of some concepts, like disordered systems and entropy. Originally introduced in termodynamics and statistical mechanics, they emerged in information theory as well, and showed up in some applications to stock market prediction, too [4]. In the early 90ties Nicolas Sourlas pioneered an effort to translate information theory problems to equivalent statistical mechanic setups [19], with the aim to recover for the formers the sensibility and the techniques developed for the latters through the years by physicists, as for instance the replica trick. Such an approach was extended and detailed in the case of the random coding problem [20] with a particular attention to turbo codes [17] and LDPC codes [16].

The scene is completed by the upcoming maturity of the probabilistic theory of large deviations. Many of the techniques adopted in channel coding analysis, as for instance the computation of error exponents or the very "method of types" devised by Csiszar [5], may be seen as special cases of large deviation computations.

Following this observation, and adopting the Sourlas approach, the present work was initiated. Our effort is to rewrite the channel coding problem in a more formal and general setup, recovering physicists' expertise and precising the mathematical problems involved. Our aim is to understand more deeply the phenomena involved in channel coding and to takle its remaining open questions.

As a test model, we choose the very same channel model originally studied by Shannon for its channel coding theorem. For this reason we refer to it at the Shannon problem. We then proceed to describe it with the language of statistical mechanics according to the Sourlas approach, identifying the critical quantities for the computation of capacity and error exponent. Such quantities have the form of free energy densities. Finally, we develop our large deviation analysis. We choose a general-to-particular approach, developing a rather general theorem which we then rephrase with simpler assumptions and eventually we it apply to the analysis of the Shannon problem in our free energy density formulation. What we obtain is a new version of the coding theorem, for both capacity and error exponent, as a function of a parameter β called inverse temperature. This inverse temperature is introduced by the Sourlas approach and it

parametrises the decoding strategy used in the problem. As $\beta \to \infty$ we recover MAP wordwise decoding, while for β equal to a particular value β_{Crit} we obtain MAP bitwise decoding. Although the results for these two particular decoding strategies are already known in the literature through dedicated analysis, we get a new channel coding theorem for a family of "intermediate" decoding strategies and we are able to make some new observations on the stability of MAP bitwise decoding against poor knowledge of the channel parameter.

Our aims for future research include a tightness result for the Shannon problem in the high temperature regime. A generalisation of the channel model or the introduction of a linear constraint for the random encoder will require a generalisation of the large deviation theorem here proved.

The layout of this work is as follows. Every chapter is written according to a general-to-particular approach: first wide discussions, abstract concepts or results, then a second part focused on the Shannon problem. After the present introduction, chapter 2 offers an introduction to coding theory and to the Shannon problem in particular. Chapter 3 is instead dedicated to the Sourlas approach. Chapter 4 contains our large deviation theorem and its application. The conclusions in chapter 5 will finally resume the work done and the results obtained.

2 Coding theory

2.1 Introduction

This chapter starts with a rather general introduction to coding theory. In order not to bore the expert, the first paragraph is devised as an optional reading: all the spanning definitions are duplicated later on. Afterwards we dedicate three paragraphs respectively to a definition of the channel coding problem, an outline of its decision formulation, and an introduction to the idea of random (channel) coding. Finally, in the last paragraph we restrict ourselves to our focus problem, i.e.what we call the Shannon problem.

In the subsequent chapters we will describe the approaches and techniques adopted in our analysis, i.e. the statistical-mechanic ansatz and large deviation theory, along with their application to the binary random coding problem. Then we will proceed to the final results and conclusions.

2.2 The general problem and its implementations

Abstractly speaking, coding theory studies maps between given couples of spaces, named the source space S and the encoded space X. Their elements are called respectively the sourcewords s and the codewords x. A code c can be defined as an ordered couple of maps t (for "transmit") and r (for "receive"), called respectively the encoder and the decoder, as follows:

$$c = (t, r)$$

$$t : S \to X$$

$$r : X \to S$$

The general problem of coding theory is the search for the code c_k minimizing a given cost function k defined on

$$C := X^S \times S^X$$

the space of all codes between S and X, and complying to some eventual further conditions that can be formalised as a subspace C' of C:

$$k : C \to \mathbb{R}$$

$$c_k := \arg\min_{c \in C'} k(c)$$

We now present the technological intuition behind coding theory along with some relevant modeling and mathematical issues. Follows a serie of introductions to the contestualised versions of the general problem.

2.2.1 Modularity in information systems

Modern coding theory was developed with a focus on direct applications in electronics and information technology.

According to operative engineering criteria, an information system is split into several sub-systems dedicated to separate tasks. This allows the design and realisation of arbitrarily large systems, like the Internet. A first example of a subsystem is the one dedicated to acquiring and dually reproducing physical world stimula. It is in other words a converter between real world signals (quantities, processes, fields) and electronic signals. Then these electronic signals can be processed, stored or remotely transmitted by other dedicated subsistems.

Today the mainstream technology for electronic signals is digital, which means that such signals are interpreted as being discrete and quantised, and that they are typically associated to binary data (binary vectors, binary strings, binary processes). Digital technology sports important features like regeneration (de-noising), redundancy control, low power consuption and low infrastructure and maintenance costs. Actually the first two mentioned features, i.e. regeneration and redundancy control, are the core topics of two sub-areas of coding theory, respectively channel coding and source coding. We will present them in the next sections.

This modular approach and this technology choice have direct modeling consequences in coding theory. First of all, the coding problem splits into subproblems: rate distortion theory for the problem of converting real world signals into electronic data and vice versa, source coding for data compression (an elaboration to use less storage or communication resources), cryptography for data protection (from unauthorised use in an open environment), channel coding for data communication.

Secondly, in any of such sub-problems at least one of the two spaces S and X is taken to be discrete, as for example the space of n-dimensional binary vectors B^n (which is finite, being $B := \{0, 1\}$) or the space of finite binary strings $D := \bigcup_{n\in\mathbb{N}} B^n$ (which is infinite but countable). It has been proven that using such spaces, "optimal" solutions in information sub-systems join up to "optimal" solutions in the complete system, that is, we have the optimality of modular solutions. This very fact justifies the engineering modular approach.

It is rather easy to characterise the elements of D and B^n with an intuitive measure of their information content (or complexity, or "self-entropy"), i.e. their length, which will be variable in the case of D and fixed and equal to n in the case of B^n. Starting from such a crude measure, information theorists introduced probabilistic descriptions of the spaces and rediscovered the concept of entropy as the proper measure of information content. So in some approaches the source space S is enriched to a probability triple (S, Σ, λ).

Sometimes the space of endless binary sequences $L := B^{\mathbb{N}}$ (which is uncountable) is also considered for digital data. This may happen for modeling reasons, as for instance the effort of describing real-time, long-operation systems. Or it may be for mathematical reasons. This choice, though, brings along its own problems: first of all, the optimality of modular solutions may break down in

some pathologic cases. Anyway, we do not plan to use such a device in this thesis. We will use the scaling of finite discrete spaces as an alternative workaround to prove the described asymptotic results.

Finally, let us spend some words on the meaning of "optimality" in information engineering. The lack of uniqueness and a kind of weak notion of existence for some crucial theorem as the Shannon's one (existence in any neighbourhood of the prescribed conditions) along with issues of physical implementability of solutions, leads to a dual concept of "weak optimality": existence and implementability of a solution arbitrarily adherent to the prescribed conditions.

Let us now proceed to a series of presentations of the main sub-problems in coding theory.

2.2.2 Rate distortion theory

Rate distortion theory is the branch of coding theory that looks for "good" finite discrete representations of real signals, where the "goodness" is formalized through the definition of a so-called distortion measure. This is the theoretical core of all technologies that translate real-world stimula into digital data, like digital cameras or portable CD players.

In rate distortion theory S represent the "real world" and is tipically a subspace of $\mathbb{R}^n$, while X is discrete and finite to be easily mapped into digital data. A distortion function d is a map like the following:

$$d : S \times S \to \mathbb{R}_0^+$$

Let us consider a code $c_d = (c_{d,S}, c_{d,X})$ which minimizes the distortion measure $d_{\max}$:

$$d_{\max} := \max_{s \in S} \; d\left[\; s, c_{d,X} \circ c_{d,S}\left(s\right) \; \right]$$

If it exists, such a code is the answer to the problem of finding a good representation of S in X according to the distortion measure $d_{\max}$, which plays the role of our aforementioned cost function f.

Another approach introduces some more knowledge of the source space S through a probability triple (S, Σ, λ), with associated expectation $\mathbb{E}$. This offer a new distortion measure $\mathbb{E}\left\{d\right\}$:

$$\mathbb{E}\left\{d\right\} := \mathbb{E}\left\{ \; d\left[\; s, c_{d,X} \circ c_{d,S}\left(s\right) \; \right] \; \right\}$$

Often S and d are empirically chosen according respectively to the "tipical" domain of the real world stimulus s and a "qualitative perception" of the final reproduction $c_{d,X} \circ c_{d,S}\left(s\right)$. Such is the case of the MP3 format, for instance.

2.2.3 Source coding theory

Source coding theory focuses on reducing the load on communication and storage resources by reducing the size of the digital data. The typical setup in source coding theory is:

$$S = X = D$$

and the main problem is the search for codes c so that each element s' on a subset S' of S is mapped biunivocally to an element s'' on S:

$$t\left(s'\right)=s''\ ,\ r\left(s''\right)=s'\ \ \forall s'\in S'$$

so that the following property holds:

$$l\left(s''\right)\leq l\left(s'\right)$$

where the function $l\left(s\right)$ returns the length of the string s. This requirement amounts to select codes that compress the strings in S' (map them into shorter strings) without losing their information (mantaining biunivocity in S'). The applications of this problem live mostly in compression software, i.e. computer programs that "reduce" the size of digital data files so that they occupy less storage or communication resources in a (limited) digital system.

There are two main approaches to operative source coding: one algorithmic and the other probabilistic. We will outline them briefly in a simple, fully bijective setup. A relaxation of this bijectivity hypothesis will be considered in the following subsection on lossy compression.

The first approach looks for deterministic algorithms which map finite sequences of bits (elements of D) into couples consisting of an incremental vocabolary and a translation of the data accoding to the same vocabulary. A concatenation of any of such couples' elements would live in D as well. The encoder t and the decoder r are in this case both bijective, but only a subset of D would be encoded into shorter codewords. This subset is specified operatively by the algorithm and is represented by S' in our model. A theoretical study of this operative encoding strategy was pioneered by Kolmogorov [15]. In analogy with our presentation of rate distortion theory, we could define a worst-case residual occupation ratio $\gamma_{\max}$:

$$\gamma_{\max}:=\max_{s\in S'}\gamma\left(s\right)$$

$$\gamma\left(s\right):=\frac{l\left(s''\right)}{l\left(s'\right)}$$

and look for codes that minimize it, or equivalently maximize the worst-case compression ratio $\rho_{\min}$:

$$\rho_{\min}:=\min_{s\in S'}\rho\left(s\right)$$

$$\rho\left(s\right):=\frac{1}{\gamma\left(s\right)}=\frac{l\left(s'\right)}{l\left(s''\right)}$$

A second approach to source coding comes from probability theory. If we introduce a probabilistic description of the source (S,Σ,λ) with expectation $\mathbb{E}$ we may choose to minimize -instead of $\rho_{\min}$- one of the following (generally different) cost functions:

$$\gamma_{\lambda}:=\mathbb{E}\left\{\gamma\left(s\right)\right\}$$

$$\rho_\lambda := \frac{1}{\mathbb{E}\{\rho(s)\}}$$

over the set of bijective codes

$$C' := \left\{ \; c \in C : r \circ t(s) = s \quad \forall s \in S \; \right\}$$

Of course, if γ_λ or ρ_λ are less than one then there exists a subset S' of S such that $\forall s \in S'$ we have $\rho(s) < 1$ or $\gamma(s) > 1$ respectively - this means that we can consider the probabilistic approach within the general frame we described at the beginning of this section. Theoretical studies of this approach identified self-information and entropy as the critical quantities in gauging the "compressibility" of elements of S and the achievable compression rate $\mathbb{E}\{\rho(s)\}$. One of the landmarks of this research effort is Shannon's paper on coding theory [18].

Lossy compression

By lossy compression we typically intend a source coding problem where bijectivity is not strictly required.

In some sense this problem can be seen as a fusion of source coding and rate distortion theory, with $S = X = D$. The cost function would then be a combination of a distortion measure and residual occupation ratio. The tradoff between these two shall then be empirically adapted to the problem considered - this is exactly the case in MP3 audio or MPEG4 video compression.

The relaxation of the bijectivity constraint can deliver benefits even without introducing a distortion measure. From an algorithmic source coding perspective, we could decide to identify a superset S'' of S' such that the function $r \circ t$ needs not to equal to the identity in its complement $S \backslash S''$. This of course can lead to simpler algorithms. From a probabilistic source coding point of view, we could impose a small arbitrary bound P_{loss} on the probability of losing bijectivity:

$$\mathbb{P}\{s \in S : r \circ t(s) \neq s\} \leq P_{loss}$$

This bound would then identify a subset of acceptable codes:

$$C'' := \{c \in C : \mathbb{P}\{s \in S : r \circ t(s) \neq s\} \leq P_{loss}\}$$

which clearly contains the set of bijective codes:

$$C'' \supset C'$$

So the domain within which we look for a minimisation of γ_λ or ρ_λ is larger, potentially allowing better codes.

2.2.4 Cryptography

Cryptography is similar to source coding theory in the fact that it typically takes

$$S = X = D$$

and that code bijectivity is required. However the aim is totally different than compression, and requires the concept of algorithmic complexity to be made precise. The task is in fact to provide a code so that encoding is "easy" but decoding is "difficult" unless we know a secret "key", so that an encoded (encrypted) message can be considered secure (unreadable) from everybody missing such a key. This can be formalised into a quest for a code c so that the problem of reconstruction of the decoder r knowing the encoder t lives in NP relatively to such a key $k \in D$. Since the description of cryptography through a cost function would require to discuss complexity theory, and since we plan not to further explore this domain in this thesis, we choose to end this paragraph here.

2.3 Channel coding theory

Channel coding theory studies the problem of communication through a noisy medium, that is a medium which can unpredictably (randomly) alter the information exchanged. Such a phenomenon is of course unwelcome.

The physical elements involved in the problem are a transmitter, a receiver and a channel (which may be a cable, or the space). In the engineering literature, a formal setup is offered by the notions of coding scheme, noise model and decoding strategy. Let us briefly recapitulate the abstract coding setup introduced in the previous paragraph, and then proceed with an identification between elements of the latter, of the engineering setup and of the physical world. Given two spaces S and X, a code c can be defined as an ordered couple of maps:

$$c = (t, r)$$

$$t : S \to X$$

$$r : X \to S$$

The coding scheme corresponds in our framework to the $t : S \to X$ encoder map and is implemented by the physical transmitter. It translates electronic data from the space S to real world signals in X. In channel coding we typically have:

$$S = B^m \quad X = \mathbb{R}^n$$

where $B := \{0, 1\}$. The decoding strategy is the $r : X \to S$ decoder map, associated to the physical receiver. It translates back real world signals from X to electronic data in S. The noise model represents the channel, that is the effect of unpredictable alteration of the transmitted signal. It can be described by a random function b (for "bruit", french word for noise), called the random noise:

$$b : (\Omega, \mathcal{F}, \mathbb{P}) \to X^X$$

where the triple $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space with associated expectation $\mathbb{E}$ and X^X is the space of the functions from X to X itself. There are, of course, measurability issues to take care of. Since the channel coding model will be progressively enriched throughout this paragraph, we choose to precise these

requirements only at the end of it. The channel coding problem can then be described as the search for a code c that minimizes some error functional $E_{f,\max}$ on a set $W \subseteq C$:

$$E_{f,\max} := \max_{s \in W} E_f$$

$$\begin{aligned} E_f &: = \mathbb{E}\left\{f\left(r \circ b\left(\omega\right) \circ t\left(s\right), s\right)\right\} \\ &: = \int_{\Omega} f\left(r \circ b\left(\omega\right) \circ t\left(s\right), s\right) \mathbb{P}\left(d\omega\right) \end{aligned}$$

given the noise n and an arbitrary measurable function f:

$$f : S \times S \to \mathbb{R}$$

Notice that E_f depends from s, t and r., while $E_{f,\max}$ only from t and r, which is of course consistent with our minimisation effort as $c = (t, r)$ varies on W. This latter subspace of C describes an eventual constraint, as for instance an "energy" constraint specified by an energy paramenter $w \in [0, \infty]$:

$$W := \left\{ c \in C : \max_{s \in S} \left\| t\left(s\right) \right\|_2 \leq w \right\}$$

If we take some extra hypotheses, as for example discreteness and finiteness of S, we can take

$$f\left(t, s\right) := \chi_{\ t \neq s}$$

where χ_A is the indicator function of the set A, and define the maximal (source)word error probability $\mathbb{P}_{\max}\left(E_{err}\right)$:

$$\mathbb{P}_{\max}\left(E_{err}\right) := \max_{s \in S} \mathbb{P}\left\{E_{err}\right\}$$

$$E_{err} := \left\{\omega \in \Omega : r \circ b \circ t\left(s\right) \neq s\right\}$$

the link between the two descriptions being the passage:

$$\begin{aligned} \mathbb{P}\left\{E_{err}\right\} &: = \mathbb{E}\left\{ \chi_{\ \left(r \circ b\left(\omega\right) \circ t\left(s\right) \neq s\right)\ } \right\} \\ &= \mathbb{E}\left\{f\left(r \circ b\left(\omega\right) \circ t\left(s\right), s\right)\right\} \\ &= E_f \end{aligned}$$

Clearly, $\mathbb{P}\left\{E_{err}\right\}$ depends from s, t and r., while $\mathbb{P}_{\max}\left(E_{err}\right)$ only from t and r. Again, we can introduce a probabilistic description of the source (S, Σ, λ). In this enriched setup we are faced with the problem of relating the two probability spaces (S, Σ, λ) of the source and $(\Omega, \mathcal{F}, \mathbb{P})$ of the noise. Since the entities so described (the unpredictability in the source and the one in the channel) are unrelated in the model, it seems natural to assume mutual independence between them. In order to have a unified probability description, we can choose to define a product probability triple $(\Omega \times S, \mathcal{F} \otimes \Sigma, \mathbb{P} \otimes \lambda)$ where $\mathcal{F} \otimes \Sigma$ is the

σ-algebra generated by the union of the two σ-algebrae $\mathcal{F}$ and Σ extended to the product space $\Omega \times S$ while $\mathbb{P} \otimes \lambda$ is of course the product measure. With this choice we can introduce a new cost function, the average (source)word error probability $\mathbb{P}_{werr} := \mathbb{P} \otimes \lambda (E_{err})$.

Alternatively, we can choose one of the two probability spaces (S, Σ, λ) and $(\Omega, \mathcal{F}, \mathbb{P})$ as the base space and describe the other as a random variable induced probability space. Since the channel noise is already described in our setup through the random variable $b : \Omega \to X^X$, it seems more natural to choose $(\Omega, \mathcal{F}, \mathbb{P})$ as the base triple and define a $(\mathcal{F}, \Sigma)$-measurable random variable

$$s_0 : (\Omega, \mathcal{F}, \mathbb{P}) \to (S, \Sigma)$$

independent of b and such that λ the measure induced by $\mathbb{P}$ through s_0 in S:

$$\lambda (B) := \mathbb{P} \left(\ s_0^{-1} (B) \ \right) \forall B \in \Sigma$$

This new random variable s_0 can be called the random message.

Of course it is not always possible to add independent random variables on a predefined space, so formally this is an assumption. Now we can also precise the measurability requirements for t, b, r. Letting χ be a given sigma-algebra on X, we assume:

$$t : (S, \Sigma) \to (X, \chi)$$
$$b : (\Omega \times X, \mathcal{F} \otimes \chi) \to (X, \chi)$$
$$r : (X, \chi) \to (S, \Sigma)$$

where we have interpreted b s a measurable function from a product space, instead as a random variable in a function space. Of course, the probability measure $\mathbb{P}$ defined on $(\Omega, \mathcal{F})$ makes the measurable function b a "random function".

The above choice lead us to redefine the error event, which we do through a small abuse of notation:

$$E_{err} := \{\omega \in \Omega : r \circ b \circ t (s_0) \neq s_0\}$$

Notice that now E_{err} depends from t and r, while b and s_0 are the random variables defining it. The average word error probability $\mathbb{P}_{werr}$.can now be formalised as follows:

$$\begin{aligned} \mathbb{P}_{werr} \quad &: \quad = \mathbb{P} (E_{err}) \\ &= \quad \mathbb{E} \left\{ \chi_{E_{err}} \right\} \\ &= \quad \mathbb{E} \left(\ \mathbb{E} \left(\chi_{E_{err}} \middle| \, s_0^{-1} (\Sigma) \right) \ \right) \\ &= \quad \mathbb{E} \left(\ \mathbb{P} \left\{ E_{err} \middle| \, s_0^{-1} (\Sigma) \right\} \ \right) \end{aligned}$$

From now on we will stick to this choice, that is we choose to employ the enriched probabilistic description of the source (S, Σ, λ) that we suppose induced by the random message $s_0 : (\Omega, \mathcal{F}, \mathbb{P}) \to S$ independently from the random noise $b : (\Omega, \mathcal{F}, \mathbb{P}) \to X^X$, with the error event defined as $E_{err} := \{\omega \in \Omega : r \circ b \circ t (s_0) \neq s_0\}$.

2.4 Channel decoding in detail: the decision setup

In the previous paragraph we have identified the channel coding problem as a quest for a cost function minimizer in a constrained subspace W of the codes' space C. Here we explore in more detail the formal setup of the channel coding problem, with a particular focus on the decoder part.

We take the further assumptions that the source space S is finite and that the induced measure λ is uniform on it. This is a typical setup taken in telecommunication engineering, the finiteness hypothesis being crucial to implement a decoding decision.

Let us now observe that, given the encoder $t : S \to X$, the choice of the cost function k identifies an optimal decoder $r : X \to S$ in the following way:

$$r_t := \underset{\gamma \in S^X}{\arg\min}\ k\left(\gamma, t\right)$$

provided, of course, that existence and unicity issues are satisfied in some sense. This problem is covered by decision theory.

The consequence of such a notion is that we can focus on looking for a good encoder after we specify a cost function and set for the decision theory-optimal decoder r. This leads to an interesting engineering approach: first set a cost function and look for a simple encoder so that, associated with the decision theory-optimal decoder, it constitutes a "good" code in such a cost-defined sense. Afterwards try to approximate such a decoder with a suboptimal one, ideally with an iterative algorithm to be artificially stopped after the decoding process is empirically good enough.

In the next two sections we explore the decoding consequences for two popular cost function choices. The first one is the already introduced average word error probability $\mathbb{P}_{werr}$ which induces wordwise maximum a posteriori probability decoding. Because of the uniformity assumption on λ, this is equivalent to maximum likelyhood decoding. This is the reason for being labelled "ML decoding" in turbo coding literature (turbo codes are a particular case of random codes). The second one is the so called average bit error probability $\mathbb{P}_{berr}$, to be defined later on, which induces bitwise maximum a posteriori probability decoding. The word bitwise relates to the projection of the messageword into the digits of its binary representation, seen as the messageword components in an appropriately defined vector space, and to the subsequent decoding of such binary digits (bits) separately from each other. Such idea of an independent reconstruction of the bit components of a sourceword vector suggests also another name for the associated decoding strategy, that is marginal decoding, as opposed to the dual label of global decoding employable for the already described wordwise decoding. Equivalence with maximum likelyhood decoding holds in this bitwise case, too. Perhaps surprisingly, this strategy is called MAP decoding in the aforementioned "turbo literature".

2.4.1 wordwise decoding - "ML decoding"

Let us take the average word error probability $\mathbb{P}_{werr}$ as our cost function. The decision theory-optimal decoder is then defined by:

$$\begin{aligned} r_t \quad : \quad &= \arg\min_{\gamma \in S^X} \mathbb{P}_{werr}(\gamma, t) \\ &= \arg\min_{\gamma \in S^X} \mathbb{P}(\gamma \circ b \circ t(s) \neq s) \\ &= \arg\min_{\gamma \in S^X} \mathbb{E}\left\{ \chi_{\ \gamma \circ b \circ t(s) \neq s} \right\} \end{aligned}$$

Of course, we need to consider existence and uniqueness conditions. Uniqueness, for instance, can be achieved only almost surely, given the presence of the expectation operator inside the arg min. In the typical cases S is discrete and finite while X is a real finite-dimensional vector space, so any decoder actually partitions the codeword space X in a number of subsets, each of which is associated to a different messageword. In case existence is granted, the optimal partitioning subsets take the name of Voronoi regions and one can easily prove that:

$$r_t(y) = \arg\max_{s \in S} \mathbb{P}\{ s | y, t \}$$

2.4.2 bitwise decoding - "MAP decoding"

Aim of this section is to introduce the so-called MAP decoding, a decision theory-optimal decoding strategy induced by the average bit error probability cost function $\mathbb{P}_{berr}$, which we are going to define.

The idea behind bitwise decoding comes from the physical binary representation (in binary digits or "bits") of the electronic data in communication systems. This representation offer the possibility to refer the probability of a transmission error to the single bits that make up a sourceword representation, instead that to the sourceword itself. In othero words, we might be interested in the probability that a single bit of the reconstructed sourceword is actually correctly received or not, instead of just considering the correctness of the entire sourceword reconstruction.

The reason for this paradigm shift stems from an engineering perspective. As already mentioned in the introduction on modularity in information systems, sometimes it is interesting to consider systems to operate indefinitely, or at least without a perceivable time frame. In such a case sourcewords and codewords are ideally unending sequences, and formalisation of codes becomes more delicate. However, under the usual engineering hyotheses of ergodicity, (source)words' errors actually happen asymptotically with probability one, and moreover their probability brings no useful information on the local behaviour of the system, where the locality can be intended in space, time or more abstract measurable resources. So the idea is to shift the computation to inherently local quantities, like bits, which can provide a meaningful and operative definition of an error probability.

One last thing to stress before delving into the formal details is that the choice of a binary representation does actually have an impact on performance. At parity of other conditions, in fact, different representations can induce different average bit error probabilities. For an intuition of this effect, let us think of associating close sourceword representations to codewords that can be easiliy be mistaken for one another: in this way, a more likely error will corrupt a smaller number of bits.

Assume that S is a finite set (not necessarily equal to B^m) and let $m \in \mathbb{N}$ be such that

$$\#\{B^m\} = 2^m \geq \#\{S\}$$

Then we can set a bijection Ψ between S and a subset S' of B^m:

$$\begin{aligned} \Psi &: S \to S' \\ \Psi^{-1} &: S' \to S \end{aligned}$$

which we can call a binary representation of S. If we consider Ψ as a random variable, then it induces an entire probability space $\left(S', \Sigma', \lambda'\right)$ as follows:

$$\begin{aligned} \Sigma' &: = \left\{C \subseteq S' : \Psi^{-1}(C) \in \Sigma\right\} \\ &= \left\{C \subseteq S' : s_0^{-1} \circ \Psi^{-1}(C) \in \mathcal{F}\right\} \end{aligned}$$

$$\begin{aligned} \lambda'(C) &: = \lambda\left(\ \Psi^{-1}(C)\ \right) \\ &= \mathbb{P}\left(\ s_0^{-1} \circ \Psi^{-1}(C)\ \right) \quad \forall C \in \Sigma' \end{aligned}$$

Since the general element

$$s' := \{s'_1, \ldots, s'_m\}$$

of S' is a m-dimensional vector, we can apply to it the k-th projection operator $\pi_k(\cdot)$:

$$\pi_k(s') := s'_k$$

for all $k \in \{1, \ldots, m\}$. With the position

$$s' = \Psi(s_0)$$

we can now define the event $E_{err,k}$ as follows:

$$\begin{aligned} E_{err,k} &: = \left\{\omega \in \Omega : \pi_k \circ \Psi \circ r \circ b \circ t \circ \Psi^{-1}(s') \neq \pi_k(s')\right\} \\ &= \dot{\bigcup_{\substack{a,b \in S \\ \pi_k(a) \neq \pi_k(b)}}} E_{a,b} \end{aligned}$$

where the symbol $\dot{\cup}$ is used to denote disjoint union, since with the position

$$E_{a,b} := \left\{\omega \in \Omega : s' = a, \Psi \circ r \circ b \circ t \circ \Psi^{-1}(s') = b\right\} \quad \forall a, b \in S$$

we have:

$$E_{a,b} \cap E_{c,d} = \emptyset \ \forall (a,b), (c,d) \in S^2 : (a,b) \neq (c,d)$$

Notice furthermore that such event $E_{err,k}$ is a function of k and Ψ besides t, r. In total analogy with the average word error probability $\mathbb{P}_{werr}$, we can now define the average k-th bit error probability $\mathbb{P}_{berr,k}$:

$$\begin{aligned}
\mathbb{P}_{berr,k} &: = \mathbb{P}\left(E_{err,k}\right) \\
&= \mathbb{E}\left\{\chi_{E_{err,k}}\right\} \\
&= \mathbb{E}\left(\ \mathbb{E}\left(\chi_{E_{err,k}} \middle| s_0^{-1} \circ \Psi^{-1}\left(\Sigma'\right)\right)\ \right) \\
&= \mathbb{E}\left(\ \mathbb{P}\left\{E_{err,k} \middle| s_0^{-1} \circ \Psi^{-1}\left(\Sigma'\right)\right\}\ \right) \\
&= \sum_{a \in S'} \mathbb{P}\left\{E_{err,k} \middle| s' = a\right\} \mathbb{P}\left\{s' = a\right\}
\end{aligned}$$

the last passage being possible because of the countable nature of S'. We can manipulate some more the $\mathbb{P}\left\{E_{err,k} \middle| s' = a\right\}$ term:

$$\begin{aligned}
\mathbb{P}\left\{E_{err,k} \middle| s' = a\right\} &= \mathbb{P}\left\{ \dot{\bigcup_{\substack{a,b \in S \\ \pi_k(a) \neq \pi_k(b)}}} E_{a,b} \middle| s' = a \right\} \\
&= \sum_{\substack{b \in S' \\ \pi_k(a) \neq \pi_k(b)}} \mathbb{P}\left\{E_{a,b} \middle| s' = a\right\} \\
&= \sum_{\substack{b \in S' \\ \pi_k(a) \neq \pi_k(b)}} \mathbb{P}\left\{\Psi \circ r \circ b \circ t \circ \Psi^{-1}\left(s'\right) = b \middle| s' = a\right\} \\
&= \sum_{\substack{b \in S' \\ \pi_k(a) \neq \pi_k(b)}} \mathbb{P}\left\{\Psi \circ r \circ b \circ t \circ \Psi^{-1}\left(a\right) = b\right\}
\end{aligned}$$

the last passage being a consequence of independence between the random source and the random noise:

$$\begin{aligned}
\mathbb{P}\left\{\Psi \circ r \circ b \circ t \circ \Psi^{-1}\left(s'\right) = b \middle| s' = a\right\} &= \mathbb{P}\left\{\Psi \circ r \circ b \circ t \circ \Psi^{-1}\left(s'\right) = b, s' = a \middle| s' = a\right\} \\
&= \mathbb{P}\left\{\Psi \circ r \circ b \circ t \circ \Psi^{-1}\left(a\right) = b, s' = a \middle| s' = a\right\} \\
&= \mathbb{P}\left\{\Psi \circ r \circ b \circ t \circ \Psi^{-1}\left(a\right) = b \middle| s' = a\right\} \\
&= \mathbb{P}\left\{\Psi \circ r \circ b \circ t \circ \Psi^{-1}\left(a\right) = b\right\}
\end{aligned}$$

In order to get rid of the dependence on k we can now take the following average:

$$\mathbb{P}_{berr} := \frac{1}{m} \sum_{k=1}^{m} \mathbb{P}_{berr,k}$$

which we call the average bit error probability. This is of course a function of Ψ, t and r.

We can now introduce the associated decision theory-optimal decoder as follows:

$$\begin{aligned} r_t \quad : \quad &= \arg\min_{\gamma \in S^X} \mathbb{P}_{berr}(\gamma, t) \\ &= \arg\min_{\gamma \in S^X} \frac{1}{m} \sum_{k=1}^{m} \mathbb{P}_{berr,k} \\ &= \arg\min_{\gamma \in S^X} \frac{1}{m} \sum_{k=1}^{m} \sum_{\substack{a,b \in S' \\ \pi_k(a) \neq \pi_k(b)}} \mathbb{P}\left\{\Psi \circ \gamma \circ b \circ t \circ \Psi^{-1}(a) = b\right\} \mathbb{P}\{s' = a\} \end{aligned}$$

If we map the decoder codomain to its binary representation we can set each of the bit decoders separately and achieve global optimality:

$$r_t = \Psi^{-1} \circ \{r_{1,t}, ..., r_{m,t}\}$$

$$\begin{aligned} r_{k,t} \quad : \quad &= \pi_k \circ \Psi \circ r_t \\ &= \arg\min_{\gamma \in \{0,1\}^X} \sum_{\substack{a,b \in S' \\ \pi_k(a) \neq \pi_k(b)}} \mathbb{P}\left\{\gamma \circ b \circ t \circ \Psi^{-1}(a) = \pi_k(b)\right\} \mathbb{P}\{s' = a\} \\ &= \arg\min_{\gamma \in \{0,1\}^X} \sum_{a \in S'} \mathbb{P}\left\{\gamma \circ b \circ t \circ \Psi^{-1}(a) \neq \pi_k(a)\right\} \mathbb{P}\{s' = a\} \\ &= \arg\min_{\gamma \in \{0,1\}^X} \left\{ \begin{array}{ll} \sum\limits_{\substack{a \in S' \\ \pi_k(a) = 0}} & \mathbb{P}\left\{\gamma \circ b \circ t \circ \Psi^{-1}(a) = 1\right\} \mathbb{P}\{s' = a\} + \\ \sum\limits_{\substack{a \in S' \\ \pi_k(a) = 1}} & \mathbb{P}\left\{\gamma \circ b \circ t \circ \Psi^{-1}(a) = 0\right\} \mathbb{P}\{s' = a\} \end{array} \right\} \end{aligned}$$

As in the wordwise case, we need to take existence and uniqueness conditions into account. Similar observations to the ones done for the wordwise decoding

can be applied here too. Once these conditions are taken care of, we get:

$$\begin{aligned} r_k(y,t) &= \chi\ \{\ \mathbb{P}\{\pi_k(s')=1|\,y\} > \mathbb{P}\{\pi_k(s')=0|\,y\}\ \} \\ &= \chi\left\{ \sum_{\substack{a\in S' \\ \pi_k(a)=1}} \mathbb{P}\{s'=a|\,y\} > \sum_{\substack{a\in S' \\ \pi_k(a)=0}} \mathbb{P}\{s'=a|\,y\} \right\} \\ &= \chi\left\{ \sum_{\substack{a\in S' \\ \pi_k(a)=1}} \mathbb{P}\{s'=a|\,y\} > \tfrac{1}{2} \right\} \end{aligned}$$

where χ is of course the indicator function of the subscripted set. Notice that the a posteriori probabilities $\mathbb{P}\{s'=a|\,y\}$ depend from t. Let us stress that this bitwise decoding strategy is different from the wordwise strategy, and that, in the same conditions, can lead to different decoded messages. In this sense it is sometimes considered "suboptimal", the optimal decoding strategy being considered the wordwise, global one. Of course this consideration is not really meaningful - optimality is by definition associated to the minimisation of a cost function, and it is perfectly natural that two different cost functions lead to different strategies.

Let us finally observe that the probability of error for the entire codeword, that is, the probability that at least one bit of the representation is incorrectly decoded, is bounded between $\mathbb{P}_{berr}$ and $m\mathbb{P}_{berr}$. Notice that this probability range needs not to be related with the error probability we get by wordwise ML decoding, since it comes from a different decoding strategy, i.e. bitwise MAP decoding.

2.5 Random Coding

In the two previous paragraphs we have generally summarised coding theory and channel coding. Here we introduce random coding, a particular strategy for producing "good" encoders.

We will first give a new geometric notion, or interpretation, of "goodness" for the encoders, followed by the introduction to the random coding idea and a brief historical sketch of coding theory's focus. Then we present the "annealed" approach to the computation of the cost functions, and its consequences for the bitwise and wordwise error probabilities.

2.5.1 A geometric notion of a code's "goodness"

As introduced in the first two paragraphs, the "goodness" of a channel code may be abstractedly described through a cost function, as for example the described quantities $\mathbb{P}_{werr}$ and $\mathbb{P}_{berr}$.

Anyway, the model engineers have usually in mind is much specific and more detailed than even the probabilistically-enriched setup presented in the last section of first paragraph. First of all, let us import the extra properties introduced in the second paragraph: let the source measure λ induced by the random message s_0 be uniform in the discrete finite source space. The noise model

$$b : (\Omega, \mathcal{F}, \mathbb{P}) \to X^X$$

is too general to be considered broadly: in many cases the noise is modeled as a "noiseword", a random variable on X which is summed to the transmitted codeword to produce the received word:

$$b' : (\Omega, \mathcal{F}, \mathbb{P}) \to X$$

$$b\,(x) := b' + x$$

Such a case is called linear or additive noise - of course, this requires the hypothesis of X being a linear space. Moreover, the distribution of b' is typically taken to be zero mean (a zero value existing because of the linear space requirement on X), single-peaked around zero and symmetric, this last hypothesis requiring a further topologic space assumption on X. Finally, the codeword space X is often described as a finite-dimensional vector space. This allows us to take further hypotheses on the mutual properties of the j-th components b'_j of the random noiseword b':

$$b'_j := \pi_j\,(b')$$

with the natural definition of the j-th component projection operator $\pi_j\,(\cdot)$. The simplest, strongest and most common hypotheses are equidistribution and mutual independence (IID) between all the components.

The basic idea behind this furtherly enriched model is that all the messages are equally likely and that the noise acts limitedly and locally. Limitedly in the sense of the linearity and the zero-mean single-peak hypotheses. Locally in the sense of the IID hypothesis.

This idea suggests a geometric intuition of the noise and a geometric interpretation the goodness of an encoder: the noise "blurs" the codewords into fuzzy spots, so it is possible for a codeword to be "mistaken" for another nearby at decoding. An encoder will then be good when the codewords are as symmetrically far apart as possible from each other, so that the fuzzy spots created by the noise will not overlap, or will do less so.

Considering now the decoding aspect, this geometric interpretation appears to couple more closely with the wordwise than with the bitwise strategy. In fact, this geometric solution minimizes the wordwise error probability (when coupled with a wordwise MAP decoder) if the hypotheses introduced in this section hold - the noise being gaussian on $\mathbb{R}^n$ for instance. But let us notice that if furthermore the cardinality of S is 2^m for a certain $m \in \mathbb{N}$, then the bit representations are equally distributed and all bits are so as well. In this case the optimal symmetric solution for bitwise MAP decoding would again be the same. However, the mapping between the bit representations and the codewords can

have an impact on performance, and in some cases there are non-symmetrical solutions that perform better.

2.5.2 Random coding's fundamental idea

Random coding is a perhaps misleading name for a strategy to find geometrically good codes that dates back to Claude E. Shannon and its foundational coding theory paper of the late '40s.

Let S_n be a finite discrete space such that $|S_n| = 2^{\lfloor Rn \rfloor}$ and consider a sequence $\{x_s^n; s \in S_n; \}_{n\in\mathbb{N}}$ of families of independent and identically distributed (IID) random variables

$$x_s^n : (\Omega, \mathcal{F}, \mathbb{P}) \to (X_n, \Xi_n) = (\mathbb{R}^n, \mathcal{B}^n)$$

so that they are mutually independent with respect to a sequence of random messages $s_{0,n} : (\Omega, \mathcal{F}, \mathbb{P}) \to S_n$ and a sequence of noisewords $b_n : (\Omega, \mathcal{F}, \mathbb{P}) \to X_n^{X_n}$, and moreover that the distribution of the x_s^n is uniform in some subset X_n' of the codeword space X_n. Such a region X_n' could be the projection onto X_n of a power constraint W' parametrised by $w \in [0, \infty]$:

$$X_n' := \left\{ x \in X_n : \frac{\|x\|_2}{n} \leq w \right\}$$

$$W' := \left\{ c \in C : \max_{s\in S} \frac{\|t(s)\|_2}{n} \leq w \right\}$$

defined analogously to the energy constraint $W \subseteq C$ introduced in the last section of the first paragraph:

$$W := \left\{ c \in C : \max_{s\in S} \|t(s)\|_2 \leq w \right\}$$

where $w \in [0, \infty]$ was an energy parameter instead of a power one. Shannon's fundamental observation is that the empirical measure defined by a set of $|S_n|$ realisations of such codewords approaches almost surely a Poisson measure with density

$$\frac{2^{\lfloor Rn \rfloor}}{B^n(nw)}$$

on X_n' as $n \to \infty$, where $B^n(nw)$ is the hypervolume of the n-dimensional hyperball of radius nw. This means that, as n increases, a random encoder t_ω^n (which we call simply, albeit improperly, a random code) consituted by $|S_n|$ realisations of the random variables x_s:

$$\begin{aligned} t_\omega^n : \quad & S_n \to X_n \\ & s \mapsto x_s^n(\omega) \end{aligned}$$

$$\begin{aligned} t^n : \quad & (\Omega, \mathcal{F}, \mathbb{P}) \to X_n^{S_n} \\ & \omega \mapsto t_\omega^n \end{aligned}$$

tends to be a good encoder in the geometric sense - meaning that the codewords tend to be equally spaced. Using this observation in his homonimous theorem, Shannon was able to prove that the wordwise error probability decays asymptotically to zero almost surely according to the ensemble $\left(X_n^{S_n}, \mathcal{B}^{n|S_n|}, \gamma_n\right)$ of random codes, defined as the probability space induced by the random code t onto the encoder measurable space $\left(X_n^{S_n}, \mathcal{B}^{n|S_n|}\right)$:

$$\gamma_n(A) := \mathbb{P}\left(\ t^{-1}(A)\ \right) \quad \forall A \in \mathcal{B}^{n|S_n|}$$

Such an ensemble is taken to be independent from the noise and message: the three random variables

$$s_0^n : (\Omega, \mathcal{F}, \mathbb{P}) \to S_n$$
$$b^n : (\Omega, \mathcal{F}, \mathbb{P}) \to X_n$$
$$t^n : (\Omega, \mathcal{F}, \mathbb{P}) \to X_n^{S_n}$$

are mutually independent from each other. We point out that, in his theorem, Shannon used the probability of codeword error as a cost function, but did not choose wordwise MAP decoding as the decoding strategy, opting instead for something called "group decoding" which, although clearly suboptimal when coupled with wordwise error probability, appeared to be more simply manageable in the proof and capable of approaching wordwise MAP decoding itself. This is a kind of theoretical dual of the operative engineering strategy, hinted in the last section of the first paragraph, that looks for a simple suboptimal iterative decoding algorithm instead of implementing the sometimes computationally intensive wordwise MAP decoder.

2.5.3 Algebraic structure versus randomness

Random codes present a batch of indesiderable properties like the lack of internal symmetry (their codewords have not a group property in general) which makes the probability of a codeword detection error depending on the codeword itself (an unwelcome effect), or the presence of a quota of "bad codewords" (that is, couple of codewords close to each other) in sufficiently crowded codes, or the problem to actually choose an encoder from the ensemble avoiding bad picks. Moreover, real codes are finite lenght codes, so the random approach to find good but asymptotically long codes has not been considered of practical importance for quite some time after Shannon's founding paper. Instead, the engineering community focused on other ideas to construct finite geometrically good codes, the most popular being of algebraic descent: the codewords of a code are viewed as the elements of a discrete linear group to be chosen so that the minimal hamming distance between any two codewords is as high as possible. Of course, this strategy is not asymptotic and allows to produce finite codes of arbitrary lenght. This line of research flourished for 40 years, from the '50s to the '90s, but without producing codes with a performance really close the theoretical bound proven by Shannon. Then in the early '90s Claude Berrou and Alain Glavieux published a groundbreaking article on the so called "turbo codes", a very simple

way to produce pseudorandom codes of arbitrary length n which are very good and, as n increase, approach exponentially fast the asymptotic Shannon's bound. These turbo codes introduce an algebraic structure in a pseudorandom encoder, solving many of the issues concerning pure, unstructured random codes. The article sparkled new interest in random coding which continues to the present, along with the "rediscovery" of some previous results by Robert Gallager who devised in the early '50s some pseudorandom encoders similar to the turbo codes, called Low Density Parity Check (LDPC) codes, which can be more easily analised mathematically and in some cases perform even better than the turbo codes themselves.

2.5.4 "Annealed" cost functions

As already hinted in the second section of this paragraph, Shannon was able to prove his theorem not for single codes but for ensemble of them, an ensemble of random codes being the probability space induced by the random code onto the encoder domain X^S. From this theorem he obtained a corollary on the existence of good codes in the asymptotic ensemble, and a proof that asymptotically almost every code in the ensemble is good (i.e., for $n \to \infty$ a random encoder is almost surely good).

Shannon's approach considers the "annealed" cost functions:

$$\begin{aligned} \mathbb{P}_{werr} \quad &: \quad = \mathbb{P}\left\{\omega \in \Omega : r \circ b \circ t\left(s_0\right) \neq s_0\right\} \\ &= \quad \sum_{a \in S} \mathbb{P}\left\{ E_{err} \middle| s_0 = a\right\} \mathbb{P}\left\{s_0 = a\right\} \\ &= \quad \sum_{a \in S} \mathbb{P}\left\{r \circ b \circ t\left(a\right) \neq a\right\} \mathbb{P}\left\{s_0 = a\right\} \end{aligned}$$

$$\begin{aligned} \mathbb{P}_{berr} \quad &: \quad = \frac{1}{m} \sum_{k=1}^{m} \mathbb{P}_{berr,k} \\ &: \quad = \frac{1}{m} \sum_{k=1}^{m} \mathbb{P}\left\{\omega \in \Omega : \pi_k \circ \Psi \circ r \circ b \circ t \circ \Psi^{-1}\left(s'\right) \neq \pi_k\left(s'\right)\right\} \\ &= \quad \frac{1}{m} \sum_{k=1}^{m} \sum_{a \in S'} \mathbb{P}\left\{ E_{err,k} \middle| s'_0 = a\right\} \mathbb{P}\left\{s'_0 = a\right\} \\ &= \quad \frac{1}{m} \sum_{k=1}^{m} \sum_{\substack{a, b \in S' \\ \pi_k(a) \neq \pi_k(b)}} \mathbb{P}\left\{\Psi \circ r \circ b \circ t \circ \Psi^{-1}\left(a\right) = b\right\} \mathbb{P}\left\{s'_0 = a\right\} \end{aligned}$$

where of course s'_0 is the binary representation of $s_{0,}$ instead of the "quenched" ones:

$$\begin{aligned} \mathbb{P}_{werr}(t) &\overset{a.s.}{:=} \sum_{a \in S} \mathbb{P}\{E_{err} | s_0 = a, t\} \mathbb{P}\{s_0 = a\} \\ &= \sum_{a \in S} \mathbb{P}\{r \circ b \circ t(a) \neq a | t\} \mathbb{P}\{s_0 = a\} \end{aligned}$$

$$\mathbb{P}_{berr}(t) \overset{a.s.}{:=} \frac{1}{m} \sum_{k=1}^{m} \mathbb{P}_{berr,k}(t)$$

$$= \frac{1}{m} \sum_{k=1}^{m} \sum_{a \in S'} \mathbb{P}\{E_{err,k} | s'_0 = a, t\} \mathbb{P}\{s'_0 = a\}$$

$$= \frac{1}{m} \sum_{k=1}^{m} \sum_{\substack{a, b \in S' \\ \pi_k(a) \neq \pi_k(b)}} \mathbb{P}\left\{\Psi \circ r \circ b \circ t \circ \Psi^{-1}(a) = b \middle| t\right\} \mathbb{P}\{s'_0 = a\}$$

which should be used to prove results for single encoders (that is, particular instances of the random code t). The akward notation $\mathbb{P}\{\cdot | s_0 = a, t\}$ is an abuse of notation that improves readability - instad of the proper notation of a realisation $\mathbb{P}\{\cdot | s_0 = a, t = t^*\}$ of the conditional probability $\mathbb{P}\{\cdot | s_0, t\}$ where t^* is a parameter. We have, of course:

$$\mathbb{P}_{werr} = \mathbb{E}\{\ \mathbb{P}_{werr}(t)\ \}$$

$$\mathbb{P}_{berr} = \mathbb{E}\{\ \mathbb{P}_{berr}(t)\ \}$$

We borrowed the terms "quenched" and "annealed" from the statistical mechanic literature, where they indicate if certain quantities are averaged also to the respect of a random environment, or are instead computed conditioning to a realization of the environment itself.

The use of the annealed cost functions induces some simplifications in the mathematical analysis of the cost functions themselves. In particular, the fact that the codeword are IID means that the quantities

$$\mathbb{P}\{E_{err,k} | s' = a\}$$

$$\mathbb{P}\{E_{err} | s = a\}$$

are actually independent of a, since any two realisations that can be mapped into one another through a permutation of the codewords are equally likely. Analogously, if $|S| = |S'| = 2^m$ then

$$\mathbb{P}\{E_{err,k} | s' = a\}$$

is also independent from k, because a permutation of the representation's bits can be mapped into a permutation of the codewords of the associated encoder. This leads to the forms:

$$\mathbb{P}_{werr} = \mathbb{P}\left\{ E_{err} | \, s = a \right\} = \mathbb{P}\left\{ r \circ b \circ t \left(a \right) \neq a \right\}$$

where $a \in S$ is arbitrarily chosen, and

$$\begin{aligned} \mathbb{P}_{berr} &= \mathbb{P}_{berr,k} = \mathbb{P}\left\{ E_{err,k} | \, s' = a \right\} \\ &= \sum_{\substack{b \in S' \\ \pi_k(a) \neq \pi_k(b)}} \mathbb{P}\left\{ \Psi \circ r \circ b \circ t \circ \Psi^{-1}(a) = b \right\} \end{aligned}$$

where $a \in S'$ and $k \in \{1, \ldots, m\}$ are arbitrarily chosen as well.

In the next chapters we will give a new proof of Shannon's theorem based on a statistical-mechanic approach and formalised through large deviation theory. We will show how a generalised error probability $\mathbb{P}_\beta$ exhibits an asymptotically exponential behaviour toward zero or one depending on some system parameters. Since, analogously to $\mathbb{P}_{berr}$ and $\mathbb{P}_{werr}$, we will have

$$\mathbb{P}_\beta = \mathbb{E}\left\{ \ \mathbb{P}_\beta(t) \ \right\}$$

we will be able to deduce that $\mathbb{P}_\beta(t)$ will asymptotically be zero or one almost surely.

2.6 The binary noise - binary random coding problem

In this last paragraph of the current chapter we will describe the binary noise -binary random coding problem, which we will call the Shannon problem for short. After some preliminary notation, we proceed to the definition of the constituent elements of the problem: the source space S and random message s_0, the random encoder t and the noise b. We will finally compute the expression of the annealed bitwise error probability, adopting the bitwise MAP decoding strategy. Since we plan to compute a scaling limit (Shannon's theorem being an asymptotic result), all the relevant elements will be dependant on some scaling parameter $n \in \mathbb{N} \setminus \{0\}$, which will afterwards go to ∞.

In the next chapter we will adopt a statistical-mechanic approach to the coding problem, introducing a generalisation of the bitwise MAP decoding strategy which we will show to contain also the wordwise MAP decoding one.

2.6.1 Some notations

Let $p \in [0,1]$ and let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability triple. Let us agree that a real random variable x is distributed according to a $Bernoulli(\pm 1, p)$ law iff

$$\begin{aligned} \mathbb{P}(x = +1) &= 1 - p \\ \mathbb{P}(x = -1) &= p \end{aligned}$$

Let $x^n := (x_1^n, \ldots, x_n^n)$, $y^n := (y_1^n, \ldots, y_n^n)$ be two real vectors of length n. Let $\langle x^n, y^n \rangle$ denote the standard scalar product of the two vectors and $d_1(x^n, y^n) := \sum_i |x_i^n - y_i^n|$ their l_1 distance. When x_i^n and y_i^n take values on $\{-1, +1\}$, the operators d_1, $\langle \cdot, \cdot \rangle$ and the Hamming distance d_H are linked as follows:

$$\begin{aligned} d_H(x^n, y^n) &:= \sum_{i=1,\ldots,n} 1_{x_i^n \neq y_i^n} \\ &= \frac{1}{2} d_1(x^n, y^n) = \frac{n - \langle x^n, y^n \rangle}{2} \end{aligned}$$

2.6.2 The source space

Let $R \in [0, 1]$ be a parameter denoting the *rate* of the code.

Let $S_n = \{0, 1\}^{\lfloor Rn \rfloor}$ be the source space and let Σ_n be the associated power set, with λ_n denoting the uniform probability on it. The source probability triple will then be $(S_n, \Sigma_n, \lambda_n)$. The encoded message s_0 is a random variable

$$s_0 : (\Omega, \mathcal{F}, \mathbb{P}) \to S_n$$

which induces Σ_n and λ_n on S_n:

$$\Sigma_n := \left\{ B \subseteq S_n : s_0^{-1}(B) \in \mathcal{F} \right\}$$

$$\lambda(B) := \mathbb{P}\left(\ s_0^{-1}(B) \ \right) \quad \forall B \in \Sigma_n$$

Since S_n is a space of binary vectors of length $\lfloor Rn \rfloor$, we have a natural binary representation S_n' of S_n induced by the projections of the vectors onto the $\lfloor Rn \rfloor$ coordinates. For this reason we will make from now on no distintion between S_n' and S_n and will always use the notation S_n.

2.6.3 The binary random code

Let $X_n = \{-1, 1\}^n$ be the encoded space. The encoder space will then be $X_n^{S_n}$.

For all $s \in S_n$ and $j = 1, \ldots, n$. let

$$x_{s,j} : (\Omega, \mathcal{F}, \mathbb{P}) \to \{-1, 1\}$$

be independent and identically distributed (IID) random variables with a $Bernoulli(\pm 1, 1/2)$ law, independent from s_0. For any $s \in S_n$, we identify the vector $x_s^n := (x_{s,1}, \ldots, x_{s,n})$ of lenght n with the codeword associated to the message s by the binary random code

$$t(R, n) := \left(x_1^n, \ldots, x_{2^{\lfloor nR \rfloor}}^n \right)$$

$$t(R, n) : (\Omega, \mathcal{F}, \mathbb{P}) \to X_n^{S_n}$$

The probability triple

$$\left(X_n^{S_n}, \Xi_n, \gamma_n \right)$$

endowed by $t(R, n)$ on $X_n^{S_n}$ is denoted by $SRE(R, n)$ and is called the *binary Shannon Random Ensemble* of rate R and length n.

2.6.4 The memoryless binary symmetric noise

For $j = 1, \dots, n$, let

$$w_j : (\Omega, \mathcal{F}, \mathbb{P}) \to \{-1, 1\}$$

be IID random variables with a $Bernoulli(\pm 1, p)$ law, independent from s_0 and $t\,(R, n)$. The noise word

$$w^n := (w_1, \dots, w_n)$$

$$w^n : (\Omega, \mathcal{F}, \mathbb{P}) \to X_n$$

represents the random contribution of a (memoryless) binary symmetric channel on a codeword. For all choices of s, j, let

$$y_{s,j} := x_{s,j} w_j$$

$$y_s^n := (y_{s,1}, \dots, y_{s,n})$$

where y_s^n is the corrupted (noisy) version of the codeword x_s^n, i.e. the received codeword when s is transmitted. The noise here introduced is multiplicative, not linear, but it is clear that with the map

$$a = (-1)^b$$

it becomes additive at the exponent, eventually with a modulo 2 for a bijection.

2.6.5 The annealed bitwise error probability

In this last section of the paragraph we compute the annealed bit error probability

$$\mathbb{P}_{berr} = \mathbb{P}_{berr,k} = \mathbb{P}\left\{ E_{err,k} \middle|\, s = a \right\}$$

associated to the bitwise MAP decoding strategy for the Shannon problem.

We can write the law of $y_{s_0}^n$ given the code $t\,(R, n)$ and the message s_0. For any choice of $z^n := \{z_1, \dots, z_n\} \in \{-1, +1\}^n$, we have that $\mathbb{P}$-a.s.

$$\begin{aligned}
\mathbb{P}\left\{ y_{s_0}^n = z^n | s_0, t\,(R, n) \right\} &= \mathbb{P}\left\{ y_{s_0}^n = z^n | x_{s_0}^n \right\} \\
&= \prod_{j=1,\dots,n} \mathbb{P}\left\{ y_{s_0,j} = z_j | x_{s_0,j} \right\} \\
&= \prod_{j=1,\dots,n} \left[\sqrt{p\,(1-p)} \left(\sqrt{\frac{1-p}{p}} \right)^{z_j x_{s_0,j}} \right] \\
&= (p\,(1-p))^{n/2} \left(\sqrt{\frac{1-p}{p}} \right)^{n - 2 d_H\left(z^n, x_{s_0}^n \right)}
\end{aligned}$$

Bayes formula yields the conditional law of s_0, given the code $t\,(R, n)$ and the received message $y_{s_0}^n$; for any $s \in S_n$ we have:

$$\mathbb{P}\left\{ s_0 = s | y_{s_0}^n, t\,(R, n) \right\} = K \exp\left\{ -\beta d_H\left(y_{s_0}^n, x_{s_0}^n \right) \right\}, \quad \mathbb{P}\text{-a.s.}$$

where K does not depend on s because all source messages are equally likely, and

$$\beta := \left(\log \frac{1-p}{p} \right)$$

For $k = 1, \ldots, \lfloor nR \rfloor$, let $\pi_k^n : S_n \to \{0,1\}$ be the map associating each message s to the k-th binary digit $\pi_k^n(s)$ of its representation. An error occurs at the bit k, if the decoded digit is different from $s_{0,k}$. By definition of bitwise MAP decoding, the decoded digit is 1 iff

$$\Delta_k^n \leq 0$$

where

$$\begin{aligned} \Delta_k^n &:= \sum_{s:\pi_k^n(s)=0} \mathbb{P}\left\{ s_0 = s | y_{s_0}^n, t(R,n) \right\} - \sum_{s:\pi_k^n(s)=1} \mathbb{P}\left\{ s_0 = s | y_{s_0}^n, t(R,n) \right\} \\ &= K \left[\sum_{s:\pi_k^n(s)=0} \exp\left\{ -\beta d_H\left(y_{s_0}^n, x_{s_0}^n \right) \right\} - \sum_{s:\pi_k^n(s)=1} \exp\left\{ -\beta d_H\left(y_{s_0}^n, x_{s_0}^n \right) \right\} \right] \end{aligned}$$

otherwise 0 is decoded. We have

$$E_{err,k} = \left\{ \omega \in \Omega : (-1)^{\pi_k^n(s)} \Delta_k^n \leq 0 \right\}$$

Let us now take $s_0 = 0$, $k = 1$ (recall that we can arbitrarily fix them in the annealed computation) and write the formula for $\mathbb{P}_{berr}$:

$$\begin{aligned} \mathbb{P}_{berr}(R,n) &= \mathbb{P}(\Delta_1^n \leq 0 | s_0 = 0) \\ &= \mathbb{P}\left(\sum_{s:\pi_k^n(s)=0} \exp\left\{ -\beta d_H\left(y_{s_0}^n, x_{s_0}^n \right) \right\} \leq \sum_{s:\pi_k^n(s)=1} \exp\left\{ -\beta d_H\left(y_{s_0}^n, x_{s_0}^n \right) \right\} | s_0 = 0 \right) \end{aligned}$$

3 The ansatz of statistical-mechanics

3.1 Introduction

In this chapter we plan to describe a recent statistical-mechanic description of random channel codes pioneered by Nicolas Sourlas.

After some general definitions, we introduce some statistical-mechanic concepts. Then we describe how to use the statistical-mechanic language to characterise random channel codes, and most importantly to identify their relevant quantities. Finally we apply this new point of view to our focus problem: the Shannon problem, as defined in the previous chapter.

In the next chapter we will develop some large deviation theorems to obtain mathematically rigorous results on the behaviour of the quantities identified through this statistical-mechanic ansatz.

3.2 Definitions

Let $\mathbb{R}^+ :=]0, \infty[$, $\mathbb{R}_0^+ := [0, \infty[$, $\mathbb{R}_\infty^+ :=]0, \infty]$ and $\mathbb{R}_{0,\infty}^+ := [0, \infty]$ in the natural real extended topology. Let us take $0 \log 0 = 0$ throughout the chapter and the rest of the thesis, so that the function $x \log x$ is well defined and continuous also at $x = 0$.

3.2.1 Relative entropy or Kullbach-Leibler distance

Let $(X, \mathcal{X})$ be a measurable space, with two finite measures μ and λ. Let furthermore λ be absolutely continuous with respect to μ, that is:

$$\mu(A) = 0 \implies \lambda(A) = 0 \quad \forall A \in \mathcal{X}$$

In this case Radon-Nikodym's theorem gives us the existence of a measurable non-negative function

$$f : (X, \mathcal{X}) \to \mathbb{R}_0^+$$

so that:

$$\lambda(A) := \int_A d\lambda = \int_A f d\mu \quad \forall A \in \mathcal{X}$$

We can then use the notation:

$$f =: \frac{d\lambda}{d\mu}$$

and define the relative entropy of the measure λ respect to the measure μ, also called the Kullbach-Leibler distance of the measure λ from the measure μ:

$$D_b(\lambda||\mu) := \int_X \log_b \frac{d\lambda}{d\mu} d\lambda$$

where $b \in \mathbb{R}^+$ is the base of the logarithm and is intended equal to the Neper number e when not explicitely indicated. Of course, we have:

$$D_b(\lambda||\mu) = \frac{1}{\log b} D(\lambda||\mu)$$

Note furthermore that in general the relative entropy is not a distance in the topological sense because it lacks symmetry, even in the case when both measures are absolutely continuous with respect to each other:

$$D_b(\lambda||\mu) \neq D_b(\mu||\lambda)$$

When μ is a probability measure then the relative entropy has the property of non-negativity, with the zero value achieved if and only if $f = 1$ μ-a.s., because of the strict concavity of the logarithm:

$$\begin{aligned} -D_b(\lambda||\mu) &= \int_X \left(\log_b \frac{1}{f}\right) d\lambda \\ &\leq \log_b \int_X \frac{1}{f} d\lambda \\ &= \log_b \int_X \frac{d\mu}{d\lambda} d\lambda = 0 \end{aligned}$$

3.2.2 Binary relative entropy and Gilbert-Varshamov distance

Let $X := \{x_1, x_2\}$ be a binary space with μ and λ two probability measures so that:

$$\begin{aligned} \lambda(x_1) = z \quad \lambda(x_2) = 1 - z \quad z \in [0,1] \\ \mu(x_1) = y \quad \mu(x_2) = 1 - y \quad y \in]0,1[\end{aligned}$$

Then we can write the following form for the relative entropy of λ with respect to μ:

$$D_b(z||y) := \begin{cases} z\log_b\left(\frac{z}{y}\right) + (1-z)\log_b\left(\frac{1-z}{1-y}\right) & \text{if} \quad z \in]0,1[\\ \log_b\left(\frac{1}{1-y}\right) & \text{if} \quad z = 0 \\ \log_b\left(\frac{1}{y}\right) & \text{if} \quad z = 1 \end{cases}$$

where we identified the measures by their relevant probability ratios z and y Notice that $D_b(z||y)$ is strictly convex, lower-semicontinuous, always non negative and equal to zero if and only if $z = y$. For any $R \in [0,1]$, let the *Gilbert-Varshamov distance* $\delta_{GV}(R)$ be the real number such that

$$D_b\left(\delta_{GV}(R) \middle\| \frac{1}{2}\right) = R\log_b 2$$

or equivalently

$$D_2\left(\delta_{GV}(R) \middle\| \frac{1}{2}\right) = R$$

with $\delta_{GV}(R) \in \left[0, \frac{1}{2}\right]$. In order to prove existence and unicity of the Gilbert-Varshamov distance, we just need to notice that $f(\delta) := D_2\left(\delta \middle\| \frac{1}{2}\right)$ is continuous

for $\delta \in \left[0, \frac{1}{2}\right]$ and

$$\delta \in \left[0, \frac{1}{2}\right] \Longrightarrow \frac{\partial}{\partial \delta} f(\delta) = \log_b \frac{\delta}{1-\delta} \leq 0$$

$$f\left(\left[0, \frac{1}{2}\right]\right) = [0, \log_b 2]$$

3.2.3 Discrete entropy

Let $(X, \mathcal{X}, \lambda)$ be a discrete finite probability space with expectation $\mathbb{E}_\lambda$. We define the entropy of the probability measure λ as follows:

$$H_b(\lambda) := \mathbb{E}_\lambda \{-\log_b \lambda(\cdot)\} := -\sum_{x \in X} \lambda(x) \log_b \lambda(x)$$

Entropy is a non-negative quantity:

$$0 \leq \lambda(x) \leq 1 \quad \forall x \in X \implies H_b(\lambda) \geq 0$$

Entropy is equal to zero if the probability mass is degenerate:

$$\delta_a(x) := \left\{ \begin{array}{lll} 1 & \text{iff} & x = a \\ 0 & \text{iff} & x \neq a \end{array} \right. \implies H_b(\delta_a) = 0$$

Let $\alpha \in [0,1]$ and λ_1, λ_2 be two probability measures on $(X, \mathcal{X})$. Let $\lambda_3 := \alpha\lambda_1 + (1-\alpha)\lambda_2$ be their linear combination probability measure:

$$\lambda_3(x) := \alpha\lambda_1(x) + (1-\alpha)\lambda_2(x) \quad \forall x \in X$$

Let us notice that $\lambda_1$and λ_2 are absolutely continuous with respecto to λ_3 for $\alpha \in (0,1)$, and that:

$$\begin{aligned} H_b(\lambda_3) \quad : \quad & = -\sum_{x \in X} \lambda_3(x) \log_b \lambda_3(x) \\ & = \alpha \sum_{x \in X} \lambda_1(x) \log_b \frac{1}{\lambda_3(x)} + (1-\alpha) \sum_{x \in X} \lambda_2(x) \log_b \frac{1}{\lambda_3(x)} \\ & = \alpha \sum_{x \in X} \lambda_1(x) \log_b \frac{\lambda_1(x)}{\lambda_3(x)} - \alpha \sum_{x \in X} \lambda_1(x) \log_b \lambda_1(x) \\ & \quad + (1-\alpha) \sum_{x \in X} \lambda_2(x) \log_b \frac{\lambda_2(x)}{\lambda_3(x)} - (1-\alpha) \sum_{x \in X} \lambda_2(x) \log_b \lambda_2(x) \\ & = : \alpha D_b(\lambda_1 || \lambda_3) + \alpha H_b(\lambda_1) \\ & \quad + (1-\alpha) D_b(\lambda_2 || \lambda_3) + (1-\alpha) H_b(\lambda_2) \\ & \geq \alpha H_b(\lambda_1) + (1-\alpha) H_b(\lambda_2) \end{aligned}$$

That is, entropy is a strictly concave functional, since

$$\lambda_1 = \lambda_2 = \lambda_3 \Leftrightarrow \left\{ \begin{array}{l} D_b(\lambda_1 || \lambda_3) = 0 \\ D_b(\lambda_2 || \lambda_3) = 0 \end{array} \right.$$

This last result implies that entropy is equal to zero if and only if $\lambda = \delta_a$, for some $a \in X$. Let now μ be the uniform probability measure on X:

$$\mu(x) = \frac{1}{|X|} \quad \forall x \in X$$

Clearly λ is absolutely continuous with respect to μ, and furthermore we have:

$$D_b(\lambda||\mu) = \sum_{x \in X} \lambda(x) \log_b \frac{\lambda(x)}{\frac{1}{|X|}} = \log_b |X| - H_b(\lambda)$$

Since $D_b(\lambda||\mu) \geq 0$ and strictly convex, we obtain:

$$H_b(\lambda) \leq \log_b |X|$$

$$H_b(\lambda) = \log_b |X| \Longleftrightarrow \lambda = \mu$$

Let us resume the just proven properties of the entropy $H_b(\lambda)$ of a discrete finite probability measure λ: $H_b(\lambda)$ is a strictly concave functional bounded between 0 and $\log_b |X|$, and achieves these extremal values if and only if λ is degenerate or uniform respectively.

One final observation, of which we will not give proof here, is that entropy is the only functional with all the aforementioned properties.

3.3 Statistical mechanics at equilibrium: the finite canonical ensemble

We introduce now some general concepts from statistical mechanics. Let S be a discrete finite "state space" and $E(s)$ a given "state energy" function:

$$\begin{array}{rl} E: & S \to \mathbb{R}_0^+ \\ & s \mapsto E(s) \end{array}$$

Let us imagine that S represents the space of possible configurations of a statistical system. Let us introduce an external reservoir (or universe) completely characterised by its given and constant normalized temperature

$$T := \frac{1}{\beta}, \quad T \in \mathbb{R}^+, \quad \beta \in \mathbb{R}^+$$

Then the fundamental "canonical ensemble" principle in statistical mechanics says that the probability $\mu_\beta(s)$ of a configuration s to be realised by the system at equilibrium with the reservoir has the form

$$\mu_\beta(s) := \frac{1}{Z(\beta)} e^{-\beta E(s)}$$

$$Z(\beta) := \sum_{s \in S} e^{-\beta E(s)}$$

where the normalizing constant $Z(\beta)$ is named the systems' partition function. The probability measure μ_β is defined on the power set $P(S)$ of S, inducing the probability triple

$$\left(S, P(S), \mu_\beta\right)$$

with expectation $\mathbb{E}_{\mu_\beta}$. We do not dare to give an external definition of equilibrium here, since we should introduce a description of the dynamics according to which the equilibrium is a stable state. Let it just be defined by the canonical ensemble principle itself, changed into a convenient definition. We just add that the term canonical relates to the possible exchanges of energy between the system and the reservoir. If no exchanges altogheter were possible then the system would be microcanonical, while if the cardinality of the state space were allowed to fluctuate then the system would be grand canonical.

Note that $Z(\beta)$ is a continuous and infinitely differentiable function, since the cardinality of S is by hypothesis finite and the exponential function is continuous and infinitely differentiable. Now, most of the interesting quantities of the system are obtained through expectations, as for instance the average energy $E(\beta)$:

$$\begin{aligned} E(\beta) \quad : \quad & = \mathbb{E}_{\mu_\beta}\{E(s)\} \\ & = \sum_{s\in S} E(s)\,\mu_\beta(s) \\ & = \frac{1}{Z(\beta)}\sum_{s\in S} E(s)\,e^{-\beta E(s)} \\ & = \frac{1}{Z(\beta)}\sum_{s\in S}\left(-\frac{\partial}{\partial\beta}e^{-\beta E(s)}\right) = -\frac{1}{Z(\beta)}\frac{\partial}{\partial\beta}Z(\beta) = -\frac{\partial}{\partial\beta}\log Z(\beta) \end{aligned}$$

or the average entropy $H(\beta)$:

$$\begin{aligned} H(\beta) \quad : \quad & = \mathbb{E}_{\mu_\beta}\left\{-\log\mu_\beta(s)\right\} \\ & = \mathbb{E}_{\mu_\beta}\{\log Z(\beta) + \beta E(s)\} = \log Z(\beta) + \beta E(\beta) \\ & = \log Z(\beta) - \beta\frac{\partial}{\partial\beta}\log Z(\beta) \end{aligned}$$

These simple computations give us a hint on why in statistical mechanic theory it is usually considered that the knowledge of the partition function $Z(\beta)$ equates to the knowledge of the system itself: although the system is identified at equilibrium by the energies $E(s)$ or equivalently by the probability measure μ_β, the partition function alone is sufficient to compute E and H. The logarithm of the partition function $\log Z(\beta)$ is also equally meaningful and, when divided by β, it is called the free energy of the system.

In the following sections of the current paragraph, after briefly studying the limiting behaviour of the probability measure μ_β, we outline some general properties of the partition function and the free energy. Finally we will introduce the concepts of scaling limit and random system.

3.3.1 The limiting behaviour of the measure μ_β

Let us now study the behaviour of μ_β as β goes to 0 or ∞. We have

$$\lim_{\beta\to 0}\mu_\beta(s) = \lim_{\beta\to 0}\frac{e^{-\beta E(s)}}{\sum\limits_{a\in S} e^{-\beta E(a)}} = \tfrac{1}{|S|} \quad \forall s\in S$$

and since $|S|<\infty$ we can say that μ_β converges uniformly toward the uniform probability distribution on S as $\beta\to 0$. Let now:

$$E_{\min} := \min_{s\in S} E(s)$$

$$S_{\min} := \{s\in S : E(s) = E_{\min}\}$$

$$\emptyset \neq S_{\min} \subseteq S$$

we have:

$$\mu_\beta(s) = \frac{e^{-\beta[E(s)-E_{\min}]}}{\sum\limits_{a\in S} e^{-\beta[E(a)-E_{\min}]}} = \frac{e^{-\beta[E(s)-E_{\min}]}}{|S_{\min}| + \sum\limits_{a\in S\setminus S_{\min}} e^{-\beta[E(a)-E_{\min}]}} \quad \forall s\in S$$

and since:

$$\lim_{\beta\to\infty}\sum_{a\in S\setminus S_{\min}} e^{-\beta[E(a)-E_{\min}]} = 0$$

$$\begin{aligned}\lim_{\beta\to\infty} e^{-\beta[E(s)-E_{\min}]} &= \begin{cases} 1 & \text{if} \quad s\in S_{\min} \\ 0 & \text{if} \quad s\notin S_{\min}\end{cases} \\ &= : \chi_{S_{\min}}(s)\end{aligned}$$

we obtain:

$$\lim_{\beta\to\infty}\mu_\beta(s) = \frac{1}{|S_{\min}|}\chi_{S_{\min}}(s)$$

that is, μ_β coverges (again uniformly) to the uniform probability measure on $S_{\min}$ as $\beta\to\infty$.

We can give a physical interpretation of these observations, based on some intuition of energy exchanges associated to state transitions (in order to be precise this would require of course some words on the dynamics of the system, which we choosed to omit). At low temperature ($\beta\to\infty \iff T\to 0$) the system "collapses" on its low energy states $S_{\min}$, which are equally likely. There is not enough reservoir energy, whose density is the temperature T, to be absorbed by the system allowing a more energetic state to be occupied. On the other hand, at high temperature ($T\to\infty \iff \beta\to 0$) the state energies become irrelevant and all states are uniformly occupied, since the reservoir temperature is so high that the finite state energy differences become negligible when energy is exchanged.

3.3.2 The partition function $Z(\beta)$

We already observerd that the partition function

$$Z(\beta) := \sum_{s \in S} e^{-\beta E(s)}$$

is continuous and infinitely differentiable in β. Let us now notice that:

$$Z(\beta) > 0 \forall \beta \in \mathbb{R}^+$$

$$\lim_{\beta \to 0} Z(\beta) = |S| =: Z(0)$$

$$\lim_{\beta \to \infty} Z(\beta) = |E_0| =: Z(\infty)$$

with the definition:

$$E_0 := \{s \in S : E(s) = 0\}$$

and that:

$$\frac{\partial}{\partial \beta} Z(\beta) = \sum_{s \in S} \frac{\partial}{\partial \beta} e^{-\beta E(s)} = -Z(\beta) E(\beta) \leq 0 \forall \beta \in \mathbb{R}^+$$

$$\frac{\partial}{\partial \beta} Z(\beta) = 0 \Longleftrightarrow E \equiv 0 \Longleftrightarrow \left\{ \ E(s) = 0 \quad \forall s \in S \ \right\}$$

$$\begin{aligned} \frac{\partial^2}{\partial \beta^2} Z(\beta) &= \sum_{s \in S} \frac{\partial^2}{\partial \beta^2} e^{-\beta E(s)} \\ &= \sum_{s \in S} E^2(s) e^{-\beta E(s)} \geq 0 \ \ \forall \beta \in \mathbb{R}^+ \end{aligned}$$

$$\frac{\partial^2}{\partial \beta^2} Z(\beta) = 0 \Longleftrightarrow \left\{ \ E(s) = 0 \quad \forall s \in S \ \right\}$$

So if we discard the trivial case $E(s) = 0 \ \forall s \in S$ then the partition function $Z(\beta)$ is a positive, convex and decreasing function bounded between the two limiting values $Z(0)$ and $Z(\infty)$. In the trivial case it is clearly constant in β.

Let us spend some words on the definition on E_0. Since the energies $E(s)$ are by hypothesis non negative and S is discrete and finite, we could think of translating uniformly the energies $E(s)$ so that $E_{\min} = 0$ and consequently $E_0 \neq \emptyset$. But we prefer not to do such a normalisation, because we will later let the energies $E(s)$ be picked according to random variables and we will consider the limit $|S| \to \infty$. In such a situation, the normalisation would be random as well and varying with the cardinality of S. Moreover, the limiting energy distribution $\{E(s)\}_{s \in S}$ may tends to zero without achieving the value. For all these reasons we stick to a setup where E_0 can also be empty.

3.3.3 The free energy $F(\beta)$

Analogously to the properties of the partition function $Z(\beta)$ just shown, the quantity $\log Z(\beta)$ is continuous, infinitely differentiable in β and

$$\lim_{\beta\to 0} \log Z(\beta) = \log Z(0)$$

$$\lim_{\beta\to \infty} \log Z(\beta) = \log Z(\infty)$$

$$\tfrac{\partial}{\partial\beta} \log Z(\beta) = -E(\beta) \leq 0 \quad \forall \beta \in \mathbb{R}^+$$

So again, discarding the trivial case $E(s) = 0\ \forall s \in S$, $\log Z(\beta)$ is a decreasing function bounded between the two limiting values $\log Z(0)$ and $\log Z(\infty)$. Notice now that if the minimal energy states $S_{\min}$ have energy $E_{\min} = 0$ then $Z(\infty) = |S_{\min}| \geq 1$ and so $\lim_{\beta\to\infty} \log Z(\beta) \geq 0$, otherwise $\lim_{\beta\to\infty} \log Z(\beta) = -\infty$.

Let us now make three observations. First, let us notice that:

$$\begin{aligned}
\lim_{\beta\to\infty} E(\beta) &= \lim_{\beta\to\infty} \sum_{s\in S} E(s)\,\mu_\beta(s) \\
&\overset{\text{d.c.t.}}{=} \sum_{s\in S} E(s) \lim_{\beta\to\infty} \mu_\beta(s) \\
&= \sum_{s\in S} E(s)\,\frac{1}{|S_{\min}|}\chi_{S_{\min}}(s) \\
&= E_{\min}
\end{aligned}$$

so we have

$$\lim_{\beta\to\infty} \frac{\partial}{\partial\beta} \log Z(\beta) = -E_{\min}$$

Secondly, from the entropy bounds we get:

$$0 \leq H \leq \log|S| \Longrightarrow \frac{\partial}{\partial\beta}\log Z(\beta) \leq \frac{\log Z(\beta)}{\beta} \leq \frac{\partial}{\partial\beta}\log Z(\beta) + \frac{\log|S|}{\beta}$$

So since $\log|S|$ is finite we have that

$$\lim_{\beta\to\infty} \frac{\frac{\partial}{\partial\beta}\log Z(\beta)}{\frac{\log Z(\beta)}{\beta}} = 1$$

In third place, observe that a uniform shift in the state energy function

$$\begin{aligned}
E_\epsilon : \quad & S \to \mathbb{R}_0^+ \\
& s \mapsto E(s) + \epsilon
\end{aligned}$$

$$\epsilon \in \mathbb{R}^+$$

induces only a rescaling in the partition function:

$$Z_\epsilon(\beta) := \sum_{s\in S} e^{-\beta E_\epsilon(s)} = e^{-\beta\epsilon} Z(\beta)$$

and equivalently a traslation of the free energy:

$$\frac{1}{\beta}\log Z_\epsilon(\beta) = \frac{1}{\beta}\log\sum_{s\in S} e^{-\beta E_\epsilon(s)} = \frac{1}{\beta}\log Z(\beta) - \epsilon$$

These three facts suggest to use a negative, β-normalised version of $\log Z(\beta)$ for the free energy F, in case we are interested in analysing the behaviour of the system at all temperatures including $T = 0 \iff \beta = \infty$ and we are not sure that the minimal state energy $E_{\min}$ is actually zero (for lack of knowledge or inherent impredictability of the system's nature, see the section "random systems" of the present paragraph):

$$F(\beta) := -\frac{1}{\beta}\log Z(\beta)$$

Of course, we have:

$$\lim_{\beta\to\infty} F(\beta) = E_{\min}$$

$$\lim_{\beta\to 0} F(\beta) = -\infty$$

Let us now study its convexity:

$$\begin{aligned}
\frac{\partial^2}{\partial\beta^2}F(\beta) &= \frac{\partial}{\partial\beta}\left\{\frac{\partial}{\partial\beta}\left[-\frac{1}{\beta}\log Z(\beta)\right]\right\} \\
&= -\frac{2}{\beta^3}\log Z(\beta) + \frac{2}{\beta^2}\frac{\partial}{\partial\beta}\log Z(\beta) - \frac{1}{\beta}\frac{\partial^2}{\partial\beta^2}\log Z(\beta) \\
&\leq -\frac{1}{\beta}\frac{\partial^2}{\partial\beta^2}\log Z(\beta) \\
&= \frac{\left[\frac{\partial}{\partial\beta}Z(\beta)\right]^2 - Z(\beta)\frac{\partial^2}{\partial\beta^2}Z(\beta)}{\beta Z^2(\beta)} \\
&= \frac{\sum_{s,t\in S}\left[E(s)E(t) - E^2(t)\right]e^{-\beta[E(s)+E(t)]}}{\beta Z^2(\beta)} \\
&= \frac{1}{2}\frac{\sum_{s,t\in S}\left[2E(s)E(t) - E^2(t) - E^2(s)\right]e^{-\beta[E(s)+E(t)]}}{\beta Z^2(\beta)} \leq 0
\end{aligned}$$

where the first inequality follows from the entropy inequality:

$$0 \leq H \Longrightarrow \frac{\partial}{\partial\beta}\log Z(\beta) \leq \frac{\log Z(\beta)}{\beta}$$

while the last equality follows from symmetry considerations over the mute parameters s and t. The definition of $F(\beta)$ induces also new forms for the entropy and energy equalities. Recalling that:

$$T = \frac{1}{\beta} \quad \frac{\partial T}{\partial \beta} = -\frac{1}{\beta^2}$$

we have that the equality:

$$\begin{aligned} H(\beta) &= \log Z(\beta) - \beta \frac{\partial}{\partial \beta} \log Z(\beta) \\ &= \beta^2 \frac{\partial}{\partial \beta} \left(-\frac{1}{\beta} \log Z(\beta) \right) \end{aligned}$$

implies:

$$H(T) = -\frac{\partial}{\partial T} F(T)$$

and analogously:

$$\begin{aligned} \beta E(\beta) &= -\beta \frac{\partial}{\partial \beta} \log Z(\beta) \\ &= \beta^2 \frac{\partial F(\beta)}{\partial \beta} + \beta F(\beta) \end{aligned}$$

translates into:

$$\begin{aligned} E(T) &= F(T) - T \frac{\partial}{\partial T} F(T) \\ &= F(T) + TH(T) \end{aligned}$$

3.3.4 The thermodynamic limit: the free energy density $f_n(\beta)$, phase transitions

In statistical mechanics, the scaling limit is the computation of the "macroscopic" properties of the system as the "size" of the system diverges. We need of course to define what we mean by macroscopic and by the size of a system, and in doing so we will take some hypotheses on the structure of the system itself.

The definition of a macroscopic property is the more tricky, since originally the term comes from empirical physics. In this world a macroscopic quantity is a quantity of the system as a whole and as such can be locally measured, producing the same value in any observation independently from the place. Let us keep this intuition in mind for later in this section, when we will be able to bridge it with mathematical definitions.

Let us now suppose that the system consists of n different indentical "subsystems" or "elements", each with a possible state space A. This means that the state space will be $S = A^n$. Let us now recall that the entropy of the system is bounded as follows:

$$0 \leq H(\beta) \leq \log |S| = n \log |A|$$

So it grows at most linearly with n, which we call the size of the system. From now on we append the subscript n to the quantities associated to a system of the same size: $Z_n(\beta)$, $H_n(\beta)$, $E_n(\beta)$, $F_n(\beta)$. Let us also recall that the free energy $F_n(\beta)$ has the following bounds:

$$\begin{array}{ccccc} -\frac{1}{\beta}\log Z(0) & \leq & F_n(\beta) & \leq & -\frac{1}{\beta}\log Z(\infty) \\ || & & & & || \\ -n\frac{1}{\beta}\log|A| & & & & -\frac{1}{\beta}\log|E_0| \end{array}$$

so this one too grows at most linearly in n. In many physical problems the energy $E_n(\beta)$ grows as well asymptotically linearly with n. In the statistical mechanic jargon this linear fenomenon is expressed by saying that these are extensive properties of the system. It becomes then natural to introduce the spatial densities $E_n(\beta)/n$, $H_n(\beta)/n$ and most importantly the free energy density $f_n(\beta)$:

$$f_n(\beta) := \frac{F_n(\beta)}{n} = -\frac{1}{\beta n}\log Z_n(\beta)$$

Analogously to the free energy $F(\beta)$, we have for the density $f_n(\beta)$ the three properties:

$$\lim_{\beta\to\infty} -\frac{\partial}{\partial\beta}\frac{\log Z_n(\beta)}{n} = \frac{E_{\min,n}}{n}$$

$$-\frac{\partial}{\partial\beta}\frac{\log Z_n(\beta)}{n} \geq f_n(\beta) \geq -\frac{\partial}{\partial\beta}\frac{\log Z_n(\beta)}{n} - \frac{\log|A|}{\beta}$$

$$f_{n,\epsilon}(\beta) := -\frac{1}{\beta n}\log Z_{n,\epsilon}(\beta) = f_n(\beta) + \frac{\epsilon}{n}$$

and the relations:

$$\frac{H_n}{n} = -\frac{\partial}{\partial T} f_n(T)$$

$$\begin{aligned} \frac{E_n}{n} &= f_n(T) - \frac{\partial}{\partial T} f_n(T) \\ &= f_n(T) + T\frac{H_n}{n} \end{aligned}$$

To sum it all we now have a quantity, the free energy density $f_n(\beta)$, which does not diverge as β or n diverge, and which enables us to compute the (densities of the) macroscopic properties of the system, entropy and energy in particular. The main problem that can eventually arise in thermodynamical limits is that the free energy may converge to a discontinuous function in β:

$$\limsup_{n\to\infty}\frac{1}{\beta n}\log\sum_{s\in A^n} e^{-\beta E(s)} \neq \liminf_{n\to\infty}\frac{1}{\beta n}\log\sum_{s\in A^n} e^{-\beta E(s)}$$

This would make the energy and entropy densities discontinuous too, of course. Such a discontinuitiy phenomenon is understandably very important, and is

called a phase transition of the system. Notice that its eventual occurrence does not affect the bounded nature of the free energy density, which thus remains a good tool to study the behaviour of the system.

Let us now share some final folklore words on the scaling of the probability μ_β: one of the basic tenets in statistical mechanics is that in real phisical systems the actual realisations of energy density $E_n(s)/n$ or self information density $-\log \mu_\beta(s)/n$ tend, in the absence of phase transitions, to concentrate (in the sense of the Law of Large Numbers or large deviation theory) around their expected values E and.H. This happens, for instance, when the system is composed of independent or weakly dependent subsystems (we do not prove this claim here). When such a concentration happens then the system becomes almost surely deterministic in the scaling limit, as is the case with the thermodynamic model of the ideal gas. Moreover, in the case of weak or no dependence of the subsystems it can can be argued that even arbitrarily small parts of the system (compared to the whole) exibit the concentration property. In this case, the realisation of a local quantity $E'_m(s')/m$ (where $E'_m(s')$ is the energy of a subsystem of size proportional to m, with $m \to \infty$ but $m/n \to 0$) would give information on the global quantity $E_n(s)/n$. Recalling the opening remark on macroscopic quantities, we now can identify these with such concentrating local random variables, or directly with the expected densities toward which they converge. According to this last interpretation, E_n/n and H_n/n are in the limit macroscopic quantities, while f_n is not because it is not directly an expectation of some random variable (it is a logarithm of an expectation). A mathematical computation and analysis of the free energy density would yeld in this framework precise and exhaustive information on the macroscopic behaviour of the system.

Before concluding the section, let us notice that in this minimal and qualitative description we omitted some further criteria used in statistical mechanics to define macroscopic quantities, such some sort of invariance under transformation which is required for physical measurability.

3.3.5 Random systems

As described in the previous parts of the current paragraph, the properties of a canonical statistical-mechanic system at equilibrium are defined by the state energy function $E(s)$.

An interesting way to generalise this model is to let the energy function $E : S \to \mathbb{R}_0^+$ be the realisation of a random variable:

$$\begin{aligned} \mathbf{E}: \qquad & \Omega \to \left(\mathbb{R}_0^+\right)^S \\ & \omega \mapsto \{E_\omega(s), s \in S\} \end{aligned}$$

where

$$(\Omega, \mathcal{F}, \mathbb{P})$$

is the original probability triple with associated expectation $\mathbb{E}$. We can dually

write:

$$\begin{aligned} \mathbf{E}_\omega : \quad & S \to \mathbb{R}_0^+ \\ & s \mapsto E_\omega(s) \quad \forall s \in S \end{aligned}$$

We have now to investigate the relationship between the new probability triple $(\Omega, \mathcal{F}, \mathbb{P})$ and $\left(S, P(S), \mu_\beta\right)$, the one previously defined and associated to the internal state of the system itself. One way of doing this is to assume the existence of the following probability kernel

$$\mathbb{P} : (\mathbf{E}, A) \to \mathbb{P}(A|\,\mathbf{E}) = \frac{\sum_{s\in A} e^{-\beta E_\omega(s)}}{\sum_{s\in S} e^{-\beta E_\omega(s)}}$$

so that,for all $\beta \geq 0$,

$$\mu_{\beta,\omega} := \mathbb{P}(\cdot|\,\mathbf{E})$$

is $\mathbb{P}$-a.s. a probability measure on $(S, P(S))$ and $\mathbb{P}(A|\cdot)$ is $\mathcal{F}$-measurable on Ω for all $A \in P(S)$. Let $\mathbb{E}\{\cdot|\,\mathbf{E}\}$ be the (conditional) expectation associated to $\mathbb{P}(\cdot|\,\mathbf{E})$. In this setup all the previously defined quantities become random variables:

$$E_\omega \overset{\mathbb{P}-a.s.}{:=} \mathbb{E}\{E(s)|\,\mathbf{E}\}$$

$$H_\omega \overset{\mathbb{P}-a.s.}{:=} \mathbb{E}\{-\log \mathbb{P}(s|\,\mathbf{E})|\,\mathbf{E}\}$$

$$F_\omega \overset{\mathbb{P}-a.s.}{:=} \frac{1}{\beta} \log \sum_{s\in S} e^{-\beta E_\omega(s)}$$

and are called the quenched energy, the quenched entropy and the quenched free energy respectively. The mutual relations hold $\mathbb{P}$-almost surely as well:

$$H_\omega \overset{\mathbb{P}-a.s.}{:=} \frac{\partial F_\omega}{\partial T}$$

$$E_\omega \overset{\mathbb{P}-a.s.}{:=} T\frac{\partial F_\omega}{\partial T} - F_\omega$$

Random system models are typically studied in the thermodynamic limit, when a further hypothesis on the concentration of the measure μ holds as well. This goes hand in glove with the qualitative discussion at the end of the previous sections, and constitutes a kind of generalisation of it. Following this line of thoughts, it is meaningful to compute the unconditioned expectations

$$E := \mathbb{E}\{E_\omega\}$$

$$H := \mathbb{E}\{H_\omega\}$$

$$F := \mathbb{E}\{F_\omega\}$$

and to check if and when the unaveraged quantities E_ω, H_ω, F_ω approximate them in some sense in the scaling limit, an indication that the further concentration hypothesis holds. We will see in the next chapter on large deviation

theory that such theory provides the perfect tool for this approach. Notice now that by linearity of the expectation operator we get:

$$H = \frac{\partial F}{\partial T}$$

$$E = T\frac{\partial F}{\partial T} - F$$

The averaged quantities E, H and F are called now the annealed energy, the annealed entropy and the annealed free energy density respectively.

3.4 The statistical mechanic ansatz to random coding

In this paragraph we will outline the duality between channel decoding and statistical-mechanic systems. We will consider random codes, and associating to random systems in the sense descibed above. We will study their limit as the size of the code increases, which translates into the scaling limit of the statistical-mechanic description. From now on let

$$m := \lfloor nR \rfloor$$

$$R \in [0, 1]$$

In particular, we will be interested in studing the conditions so that the probability of error vanishes in the limit:

$$\lim_{n\to\infty} \mathbb{P}_{err} = 0$$

and in computing in such a case the limiting density of the logarithm of the probability of error

$$\lim_{n\to\infty} \frac{1}{n} \log \mathbb{P}_{err}$$

which we will call the error exponent of the system.

After briefly resuming channel coding theory, we will show how one can associate this coding system to a statistical mechanic and we will study the distribution of the associated state-space energies.

3.4.1 Random coding summary

In this section we briefly recapitulate the random channel coding - MAP bitwise decoding introduced in the previous chapter. We described a generic code as a couple $c = (t, r)$ of maps between a source space S and a codeword space X (please note that, fow now, there is no relationship between the source space S in coding theory and the state space S in the statistical mechanic approach):

$$t : S \to X \quad r : X \to S$$

We introduced a probabilistic description of the channel coding problem built on a probability triple $(\Omega, \mathcal{F}, \mathbb{P})$. This description consists of a random message and a random noise. The former is a random variable

$$s_0 : (\Omega, \mathcal{F}, \mathbb{P}) \to (S, \mathbb{P}(S))$$

inducing a probability measure λ on $(S, \mathbb{P}(S))$.with the further assumptions that the source space S is finite, $|S| = 2^m = 2^{\lfloor nR \rfloor}$ and that the s_0-induced measure λ is uniform on it. The random noise is a second random variable

$$b : (\Omega, \mathcal{F}, \mathbb{P}) \to X^X$$

independent from the random message. To get a random coding setup we have to introduce a further random variable, the random code:

$$\begin{array}{rc} t: & (\Omega, \mathcal{F}, \mathbb{P}) \to X^S \\ & \omega \mapsto t_\omega \end{array}$$

where

$$\begin{array}{rc} t_\omega: & S \to X \\ & s \mapsto t_\omega(s) \end{array}$$

with the hypothesis that the random variables $\omega \to t_\omega(s)$, $s \in S$ are independent and identically distributed. In this model the received word (living in the codeword space, but not necessarily a codeword) is the random variable $y = r \circ b \circ t(s_0)$. Let us notice that we could have introduced random coding in a more general framework, without the hypotheses of uniformity of λ over S and of the identical distribution of the $\omega \to t_\omega(s)$. But since these hypotheses are typical in coding problems and critical in order to simplify certain computations, we choosed to assume them from the beginning.

We have then deployed decision theory arguments to find optimal decoding strategies for r, in the sense of minimisation of the quenched average wordwise error probability given the code t:

$$\mathbb{P}_{werr}(t) = \sum_{a \in S} \mathbb{P}(s_0 = a) \sum_{b \in S \setminus \{a\}} \mathbb{P}\{r \circ b \circ t(b) = a | t\}$$

obtaining the MAP wordwise decoder:

$$r_t = \arg\max_{s \in S} \mathbb{P}\{s_0 = s | y, t\}$$

or analogously in the sense of minimisation of the quenched average bitwise error probability:

$$\mathbb{P}_{berr}(t) = \sum_{k=1}^{m} \frac{1}{m} \sum_{a \in S} \mathbb{P}(s_0 = a) \sum_{\substack{b \in S \\ \pi_k(b) \neq \pi_k(a)}} \mathbb{P}\{r \circ b \circ t(b) = a | t\}$$

where we willingly equivocated between an element b of S and its binary representation. We obtained in this way the MAP bitwise decoder:

$$r_t(y) = \{r_{1,t}(y), ..., r_{m,t}(y)\}$$

$$r_{k,t}(y) = \chi \left\{ \begin{array}{cc} \sum_{\substack{s \in S \\ \pi_k(s)=1}} \mathbb{P}\{s_0 = s | y, t\} > & \sum_{\substack{s \in S \\ \pi_k(s)=0}} \mathbb{P}\{s_0 = s | y, t\} \end{array} \right\}$$

We then introduced the annealed average wordwise error probability:

$$\mathbb{P}_{werr} = \mathbb{E}\{\mathbb{P}_{werr}(t)\} = \mathbb{P}\{r \circ b \circ t(a) \neq a\}$$

and bitwise error probability:

$$\mathbb{P}_{berr} = \mathbb{E}\{\mathbb{P}_{berr}(t)\} = \sum_{\substack{b \in S \\ \pi_k(b) \neq \pi_k(a)}} \mathbb{P}\{r \circ b \circ t(a) = b\}$$

where now k and a are arbitrarily chosen, not random. The reason for introducing these new quantities is that they are easier to study mathematically. The drastic simplifications were made possible by the aforementioned hypotheses of uniformity of λ over S and of identical distribution of the $\omega \to t_\omega(s)$.

Once a concentration of a quenched quantity is proved (in our setup this will be a large deviation concentration property), we can obtain significant information on its typical behaviour.from the behaviour of the annealed one.

3.4.2 The fundamental identification

In order to describe the machinery just recapitulated with the statistical mechanic concepts, we need to identify entities analogous to a state space and a state energy function. We do this in the present section. Let us take the two following identification:

- the state space S of the statistical mechanic description with the source space S of the channel coding setup;

- the random energies $E_\omega(s)$ of the statistical mechanic description as follows:

$$E_\omega(s) := -\log \mathbb{P}\{s | y, t\}$$

that is, we associate the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ inducing the energies of the statistical mechanic system with the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ which generates the random source, code and noise in the channel coding setup. Of course the associated probabilities and partition function will be:

$$\mu_{\beta,\omega}(s) := \frac{1}{Z_\omega} \left[\ \mathbb{P}\{s | y, t\} \ \right]^\beta$$

$$Z_\omega := \sum_{s \in S} e^{-\beta E_\omega(s)} = \sum_{s \in S} \left[\ \mathbb{P}\{ s | \, y, t\} \ \right]^\beta$$

where $\beta \in [0, \infty]$.

3.4.3 The "spin magnetisation"

Now that we have the basic indentification in place, let us consider the following quantity $\Delta_k(\beta)$, for $k \in \{1, ..., m\}$:

$$\begin{aligned} \Delta_k(\beta) &: = \mathbb{E}_{\beta,\omega} \left\{ \ 2\pi_k(s) - 1 \ \right\} \\ &= 2 \sum_{s \in S} \pi_k(s) \, \mu_{\beta,\omega}(s) - 1 \\ &= \sum_{\substack{s \in S \\ \pi_k(s) = 1}} \mu_{\beta,\omega}(s) - \sum_{\substack{s \in S \\ \pi_k(s) = 0}} \mu_{\beta,\omega}(s) \end{aligned}$$

If we consider $\mu_{\beta,\omega}$ to be a kind of parametric "a posteriori" probability, the associated parametric MAP bitwise decoder (for the first bit) would have the form

$$\begin{aligned} r_{\beta,k}(y,t) &= \chi \left\{ \sum_{\substack{s \in S \\ \pi_k(s) = 1}} \mu_{\beta,\omega}(s) > \sum_{\substack{s \in S \\ \pi_k(s) = 0}} \mu_{\beta,\omega}(s) \right\} \\ &= \chi \ \{\Delta_k(\beta) > 0\} \end{aligned}$$

so actually the sign of $\Delta_k(\beta)$ is the discriminant information in the bitwise extimation. In the statistical mechanic literature, a quantity like $\Delta_k(\beta)$ is called a spin magnetisation of the system, that is a global quantity resulting by the state average of a binary quantity associated to each state. The (annealed) error probability $\mathbb{P}_{berr,k}(\beta, R)$ in this estimation is

$$\begin{aligned} \mathbb{P}_{berr,k}(\beta, R) &: = \mathbb{P}\{r_{\beta,k}(y,t) \neq \pi_k(s_0)\} \\ &= \mathbb{P}\left\{ (-1)^{\ \pi_k(s_0)} \ \Delta_k(\beta) > 0 \right\} \end{aligned}$$

while the messageword estimator $r_\beta(y,t)$ built on these bit estimators can be defined as

$$r_\beta(y,t) := \{r_{\beta,1}(y,t), ..., r_{\beta,m}(y,t)\}$$

Let us now do an analysis of the effect of different values of β. If $\beta = 1$ we immediately get the bitwise MAP decoding. If instead we let $\beta \to \infty$, we can

observe that we are weighting more and more so in the average the most likely term (in the case there is only one of such terms):

$$\lim_{\beta\to\infty}\mu_{\beta,\omega}\left(s\right)=\lim_{\beta\to\infty}\frac{1}{Z_{\omega}}\left[\ \mathbb{P}\left\{\,s|\,y,t\right\}\ \right]^{\beta}=\chi_{\ \arg\max_{a\in S}\mathbb{P}\left\{\,a|\,y,t\right\}}$$

This means that, in the limit, the only significant contribution will be the one of the most likely message-word, i.e. we would be decoding the bit corresponding to such a messageword, for all bits: this is precisely wordwise MAP decoding. Clearly, if there were more that one term messageword associated with the highest a posteriori $\mu_{\beta,\omega}$-probability, then of course the wordwise MAP decoding strategy would fail, too. We just proved that, through the analysis for any $\beta\in[0,\infty]$ of our generalised setup, we are able to recover the bitwise and wordwise decoding strategies as particular cases.

3.4.4 The error exponent

In words, we are interested in studing the conditions for which the error probability vanishes in the thermodynamic limit for the ensemble, and how fast does this eventually happen. In order to make this statement rigorous, we have to make some definitions.

The "natural" approach and its problem

Our first idea would be to define the codeword error exponent $E_{err}(t,\beta,R)$ for all possible random codes realizations t and the decoder $r_{\beta}\left(y,t\right)$ as follows:

$$E_{err}(t,\beta,R)\overset{a.s.}{:=}\liminf_{n\to\infty}-\frac{1}{n}\log\mathbb{P}_{werr}\left(t,\beta,R,n\right)$$

where

$$\mathbb{P}_{werr}\left(t,\beta,R,n\right)\overset{a.s.}{:=}\mathbb{P}\left\{\ \ r_{\beta}\circ b\circ t\left(s_{0}\right)\neq s_{0}|\,t\ \right\}$$

Of course, $E_{err}(R,r,t)$ is non-negative. For all the values of β and R so that $E_{err}(R,r,t)$ is strictly positive, the error probability $\mathbb{P}_{werr}\left(t,\beta,R,n\right)$ vanishes exponentially fast. This gives us a first formalisation of the aim stated at the beginning of this section, but does not tell us yet what happens for values of β and R external to the domain of strict positivity of $E_{err}(R,r,t)$.

Now let us now notice that our definition of $\mathbb{P}_{werr}\left(t,\beta,R,n\right)$ amount to an error average on the codewords $a\in S$:

$$\begin{aligned}\mathbb{P}_{werr}\left(t,\beta,R,n\right)&\overset{a.s.}{=}\mathbb{P}\left\{\ \ r_{\beta}\circ b\circ t\left(s_{0}\right)\neq s_{0}|\,t\ \right\}\\&\overset{a.s.}{=}\sum_{a\,\in\,S}\mathbb{P}\left(s_{0}=a\right)\mathbb{P}\left\{\ \ r_{\beta}\circ b\circ t\left(s_{0}\right)\neq s_{0}|\,t\ \right\}\\&\overset{a.s.}{=}\sum_{a\,\in\,S}\frac{1}{2^{m}}\mathbb{P}\left\{\ \ r_{\beta}\circ b\circ t\left(a\right)\neq a|\,t,s_{0}=a\ \right\}\end{aligned}$$

Taking this average sounds reasonable when the code is a finite length one and ergodicity in its successive independent use assures that the measured error rate converges to this average error probability. But if we consider a scaling limit of the code, with the length $n \to \infty$, then we realise that the limit code is actually used only once (since a codeword has an infinite length). This fact, in conjunction with the simmetry of the code (the fact that codewords are independent and identically distributed random variables), allows us to discard the study of $\mathbb{P}_{werr}(t, \beta, R, n)$ in favor of the study of the sourceword-conditioned error probability $\mathbb{P}_{werr}(t, \beta, R, n, a)$ defined as follows:

$$\mathbb{P}_{werr}(t, \beta, R, n, a) \stackrel{a.s.}{:=} \mathbb{P}\left\{ \ r_\beta \circ b \circ t(a) \neq a | \, t, s_0 = a \ \right\}$$

Let us now define

$$\mathbb{P}_{berr,k}(t, \beta, R, n, a) \stackrel{a.s.}{:=} \mathbb{P}\left\{ \ \pi_k \circ r_\beta \circ b \circ t(a) \neq \pi_k(a) | \, t, s_0 = a \ \right\}$$

for any fixed $k \in \{1, ..., m\}$ and notice that this probability is independent on k because of the symmetry of the code and the fact that the probability measure on S is uniform and $|S| = 2^m$.

Since the probability $\mathbb{P}_{werr}(t, \beta, R, n, a)$ of error for the entire codeword, that is, the probability that at least one bit is incorrectly decoded, is bounded between $\mathbb{P}_{berr,k}(t, \beta, R, n, a)$ and the sum for $k \in \{1, ..., m\}$ of $\mathbb{P}_{berr,k}(t, \beta, R, n, a)$, we get the following passages:

$$\begin{aligned}
& -\frac{1}{n} \log \mathbb{P}_{berr,k}(t, \beta, R, n, a) \\
\geq \ & -\frac{1}{n} \log \mathbb{P}_{werr}(t, \beta, R, n, a) \\
\geq \ & -\frac{1}{n} \log \sum_{k=1}^{m} \mathbb{P}_{berr,k}(t, \beta, R, n, a) \\
\geq \ & -\frac{1}{n} \log \left[\ m \sup_{k \in \{1,...,m\}} \mathbb{P}_{berr,k}(t, \beta, R, n, a) \ \right]
\end{aligned}$$

which lead to

$$\begin{aligned}
& \inf_{k \in \{1,...,m\}} \left[-\frac{1}{n} \log \mathbb{P}_{berr,k}(t, \beta, R, n, a) \right] \\
\geq \ & -\frac{1}{n} \log \mathbb{P}_{werr}(t, \beta, R, n, a) \\
\geq \ & \inf_{k \in \{1,...,m\}} \left[-\frac{1}{n} \log \mathbb{P}_{berr,k}(t, \beta, R, n, a) \right] - \frac{1}{n} \log \lfloor Rn \rfloor
\end{aligned}$$

and finally

$$\begin{aligned}
& \liminf_{n \to \infty} \inf_{k \in \{1,...,m\}} \left[-\frac{1}{n} \log \mathbb{P}_{berr,k}(t, \beta, R, n, a) \right] \\
= \ & \liminf_{n \to \infty} -\frac{1}{n} \log \mathbb{P}_{werr}(t, \beta, R, n, a) \stackrel{a.s.}{=:} E_{err}(R, r, t)
\end{aligned}$$

where we have abused, reusing it, the notation for the error exponent.

We would now aim for some concentration result for the random variable sequence

$$\inf_{k\in\{1,...,m\}} -\frac{1}{n}\log \mathbb{P}\left\{ \ \pi_k \circ r_\beta \circ b \circ t\,(a) \neq \pi_k\,(a) |\, b,t,s_0 = a \ \right\}$$

but the infimum operation hinders us in the computations, since the random variables

$$\mathbb{P}\left\{ \ \pi_k \circ r_\beta \circ b \circ t\,(a) \neq \pi_k\,(a) |\, b,t,s_0 = a \ \right\} \overset{a.s.}{=} \chi_{\left\{ \ \pi_k \circ r_\beta \circ b \circ t\,(a) \neq \pi_k\,(a) |\, b,t,s_0 = a \ \right\}}$$

are not independent - it is actually easy to find qualitative reasons for their strong correlation. This is the reason behind our adoption of Shannon's annealed approach.

The Shannon annealed approach

The Shannon's approach, as explained in the previous chapter, consists in computing the annealed error probability

$$\mathbb{P}_{werr}\,(\beta,R,n) \overset{a.s.}{:=} \mathbb{P}\left\{r_\beta \circ b \circ t\,(s_0) \neq s_0\right\}$$

and its associated error exponent

$$E_{err}(\beta,R) \overset{a.s.}{:=} \liminf_{n\to\infty} -\frac{1}{n}\log \mathbb{P}_{werr}\,(\beta,R,n)$$

Recalling properties already observed in the first chapter, with the definition

$$\mathbb{P}_{berr}\,(\beta,R,n) := \sum_{k=1}^{m} \frac{1}{m}\mathbb{P}\left\{\pi_k \circ r_\beta \circ b \circ t\,(s_0) \neq \pi_k\,(s_0)\right\}$$

we have

$$\mathbb{P}_{berr}\,(\beta,R,n) \leq \mathbb{P}\left\{r_\beta \circ b \circ t\,(s_0) \neq s_0\right\} \leq m\mathbb{P}_{berr}\,(\beta,R,n)$$

so there is no difference in the vanishing exponent for bit or word error probability. Furthermore

$$\mathbb{P}_{berr}\,(\beta,R,n) = \mathbb{P}\left\{\,\pi_k \circ r_\beta \circ b \circ t\,(a) \neq \pi_k\,(a) |\, s_0 = a\right\}$$

for any $k \in \{1,...,m\}$ and $a \in S$, so we can fix them as parameters. We can for instance conveniently take $k = 1$ and $a = 0$, the messageword corresponding to the all zero representation.

According to the definitions in the section dedicated to spin magnetisation, we now have that

$$\mathbb{P}\left\{\,\pi_1 \circ r_\beta \circ b \circ t\,(a) \neq 0 |\, s_0 = 0\right\} = \mathbb{P}\left\{\,\Delta_1^n\,(\beta) < 0 |\, s_0 = 0\right\}$$

$$= \mathbb{P}\left\{ \sum_{\substack{s \in S \\ \pi_1(s)=1}} \mu_{\beta,\omega}(s) > \sum_{\substack{s \in S \\ \pi_1(s)=0}} \mu_{\beta,\omega}(s) \,\middle|\, s_0 = 0 \right\}$$

$$= \mathbb{P}\left\{ \sum_{\substack{s \in S \\ \pi_1(s)=1}} e^{-\beta E_\omega(s)} > \sum_{\substack{s \in S \\ \pi_1(s)=0}} e^{-\beta E_\omega(s)} \right\}$$

where we have again slightly abused the notation by redefining the energies $E_\omega(s)$ as follows:

$$\begin{aligned} E_\omega(s) &:= -\log \mathbb{P}\{s|\, y, t; s_0 = a\} \\ &= -\log \frac{\mathbb{P}\{s_0 = a|\, t\}}{\mathbb{P}\{y|\, t\}} \mathbb{P}\{y|\, s_0 = a, t\} \\ &= -\log \mathbb{P}\{y|\, s_0 = a, t\} + K \end{aligned}$$

the second passage being a consequence of Bayes theorem, and the constant K being

$$\begin{aligned} K &:= -\log \frac{\mathbb{P}\{s_0 = a|\, t\}}{\mathbb{P}\{y|\, t\}} \\ &= -\log \mathbb{P}\{s_0 = a\} + \log \mathbb{P}\{y|\, t\} \\ &= \log m + \log \mathbb{P}\{y|\, t\} \end{aligned}$$

independent from s. Finally we obtain

$$E_{err}(\beta, R) \stackrel{a.s.}{:=} \liminf_{n\to\infty} -\frac{1}{n} \log \mathbb{P}\left\{ \sum_{\substack{s \in S \\ \pi_1(s)=1}} e^{-\beta E_\omega(s)} > \sum_{\substack{s \in S \\ \pi_1(s)=0}} e^{-\beta E_\omega(s)} \right\}$$

If we are then able to prove a concentration of the argument of the above probability in a large deviation sense - in the randomness induced by both the noise and the random encoder - then we will be able to see when the error probability in question vanishes almost surely and, if so, give an indication and a measure of the exponential nature of such a behaviour. This will be explained in more detail in the next chapter on large deviations.

3.5 The Shannon problem as a statistical mechanic system

In this last paragraph we apply the point of view developed in the present chapter to the Shannon problem defined at the end of the previous one. After

a brief summary of the quantities of interest in such a problem, we proceed to the statistical mechanic interpretation and to the analysis of the properties of the messagewords' associated energies.

3.5.1 The Shannon problem

Let us recall that in the Shannon problem we obtained the following expression for the annealed bitwise error probability:

$$\mathbb{P}_{berr}(R,n) = \mathbb{P}(\Delta_1^n \leq 0 | s_0 = 0)$$

where

$$\Delta_1^n := \sum_{s:\pi_1^n(s)=0} \mathbb{P}\{s|y^n, t(R,n)\} - \sum_{s:\pi_1^n(s)=1} \mathbb{P}\{s|y^n, t(R,n)\}$$

and

$$\mathbb{P}\{s|y^n, t(R,n)\} \stackrel{a.s.}{=} K \exp\{-\beta_{trans} d_H(y^n, x_0^n)\}$$

$$\beta_{trans} := \left(\log \frac{1-p}{p}\right)$$

3.5.2 The statistical mechanic identification

In order to adopt the statistical mechanic point of view for the Shannon problem, we just need to let a parameter β take the place of the constant β_{trans}, and in defining the messageword energies as follows:

$$E_\omega^n(s) := d_H(y^n(\omega), x_s^n(\omega))$$

obtaining

$$\mathbb{P}\{s|y^n, t(R,n)\} \stackrel{a.s.}{=} Ke^{-\beta E_\omega^n(s)}$$

We have left the constant K, which will become irrelevant within the condition $\Delta_1^n(\beta) < 0$.

3.5.3 A code rotation

In this section we study the behaviour of the energies $E_\omega^n(s)$. Let us define the following quantities:

$$\widetilde{x}_{s,j} := x_{s,j} x_{0,j} w_j$$

$$\widetilde{x}_s^n := \{\widetilde{x}_{s,1}, \ldots, \widetilde{x}_{s,n}\}$$

$$\widetilde{t}(R,n) := \{\widetilde{x}_s^n\}_{s \in S_n}$$

We can think of $\widetilde{t}(R,n)$ as the encoder rotated by the transmitted codeword and the noise realization. It is easy to show that:

1. for $s \neq 0$ the $\widetilde{x}_{s,j}$ are $Bernoulli\left(\pm 1, \frac{1}{2}\right)$;

2. $\widetilde{x}_{0,j} = w_j \sim Bernoulli(\pm 1, p)$.

It is less immediate (but extremely important) to show that the $\widetilde{x}_s^n$ are mutually independent. In order to do so we notice that for $b, c \in \{-1, 1\}$ we have

$$P\left(\widetilde{x}_{s,j} = b, \widetilde{x}_{s',j} = c\right) = P\left(x_{s,j} = bx_{0,j}w_j, x_{s',j} = cx_{0,j}w_j\right)$$

$$\begin{aligned}
&= P\left(x_{s,j} = bx_{0,j}w_j, x_{s',j} = cx_{0,j}w_j | x_{0,j}w_j = +1\right) P\left(x_{0,j}w_j = +1\right) \\
&+ P\left(x_{s,j} = bx_{0,j}w_j, x_{s',j} = cx_{0,j}w_j | x_{0,j}w_j = -1\right) P\left(x_{0,j}w_j = -1\right) \\
&= \frac{1}{4} = P\left(\widetilde{x}_{s,j} = b\right) P\left(\widetilde{x}_{s',j} = c\right)
\end{aligned}$$

for all $s \neq s'$, both different from 0, the all zero codeword that is transmitted. Moreover

$$\begin{aligned}
P\left(\widetilde{x}_{s,j} = b, \widetilde{x}_{0,j} = c\right) &= P\left(x_{s,j}x_{0,j} = bw_j, w_j = c\right) \\
&= P\left(x_{s,j}x_{0,j} = b\right) 1_{(c=1)} P\left(w_j = +1\right) \\
&+ P\left(x_{s,j}x_{0,j} = -b\right) 1_{(c=-1)} P\left(w_j = -1\right) \\
&= \frac{1}{2}\left[1_{(c=1)}\left(1-p\right) + 1_{(c=-1)}p\right] \\
&= P\left(\widetilde{x}_{s,j} = b\right) P\left(\widetilde{x}_{0,j} = c\right)
\end{aligned}$$

for all $s \neq 0$. The independence for different j's is immediate to see. Let us notice that this mutual independence is a peculiar property of the binary $SRE(R, n)$ together with the BSC – this would not work in a general non-binary setting. Notice furthermore that, after the rotation, the internal relative geometry of the code is maintained – the Hamming distances between all the couples of codewords are maintained.

It is now immediate to see that the

$$E_\omega^n\left(s\right) := d_H\left(y^n(\omega), x_s^n(\omega)\right) = \frac{1}{2}\left[n - \sum_{j=1,\ldots,n} \widetilde{x}_{s,j}(\omega)\right]$$

are mutually independent, and for all $s \neq 0$ are identically distributed as well.

3.5.4 The associated statistical-mechanic model

In the end, we rephrase the computation of the error exponent in the Shannon problem as follows:

$$E_{err}\left(\beta, R\right) := \liminf_{n\to\infty} -\frac{1}{n}\log \mathbb{P}\left[\sum_{s:\pi_1^n(s)=0} e^{-\beta E_\omega^n(s)} - \sum_{s:\pi_1^n(s)=1} e^{-\beta E_\omega^n(s)} < 0\right]$$

where we have

$$\left|\{s \in S : \pi_1^n\left(s\right) = 0\}\right| = \left|\{s \in S : \pi_1^n\left(s\right) = 1\}\right| = 2^{m-1} = 2^{\lfloor Rn \rfloor - 1}$$

$$E_{\omega}^{n}(s) := d_H\left(y^n(\omega), x_s^n(\omega)\right) = \frac{1}{2}\left[n - \sum_{j=1,\ldots,n} \widetilde{x}_{s,j}(\omega)\right]$$

independent, and identically distributed for $s \neq 0$. With the further definitions

$$\begin{aligned} Z_{\beta,n}^{*}(\omega) &= e^{-\beta E_{\omega}^{n}(0)} \\ \hat{Z}_{\beta,n}^{0}(\omega) &= \sum_{\substack{s\,:\,\pi_1^n(s)=0 \\ s\neq 0}} e^{-\beta E_{\omega}^{n}(s)} \\ Z_{\beta,n}^{1}(\omega) &= \sum_{s\,:\,\pi_1^n(s)=1} e^{-\beta E_{\omega}^{n}(s)} \end{aligned}$$

we get

$$E_{err}(\beta, R) = \liminf_{n\to\infty} -\frac{1}{n}\log P\left[\; Z_{\beta,n}^{1}(\omega) \geq Z_{\beta,n}^{*}(\omega) + \hat{Z}_{\beta,n}^{0}(\omega) \;\right]$$

where the condition for errorfree coding in terms of (β, R) is of course

$$\lim_{n\to\infty} P\left\{\omega \in \Omega :\; Z_{\beta,n}^{1}(\omega) < Z_{\beta,n}^{*}(\omega) + \hat{Z}_{\beta,n}^{0}(\omega) \;\right\} = 1$$

We will proceed to compute the error exponent $E_{err}(\beta, R)$ through large deviation techniques detailed in the following chapter.

4 Large Deviations

4.1 Introduction

In this chapter we concern ourselves with the large deviation analysis of sequences of random variables of the general form

$$\frac{1}{n}\log\int_{S_n} e^{-\beta H_n(s,\omega)}d\lambda_n(s)$$

where $\{\lambda_n(s)\}_{n\in\mathbb{N}}$ is a sequence of concentrating measures with a given rate function and $\{H_n(s,\omega)\}_{n\in\mathbb{N}}$ are random energy functions.

In the first paragraph will prove a large deviation theorem for such objects, identifying a set of rather general conditions for it to hold. This theorem could be seen as a generalization of Varadhan lemma for random log-Laplace integrals. Then in the second paragraph we will specialise the conditions to a simpler, albeit important case, which we call "extended REM", not to be confused with the "generalised REM" of the statistical mechanics' literature. Finally in the third paragraph we will show how this simpler version can be applied to the Shannon problem, described according ot the approach of statistical mechanics presented in the previous chapter. As a result we will be able to compute the random coding error exponent. The relevant discussion will take place in the next, conclusive chapter.

4.2 A large deviation principle for random log-Laplace integrals

Random measures of Gibbs type appear in various fields including disordered systems of statistical mechanics and random coding theory. For the purpose of this paragraph, let us fix some notations. Let $(\Omega,\mathcal{F},P)$ be a probability space with expectation E ($\omega\in\Omega$ will be the random parameter of the Gibbs measures). For every $n\in\mathbb{N}$ let $(S_n,\mathcal{B}_n,\lambda_n)$ be a measure space (space of configurations) with λ_n being a finite measure, and let $H_n : S_n\times\Omega\to\mathbb{R}$ be a measurable function (the energy). Given $\beta>0$ (inverse temperature) we set

$$Z_{\beta,n}^{\omega} := \int_{S_n} e^{-\beta H_n(s,\omega)}d\lambda_n(s)$$

and introduce the Gibbs measure on $(S_n,\mathcal{A}_n)$, depending on $\omega\in\Omega$, defined as

$$d\nu_n^{\omega}(s) = \frac{1}{Z_{\beta,n}^{\omega}}e^{-\beta H_n(s,\omega)}d\lambda_n(s)$$

Finally, let us introduce the *random free energy density*

$$f_{\beta,n}^{\omega} = \frac{1}{\beta n}\log Z_{\beta,n}^{\omega}$$

We want to study its limit as $n \to \infty$ and its fluctuations, in particular its large deviations. One natural approach to analyze the LD of $f^{\omega}_{\beta,n}$ is the application of the Gärtner-Ellis theorem, since

$$E\left[e^{n\alpha f^{\omega}_{\beta,n}}\right] = E\left[\left(Z^{\omega}_{\beta,n}\right)^{\frac{\alpha}{\beta}}\right]$$

For positive integers $\frac{\alpha}{\beta}$ this approach is often feasible and resembles the replica approach. However, it is necessary to compute the previous expected value also for non-integer $\frac{\alpha}{\beta}$, and this is seldom an easy task. This difficulty can be bypassed by another approach, although somewhat longer in principle. It has the conceptual appeal to be based on properties of the *number of states with a given energy*, a well known quantity both in statistical mechanics and information sciences (where it is related to the so-called *spectral shape function*).

Let us introduce the (random) occupation measure of $H_n(x,\omega)$:

$$\rho^{\omega}_n = \text{law of } H_n(\cdot,\omega) \text{ under } \lambda_n$$

which, in a finite space with uniform measure $(S_n, \mathcal{B}_n, \lambda_n)$, is proportional to the number of configurations $s \in S_n$ with a given energy.

Let $O_a{:}\mathbb{R} \to \mathbb{R}$ be the homothety $O_a(b) = \frac{b}{a}$. We may think that O_n shrinks sets and measures by a factor n. Then

$$\begin{aligned} Z^{\omega}_{\beta,n} &= \int_{\mathbb{R}} e^{-\beta\sigma} d\rho^{\omega}_n(\sigma) \\ &= \int_{\mathbb{R}} e^{-\beta n\sigma} d\left(O_n\rho^{\omega}_n\right)(\sigma) \end{aligned}$$

where

$$\left(O_n\rho^{\omega}_n\right)([a,b]) = \rho^{\omega}_n([na,nb])$$

We see that the limit properties of $f^{\omega}_{\beta,n}$ become the analogous properties of

$$\frac{1}{\beta n}\log\int_{\mathbb{R}} e^{-\beta n\sigma} d\left(O_n\rho^{\omega}_n\right)(\sigma)$$

Without randomness, the asymptotic of this quantity can be studied, in principle, by means of Varadhan lemma on log-Laplace integrals. Therefore the theorem we are about to prove may be seen as a generalization of Varadhan lemma to random measures. Such an approach has been particularly inspired by [9].

Notice that, when λ_n is not a probability measure, also $O_n\rho^{\omega}_n$ is not a probability measure, so we have to normalize it in order to work in the usual framework of Varadhan lemma. Denote by μ^{ω}_n the normalization of $O_n\rho^{\omega}_n$. If the scaling factor between $O_n\rho^{\omega}_n$ and μ^{ω}_n is deterministic and exponential in n, it is sufficient to understand the LD of $\frac{1}{\beta n}\log\int_{\mathbb{R}} e^{-\beta n\sigma} d\mu^{\omega}_n(\sigma)$ and then translate the result.

With the previous motivations in mind, in this paragraph we proceed to the study of a general large deviation theorem.

We consider a regular topological space Σ, a sequence of random measures μ_n^ω on it and a continuous function $\phi : \Sigma \to \mathbb{R}$. We will study a large deviation principle for a sequence of random variables

$$\frac{1}{n} \log \int_\Sigma e^{n\phi(\sigma)} \mu_n^\omega (d\sigma)$$

under suitable assumptions on μ_n^ω. What we will show is that, under some particular conditions, there exists a number $\widetilde{f}$ such that the above integral concentrates around it and there exists a function $\widetilde{I} : \mathbb{R} \to \mathbb{R}$ such that it satisfies a LDP with it as a rate function. In order to carry out our proof, we will need to distinguish between positive and negative fluctuations:

$$\begin{aligned} \frac{1}{n} \log \int_\Sigma e^{n\phi(\sigma)} \mu_n^\omega (d\sigma) &> \widetilde{f} \\ \frac{1}{n} \log \int_\Sigma e^{n\phi(\sigma)} \mu_n^\omega (d\sigma) &< \widetilde{f} \end{aligned}$$

since their probabilities are governed by different principles. In the following sections we will consider the two cases separately, and at the end of the paragraph will glue them together into a single theorem. We will also underline a necessary "tightness" tail condition hypothesis, we will consider a simplified real line setup for the positive fluctuations and we will tackle the generally difficult problem of computing the LDP for negative fluctuations by identifying some easy situations which are common in real modeling problems.

4.2.1 LDP for positive fluctuations: a Varadhan-like lemma

Let Σ be a regular topological space with Borel σ-field $\mathcal{B}(\Sigma)$ (recall that a Hausdorff space is a topological space where every two distinct points admit disjoint neighborhoods, and such a space is called regular when the same property holds true between any closed set C and any point σ with $\sigma \notin C$). Let $\mathcal{B}(\mathbb{R})$ be the Borel σ-field on $\mathbb{R}$, $(\Omega, \mathcal{F}, P)$ be a probability space with expectation E, $\{\mu_n^\omega; n \in \mathbb{N}, \omega \in \Omega\}$ a sequence of random probability measures on Σ (we say that a probability measure μ^ω depending on $\omega \in \Omega$ is a random probability measure if $\int_\Sigma g(\sigma) \mu^\omega (d\sigma)$ is a random variable, for every bounded continuous function $g : \Sigma \to \mathbb{R}$).

Let $\mathcal{A}$ be a base of the topology of Σ. Let $\phi : \Sigma \to \mathbb{R}$ be a continuous function.

Let us assume the existence of a function

$$J : \Sigma \times \mathbb{R} \to [0, \infty]$$

such that for every $A \in \mathcal{A}$ and every $v \in \mathbb{R}$ we have

$$-\inf_{\sigma \in \overset{\circ}{A}, x > v} J(\sigma, x) \leq \underline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \mu_n(A) > v\right) \tag{1}$$

$$\leq \overline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \mu_n(A) > v\right) \leq -\inf_{\sigma\in\overline{A}, x\geq v} J(\sigma, x) \tag{2}$$

We also assume that J is lower semicontinuous and has the following form of the property of good rate functions: given two real numbers v and γ, the set

$$\{(\sigma, x) : x \geq v, J(\sigma, x) \leq \gamma\}$$

is compact. We could call the above assumption a global large deviation property of the random probability measure sequence $\{\mu_n^\omega; n \in \mathbb{N}, \omega \in \Omega\}$ according to the rate function J.

Remark 1 *Since* $\log \mu_n^\omega(A) \leq 0$*, we could avoid to consider* $v \geq 0$ *in this definition. With the present definition we readily have (from the lower bound)*

$$J(\sigma, x) = \infty \quad \forall (\sigma, x) \in \Sigma \times (0, \infty)$$

Remark 2 *Let* $I_\mu : \Sigma \to [0, \infty]$ *be defined as*

$$I_\mu(\sigma) = -\inf\{x : J(\sigma, x) > 0\}$$

Heuristically, I_μ *is the rate function of* μ_n^ω *for* P*-a.e.* $\omega \in \Omega$*. A rigorous statement has to involve also large deviations from below treated in Theorem B below, so it is not given here. Here we could give only an inequality, of little use. But in the applications, it is very useful to have the right guess of the rate function* I_μ *from the previous formula.*

In addition to the previous assumptions, we need to assume a tail condition, since we treat possibly unbounded functions ϕ. We assume that there exists a sequence $(K_j)_{j\in\mathbb{N}}$ of compact subsets of Σ such that

$$\lim_{j\to\infty} \overline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \int_{K_j^c} e^{n\phi(\sigma)} \mu_n(d\sigma) > L\right) = -\infty \tag{3}$$

for all $L \in \mathbb{R}$. We could call this tail condition exponential tightness of the random probability measure sequence $\{\mu_n^\omega; n \in \mathbb{N}, \omega \in \Omega\}$ according to the scaling function ϕ. We state in the next subsection a sufficient condition for it, very easy to check in the application to the REM, at least.

Theorem 3 (A) *Let* $\widetilde{I}^+ : \mathbb{R} \to [0, \infty]$ *be defined as*

$$\widetilde{I}^+(x) = \inf_{\sigma\in\Sigma} J(\sigma, x - \phi(\sigma))$$

Then, under the previous assumptions of global large deviation property and exponential tightness of $\{\mu_n^\omega; n \in \mathbb{N}, \omega \in \Omega\}$*, we have for every* $a \in \mathbb{R}$ *that*

$$\underline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \int_\Sigma e^{n\phi(\sigma)} \mu_n(d\sigma) > a\right) \geq -\inf_{(a,\infty)} \widetilde{I}^+ \tag{4}$$

$$\overline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \int_\Sigma e^{n\phi(\sigma)} \mu_n(d\sigma) > a\right) \leq -\inf_{[a,\infty)} \widetilde{I}^+ \tag{5}$$

Proof. Along the proof, we shall denote $\int_\Sigma e^{n\phi(\sigma)}\mu_n^\omega(d\sigma)$ by $\widetilde{Z}_n^\omega$.

Step 1 (lower bound (4)). We can write (4) as

$$\underline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\widetilde{Z}_n^\omega > a\right) \geq -\inf_{\sigma\in\Sigma, x>a} J(\sigma, x-\phi(\sigma))$$

For every $A\in\mathcal{A}$, setting $\phi_A^- = \inf_A\phi$, we have

$$\int_\Sigma e^{n\phi(\sigma)}\mu_n^\omega(d\sigma) \geq e^{n\phi_A^-}\mu_n^\omega(A)$$

Hence

$$\begin{aligned} P\left(\frac{1}{n}\log\widetilde{Z}_n^\omega > a\right) &\geq P\left(e^{n\phi_A^-}\mu_n(A) > e^{na}\right) \\ &= P\left(\frac{1}{n}\log\mu_n(A) > a-\phi_A^-\right) \end{aligned}$$

Therefore from our assumptions

$$\begin{aligned} &\underline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\widetilde{Z}_n^\omega > a\right) \\ &\geq \underline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\mu_n(A) > a-\phi_A^-\right) \\ &\geq -\inf_{\sigma\in\overset{\circ}{A}, x>a-\phi_A^-} J(\sigma,x) \end{aligned}$$

and thus

$$\underline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\widetilde{Z}_n^\omega > a\right) \geq -\inf_{A\in\mathcal{A}}\inf_{\sigma\in\overset{\circ}{A}, x>a-\phi_A^-} J(\sigma,x)$$

The assertion of this step follows from the identity

$$\inf_{A\in\mathcal{A}}\inf_{\sigma\in\overset{\circ}{A}, x>a-\phi_A^-} J(\sigma,x) = \inf_{\sigma\in\Sigma, x>a} J(\sigma, x-\phi(\sigma))$$

In order to prove it, let Λ be the set of all (σ,x) such that $\sigma\in\overset{\circ}{A}$ for some $A\in\mathcal{A}$ and $x > a+\left(\phi(\sigma)-\phi_A^-\right)$; and let Δ be the set $\Sigma\times(a,\infty)$. We prove that $\Lambda=\Delta$. To prove $\Lambda\subset\Delta$ we have only to prove that $(\sigma,x)\in\Lambda$ implies $x>a$; but $(\sigma,x)\in\Lambda$ implies $x > a+\left(\phi(\sigma)-\phi_A^-\right)$ and $\left(\phi(\sigma)-\phi_A^-\right)\geq 0$, so $x>a$. Conversely, let $(\sigma,x)\in\Delta$. Since $x>a$, there is $\varepsilon>0$ such that $x > a+\varepsilon$. Let $A_\varepsilon\in\mathcal{A}$ be such that $\sigma\in\overset{\circ}{A}_\varepsilon$ and $\phi_{A_\varepsilon}^+ - \phi_{A_\varepsilon}^- \leq \varepsilon$ where $\phi_{A_\varepsilon}^+ = \sup_{A_\varepsilon}\phi$, $\phi_{A_\varepsilon}^- = \inf_{A_\varepsilon}\phi$. Then $x > a+\varepsilon$ implies

$$x > a+\phi_{A_\varepsilon}^+ - \phi_{A_\varepsilon}^- \geq a+\left(\phi(\sigma)-\phi_{A_\varepsilon}^-\right)$$

The proof is complete.

Step 2 (preparation to (5)) Recall the tail condition (3). For every $j\in\mathbb{N}$, let $\mathcal{A}_j\subset\mathcal{A}$ be a finite family with the two properties:

- $K_j \subset \bigcup_{A\in\mathcal{A}_j} A$
- $\phi_A^+ - \phi_A^- \leq \frac{1}{j}$ for every $A \in \mathcal{A}_j$

where $\phi_A^+ = \sup_A \phi$, $\phi_A^- = \inf_A \phi$. Then for every sequence $(c_j)_{j\in\mathbb{N}}$ such that $\lim_{j\to\infty} c_j = \infty$ we have

$$\sup_{j\in\mathbb{N}} \left[\left(\inf_{A\in\mathcal{A}_j} \inf_{\sigma\in\overline{A}, x\geq a-\frac{1}{j}-\phi_A^+} J(\sigma, x) \right) \wedge c_j \right] \geq \inf_{\sigma\in\Sigma, x\geq a} J\left(\sigma, x-\phi\left(\sigma\right)\right) \tag{6}$$

It is sufficient to prove

$$\begin{aligned} &\overline{\lim}_{j\to\infty} \left[\inf_{A\in\mathcal{A}_j} \inf_{\sigma\in\overline{A}, x\geq a-\frac{1}{j}+\left(\phi(\sigma)-\phi_A^+\right)} J(\sigma, x-\phi\left(\sigma\right)) \right] \\ &\geq \inf_{\sigma\in\Sigma, x\geq a} J\left(\sigma, x-\phi\left(\sigma\right)\right) \end{aligned}$$

We have

$$\inf_{A\in\mathcal{A}_j} \inf_{\sigma\in\overline{A}, x\geq a-\frac{1}{j}+\left(\phi(\sigma)-\phi_A^+\right)} J(\sigma, x-\phi\left(\sigma\right)) \geq \inf_{\sigma\in\Sigma, x\geq a-\frac{2}{j}} J\left(\sigma, x-\phi\left(\sigma\right)\right)$$

because the set Λ of all (σ, x) such that $\sigma \in \overline{A}$ for some $A \in \mathcal{A}_j$ and $x \geq a - \frac{1}{j} + \left(\phi\left(\sigma\right) - \phi_A^+\right)$ is smaller than the set $\Delta = \Sigma \times [a - \frac{2}{j}, \infty)$. Let us check that $\Lambda \subset \Delta$: if $(\sigma, x) \in \Lambda$ then

$$x \geq a - \frac{1}{j} + \left(\phi\left(\sigma\right) - \phi_A^+\right) \geq a - \frac{1}{j} - \frac{1}{j} = a - \frac{2}{j}$$

From the previous inequality, it is sufficient to prove

$$\overline{\lim}_{j\to\infty} \inf_{\sigma\in\Sigma, x\geq a-\frac{2}{j}} J\left(\sigma, x-\phi\left(\sigma\right)\right) \geq \inf_{\sigma\in\Sigma, x\geq a} J\left(\sigma, x-\phi\left(\sigma\right)\right)$$

This is a consequence of the lower semicontinuity of J. Indeed, given $\varepsilon > 0$, for every j let $(\sigma_j, x_j) \in \Sigma \times [a - \frac{2}{j}, \infty)$ be such that

$$J\left(\sigma_j, x_j - \phi\left(\sigma_j\right)\right) \leq \inf_{\sigma\in\Sigma, x\geq a-\frac{2}{j}} J\left(\sigma, x-\phi\left(\sigma\right)\right) + \varepsilon$$

Then

$$\overline{\lim}_{j\to\infty} \inf_{\sigma\in\Sigma, x\geq a-\frac{2}{j}} J\left(\sigma, x-\phi\left(\sigma\right)\right) \geq \overline{\lim}_{j\to\infty} J\left(\sigma_j, x_j - \phi\left(\sigma_j\right)\right) - \varepsilon$$

We claim that $x_j \geq a$ eventually. If not, there is a subsequence (x_{j_k}) of (x_j) such that $x_{j_k} \in [a - \frac{2}{j}, a)$. We have $x_{j_k} \to a$. From the assumption on the union of the level sets of $\sigma \mapsto J(\sigma, x)$ for x in a positive half-line, there is a subsequence of (σ_j), say (σ_{j_k}) for simplicity of notations, such that $\sigma_{j_k} \to \sigma^*$

for some σ^*. Indeed, the sequence $(\sigma_j, x_j - \phi(\sigma_j))$ has to live in a compact set, so in particular this is true for the first component. But then

$$\overline{\lim}_{j\to\infty} J(\sigma_j, x_j - \phi(\sigma_j)) \geq \underline{\lim}_{k\to\infty} J(\sigma_{j_k}, x_{j_k} - \phi(\sigma_{j_k}))$$

which is $\geq J(\sigma^*, a - \phi(\sigma^*))$, by the lower semicontinuity. The proof is complete.

Step 3 (upper bound (5)) We can rewrite (5) as

$$\overline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\log \widetilde{Z}_n^\omega > a\right) \leq -\inf_{\sigma\in\Sigma, x\geq a} J(\sigma, x - \phi(\sigma))$$

For every fixed $j \in \mathbb{N}$, the following implication is true:

$$\frac{1}{n}\log \widetilde{Z}_n^\omega > a \Rightarrow \int_{K_j^c} e^{n\phi(\sigma)}\mu_n^\omega(d\sigma) > e^{na}\frac{1}{N} \text{ or}$$
$$\exists A \in \mathcal{A}_j : e^{n\phi_A^+}\mu_n^\omega(A) > e^{na}\frac{1}{N}$$

where N is equal to the cardinality of $\mathcal{A}_j$ plus one. This implication is true because, by contradiction, if $e^{n\phi_A^+}\mu_n^\omega(A) \leq e^{na}\frac{1}{N}$ for every $A \in \mathcal{A}_j$ and $\int_{K_j^c} e^{n\phi(\sigma)}\mu_n^\omega(d\sigma) \leq e^{na}\frac{1}{N}$, then

$$\begin{aligned}\int_\Sigma e^{n\phi(\sigma)}\mu_n^\omega(d\sigma) &\leq \sum_{A\in\mathcal{A}_j}\int_A e^{n\phi(\sigma)}\mu_n^\omega(d\sigma) + \int_{K_j^c} e^{n\phi(\sigma)}\mu_n^\omega(d\sigma)\\ &\leq \sum_{A\in\mathcal{A}_j} e^{n\phi_A^+}\mu_n^\omega(A) + \int_{K_j^c} e^{n\phi(\sigma)}\mu_n^\omega(d\sigma) \leq e^{na}\end{aligned}$$

Therefore

$$\begin{aligned}&P\left(\frac{1}{n}\log \widetilde{Z}_n^\omega > a\right)\\ &\leq \sum_{A\in\mathcal{A}_j} P\left(e^{n\phi_A^+}\mu_n(A) > e^{na}\frac{1}{N}\right)\\ &+ P\left(\int_{K_j^c} e^{n\phi(\sigma)}\mu_n(d\sigma) > e^{na}\frac{1}{N}\right)\end{aligned}$$

$$\begin{aligned}&= \sum_{A\in\mathcal{A}_j} P\left(\frac{1}{n}\log\mu_n(A) > a - \phi_A^+ - \frac{1}{n}\log N\right)\\ &+ P\left(\frac{1}{n}\log\int_{K_j^c} e^{n\phi(\sigma)}\mu_n(d\sigma) > a - \frac{1}{n}\log N\right)\end{aligned}$$

and therefore

$$\leq \sum_{A\in\mathcal{A}_j} P\left(\frac{1}{n}\log\mu_n(A) > a - \frac{1}{j} - \phi_A^+\right)$$
$$+ P\left(\frac{1}{n}\log\int_{K_j^c} e^{n\phi(\sigma)}\mu_n(d\sigma) > a-1\right)$$

for sufficiently large n. By our assumptions, for every $A\in\mathcal{A}_j$

$$\overline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\mu_n(A) > a - \frac{1}{j} - \phi_A^+\right)$$
$$\leq -\inf_{\sigma\in\overline{A}, x\geq a-\frac{1}{j}-\phi_A^+} J(\sigma,x)$$

and

$$\overline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\int_{K_j^c} e^{n\phi(\sigma)}\mu_n(d\sigma) > a-1\right) = -c_j$$

with $\lim_{j\to\infty} c_j = \infty$. Therefore

$$\overline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\widetilde{Z}_n^\omega > a\right)$$
$$\leq -\sup_{j\in\mathbb{N}}\left[\left(\inf_{A\in\mathcal{A}_j}\inf_{\sigma\in\overline{A}, x\geq a-\frac{1}{j}-\phi_A^+} J(\sigma,x)\right)\wedge c_j\right]$$

By (6) we have the result of this step. The proof is complete. ■

4.2.2 On the exponential tightness condition

Lemma 4 *Assume that for some $\lambda > 1$ we have*

$$\lim_{L\to\infty}\overline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\int_\Sigma e^{n\lambda\phi(\sigma)}\mu_n(d\sigma) > L\right) = -\infty \tag{7}$$

and

$$\lim_{j\to\infty}\inf_{\sigma\in\overline{K_j^c}x\geq L} J(\sigma,x) = +\infty \tag{8}$$

for every $L\in\mathbb{R}$, where J and K_j^c are given in the previous section. Then the tail condition (3) holds true.

Proof. Since

$$\int_{K_j^c} e^{n\phi(\sigma)}\mu_n^\omega(d\sigma) \leq \left(\mu_n^\omega\left(K_j^c\right)\right)^{\frac{\lambda-1}{\lambda}}\cdot\left(\int_\Sigma e^{n\lambda\phi(\sigma)}\mu_n^\omega(d\sigma)\right)^{\frac{1}{\lambda}}$$

we have

$$P\left(\frac{1}{n}\log\int_{K_j^c} e^{n\phi(\sigma)}\mu_n\left(d\sigma\right) > L\right)$$
$$\leq P\left(\frac{\lambda-1}{\lambda}\frac{1}{n}\log\left(\mu_n\left(K_j^c\right)\right) + \frac{1}{\lambda}\frac{1}{n}\log\left(\int_\Sigma e^{n\lambda\phi(\sigma)}\mu_n\left(d\sigma\right)\right) > L\right)$$

Taken any two real numbers L_1 and L_2 such that $L_1 + L_2 = L$, we have

$$P\left(\frac{1}{n}\log\int_{K_j^c} e^{n\phi(\sigma)}\mu_n\left(d\sigma\right) > L\right)$$
$$\leq P\left(\frac{\lambda-1}{\lambda}\frac{1}{n}\log\left(\mu_n\left(K_j^c\right)\right) > L_1\right)$$
$$+ P\left(\frac{1}{\lambda}\frac{1}{n}\log\left(\int_\Sigma e^{n\lambda\phi(\sigma)}\mu_n\left(d\sigma\right)\right) > L_2\right)$$

Therefore

$$\overline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\int_{K_j^c} e^{n\phi(\sigma)}\mu_n\left(d\sigma\right) > L\right)$$
$$\leq \overline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\left(\mu_n\left(K_j^c\right)\right) > L_1\frac{\lambda}{\lambda-1}\right)$$
$$\vee\, \overline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\left(\int_\Sigma e^{n\lambda\phi(\sigma)}\mu_n\left(d\sigma\right)\right) > \lambda L_2\right)$$

$$\leq -\inf_{\sigma\in\overline{K_j^c}, x\geq L_1\frac{\lambda}{\lambda-1}} J\left(\sigma, x\right)$$
$$\vee\, \overline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\left(\int_\Sigma e^{n\lambda\phi(\sigma)}\mu_n\left(d\sigma\right)\right) > \lambda L_2\right)$$

The result is now a direct consequence of the two assumptions. ■

Theorem 5 *Consider the sequence of probability measures*

$$\nu_n = E\mu_n$$

and assume that

$$\overline{\lim}_{n\to\infty}\frac{1}{n}\log\int_\Sigma e^{n\lambda\phi(\sigma)}\nu_n\left(d\sigma\right) < \infty \tag{9}$$

for some $\lambda > 1$. Then the condition (7) is verified. Therefore, if (8) is also verified, the tail condition (3) holds true.

Proof. The assertion follows easily from the inequality

$$\begin{aligned} & P\left(\frac{1}{n}\log\int_{\Sigma}e^{n\lambda\phi(\sigma)}\mu_n\left(d\sigma\right)>L\right) \\ & =P\left(\int_{\Sigma}e^{n\lambda\phi(\sigma)}\mu_n\left(d\sigma\right)>\exp nL\right) \\ & \leq\exp\left(-nL\right)E\left[\int_{\Sigma}e^{n\lambda\phi(\sigma)}\mu_n\left(d\sigma\right)\right] \\ & =\exp\left(-nL\right)\int_{\Sigma}e^{n\lambda\phi(\sigma)}\nu_n\left(d\sigma\right) \end{aligned}$$

■

Remark 6 *In many applications one has an easy control of ν_n,which allows to prove the tail condition (9).*

4.2.3 A simplified version on the real line

It would be very useful to have a simplified version, less general but easier to handle, for assumptions 1 and 2 of what we called the global large deviation property. So that it would become easier to check the hypotheses of our Varadhan-like lemma in certain common models. For this reason we present, in the present section, a simpler condition on the real line. In the next paragraph we will use it to check the conditions of our large deviation principle applied to what we call an "extended REM" model.

Consider the case when $\Sigma=\mathbb{R}$, with the Euclidean topology. Let $\mathcal{A}$ be a base of the topology given by all open intervals (a,b) with a and b taken in a dense subset of Σ (in the application we shall exclude a few isolated critical values). Let us recall the global large deviation property in this case: $J:\Sigma\times\mathbb{R}\to[0,\infty]$ is a lower semicontinuous function such that, given two real numbers v and γ, the set

$$\{(\sigma,x):x\geq v,J(\sigma,x)\leq\gamma\}$$

is compact, and that for every $(a,b)\in\mathcal{A}$ and every $v\in\mathbb{R}$ the following inequalities hold

$$-\inf_{\sigma\in(a,b),x>v}J(\sigma,x)\leq\underline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\mu_n\left(a,b\right)>v\right) \tag{10}$$

$$\leq\overline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\mu_n\left(a,b\right)>v\right)\leq-\inf_{\sigma\in[a,b],x\geq v}J(\sigma,x) \tag{11}$$

We want to give a simple set of conditions implying these two last inequalities. The following propositions are, in a sense, well known in LD theory, but here we deal with a functional $J(\sigma,x)$ of two variables where σ has no classical meaning in LD theory, so we state and prove these simple facts for completeness.

Lemma 7 *Assume that for every $\sigma \in \Sigma$, the function $x \mapsto J(\sigma, x)$ is non decreasing; and for every x, the function $\sigma \mapsto J(\sigma, x)$ is non decreasing (resp. non increasing) on some interval $A \subset \Sigma$. If, for a triple (a, b, v) with $[a, b] \subset A$, we have*

$$\lim_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \mu_n(a, b) > v\right) = -J(a, v)$$

(resp. $= -J(b, v)$), then for this triple we have (10) and (11).

Proof. We have

$$\inf_{\sigma\in[a,b], x\geq v} J(\sigma, x) = J(a, v) \text{ and } \inf_{\sigma\in(a,b), x>v} J(\sigma, x) \geq J(a, v)$$

The claim easily follows from this fact. ■

Lemma 8 *Assume that, for an interval $(a, b) \in \mathcal{A}$, (10) holds true for all v except those of a discrete set. Then (10) holds true for every v. The same property applies to (11).*

Proof. Step 1. Let v be a value in the exceptional set and let $\varepsilon^* > 0$ be a value such that $v \pm \varepsilon^*$ is in a connected neighborhood of v without others exceptional points. For $\varepsilon \in (0, \varepsilon^*)$ we have

$$\begin{aligned}
&\overline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \mu_n(a, b) > v\right) \\
&\leq \overline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \mu_n(a, b) > v - \varepsilon\right) \leq -\inf_{\sigma\in[a,b], x\geq v-\varepsilon} J(\sigma, x)
\end{aligned}$$

so

$$\begin{aligned}
&\overline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \mu_n(a, b) > v\right) \\
&\leq -\sup_{\varepsilon\in(0,\varepsilon^*)} \inf_{\sigma\in[a,b], x\geq v-\varepsilon} J(\sigma, x)
\end{aligned}$$

The proof is complete if we show that

$$\sup_{\varepsilon\in(0,\varepsilon^*)} \inf_{\sigma\in[a,b], x\geq v-\varepsilon} J(\sigma, x) \geq \inf_{\sigma\in[a,b], x\geq v} J(\sigma, x)$$

By contradiction, if there is $\delta > 0$ such that

$$\sup_{\varepsilon\in(0,\varepsilon^*)} \inf_{\sigma\in[a,b], x\geq v-\varepsilon} J(\sigma, x) \leq \inf_{\sigma\in[a,b], x\geq v} J(\sigma, x) - \delta$$

then, taken a sequence $\varepsilon_n \in (0, \varepsilon^*)$ with $\varepsilon_n \to 0$, there exists $\sigma_n \in [a, b]$, $x_n \geq v - \varepsilon_n$ such that

$$J(\sigma_n, x_n) \leq \inf_{\sigma\in[a,b], x\geq v} J(\sigma, x) - \frac{\delta}{2}$$

but since

$$x_n \geq v \Rightarrow J(\sigma_n, x_n) \geq \inf_{\sigma\in[a,b], x\geq v} J(\sigma, x)$$

then it must be $x_n \in [v - \varepsilon_n, v[$ and so there is a subsequence $(\sigma_{n_k}, x_{n_k}) \to (\sigma^*, v)$ for some $\sigma^* \in [a, b]$. By the lower semicontinuity

$$\underline{\lim}_{k\to\infty} J(\sigma_{n_k}, x_{n_k}) \geq J(\sigma^*, v) \geq \inf_{\sigma\in[a,b], x\geq v} J(\sigma, x)$$

This contradicts the previous inequality, so the proof is complete.

Step 2. Let v, ε^* and ε as above. We have

$$\begin{aligned}&\underline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\log \mu_n(a,b) > v\right)\\ &\geq \underline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\log \mu_n(a,b) > v+\varepsilon\right) \geq -\inf_{\sigma\in(a,b), x>v+\varepsilon} J(\sigma, x)\end{aligned}$$

so

$$\begin{aligned}&\underline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\log \mu_n(a,b) > v\right)\\ &\geq -\inf_{\varepsilon\in(0,\varepsilon^*)} \inf_{\sigma\in(a,b), x>v+\varepsilon} J(\sigma, x)\end{aligned}$$

The proof is complete because

$$\inf_{\varepsilon\in(0,\varepsilon^*)} \inf_{\sigma\in(a,b), x>v+\varepsilon} J(\sigma, x) = \inf_{\sigma\in(a,b), x>v} J(\sigma, x)$$

■

4.2.4 LDP for negative fluctuations: studying joint events

We always assume the tail condition (3); $(K_j)_{j\in\mathbb{N}}$ will be the sequence of compact subsets of Σ described there.

Moreover, we assume to have a functional $\widetilde{I}^- : \mathbb{R} \to [0, \infty]$, a positive integer j_0, a finite family $\mathcal{B}_j \subset \mathcal{B}(\Sigma)$, for every $j \geq j_0$ $(j \in \mathbb{N})$, such that $\Sigma = K_j^c \cup \left(\bigcup_{B\in\mathcal{B}_j} B\right)$ and two sequences $\varepsilon_j \to 0$ and $R_j \to -\infty$ such that

$$-\inf_{y<x} \widetilde{I}^-(y) - \varepsilon_j \leq \underline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\log \mu_n(B) < (x-\phi)_B^- \quad \forall B \in \mathcal{B}_j\right) \tag{12}$$

$$\overline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\log \mu_n(B) < (x-\phi)_B^+ \quad \forall B \in \mathcal{B}_j\right) \leq \left(-\inf_{y\leq x} \widetilde{I}^-(y) + \varepsilon_j\right) \vee R_j \tag{13}$$

for every $x \in \mathbb{R}$ and $j \geq j_0$.

Remark 9 *The previous conditions correspond to the intuitive requirement that for every $x \in \mathbb{R}$ we have (we cannot speak of $\mu_n(\sigma)$)*

$$
\begin{aligned}
&- \inf_{y<x} \widetilde{I}^{-}(y) - \varepsilon_j \\
&\leq \underline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\log \mu_n(\sigma) \leq x - \phi(\sigma) \ \text{ for all } \sigma \in \Sigma\right) \\
&\leq \overline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\log \mu_n(\sigma) \leq x - \phi(\sigma) \ \text{ for all } \sigma \in \Sigma\right) \\
&\leq \left(- \inf_{y\leq x} \widetilde{I}^{-}(y) + \varepsilon_j\right) \vee R_j
\end{aligned}
$$

When Σ is finite, this formulation is meaningful and may help to understand the proof of Theorem B below.

Remark 10 *Let $T_\phi = -\phi + \mathbb{R}$ be the set of all functions of the form $-\phi + c$, $c \in \mathbb{R}$. Since T_ϕ is isomorphic to $\mathbb{R}$, the following formulation is equivalent to the main assumption (the correspondence is $\widetilde{I}^{-}(x) = J^*(x - \phi(.))$): assume to have a functional $J^* : T_\phi \to [0,\infty]$, a positive integer j_0 and a finite family $\mathcal{B}_j \subset \mathcal{B}(\Sigma)$, for every $j \geq j_0$, such that $\Sigma = K_j^c \cup \left(\bigcup_{B\in\mathcal{B}_j} B\right)$ and*

$$
\begin{aligned}
&- \inf_{g\in T_\phi, g<f} J^* - \varepsilon_n \leq \underline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\log \mu_n(B) < f_B^- \quad \forall B \in \mathcal{B}_j\right) \\
&\leq \overline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\log \mu_n(B) < f_B^+ \quad \forall B \in \mathcal{B}_j\right) \\
&\leq \left(- \inf_{g\in T_\phi, g\leq f} J^* + \varepsilon_n\right) \vee R_n
\end{aligned}
$$

for every $f \in T_\phi$ and $j \geq j_0$. This more cumbersome scheme underlines the functional dependence on ϕ of the rate $\widetilde{I}^{-}(x)$, which is somewhat similar to the one of the theorem for positive fluctuations.

Theorem 11 (B) *Under the previous assumptions, for every $b \in \mathbb{R}$ we have*

$$
\underline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\log \int_\Sigma e^{n\phi(\sigma)} \mu_n(d\sigma) < b\right) \geq - \inf_{x<b} \widetilde{I}^{-}(x) \tag{14}
$$

$$
\overline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\log \int_\Sigma e^{n\phi(\sigma)} \mu_n(d\sigma) < b\right) \leq - \inf_{x\leq b} \widetilde{I}^{-}(x) \tag{15}
$$

Proof. Again, in the proof, we denote $\int_\Sigma e^{n\phi(\sigma)} \mu_n^\omega(d\sigma)$ by $\widetilde{Z}_n^\omega$.

Step 1 (upper bound (15)). For every $j \geq j_0$, we have

$$\begin{aligned} P\left(\frac{1}{n}\log \widetilde{Z}_n^{\omega} < b\right) &= P\left(\int_{\Sigma} e^{n\phi(\sigma)}\mu_n\left(d\sigma\right) < \exp\left(nb\right)\right) \\ &\leq P\left(\int_{B} e^{n\phi(\sigma)}\mu_n\left(d\sigma\right) < \exp\left(nb\right) \quad \forall B \in \mathcal{B}_j\right) \\ &\leq P\left(e^{n\phi_B^-}\mu_n\left(B\right) < \exp\left(nb\right) \quad \forall B \in \mathcal{B}_j\right) \\ &= P\left(\frac{1}{n}\log \mu_n\left(B\right) < b - \phi_B^- \quad \forall B \in \mathcal{B}_j\right) \end{aligned}$$

Therefore, from our assumptions (notice that $b - \phi_B^- = (b-\phi)_B^+$)

$$\begin{aligned} &\overline{\lim}_{n\to\infty} \frac{1}{n}\log P\left(\frac{1}{n}\log \widetilde{Z}_n^{\omega} < b\right) \\ &\leq \overline{\lim}_{n\to\infty} \frac{1}{n}\log P\left(\frac{1}{n}\log \mu_n\left(B\right) < (b-\phi)_B^+ \quad \forall B \in \mathcal{B}_j\right) \\ &\leq \left(-\inf_{y\leq b} \widetilde{I}^- + \varepsilon_j\right) \vee R_j \end{aligned}$$

The claim easily follows.

Step 2 (lower bound (14)). If $N_j = \mathrm{card}\left(\mathcal{B}_j\right) + 1$, we have

$$P\left(\frac{1}{n}\log \widetilde{Z}_n^{\omega} < b\right) = P\left(\int_{\Sigma} e^{n\phi(\sigma)}\mu_n\left(d\sigma\right) < \exp\left(nb\right)\right)$$

$$\begin{aligned} &\geq P\left(\left\{\begin{array}{c} \int_B e^{n\phi(\sigma)}\mu_n\left(d\sigma\right) < \frac{1}{N_j}\exp\left(nb\right) \quad \forall B \in \mathcal{B}_j \\ \int_{K_j^c} e^{n\phi(\sigma)}\mu_n\left(d\sigma\right) < \frac{1}{N_j}\exp\left(nb\right) \end{array}\right.\right) \\ &\geq P\left(\left\{\begin{array}{c} \frac{1}{n}\log \mu_n\left(B\right) < b - \phi_B^+ - \frac{1}{n}\log N_j \quad \forall B \in \mathcal{B}_j \\ \frac{1}{n}\log \int_{K_j^c} e^{n\phi(\sigma)}\mu_n\left(d\sigma\right) < b - \frac{1}{n}\log N_j \end{array}\right.\right) \end{aligned}$$

and therefore, given $\varepsilon > 0$, there exists n_0 (depending on j) such that for every $n \geq n_0$

$$\begin{aligned} &P\left(\frac{1}{n}\log \widetilde{Z}_n^{\omega} < b\right) \\ &\geq P\left(\left\{\begin{array}{c} \frac{1}{n}\log \mu_n\left(B\right) < b - \phi_B^+ - \varepsilon \quad \forall B \in \mathcal{B}_j \\ \frac{1}{n}\log \int_{K_j^c} e^{n\phi(\sigma)}\mu_n\left(d\sigma\right) < b - \varepsilon \end{array}\right.\right) \\ &\geq P\left(\frac{1}{n}\log \mu_n\left(B\right) < b - \phi_B^+ - \varepsilon \quad \forall B \in \mathcal{B}_j\right) \\ &- P\left(\frac{1}{n}\log \int_{K_j^c} e^{n\phi(\sigma)}\mu_n\left(d\sigma\right) \geq b - \varepsilon\right) \end{aligned}$$

Assume $\inf_{x<b} \widetilde{I}^-(x) < \infty$, otherwise the proof is complete. Let $x' < b$ be a point such that $\widetilde{I}^-(x') < \infty$ and let $\varepsilon > 0$ be such that $b - \varepsilon > x'$. We have $\inf_{x<b-\varepsilon} \widetilde{I}^-(x) < \infty$. Choose the previous ε smaller than one and choose a new j_0 if necessary such that for all $j \geq j_0$ we also have

$$\begin{aligned}&\overline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \int_{K_j^c} e^{n\phi(\sigma)} \mu_n(d\sigma) > b - \varepsilon\right) \\ &< -\inf_{x<b-\varepsilon} \widetilde{I}^-(x) - 2\end{aligned}$$

Since

$$\begin{aligned}&\underline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \mu_n(B) < b - \phi_B^+ - \varepsilon \quad \forall B \in \mathcal{B}_j\right) \\ &\geq -\inf_{x<b-\varepsilon} \widetilde{I}^-(x) - \varepsilon_j - \varepsilon\end{aligned}$$

by the previous choice of j_0 we have, for all $j > j_0$,

$$\underline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \widetilde{Z}_n^\omega < b\right) \geq -\inf_{x<b-\varepsilon} \widetilde{I}^-(x) - \varepsilon_j - \varepsilon$$

Thus

$$\underline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \widetilde{Z}_n^\omega < b\right) \geq -\inf_{x<b-\varepsilon} \widetilde{I}^-(x) - \varepsilon$$

Given $\varepsilon_0 > 0$, for all $\varepsilon < \varepsilon_0$

$$\inf_{x<b-\varepsilon} \widetilde{I}^-(x) + \varepsilon \leq \inf_{x<b-\varepsilon} \widetilde{I}^-(x) + \varepsilon_0$$

$$\begin{aligned}\inf_{0<\varepsilon<\varepsilon_0} \left(\inf_{x<b-\varepsilon} \widetilde{I}^-(x) + \varepsilon\right) &\leq \inf_{0<\varepsilon<\varepsilon_0} \inf_{x<b-\varepsilon} \widetilde{I}^-(x) + \varepsilon_0 \\ &= \inf_{x<b} \widetilde{I}^-(x) + \varepsilon_0\end{aligned}$$

hence

$$\underline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \widetilde{Z}_n^\omega < b\right) \geq -\inf_{x<b} \widetilde{I}^-(x) - \varepsilon_0$$

The proof is complete by the arbitrariety of ε_0. ■

4.2.5 How to find $\widetilde{I}^-(x)$

A direct analysis of conditions of theorem B is hard for two reasons: 1) they involve joint distributions, 2) therefore one needs an a priori guess of $\widetilde{I}^-(x)$. About 2), it is helpful to detect the region where $\widetilde{I}^-(x) = 0$. The first lemma of this section is devoted to this purpose. About 1), a lucky occurrence for some particular models is that, in the complementary region where $\widetilde{I}^-(x) \neq 0$, it is not necessary to compute joint probabilities, since just one marginal is sufficient to get $\widetilde{I}^-(x) = \infty$ (a fortiori the same is true for the joint probabilities). This is the content of the second lemma.

Lemma 12 *Assume that the conditions (1)-(2) of Theorem A holds true and let* $I_\mu : \Sigma \to [0,\infty]$ *be defined as*

$$I_\mu(\sigma) = -\inf\{x : J(\sigma, x) > 0\} \tag{16}$$

($I_\mu(\sigma) = +\infty$ *if* $\{x : J(\sigma,x) > 0\} = \emptyset$*). Let*

$$\widetilde{f} = \sup_{\sigma\in\Sigma}(\phi(\sigma) - I_\mu(\sigma)) \tag{17}$$

Assume there exists a function $\widetilde{I}^- : \left(-\infty, \widetilde{f}\right) \to [0,\infty]$*, a positive integer* j_0*, a finite family* $\mathcal{B}_j \subset \mathcal{B}(\Sigma)$*, for every* $j \geq j_0$ *(*$j \in \mathbb{N}$*), such that* $\Sigma = K_j^c \cup \left(\bigcup_{B\in\mathcal{B}_j} B\right)$ *and two sequences* $\varepsilon_j \to 0$ *and* $R_j \to -\infty$ *such that (12)-(13) hold true for every* $x < \widetilde{f}$*. Assume the following technical conditions:*

i) for every $j \geq j_0$ *and* $B \in \mathcal{B}_j$ *there is* $U_B \in \mathcal{A}$ *such that* $B \subset U_B$ *and*

$$\lim_{j\to\infty}\max_{B\in\mathcal{B}_j}\left(\phi_B^+ - \phi_{U_B}^-\right) = 0,$$

where $\mathcal{A}$ *is the base of the topology of* Σ *used in Theorem A,*

ii) for every $\varepsilon > 0$ *there exists* $j_\varepsilon \geq j_0$ *such that for every* $j \geq j_\varepsilon$ *and* $B \in \mathcal{B}_j$ *we have*

$$\inf_{\sigma\in\overline{U_B}, y\geq -I_\mu(\sigma)+\varepsilon} J(\sigma, y) > 0$$

Extend $\widetilde{I}^-$ *on* $[\widetilde{f},\infty)$ *by setting*

$$\widetilde{I}^-(x) = 0 \text{ for every } x \geq \widetilde{f}$$

Then (12)-(13) hold true also for every $x \geq \widetilde{f}$*, with the same* $\mathcal{B}_j$*,* ε_j*,* R_j *and possibly a new* j_0*.*

Proof. Step 1 (upper estimate). Since for every $x \geq \widetilde{f}$

$$\begin{aligned}&\overline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\mu_n(B) < (x-\phi)_B^+ \quad \forall B\in\mathcal{B}_j\right)\\ &\leq 0 \leq \varepsilon_j = \left(-\inf_{y\leq x}\widetilde{I}^-(y) + \varepsilon_j\right)\vee R_j\end{aligned}$$

the upper estimate is obvious.

Step 2 (lower estimate, $x > \widetilde{f}$). We have to prove that for sufficiently large $j \in \mathbb{N}$ and all $x > \widetilde{f}$,

$$\underline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\mu_n(B) < (x-\phi)_B^- \quad \forall B\in\mathcal{B}_j\right) \geq -\varepsilon_j$$

This is true if we prove that there exists j_0 such that for every $j \geq j_0$

$$\underline{\lim}_{n\to\infty} P\left(\frac{1}{n}\log\mu_n(B) < (x-\phi)_B^- \quad \forall B\in\mathcal{B}_j\right) = 1$$

which in turn is implied by

$$\lim_{n\to\infty} P\left(\frac{1}{n}\log\mu_n(B) < (x-\phi)_B^-\right) = 1$$

$\forall B \in \mathcal{B}_j$, or equivalently

$$\lim_{n\to\infty} P\left(\frac{1}{n}\log\mu_n(B) \geq (x-\phi)_B^-\right) = 0$$

Therefore it is sufficient to prove that there exists j_0 such that for every $j \geq j_0$ and $B \in \mathcal{B}_j$,

$$\lim_{n\to\infty} P\left(\frac{1}{n}\log\mu_n(U_B) \geq (x-\phi)_B^-\right) = 0$$

From (2) (recall that U_B was chosen in $\mathcal{A}$), this limit is zero in particular when

$$\inf_{\sigma\in\overline{U_B}, y\geq(x-\phi)_B^-} J(\sigma,y) > 0$$

Set $\varepsilon = \frac{x-\tilde{f}}{2}$. Let j_0 be such that for every $j \geq j_0$ we have $\phi_B^+ - \phi_{U_B}^- < \varepsilon$ for every $B \in \mathcal{B}_j$. Then for every $\sigma \in U_B$

$$\begin{aligned}(x-\phi)_B^- &= x - \phi_B^+ > x - \phi_{U_B}^- - \varepsilon \\ &\geq x - \phi(\sigma) - \varepsilon = x - (\phi(\sigma) - I_\mu(\sigma)) - I_\mu(\sigma) - \varepsilon \\ &\geq x - \widetilde{f} - I_\mu(\sigma) - \varepsilon = -I_\mu(\sigma) + \varepsilon\end{aligned}$$

Therefore

$$\inf_{\sigma\in\overline{U_B}, y\geq(x-\phi)_B^-} J(\sigma,y) \geq \inf_{\sigma\in\overline{U_B}, y\geq -I_\mu(\sigma)+\varepsilon} J(\sigma,y)$$

and the latter is positive by assumption.

Step 3 (lower estimate, $x = \widetilde{f}$). Given j, we have

$$\begin{aligned}&\underline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\mu_n(B) < \left(\widetilde{f}-\phi\right)_B^- \quad \forall B\in\mathcal{B}_j\right) \\ &\geq \sup_{\varepsilon>0}\underline{\lim}_{n\to\infty}\frac{1}{n}\log P\left(\frac{1}{n}\log\mu_n(B) < \left(\widetilde{f}-\varepsilon-\phi\right)_B^- \quad \forall B\in\mathcal{B}_j\right) \\ &\geq -\inf_{\varepsilon>0}\inf_{y<\widetilde{f}-\varepsilon}\widetilde{I}^-(y) = -\inf_{y<\widetilde{f}}\widetilde{I}^-(y)\end{aligned}$$

The proof is complete. ■

Lemma 13 *Given a number b, a positive integer j_0, a finite family $\mathcal{B}_j \subset \mathcal{B}(\Sigma)$, for every $j \geq j_0$ ($j \in \mathbb{N}$), such that $\Sigma = K_j^c \cup \left(\bigcup_{B\in\mathcal{B}_j} B\right)$ and a sequence*

$\varepsilon_j \to 0$, *assume that for every* $x < b$ *there is a sequence* $B_j \in \mathcal{B}_j$ *such that* $\lim_{j\to\infty} R_j = -\infty$ *where*

$$R_j = \overline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\log \mu_n\left(B_j\right) < (x-\phi)^+_{B_j}\right)$$

Define

$$\widetilde{I}^-(x) = \infty \text{ for every } x < b$$

Then, with this sequence R_j*, conditions (12)-(13) hold true for every* $x < b$*.*

Proof. Define $\widetilde{I}^-(x) = \infty$ for every $x < b$. In order to prove (12)-(13) we have only to prove that

$$\overline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\log \mu_n\left(B\right) < (b-\phi)^+_B \quad \forall B \in \mathcal{B}_j\right) \le R_j$$

for some $R_j \to -\infty$. This is obvious with the choice of R_j done in the statement, hence the lemma is proved. ■

Remark 14 *The good sets* B_j *to apply the previous lemma are expected to be the small sets around the maximum points of* $\phi - I_\mu$*.*

4.2.6 Full LDP

Theorem 15 *Under the assumptions of both Theorems A and B, assume that there exists a number* $\widetilde{f}$ *such that the function*

$$\widetilde{I}(x) = \begin{cases} \widetilde{I}^+(x) & \text{if} \quad x > \widetilde{f} \\ 0 & \text{if} \quad x = \widetilde{f} \\ \widetilde{I}^-(x) & \text{if} \quad x < \widetilde{f} \end{cases}$$

is strictly increasing for $x \ge \widetilde{f}$ *and strictly decreasing for* $x \le \widetilde{f}$*. Then* $\frac{1}{n}\log \int_\Sigma e^{n\phi(\sigma)} \mu_n^\omega(d\sigma)$ *satisfies a LDP with rate function* $\widetilde{I}$*.*

In the previous statement we understand that strict monotonicity has to hold in the regions where $\widetilde{I}$ is finite.

The passage from half-lines to general Borel sets under strict monotonicity of the rate function is standard (see for instance [10]), so the proof of the theorem will be omitted.

4.3 LDP for an "extended REM" model

In the previous paragraph we proved a very general large deviations theorem. In this one we will apply it to a more specific setup, which is actually a generalisation of Derrida's REM. We will call it "extended REM", not to be confused with the "generalised REM" of statistical mechanics literature. In the following paragraph we will then illustrate how our Shannon problem, described as a

statistical mechanical problem, fits into this "extended REM" framework. And we will compute the associated rate function, directly connected with the error exponent of the problem.

Now, after a quick outline of the theorem to be proved, we proceed to illustrate an equivalent description in terms of a random measure. Then we separately prove the positive and negative fluctuation LPD results.

Let $(\Omega, \mathcal{F}, P)$ be a probability space with expectation E and $\{X_n\}$ be a sequence of real random variables on $(\Omega, \mathcal{F}, P)$, with symmetric distribution, such that the following limit

$$\Lambda(\beta) = \lim_{n\to\infty} \frac{1}{n} \log E\left[e^{n\beta X_n}\right] \in [0, \infty)$$

exists (finite) for all real numbers β, and that $\Lambda(\beta)$ is a differentiable, strictly convex function. Recall that Λ is always convex (see Lemma 2.3.9 of [6]), so the *strict* convexity is the true assumption. Under these assumptions the sequence (X_n) satisfies a LDP with rate function $\Lambda^*(a)$, the Fenchel-Legendre transform of $\Lambda(\beta)$:

$$\Lambda^*(a) = \sup_{\beta\in\mathbb{R}} (a\beta - \Lambda(\beta))$$

Let $R \in (0, 1]$ be a given real number. Let $\{S_n\}$ be a sequence a finite sets such that

$$\lim_{n\to\infty} \frac{1}{n} \log \operatorname{card}(S_n) = R \log 2$$

For every $n \in \mathbb{N}$, let $(X_{s,n}; s \in S_n)$ be a family of i.i.d. random variables on $(\Omega, \mathcal{F}, P)$, distributed as X_n. We consider an extended version of the classical REM with energies $nX_{s,n}$: its partition function at inverse temperature $\beta > 0$ is defined as

$$Z^{\omega}_{\beta,n} = \sum_{s\in S_n} e^{n\beta X_{s,n}(\omega)}$$

Our main result is the following theorem.

Theorem 16 *Let*

$$a^* := \sup\{a : \Lambda^*(a) < R \log 2\}$$

$$\beta^* := \begin{cases} \frac{d\Lambda^*}{da}(a^*) & \text{if} \quad a^* \in D^{\circ}_{\Lambda^*} \\ \infty & \text{if} \quad a^* \notin D^{\circ}_{\Lambda^*} \end{cases}$$

$$f_\beta := \begin{cases} \frac{\Lambda(\beta) + R\log 2}{\beta} & \text{if} \quad \beta < \beta^* \\ a^* & \text{if} \quad \beta \geq \beta^* \end{cases}$$

. Then $f^{\omega}_{\beta,n} = \frac{1}{\beta n} \log Z^{\omega}_{\beta,n}$ satisfies a LDP with the good rate function

$$I(t) = \begin{cases} +\infty & \text{if} \quad t < f_\beta \\ 0 & \text{if} \quad t = f_\beta \\ \Lambda^*(t) - R\log 2 & \text{if} \quad t > f_\beta \end{cases}$$

Moreover, $I(t) > 0$ (possibly infinite) for $t > f_\beta$.

Since $I(t) > 0$ for all $t \neq f_\beta$, $f^{\omega}_{\beta,n}$ converges to f_β in the mean and P-a.s.

4.3.1 Preliminaries

We recall some basic facts about the Fenchel-Legendre transform which will come handy. The following statement can be found for example in [10], Theorems VI.5.3 and VI.5.6 and the subsequent discussion.

Theorem 17 *Let f be a convex lower semicontinuous function on $\mathbb{R}$ and let * denote F-L transform. Then the following conclusions hold.*

1. *f^* is a convex lower semicontinuous function on $\mathbb{R}$.*
2. *$xy \leq f(x) + f^*(y)$ for all x and y in $\mathbb{R}$.*
3. *$xy = f(x) + f^*(y)$ if and only if y is a value between the left and the right derivative of f in x.*
4. *$f^{**} = f$.*
5. *f is strictly convex on its domain if and only if f^* is essentially smooth.*

A function f is *essentially smooth* if:
i) the interior D_f° of its domain $D_f = \{x : f(x) < \infty\}$ is non empty;
ii) f is differentiable in D_f°;
iii) at finite boundary points of D_f, the lateral derivative is infinite.

Therefore by our assumptions on Λ, all these properties hold true both for Λ and Λ^*.

By elementary inspection, Λ and Λ^* are symmetric and $\Lambda(0) = \Lambda^*(0) = 0$. Hence by strict convexity both are strictly positive away from zero, increasing on the positive half-line and $\lim_{a\to\infty} \Lambda^*(a) = +\infty$ and $\lim_{\beta\to\infty} \Lambda(\beta) = +\infty$. By the same reason, they are both *good* rate functions.

The assumptions of Gärtner-Ellis theorem apply ([6], Thm. 2.3.6), so $\{X_n\}$ satisfies the LDP with rate function Λ^*. From the continuity of Λ^* in $D_{\Lambda^*}^\circ$, we have

$$\Lambda^*(a) = -\lim_{n\to\infty} \frac{1}{n} \log P(X_n \in (a, b))$$

for $0 \leq a < b$, independently of b (even infinite), possibly with the exception of the value $a = \sup D_{\Lambda^*}$.

Let us now clarify a few facts about Λ^*, a^* and f_β in our "extended REM" setup - the proofs will be in the following sections. Let $D_{\Lambda^*} = \{\beta : \Lambda^*(\beta) < \infty\}$ and $D_{\Lambda^*}^\circ$ be its interior. The set D_{Λ^*} is either a symmetric interval of the form $[-\overline{a}, \overline{a}]$ with $\overline{a} > 0$, or $D_{\Lambda^*} = \mathbb{R}$, Λ^* is differentiable in $D_{\Lambda^*}^\circ$, and it has infinite lateral derivatives at $\pm\overline{a}$ (when applicable, namely when $D_{\Lambda^*} \neq \mathbb{R}$). About the definition of a^*, since $\Lambda^*(a)$ is strictly increasing for $a > 0$, and it is lower semicontinuous, we have $|a| < a^*$ if and only if $\Lambda^*(a) < R\log 2$, $|a| > a^*$ if and only if $\Lambda^*(a) > R\log 2$ (possibly $\Lambda^*(a) = \infty$), and $\Lambda^*(a^*) \leq R\log 2$, with equality holding when $a^* \in D_{\Lambda^*}^\circ$. Finally, $f_\beta \geq a^*$ for all $\beta > 0$ because, by Theorem 17, $a\beta \leq \Lambda^*(a) + \Lambda(\beta)$ for all values of a and β, so we choose $a = a^*$, recall that $\Lambda^*(a^*) \leq R\log 2$ and look to the definition of f_β.

4.3.2 Rescaling the problem

Let

$$\mu_n^\omega := \frac{1}{\mathrm{card}(S_n)} \sum_{s \in S_n} \delta_{X_{s,n}(\omega)}$$

$$\widetilde{Z}_{\beta,n}^\omega := \int_{\mathbb{R}} e^{n\beta\sigma} \mu_n^\omega (d\sigma)$$

so we have

$$Z_{\beta,n}^\omega = \mathrm{card}(S_n) \cdot \widetilde{Z}_{\beta,n}^\omega$$

The strategy is to apply Theorems A and B to $\widetilde{Z}_{\beta,n}^\omega$ and then to get, with a suitable translation, a LDP for $Z_{\beta,n}^\omega$.

Proposition 18 *With the notations of Theorem 16, $\frac{1}{n} \log \widetilde{Z}_{\beta,n}^\omega$ satisfies a LDP with the good rate function*

$$\begin{aligned} \widetilde{I}(t) &= I\left(\frac{t + R\log 2}{\beta}\right) \\ &= \begin{cases} +\infty & \text{if} \quad t < \widetilde{f}_\beta \\ 0 & \text{if} \quad t = \widetilde{f}_\beta \\ \Lambda^*\left(\frac{t+R\log 2}{\beta}\right) - R\log 2 & \text{if} \quad t > \widetilde{f}_\beta \end{cases} \end{aligned}$$

where

$$\widetilde{f}_\beta = \beta \cdot f_\beta - R\log 2 = \begin{cases} \Lambda(\beta) & \text{if} \quad \beta < \beta^* \\ \beta a^* - R\log 2 & \text{if} \quad \beta \geq \beta^* \end{cases}$$

Moreover, $\widetilde{I}(t) > 0$ (possibly infinite) for $t > \widetilde{f}_\beta$.

To prove this, we shall apply Theorem A to prove the LDP for $\frac{1}{n} \log \widetilde{Z}_{\beta,n}^\omega$ for positive fluctuations, and we shall find the rate $\widetilde{I}^+$ equal to $\Lambda^*\left(\frac{t+R\log 2}{\beta}\right) - R\log 2$ for $t > \widetilde{f}_\beta$ (and zero otherwise). Then we shall apply Theorem B to prove the LDP for negative fluctuations, with rate $\widetilde{I}^-$ equal to ∞ for $t < \widetilde{f}_\beta$ and 0 for $t \geq \widetilde{f}_\beta$. Finally we glue the rates at $t = \widetilde{f}_\beta$ as in Theorem 15. The new global rate, which is the function $\widetilde{I}(t)$ defined above, satisfies the assumption of strict monotonicity of Theorem 15 (for the reason described next), so $\frac{1}{n} \log \widetilde{Z}_{\beta,n}^\omega$ satisfies the LDP with rate $\widetilde{I}$. Moreover, $\widetilde{I}$ is a good rate function since Λ^* is a good rate function. Finally, the property $\widetilde{I}(t) > 0$ for $t > \widetilde{f}_\beta$ is equivalent to $I(t) > 0$ for $t > f_\beta$, namely $\Lambda^*(t) > R\log 2$ for $t > f_\beta$, which is true since $t > f_\beta$ implies $t > a^*$, and $t > a^*$ implies $\Lambda^*(t) > R\log 2$.

We promised to prove that $\widetilde{I}(t)$ is strictly increasing for $t \geq \widetilde{f}_\beta$. The claim is algebraically equivalent to prove that for $t > f_\beta$, $\Lambda^*(t) > R\log 2$ and $\Lambda^*(t)$ is strictly increasing. The first fact has been proved above; the second claim is true for all positive values of t, and of course $f_\beta \geq a^* > 0$. The proof is complete.

The relation with Theorem 16 is:

Proposition 19 *Theorem 16 and Proposition 18 are equivalent.*

The equivalence is simply based on the identity

$$\frac{1}{\beta n}\log Z^{\omega}_{\beta,n} = \frac{1}{\beta}\left(\frac{1}{n}\log \widetilde{Z}^{\omega}_{\beta,n}\right) + \frac{R\log 2}{\beta} + \frac{1}{\beta}\left(\frac{1}{n}\log \mathrm{card}(S_n) - R\log 2\right)$$

and the first two claims of the following lemma.

Lemma 20 *Let $\{\alpha_n\}$ be a sequence of real numbers such that $\lim_{n\to\infty}\alpha_n = 0$. Let $\lambda \neq 0$ and η be two real numbers. Let $\{Y_n\}$ be a sequence of real random variables satisfying the LDP with the good rate function G. Then:*

i) $\{Y_n + \alpha_n\}$ satisfies the LDP with the same rate function G;

ii) $\{\lambda Y_n + \eta\}$ satisfies the LDP with the good rate function H defined as $H(x) = G\left(\frac{x-\eta}{\lambda}\right)$;

iii) if $E[Y_n]$ converges to a real number M, then $\{Y_n - E[Y_n]\}$ satisfies the LDP with the rate function $G(\cdot + M)$.

Proof. About (i) We do not give all the standard details but just notice that the estimate from above boils down to prove (easy by contradiction) that

$$\sup_n \inf_{\overline{A_{\frac{1}{n}}}} G \geq \inf_{\overline{A}} G$$

for any Borel set A, where A_ε denotes the ε-neighborhood of A; the estimate from below reduces to prove (again easy by contradiction) that

$$\inf_{U\in C}\inf_{U} G \leq \inf_{\overset{\circ}{A}} G$$

for any Borel set A with $\overset{\circ}{A} \neq \emptyset$, where C is the class of all open sets $U \subset \overset{\circ}{A}$ having positive distance from $\overline{A^c}$.

The proof of (ii) is trivial, and (iii) follows from (i) and (ii). ■

4.3.3 LDP for $\frac{1}{n}\log \widetilde{Z}^{\omega}_{\beta,n}$, positive fluctuations

Aim of this section is to apply Theorem A to

$$\frac{1}{n}\log \widetilde{Z}^{\omega}_{\beta,n} = \frac{1}{n}\log \int_{\mathbb{R}} e^{n\beta\sigma}\mu^{\omega}_n(d\sigma)$$

To verify the assumptions of Theorem A we apply Lemmas 7 and 8. In the following subsection we prove that the random measure μ^{ω}_n satisfies the assumptions of Lemma 10. Then we identify $\widetilde{I}^{+}$ as

$$\widetilde{I}^{+}(x) = \begin{cases} 0 & \text{if} \quad x \leq \widetilde{f}_\beta \\ \Lambda^*\left(\frac{x+R\log 2}{\beta}\right) - R\log 2 & \text{if} \quad x > \widetilde{f}_\beta \end{cases}$$

Check of the assumptions of Theorem A

Let $\Sigma = \mathbb{R}$ with the Euclidean topology. Let $\mathcal{A}$ be the base of the topology given by all open intervals (a, b) such that

$$a \neq 0, \quad b \neq 0$$

$$\Lambda^*(a) \neq R\log 2, \quad \Lambda^*(b) \neq R\log 2$$

Let $J : \Sigma \times \mathbb{R} \to [0, \infty]$ be defined as

$$J(\sigma, x) = \begin{cases} (\Lambda^*(\sigma) - R\log 2)^+ & \text{for} \quad x \leq -(\Lambda^*(\sigma) \wedge R\log 2) \\ \infty & \text{for} \quad x > -(\Lambda^*(\sigma) \wedge R\log 2) \end{cases}$$

It is easy to verify that J is lower semicontinuous and that for every v and γ, the set of (σ, x) such that $x \geq v$ and $J(\sigma, x) \leq \gamma$ is compact. So it satisfies the condition similar to the property of good rate function. Moreover, for every x, $\sigma \mapsto J(\sigma, x)$ is convex, symmetric and non decreasing over positive σ, and for every σ, $x \mapsto J(\sigma, x)$ is non decreasing and piecewise constant.

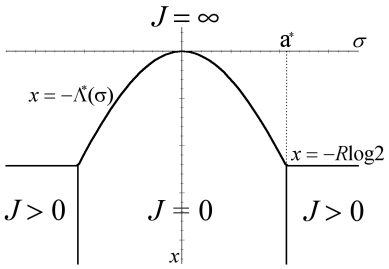

Figure 1: Values of the function $J(\sigma, x)$

The function I_μ defined by (16) takes the form

$$I_\mu(\sigma) = \begin{cases} \Lambda^*(\sigma) & \text{if} \quad \Lambda^*(\sigma) \leq R\log 2 \\ \infty & \text{if} \quad \Lambda^*(\sigma) > R\log 2 \end{cases} \tag{18}$$

It will turn out to be the rate function of μ_n^ω, for P-a.e. $\omega \in \Omega$.

Lemma 21

$$\lim_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \mu_n(a, b) > v\right) = -J(a, v)$$

holds true for every triple (a, b, v) *such that*

$$0 < a < b, \quad \Lambda^*(a) \neq R\log 2, \quad a \neq \sup D_{\Lambda^*}$$
$$v \neq 0, v \neq -\Lambda^*(a), v \neq -R\log 2$$

Proof. As we announced, we have

$$P\left(\frac{1}{n}\log\mu_n(a,b) > v\right)$$
$$= P\left(\sum_{s\in S_n} 1_{\{X_{s,n}\in(a,b)\}} > \mathrm{card}(S_n)e^{nv}\right) = P\left(B_{N_n,p_n} > x_n\right)$$

where B_{N_n,p_n} is a binomial of parameters $N_n = \mathrm{card}(S_n)$ and $p_n = P(X_n \in (a,b))$, and $x_n = N_n e^{nv}$. We have

$$\lim_{n\to\infty}\frac{1}{n}\log N_n = R\log 2$$

$$\lim_{n\to\infty}\frac{1}{n}\log p_n = -\Lambda^*(a)$$

$$\lim_{n\to\infty}\frac{1}{n}\log x_n = v + R\log 2$$

Therefore we use the result of the Appendix A with

$$\alpha = R\log 2, \quad \beta = \Lambda^*(a), \quad \gamma = v + R\log 2$$

This is the table of correspondence between the results of the Appendix A and what we need:

- $J(a,v) = 0$ when $v \le -(\Lambda^*(a) \wedge R\log 2)$ and $\Lambda^*(a) \le R\log 2$, namely when $v \le -\Lambda^*(a)$ and $\Lambda^*(a) \le R\log 2$; this case corresponds to $\gamma \le \alpha - \beta$ and $\alpha - \beta \ge 0$ and hence is covered by (31) and (36), when $v \ne -\Lambda^*(a)$, $\Lambda^*(a) \ne R\log 2$ and $v \ne -R\log 2$;
- $J(a,v) = \Lambda^*(a) - R\log 2$ if $v \le -R\log 2$ and $R\log 2 < \Lambda^*(a)$; this case is covered by (30) with $v \ne -R\log 2$;
- $J(a,v) = \infty$ when $v > -(\Lambda^*(a) \wedge R\log 2)$, that is, when $v + R\log 2 > (R\log 2 - \Lambda^*(a))^+$; letting $v \ne 0$, this case is covered by (29) if $v > 0$ and by (32) if $v < 0$.

The proof is complete. ■

Lemma 22 *With $K_j = [-j, j]$, the tail condition (3) holds true.*

Proof. We have

$$\inf_{\sigma\in\overline{K_j^c}x\ge L} J(\sigma,x) = \inf_{|\sigma|\ge j, x\ge L} J(\sigma,x) = J(j,L)$$

so, for j sufficiently large ($j > a^*$),

$$\inf_{\sigma\in\overline{K_j^c}x\ge L} J(\sigma,x) = \begin{cases} (\Lambda^*(j) - R\log 2)^+ & \text{for} \quad L \le -R\log 2 \\ \infty & \text{for} \quad L > -R\log 2 \end{cases}$$

and therefore $\lim_{j\to\infty} \inf_{\sigma\in\overline{K_j^c} x\geq L} J(\sigma,x) = +\infty$ for every L. This is condition (8). As to (9), for bounded continuous functions g we have

$$E\left[\int_\Sigma g(\sigma)\,\mu_n^\omega(d\sigma)\right] = \frac{1}{\mathrm{card}(S_n)}\sum_{s\in S_n} E\left[g(X_{s,n})\right]$$
$$= E\left[g(X_n)\right] = \int_{\mathbb{R}} g(x)\,\nu_n(dx)$$

where ν_n is the law of X_n. Thus, $E[\mu_n] = \nu_n$. Given any $\lambda > 1$, we have

$$\int_{\mathbb{R}} e^{n\lambda\beta x}\nu_n(dx) = E\left[\exp(n\lambda\beta X_n)\right]$$

so condition (9) is a consequence of our assumptions on $\Lambda(\beta)$. By theorem 5, the tail condition (3) holds true. ■

Corollary 23 *The assumptions of theorem A hold true, with J given above.*

Proof. The function J, restricted to $\sigma \in (0,\infty)$, satisfies the assumptions of lemma 7. Hence for the triples (a,b,v) of lemma 21 we have (10)-(11). The same is true for the triples (a,b,v) such that

$$a < b < 0, \quad \Lambda^*(b) \neq R\log 2, \quad b \neq \inf D_{\Lambda^*}$$
$$v \neq 0, v \neq -\Lambda^*(b), v \neq -\log 2$$

because the random variables $\mu_n(a,b)$ and $\mu_n(-b,-a)$ have the same law under P, and J is also symmetric in σ. In this way, for every interval (a,b) of $\mathcal{A}$ except those with $a < 0 < b$, (10)-(11) are satisfied for all v except three values. By lemma 8, they are satisfied for every v.

When $a < 0 < b$, the case $v > 0$ is trivial as before. When $v < 0$,

$$\inf_{\sigma\in(a,b),x>v} J(\sigma,x) = \inf_{\sigma\in[a,b],x\geq v} J(\sigma,x) = \inf_{x>v} J(0,x) = 0$$

Thus we just need to show that $P\left(\frac{1}{n}\log\mu_n(a,b) > v\right) \to 1$. Since

$$P\left(\mu_n(a,b) > \exp nv\right)$$
$$\geq P\left(\mu_n[b,\infty) < \frac{1-\exp nv}{2}, \mu_n(\infty,a] < \frac{1-\exp nv}{2}\right)$$

it is sufficient to prove that $P(\mu_n(\alpha,\infty) > C) \to 0$ for every $\alpha > 0$ and $C \in (0,1)$. This is guaranteed by (37) and completes the proof when $v < 0$. The case $v = 0$ follows again from lemma 8.

Finally, the tail condition (3) has been verified in the previous lemma. The proof of the corollary is complete. ■

Remark 24 *It is natural to ask ourselves whether the properties of μ_n just discussed are a consequence of Sanov Theorem or the k-dimensional Cramer*

theorem, since, at least in the particular case when all the $X_{i,n}$ have the same law, μ_n is an empirical measure of those treated by these theorems. Given Borel subsets $A_1, ..., A_k$ and $F_1, ..., F_k$ are of $\mathbb{R}$, these theorems allow us to compute asymptotically

$$P\left(\mu_n\left(A_1\right) \in F_1, ..., \mu_n\left(A_k\right) \in F_k\right)$$

This is not sufficient for us since we have to deal with sets $A_{j,n}$ and $F_{k,n}$ which depend on n themselves.

Computation of $\widetilde{I^+}$

We first have to show that $\tilde{f}.$ is the F-L transform of I_μ. Recall the definition of a^* and β^* given in Theorem 16.

Lemma 25 *I_μ and $\tilde{f}.$ are F-L transforms of each other.*

Proof. Convexity and F-L transform of a function f are better understood introducing the set of lines which are below f. Let

$$\mathfrak{R}(f) = \{(a,b) : ax + b \leq f(x), \quad \forall x \in \mathbb{R}\}$$

Then f is convex iff $f(x) = \sup\{ax + b : (a,b) \in \mathfrak{R}(f)\}$ and moreover its transform is given by

$$\begin{aligned} f^*(a) &= -\inf_x (f(x) - ax) \\ &= -\sup\{b : b \leq f(x) - ax, \quad \forall x \in \mathbb{R}\} \\ &= -\sup\{b : (a,b) \in \mathfrak{R}(f)\} \end{aligned}$$

We claim that $\mathfrak{R}(\tilde{f})$ is the subset of $\mathfrak{R}(\Lambda)$ with $a \leq a^*$. This is quite obvious, since $\tilde{f}$ and Λ coincide until the derivative of Λ reach a^*, then the former grows with constant slope. More formally, note that $\Lambda(\beta)$ is nowhere below $a^*\beta - \Lambda^*(a^*) \geq \beta a^* - R\log 2$ and they are tangent at $\beta = \beta^*$ (possibly infinite, with an asymptote), so that $\tilde{f} \geq \Lambda$ and hence $\mathfrak{R}(\tilde{f}) \subseteq \mathfrak{R}(\Lambda)$. By the same argument, even when $\beta^* = \infty$, if $(a,b) \in \mathfrak{R}(\tilde{f})$, then the line $a\beta + b$ must be definitively below $a^*\beta - \Lambda^*(a^*)$, so $a \leq a^*$. The other direction is even simpler. Supposing $(a,b) \in \mathfrak{R}(\Lambda)$ and $a \leq a^*$, for all β we have $a\beta + b \leq \Lambda(\beta)$. On the set $\beta < \beta^*$ we are done; then supposing $\beta^* < \infty$, at $\beta = \beta^*$ we have $a\beta^* + b \leq \Lambda(\beta^*) = a^*\beta^* - R\log 2$; since the first line starts below and has smaller slope, it will be below the second line on all $[\beta^*, \infty)$.

Perusing the claim,

$$\begin{aligned} \tilde{f}^*(a) &= -\sup\{b : (a,b) \in \mathfrak{R}(\tilde{f})\} \\ &= \begin{cases} -\sup\{b : (a,b) \in \mathfrak{R}(\Lambda)\} & \text{se } a \leq a^* \\ -\sup\emptyset & \text{se } a > a^* \end{cases} \\ &= \begin{cases} \Lambda^*(a) & \text{if } a \leq a^* \\ +\infty & \text{if } a > a^* \end{cases} \end{aligned}$$

This implies that I_μ is the F-L transform of $\tilde{f}.$. To see the converse we simply note that f_β is convex. ■

Let us now compute the function

$$\widetilde{I}^{+}(x) = \inf_{\sigma\in\Sigma} J(\sigma, x-\phi(\sigma)) = \inf_{\sigma\in\mathbb{R}} J(\sigma, x-\beta\sigma)$$

Lemma 26 *The function $\widetilde{I}^{+}$ is given by*

$$\widetilde{I}^{+}(x) = \begin{cases} \Lambda^*\left(\frac{x+R\log 2}{\beta}\right) - R\log 2 & \text{if} \quad x > f_\beta \\ 0 & \text{if} \quad x \le f_\beta \end{cases}$$

Proof. By the definition of $\tilde{I}^{+}$ and I_μ, and by Lemma 25, we have

$$\begin{aligned} \tilde{I}^{+}(x) > 0 &\Leftrightarrow J(\sigma, x-\beta\sigma) > 0, \quad \forall\sigma\in\mathbb{R} \\ &\Leftrightarrow I_\mu(\sigma) > -x+\beta\sigma, \quad \forall\sigma\in\mathbb{R} \\ &\Leftrightarrow x > f_\beta \end{aligned}$$

Suppose $x > f_\beta$, so that for all σ, $\beta\sigma - x < I_\mu(\sigma)$, and hence $\beta\sigma - x < \Lambda^*(\sigma)$ on the set $|\sigma| \le a^*$ i.e. $\{\sigma : \Lambda^*(\sigma) \le R\log 2\}$. Since in general

$$\widetilde{I}^{+}(x) = \inf_{\sigma\in D_x} \left(\Lambda^*(\sigma) - R\log 2\right)^{+}$$

where

$$D_x = \{\sigma\in\mathbb{R}: \quad \beta\sigma - x \ge \Lambda^*(\sigma) \wedge R\log 2\}$$

D_x reduces to the half-line

$$\{\sigma\in\mathbb{R}: \quad \beta\sigma - x \ge R\log 2\} = \left[\frac{x+R\log 2}{\beta}, \infty\right)$$

The claim is now proved by the monotonicity of Λ^*. ■

4.3.4 LDP for $\frac{1}{n}\log\widetilde{Z}^{\omega}_{\beta,n}$, negative fluctuations

Aim of this section is to apply Theorem B to

$$\frac{1}{n}\log\widetilde{Z}^{\omega}_{\beta,n} = \frac{1}{n}\log\int_{\mathbb{R}} e^{n\beta\sigma}\mu^{\omega}_n(d\sigma)$$

and identify $\widetilde{I}^{-}$ as

$$\widetilde{I}^{-}(x) = \begin{cases} \infty & \text{if} \quad x < \widetilde{f}_\beta \\ 0 & \text{if} \quad x \ge \widetilde{f}_\beta \end{cases}$$

To this purpose, we use Lemmas 12 and 13. In Lemma 13 we take $b = \widetilde{f}_\beta$. In Lemma 12, we have to prove that $\widetilde{f}$, defined by (17), is equal to $\widetilde{f}_\beta$. As a simple consequence of the result of the previous section we have:

Corollary 27 *The number $\widetilde{f}$, defined by (17), Fenchel-Legendre transform of I_μ, is equal to $\widetilde{f}_\beta$.*

For every $j \in \mathbb{N}$, let $K_j = [-j, j]$ as above and $\mathcal{B}_j$ be the family

$$\left[\frac{k}{j}, \frac{k+1}{j}\right], -j^2 \leq k < j^2$$

Let us verify (at the same time) the assumptions of lemma 12 and of lemma 13 with $b = \widetilde{f}_\beta$. The assumptions (1)-(2) of Theorem A have been already verified. For every $j \in \mathbb{N}$ and $B \in \mathcal{B}_j$ let $U_B \in \mathcal{A}$, $U_B \supset B$, be any open interval having Hausdorff distance from B at most $\frac{1}{2j}$. We have

$$\phi_B^+ - \phi_{U_B}^- \leq \beta \frac{2}{j}$$

hence (technical condition (i))

$$\lim_{j \to \infty} \max_{B \in \mathcal{B}_j} \left(\phi_B^+ - \phi_{U_B}^-\right) = 0$$

As to condition (ii), given $\varepsilon > 0$, for every $\sigma \in \Sigma$ we have that $\inf_{y \geq -I_\mu(\sigma)+\varepsilon} J(\sigma, y)$ is ∞ for $|\sigma| \leq a^*$ and $(\Lambda^*(\sigma) - R \log 2)$ for $|\sigma| > a^*$, so

$$\inf_{\sigma \in \overline{U_B}, y \geq -I_\mu(\sigma)+\varepsilon} J(\sigma, y) \geq \inf_{\sigma \in \Sigma, y \geq -I_\mu(\sigma)+\varepsilon} J(\sigma, y) > 0$$

Therefore, as soon as we have identified $\widetilde{I}^-(x)$ for every $x < \widetilde{f}_\beta$, the choice

$$\widetilde{I}^-(x) = 0 \text{ for every } x \geq \widetilde{f}_\beta$$

is correct.

For the complementary region, $x < \widetilde{f}_\beta$, we apply lemma 13 with $b = \widetilde{f}_\beta$. To this end, given $x < \widetilde{f}_\beta$, it is sufficient to find a sequence $B_j \in \mathcal{B}_j$ such that

$$\lim_{j \to \infty} \overline{\lim}_{n \to \infty} \frac{1}{n} \log P\left(\frac{1}{n} \log \mu_n(B_j) < (x - \phi)_{B_j}^+\right) = -\infty$$

Good sets B_j are expected to be those around the values of σ which realize $\sup V$ where $V(\sigma) = \phi(\sigma) - I_\mu(\sigma) = \beta\sigma - I_\mu(\sigma)$. Denote by σ^* the value which maximizes $V(\sigma)$. By Theorem 17 and Lemma 25, $V(\sigma) \leq \widetilde{f}_\beta$, with equality holding only for $\sigma = \sigma^* = \frac{d}{d\beta}\widetilde{f}_\beta$. Note that $\sigma^* \in [0, a^*]$ and it is a^* for $\beta \geq \beta^*$.

We take as B_j the interval, or the left one of the two intervals, which contains σ^*. In particular, with this choice, if $\sigma^* = a^*$ then $a_j := \inf B_j < a^*$.

Take $x < \widetilde{f}_\beta$, so of the form

$$x = \widetilde{f}_\beta - \varepsilon$$

for a suitable $\varepsilon > 0$. Then

$$(x-\phi)^+_{B_j} = x - \beta a_j = x - \beta\sigma^* + \beta\delta_j$$

with $|\delta_j| \leq \frac{1}{j}$

$$= \widetilde{f}_\beta - \beta\sigma^* - \varepsilon + \beta\delta_j = -I_\mu(\sigma^*) - \varepsilon + \beta\delta_j$$

It is now sufficient to prove

$$\lim_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\log \mu_n(B_j) < -I_\mu(\sigma^*) - \varepsilon + \beta\delta_j\right) = -\infty$$

for all j sufficiently large. Take the range of j such that $-\varepsilon + \beta\delta_j \leq -\frac{\varepsilon}{2}$. Write $B_j = [a_j, b_j]$. Then

$$\begin{aligned} P\left(\frac{1}{n}\log \mu_n(B_j) < -I_\mu(\sigma^*) - \varepsilon + \beta\delta_j\right) &\leq P\left(\frac{1}{n}\log \mu_n(B_j) < -I_\mu(\sigma^*) - \frac{\varepsilon}{2}\right) \\ &= P\left(\sum_{s\in S_n} 1_{\{X_{s,n}\in[a_j,b_j]\}} < \mathrm{card}(S_n) e^{-n\left(I_\mu(\sigma^*)+\frac{\varepsilon}{2}\right)}\right) \\ &= P\left(B_{N_n,p_n} < x_n\right) \end{aligned}$$

where B_{N_n,p_n} is a binomial of parameters $N_n = \mathrm{card}(S_n)$ and $p_n = P(X_n \in [a_j, b_j])$, and $x_n = N_n e^{nv}$ with $v = -\left(I_\mu(\sigma^*) + \frac{\varepsilon}{2}\right)$. We have

$$\lim_{n\to\infty} \frac{1}{n}\log N_n = R\log 2$$

$$\lim_{n\to\infty} \frac{1}{n}\log p_n = -\Lambda^*(a_j)$$

(recall from above that $a_j < a^*$)

$$\lim_{n\to\infty} \frac{1}{n}\log x_n = v + R\log 2$$

Therefore we use the result of the Appendix with

$$\alpha = R\log 2, \quad \beta = \Lambda^*(a_j), \quad \gamma = v + R\log 2$$

Since $a_j < \sigma^* \leq a^*$ and Λ^* is increasing on the positive half-line, $\Lambda^*(a_j) < \Lambda^*(\sigma^*) \leq \Lambda^*(a^*) \leq R\log 2$, so that $\alpha > \beta$. The same argument shows that $\gamma < \alpha - \beta$, since this is equivalent to

$$\Lambda^*(a_j) < I_\mu(\sigma^*) + \frac{\varepsilon}{2}$$

Equations (38) and (39) in the Appendix complete the proof as long as $\gamma \neq 0$, i.e. when $I_\mu(\sigma^*) + \varepsilon/2 \neq R\log 2$. If this is not the case, one can simply use $\frac{\varepsilon}{3}$ instead of $\frac{\varepsilon}{2}$ and get the desired result.

4.4 LDP for the Shannon problem

In this paragraph we will compute a formula for the Shannon error exponent $E_{err}(\beta, R)$ through the above LD theorem in its "extended REM" version.

We now briefly recall the Shannon problem setup in the statistical mechanical formulation introduced in the previous chapter. Let $\beta \in [0, \infty[$, $R \in [0, 1]$, $n \in \mathbb{N}$, $B := \{0, 1\}$, $S_n := B^{\lfloor Rn \rfloor}$ and let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability triple. For all $n \in \mathbb{N}$ let

$$\widetilde{x}_{s,j} : \begin{array}{c} (\Omega, \mathcal{F}, \mathbb{P}) \to B \\ \omega \mapsto \widetilde{x}_{s,j}(\omega) \end{array} \quad \text{for all } s \in S_n, 1 \leq j \leq n$$

be a sequence of families of independent random variables such that

1. for $s \neq 0$ the $\widetilde{x}_{s,j}$ are $Bernoulli\left(\pm 1, \frac{1}{2}\right)$;
2. $\widetilde{x}_{0,j} = w_j \sim Bernoulli(\pm 1, p)$.

   Let now

$$E_\omega^n(s) := \frac{1}{2}\left[n - \sum_{j=1,\ldots,n} \widetilde{x}_{s,j}(\omega)\right]$$

$$\begin{aligned} Z_{\beta,n}^*(\omega) &= e^{-\beta E_\omega^n(0)} \\ \hat{Z}_{\beta,n}^0(\omega) &= \sum_{\substack{s\,:\,\pi_1^n(s)=0 \\ s \neq 0}} e^{-\beta E_\omega^n(s)} \\ Z_{\beta,n}^1(\omega) &= \sum_{s\,:\,\pi_1^n(s)=1} e^{-\beta E_\omega^n(s)} \end{aligned}$$

   where $\pi_1^n(s)$ is the projection of s onto its first component. Then the error exponent is defined as:

$$E_{err}(\beta, R) = \liminf_{n\to\infty} -\frac{1}{n} \log P\left[\ Z_{\beta,n}^1(\omega) \geq Z_{\beta,n}^*(\omega) + \hat{Z}_{\beta,n}^0(\omega)\ \right]$$

   while the condition for errorfree coding in terms of (β, R) is

$$\lim_{n\to\infty} P\left\{\omega \in \Omega :\ Z_{\beta,n}^1(\omega) < Z_{\beta,n}^*(\omega) + \hat{Z}_{\beta,n}^0(\omega)\ \right\} = 1$$

Conceptually, behind $E_\omega^n(s)$ there is a basic energy $h_n(s, \omega)$ which corresponds to the distance between the codeword associated to s and the one associated to the zero codeword $0 \in S$. Thus this zero codeword is a *groundstate* of $h_n(\cdot, \omega)$ for each chosen code ω, namely $0 = h_n\left(s^0, \omega\right) \leq h_n(s, \omega)$ for every $s \in S_n$ and $\omega \in \Omega$. The energy $E_\omega^n(s)$ is a random perturbation of $h_n(s, \omega)$ which takes into account also the effect of the channel-noise on the codeword

associated to the zero codeword. We shall never use $h_n(s,\omega)$ explicitly, it was introduced here for explanatory reasons only.

In order to perform our large deviation analysis, we will need first of all some free energy density bounds on the error exponent. Afterwards we will proceed to their computation through the LD theorem. At the end of the chapter we will then write the final result for the error exponent.

4.4.1 Free energy density bounds of the error exponent

In the present chapter, we have developed tools for studying "free energy densities"-like quantities. But we are now interested in studying a relation between "partition function"-like quantities:

$$\liminf_{n\to\infty} -\frac{1}{n}\log P\left[\ Z^1_{\beta,n}(\omega) \geq Z^*_{\beta,n}(\omega) + \hat{Z}^0_{\beta,n}(\omega)\ \right]$$

Our aim in this section is to find a large deviation relationship between the above object and the following "free energy densities":

$$\begin{aligned} F^*_{\beta,n} &:= \frac{1}{\beta n}\log Z^*_{\beta,n} \\ F^1_{\beta,n} &:= \frac{1}{\beta n}\log Z^1_{\beta,n} \\ \hat{F}^0_{\beta,n} &= \frac{1}{\beta n}\log \hat{Z}^0_{\beta,n} \end{aligned}$$

We fulfil the stated aim through the following

Lemma 28 *Assume that:*

1) $F^*_{\beta,n}$ satisfies a LDP with a rate functional I^*_β;

2) $F^1_{\beta,n}$ and $\hat{F}^0_{\beta,n}$ satisfy LDP's with the same rate functional I^1_β, which is right-continuous at some point f_β, such that $I^1_\beta[(f_\beta] = 0$;

3) $F^*_{\beta,n}$, $F^1_{\beta,n}$ and $\hat{F}^0_{\beta,n}$ are independent;

4) defined $I_\beta(x)$ as

$$I_\beta(x) = \inf_{r\in\mathbb{R}}\left(I^1_\beta(x+r) + I^*_\beta(r)\right)$$

we assume

$$\inf_{x>0} I_\beta(x) = \inf_{x\geq 0} I_\beta(x) = I_\beta(0)$$

Then

$$E_{err}(\beta,R) = \lim_{n\to\infty} -\frac{1}{n}\log P\left[\ Z^1_{\beta,n}(\omega) \geq Z^*_{\beta,n}(\omega) + \hat{Z}^0_{\beta,n}(\omega)\ \right] = -I_\beta(0)$$

In order to prove it, we split our discussion in two parts: bounds from above and from below.

Upper bound

We obviously have

$$P\left(Z_{\beta,n}^1 \geq Z_{\beta,n}^* + \hat{Z}_{\beta,n}^0\right) \leq P\left(Z_{\beta,n}^1 \geq Z_{\beta,n}^*\right) = P\left(F_{\beta,n}^1 \geq F_{\beta,n}^*\right)$$

So we can prove:

Proposition 29 *Assume that:*

1) $F_{\beta,n}^*$ *satisfies a LDP with a rate functional* I_β^**;*

2) $F_{\beta,n}^1$ *satisfies a LDP with a rate functional* I_β^1*;*

3) $F_{\beta,n}^*$ *and* $F_{\beta,n}^1$ *are independent.*

Then

$$\overline{\lim_{n\to\infty}} \frac{1}{n} \log p_n^{err} \leq -\inf_{x\geq 0} I_\beta(x)$$

where

$$I_\beta(x) = \inf_{r\in\mathbb{R}} \{I_\beta^1(x+r) + I_\beta^*(r)\}$$

Proof. By Lemma 48 (see the Appendix), $F_{\beta,n}^1 - F_{\beta,n}^*$ satisfies a LDP with rate $I_\beta(x)$, hence

$$\overline{\lim_{n\to\infty}} \frac{1}{n} \log P\left(F_{\beta,n}^1 - F_{\beta,n}^* \geq 0\right) \leq -\inf_{x\geq 0} I_\beta(x)$$

The conclusion is now obvious. ■

We shall see that in the Shannon example these assumptions can be checked and the infimum can be explicitly computed.

Lower bound

Proposition 30 *Assume that:*

1) $F_{\beta,n}^*$ *satisfies a LDP with a rate functional* I_β^**;*

2) $F_{\beta,n}^1$ *and* $\hat{F}_{\beta,n}^0$ *satisfy LDPs with rate functionals* I_β^1 *an3)* $F_{\beta,n}^*$*,* $F_{\beta,n}^1$ *and* $\hat{F}_{\beta,n}^0$ *are independent;*

4)

$$\inf_{x>0} \inf_{r\in\mathbb{R}} \{I_\beta^1(x+r) + I_\beta^0(r)\} = 0$$

Then

$$\underline{\lim}_{n\to\infty} \frac{1}{n} \log P\left(Z_{\beta,n}^1 \geq Z_{\beta,n}^* + \hat{Z}_{\beta,n}^0\right) \geq -\inf_{x>0} I_\beta(x)$$

where $I_\beta(x)$ *has been defined in the previous proposition.*

Proof. We have

$$\left\{Z_{\beta,n}^1 \geq 2Z_{\beta,n}^*,\ Z_{\beta,n}^1 \geq 2\hat{Z}_{\beta,n}^0\right\} \subset \left\{Z_{\beta,n}^1 \geq Z_{\beta,n}^* + \hat{Z}_{\beta,n}^0\right\}$$

Since $Z^0_{\beta,n}$, $\hat{Z}^+_{\beta,n}$ and $Z^-_{\beta,n}$ are independent, we have

$$P\left(Z^1_{\beta,n} \geq 2Z^*_{\beta,n},\ Z^1_{\beta,n} \geq 2\hat{Z}^0_{\beta,n}\right) \geq P\left(Z^1_{\beta,n} \geq 2Z^*_{\beta,n}\right) \cdot P\left(Z^1_{\beta,n} \geq 2\hat{Z}^0_{\beta,n}\right)$$

Indeed, the two events are positively correlated. A complete proof is given in the Appendix. Therefore we have

$$\begin{aligned} P\left(Z^1_{\beta,n} \geq Z^*_{\beta,n} + \hat{Z}^0_{\beta,n}\right) &\geq P\left(Z^1_{\beta,n} \geq 2Z^*_{\beta,n}\right) \cdot P\left(Z^1_{\beta,n} \geq 2\hat{Z}^0_{\beta,n}\right) \\ &= P\left(F^1_{\beta,n} \geq \frac{1}{\beta n}\log 2 + F^*_{\beta,n}\right) \cdot P\left(F^1_{\beta,n} \geq \frac{1}{\beta n}\log 2 + \hat{F}^0_{\beta,n}\right) \end{aligned}$$

Hence, given $\varepsilon, \varepsilon' > 0$, eventually

$$\begin{aligned} \frac{1}{n}\log P\left(Z^1_{\beta,n} \geq Z^*_{\beta,n} + \hat{Z}^0_{\beta,n}\right) \geq \frac{1}{n}\log P\left(F^1_{\beta,n} \geq \varepsilon + F^*_{\beta,n}\right) \\ + \frac{1}{n}\log P\left(F^1_{\beta,n} \geq \varepsilon' + \hat{F}^0_{\beta,n}\right) \end{aligned}$$

By lemma 48, we have

$$\underline{\lim}_{n\to\infty} \frac{1}{n}\log P\left(F^1_{\beta,n} - F^*_{\beta,n} \geq \varepsilon\right) \geq -\inf_{x>\varepsilon} I_\beta(x)$$

$$\underline{\lim}_{n\to\infty} \frac{1}{n}\log P\left(F^1_{\beta,n} - \hat{F}^0_{\beta,n} \geq \varepsilon'\right) \geq -\inf_{x>\varepsilon'} \tilde{I}_\beta(x)$$

where

$$\tilde{I}_\beta(x) = \inf_{r\in\mathbb{R}} \left(I^1_\beta(x+r) + I^0_\beta(r)\right)$$

Therefore

$$\underline{\lim}_{n\to\infty} \frac{1}{n}\log P\left(Z^1_{\beta,n} \geq Z^*_{\beta,n} + \hat{Z}^0_{\beta,n}\right) \geq -\left(\inf_{x>\varepsilon} I_\beta(x) + \inf_{x>\varepsilon'} \tilde{I}_\beta(x)\right)$$

Taking the supremum in $\varepsilon > 0$ and $\varepsilon' > 0$ on the right-hand-side, we easily get

$$\underline{\lim}_{n\to\infty} \frac{1}{n}\log P\left(Z^1_{\beta,n} \geq Z^*_{\beta,n} + \hat{Z}^0_{\beta,n}\right) \geq -\left(\inf_{x>0} I_\beta(x) + \inf_{x>0} \tilde{I}_\beta(x)\right)$$

The conclusion of the proposition follows now immediately from assumption 4). ■

In our application we shall prove that $I^1_\beta = I^0_\beta$, hence

$$\tilde{I}_\beta(x) = \inf_{r\in\mathbb{R}} \left(I^1_\beta(x+r) + I^1_\beta(r)\right)$$

If 0 is an accumulation point for the image of I^1_β, then it is easy to show that $\inf_{x>0} \tilde{I}_\beta(x) = 0$, so assumption 4) is satisfied. We shall check also the other assumptions and compute $\inf_{x>0} I_\beta(x)$, which turns out to be equal to $\inf_{x\geq 0} I_\beta(x)$ and also simply to $I_\beta(0)$.

4.4.2 The LDP of F_n^*

Lemma 31 *F_n^* satisfies a LDP with rate functional $I^*(x) = D_e(-x||p)$.*

Proof. The result follows since

$$-nF_n^* := E_\omega^n(0) = \sum_{j=1}^n \left(\frac{1}{2} - \frac{1}{2}\widetilde{x}_{s,j}\right)$$

is just a binomial r.v. (see for example [6]). ■

4.4.3 The LDP of $F^1_{\beta,R,n}$ and $\hat{F}^0_{\beta,R,n}$

In order to study the LDP of $F^1_{\beta,R,n}$ and compute its rate functional $I^1_{\beta,R}(x)$ we use theorem 16.

Corollary 32 *$F^1_{\beta,R,n}$ satisfies a LDP with rate functional*

$$I^1_{\beta,R}(x) = \begin{cases} +\infty & \text{if} \quad x < f(\beta,R) \\ 0 & \text{if} \quad x = f(\beta,R) \\ D_e\left(-x \,\|\, \frac{1}{2}\right) - R\log 2 & \text{if} \quad x > f(\beta,R) \end{cases} \tag{19}$$

where

$$f(\beta,R) := \begin{cases} \frac{1}{\beta}\left[\log\cosh\frac{\beta}{2} - \frac{\beta}{2} + R\log 2\right] & \text{if} \quad R > R^*(\beta) \\ -\delta_{GV}(R) & \text{if} \quad R \le R^*(\beta) \end{cases} \tag{20}$$

$\delta_{GV}(\cdot)$ being the Gilbert-Varshamov distance as defined at the beginning of the previous chapter, and

$$R^*(\beta) := D_2\left(\frac{1}{e^\beta+1}\middle\|\frac{1}{2}\right) = \frac{\frac{\beta}{2}\tanh\frac{\beta}{2} - \log\cosh\frac{\beta}{2}}{\log 2} \tag{21}$$

Proof. Firstly notice that Theorem 16 applies formally unchanged if the r.v.'s h_n are substituted with some other h_n^θ that are symmetric with respect to some fixed real $\theta \neq 0$. To see this, pose $h_{s,n} = h^\theta_{s,n} - \theta$ and apply the theorem to find that (with obvious notations) $\frac{1}{\beta n}\log Z_n^\theta = \frac{1}{\beta n}\log Z_n + \theta$ and $I^\theta(t) = I(t-\theta)$.

Now we check that the assumptions of the theorem are satisfied posing

$$\Sigma_n := \{s \in S_n : \pi_1^n(s) = 1\} \qquad \text{and} \qquad h_{s,n} := \frac{1}{2n}\sum_{j=1}^n \widetilde{x}_{s,j} - \frac{1}{2}$$

We already know that $\{h_{s,n}; s \in \Sigma_n\}$ are IID; they are also symmetric with respect to $-1/2$. The set Σ_n has the required exponential behaviour, since

$|\Sigma_n| = 2^{\lfloor Rn \rfloor - 1}$; by definition $\frac{1}{\beta n} \log Z_n$ is equal to $F^1_{\beta,R,n}$. So Λ becomes:

$$\begin{aligned}
\Lambda(\beta) &= \lim_{n\to\infty} \frac{1}{n} \log E\left[\exp\left(\frac{\beta}{2} \sum_{j=1,\dots,n} \widetilde{x}_{s,j} - \frac{\beta n}{2}\right)\right] \\
&= -\frac{\beta}{2} + \lim_{n\to\infty} \frac{1}{n} \sum_{j=1}^{n} \log E\left[\exp\left(\frac{\beta}{2}\widetilde{x}_{1,j}\right)\right] \\
&= -\frac{\beta}{2} + \log\cosh\frac{\beta}{2}
\end{aligned}$$

which is convex, decreasing, smooth and finite for all $\beta \in \mathbb{R}$.

Having verified all the assumptions, we can now apply the theorem, obtaining a LDP for $F^1_{\beta,R,n}$. To get the rate functional we have to compute $\Lambda^*(a)$. By differentiation one finds that

$$a = \Lambda'(\beta) \Leftrightarrow \beta = \log\frac{1+a}{-a}$$

and hence

$$\begin{aligned}
\Lambda^*(a) &= a \log\frac{1+a}{-a} - \log\frac{1}{2(1+a)} \\
&= -a\log\frac{-a}{1/2} + (1+a)\log\frac{1+a}{1/2} \\
&= D_e\left(-a \,\middle\|\, \frac{1}{2}\right)
\end{aligned}$$

The rest follows with little computations:

$$a^*(R) = -\delta_{GV}(R)$$

$$\beta^*(R) = \begin{cases} \log\frac{1-\delta_{GV}(R)}{\delta_{GV}(R)} & \text{if} \quad R < 1 \\ \infty & \text{if} \quad R = 1 \end{cases}$$

To conclude the proof, note that the definition of R^* is such that $\beta < \beta^*(R) \Leftrightarrow R > R^*(\beta)$. ■

Remark 33 *About the regularity of $f(\cdot, R)$ at $\beta^* = \beta^*(R)$. From the definitions of a^* and β^* we have*

$$a^*\beta^* = \Lambda(\beta^*) + \Lambda^*(a^*) = \Lambda(\beta^*) + R\log 2$$

and

$$a^* = \Lambda'(\beta^*)$$

which gives us the continuity at $\beta = \beta^*(R)$ respectively of the map $\beta \mapsto \beta f(\beta, R)$, and of its derivative. It follows that the map $\beta \mapsto f(\beta, R)$ has a continuous derivative, is non-increasing and therefore always greater than or equal to $-\delta_{GV}(R)$.

Remark 34 *The rate functional $I^1_{\beta,R}(x)$ is convex (and right-continuous in $x = f(\beta,R)$) iff $R \leq R^*(\beta)$. In fact, if $R \leq R^*(\beta)$, then $f(\beta,R) = -\delta_{GV}(R)$ and*

$$\lim_{x \downarrow f(\beta,R)} I^1_{\beta,R}(x) = D_e\left(\delta_{GV}(R) \middle\| \frac{1}{2}\right) - R\log 2 = 0$$

and since the relative entropy is non-negative and strictly convex this can never happen when $f(\beta,R) \neq -\delta_{GV}(R)$.

Corollary 35 *$\hat{F}^0_{\beta,R,n}$ satisfies a LDP with the same rate functional as $F^1_{\beta,R,n}$.*

Proof. This is obvious, since the only difference with the previous case lies in the definition of $\Sigma_n := \{s \in S_n \setminus \{0\} : \pi^n_1(s) = 0\}$, which is one element smaller. ■

4.4.4 The analysis of $I_{\beta,R}(x)$

Having in mind proposition 29, we are interested in the rate:

$$I_{\beta,R}(x) = \inf_{s\in\mathbb{R}} \left(I^1_{\beta,R}(s) + I^*(s-x)\right)$$

By (19) the argument is infinite if $s < f(\beta,R)$, so we simply get

$$I_{\beta,R}(x) = \min\left\{\xi(\beta,R,x), \zeta(\beta,R,x)\right\} \tag{22}$$

where

$$\begin{aligned}
\zeta(\beta,R,x) &:= I^1_{\beta,R}(f(\beta,R)) + I^*(f(\beta,R) - x) \\
&= D_e(x - f(\beta,R) \,\|p) \qquad (23)\\
\xi(\beta,R,x) &:= \inf_{s > f(\beta,R)} \left[I^1_{\beta,R}(s) + I^*(s-x)\right] \\
&= -R\log 2 + \inf_{u < -f(\beta,R)} \xi^*(u,x) \qquad (24)
\end{aligned}$$

and

$$\xi^*(u,x) := D_e\left(u \middle\| \frac{1}{2}\right) + D_e(u + x \,\|p) \tag{25}$$

Remark 36 *Although $I^1_{\beta,R}(x)$ may be discontinuous on the set $\left\{x : I^1_{\beta,R}(x) < \infty\right\}$, $I_{\beta,R}(x)$ is always continuous on all $\{x : I_{\beta,R}(x) < \infty\}$.*

The following statement is a technical tool that we shall use often.

Lemma 37 *For any $\theta \in [0,1/2]$ and any $\beta > 0$, the equation*

$$-f(\beta,R) = \theta$$

has exactly one solution $(\beta, R) = (\beta, \varphi_\theta(\beta))$. *The map* φ_θ *is continuous, non-decreasing, hits* $R^* = R^*(\beta)$ *when*

$$\beta = \beta_\theta = \log\left(\frac{1}{\theta} - 1\right)$$

is constant for $\beta \geq \beta_\theta$, *and in fact is equal to*

$$\varphi_\theta(\beta) = \begin{cases} \frac{1}{\log 2}\left((1-2\theta)\frac{\beta}{2} - \log\cosh\frac{\beta}{2}\right) & \text{if} \quad \beta < \beta_\theta \\ D_2\left(\theta \,\middle\|\, \frac{1}{2}\right) & \text{if} \quad \beta \geq \beta_\theta \end{cases}$$

Proof. The formulae for β_θ and φ_θ are direct consequences of Equations (21) and (20) respectively. The other statements follow from an elementary study of φ_θ. ■

The following statement help us to understand the last assumption of Corollary 28.

Proposition 38 *The condition*

$$\inf_{x>0} I_{\beta,R}(x) = \inf_{x\geq 0} I_{\beta,R}(x) = I_{\beta,R}(0)$$

is satisfied if and only if R *is below the Channel capacity* $C(\beta)$, *defined by*

$$C(\beta) := \begin{cases} \frac{1}{\log 2}\left((1-2p)\frac{\beta}{2} - \log\cosh\frac{\beta}{2}\right) & \text{if} \quad \beta < \beta_{Trans} \\ D_2\left(p \| \frac{1}{2}\right) & \text{if} \quad \beta \geq \beta_{Trans} \end{cases}$$

where

$$\beta_{Trans} := \log\frac{1-p}{p}$$

Proof. We claim that both $\xi(\beta, R, x)$ and $\zeta(\beta, R, x)$ are convex in x and they attain their minimum when $x = p + f(\beta, R)$. Then recalling Equation (22), we only need to show that

$$p + f(\beta, R) \leq 0 \Leftrightarrow R \leq C(\beta)$$

but, by Lemma 37,

$$-f(\beta, R) \geq p \Leftrightarrow R \leq \varphi_p(\beta)$$

and the definition of $C(\beta)$ is exactly $\varphi_p(\beta)$. ■

Proof. We now come to the claim. The part concerning ζ easily follows from Equation (23) and the properties of D. Convexity of ξ in x comes from the same property of ξ^* and the fact that the domain of existence of the latter is a convex, compact set. ■

Proof. By the properties of D and Equation (24), the minimum of ξ^* for fixed u, is attained for $u + x = p$,

$$\min_{x\in\mathbb{R}} \xi^*(u, x) = \xi^*(u, p-u) \tag{26}$$

This also implies that the absolute minimum is at $u = 1/2$, $x = p - 1/2$. But since $-f(\beta, R) \leq \delta_{GV}(R) \leq 1/2$ (see Remark 33), the minimum of ξ^* on $\{(u, x) : u \leq -f(\beta, R)\}$ will be attained on the boundary

$$\inf_{x \in \mathbb{R}} \xi(\beta, R, x) = \min_{\substack{u \leq -f(\beta,R) \\ x \in \mathbb{R}}} \xi^*(u, x) = \min_{x \in \mathbb{R}} \xi^*(-f(\beta, R), x)$$

so that by (26), again ξ is minimum at $x = p + f(\beta, R)$. ■

4.4.5 Computation of $I_{\beta,R}(0)$

Finally, we only have to study the expression of $I_{\beta,R}(0)$ obtained above to compute the error exponent. From now on x will always be 0.

The study of $\xi(\beta, R, 0)$

Let

$$\delta_{Crit} := \frac{\sqrt{p}}{\sqrt{p} + \sqrt{1-p}}$$

then

$$\xi^*(u, 0) = \log\left(\frac{2}{1 + 2\sqrt{p(1-p)}}\right) + 2D_e(u||\delta_{Crit}) \tag{27}$$

where we have used the following lemma.

Lemma 39 *The following identity holds:*

$$D_e(u||q) + D_e(u||p)$$

$$= 2D_e\left(u \left\| \frac{\sqrt{pq}}{\sqrt{pq} + \sqrt{(1-p)(1-q)}}\right.\right) - 2\log\left(\sqrt{pq} + \sqrt{(1-p)(1-q)}\right)$$

Proof. *We just need to apply the definition of relative entropy:*

$$\begin{aligned} D_e(u||q) + D_e(u||p) &= u\log\frac{u}{q} + (1-u)\log\frac{1-u}{1-q} \\ &\quad + u\log\frac{u}{p} + (1-u)\log\frac{1-u}{1-p} \\ &= 2u\log u + 2(1-u)\log(1-u) \\ &\quad - 2u\log\sqrt{pq} - 2(1-u)\log\sqrt{(1-p)(1-q)} \end{aligned}$$

from which the thesis follows. ■

Looking to (27) we notice that the absolute minimum of $\xi^*(u, 0)$ is at $u = \delta_{Crit}$. Nevertheless:

$$\xi(\beta, R, 0) = \min_{u \leq -f(\beta,R)} \xi^*(u, 0) - R\log 2$$

so that

$$\xi(\beta, R, 0) = \xi^*(\min(\delta_{Crit}, -f(\beta, R)), 0) - R\log 2$$

The condition $-f(\beta, R) < \delta_{Crit}$ is well understood by Lemma 37, so now we compute ξ^* in the two cases.

By Equation (27),

$$\xi^*(\delta_{Crit}) = \log\left(\frac{2}{1+2\sqrt{p(1-p)}}\right)$$
$$\xi^*(-f(\beta, R)) = \log\left(\frac{2}{1+2\sqrt{p(1-p)}}\right) + 2D_e(-f(\beta, R)\,||\delta_{Crit})$$

anyway, notice that when $-f(\beta, R) = \delta_{GV}(R)$, by applying directly Equation (25), we have

$$\xi^*(\delta_{GV}(R)) = D_e(\delta_{GV}(R)\,||p) + R\log 2$$

Putting things together we have proved the following statement

Lemma 40 *Let*

$$R_{Crit}(\beta) := \begin{cases} \frac{1}{\log 2}\left((1-2\delta_{Crit})\frac{\beta}{2} - \log\cosh\frac{\beta}{2}\right) & \text{if} \quad \beta < \beta_{Crit} \\ D_2\left(\delta_{Crit}||\frac{1}{2}\right) & \text{if} \quad \beta \geq \beta_{Crit} \end{cases}$$

where

$$\beta_{Crit} := \frac{1}{2}\log\frac{1-p}{p}$$

Then $R_{Crit}(\beta) \leq C(\beta)$ *for all* β*, and*

$$-f(\beta, R) < \delta_{Crit} \Leftrightarrow R > R_{Crit}(\beta)$$

Moreover, we can distinguish three cases:

1. *If* $R \leq R_{Crit}(\beta)$ *we have*

$$\xi(\beta, R, 0) = \log\left(\frac{2}{1+2\sqrt{p(1-p)}}\right) - R\log 2$$

2. *If* $R^*(\beta) \leq R_{Crit}(\beta) < R$ *we have*

$$\xi(\beta, R, 0) = \log\left(\frac{2}{1+2\sqrt{p(1-p)}}\right) - R\log 2 + 2D_e(-f(\beta, R)\,||\delta_{Crit})$$

3. *If* $R_{Crit}(\beta) < R \leq R^*(\beta)$ *we have*

$$\xi(\beta, R, 0) = D_e(\delta_{GV}(R)\,||p)$$

The study of $\zeta(\beta, R, 0)$

We recall from (23) that

$$\zeta(\beta, R, 0) := D_e(-f(\beta, R)||p)$$

and only notice that, when $R \leq R^*(\beta)$, this expression becomes

$$\zeta(\beta, R, 0) = D_e(\delta_{GV}(R)||p)$$

4.4.6 The four cases breakdown

Let us now compute $I_{\beta,R}(0)$ in the possible situations.

Case $R_{Crit}(\beta) \leq R < R^*(\beta)$

Here the expressions for $\xi(\beta, R, 0)$ and $\zeta(\beta, R, 0)$ coincide, so

$$I_{\beta,R}(0) = D_e(\delta_{GV}(R)||p)$$

Case $R < \min\{R^*(\beta), R_{Crit}(\beta)\}$

Since $u = \delta_{Crit}$ minimizes $\xi^*(u, 0)$,

$$D_e(\delta_{Crit}||p) + D_e\left(\delta_{Crit}\left\|\frac{1}{2}\right.\right) \leq D_e(\delta_{GV}(R)||p) + D_e\left(\delta_{GV}(R)\left\|\frac{1}{2}\right.\right)$$

so that

$$D_e(\delta_{Crit}||p) + D_e\left(\delta_{Crit}\left\|\frac{1}{2}\right.\right) - R\log 2 \leq D_e(\delta_{GV}(R)||p)$$

and

$$I_{\beta,R}(0) = \log\left(\frac{2}{1+2\sqrt{p(1-p)}}\right) - R\log 2$$

Case $\max\{R^*(\beta), R_{Crit}(\beta)\} \leq R \leq R_0(\beta)$

By Remark 33 we know that

$$-f(\beta, R) \leq \delta_{GV}(R)$$

so that, by the monotonicity of $D_e\left(\cdot\left\|\frac{1}{2}\right.\right)$ on $\left[0, \frac{1}{2}\right]$,

$$D_e\left(-f(\beta, R)\left\|\frac{1}{2}\right.\right) \geq D_e\left(\delta_{GV}(R)\left\|\frac{1}{2}\right.\right)$$

and finally

$$D_e(-f(\beta, R)||p) + D_e\left(-f(\beta, R)\left\|\frac{1}{2}\right.\right) - R\log 2 \geq D_e(-f(\beta, R)||p)$$

yielding

$$I_{\beta,R}(0) = D_e\left(\frac{1}{\beta}\left[\frac{\beta}{2} - \log\cosh\frac{\beta}{2} - R\log 2\right]\middle\| p\right)$$

Case $R^*(\beta) \leq R < R_{Crit}(\beta)$

Here we will see that there is a separating curve with no explicit expression between two subregions. Recall that

$$R^*(\beta) \leq R_{Crit}(\beta) \Longleftrightarrow \beta \in [0, \beta_{Crit}]$$

with equality holding only at the border. We claim that the set of points (β, R) such that $R^*(\beta) \leq R < R_{Crit}(\beta)$ and

$$\log\left(\frac{2}{1+2\sqrt{p(1-p)}}\right) - R\log 2 - D_e\left(\frac{1}{\beta}\left[\frac{\beta}{2} - \log\cosh\frac{\beta}{2} - R\log 2\right] \middle\| p\right) = 0 \tag{28}$$

is the graph of a function of β on $[0, \beta_{Crit}]$. We shall need to extend the latter to a map $R_s(\beta)$ defined on all $[0, \infty)$, and we set arbitrarily $R_s(\beta) = 1$ if $\beta > \beta_{Crit}$. We clearly have, for $0 \leq \beta \leq \beta_{Crit}$,

$$I_{\beta,R}(0) = \begin{cases} \log\left(\frac{2}{1+2\sqrt{p(1-p)}}\right) - R\log 2 & \text{if } R^*(\beta) \leq R < R_s(\beta) \\ D_e\left(\frac{1}{\beta}\left[\frac{\beta}{2} - \log\cosh\frac{\beta}{2} - R\log 2\right] \middle\| p\right) & \text{if } R_s(\beta) \leq R \leq R_{Crit}(\beta) \end{cases}$$

The existence of such a separating curve within the region is essentially due to Rolle theorem. In order to prove that the curve $R_s(\beta)$ is a function in β (and not a generic curve), observe that differentiating the left-hand side of (28) with respect to R, we get

$$-\log 2 + \frac{1}{\beta}\log\frac{g}{1-g}\log\left(2\frac{1-p}{p}\right)$$

which is always negative in the given domain.

4.4.7 The final result

We proved the following.

Theorem 41 *If F_n^*, $F_{\beta,R,n}^1$, $\hat{F}_{\beta,R,n}^0$ are defined as in the first section of the current paragraph and R is less than the Channel capacity $C(\beta)$:*

$$C(\beta) := \begin{cases} \dfrac{(1-2p)\frac{\beta}{2} - \log\cosh\frac{\beta}{2}}{\log 2} & \text{if} \quad \beta < \beta_{Trans} \\ D_2\left(p \middle\| \frac{1}{2}\right) & \text{if} \quad \beta \geq \beta_{Trans} \end{cases}$$

where

$$\beta_{Trans} := \log\frac{1-p}{p}$$

then they satisfy hypotheses 1), 3) and 4) of Corollary 28. Moreover we have

$$I_{\beta,R}(0) = \begin{cases} \log\left(\frac{2}{1+2\sqrt{p(1-p)}}\right) - R\log 2 & \text{if} \quad (\beta, R) \in D_1 \\ D_e\left(\delta_{GV}(R) \| p\right) & \text{if} \quad (\beta, R) \in D_2 \\ D_e\left(\frac{1}{\beta}\left[\frac{\beta}{2} - \log\cosh\frac{\beta}{2} - R\log 2\right] \middle\| p\right) & \text{if} \quad (\beta, R) \in D_3 \end{cases}$$

where we defined

$$D_1 := \left\{(\beta, R) \in \mathbb{R}^2, 0 \leq R \leq \min\left\{R_s(\beta), R_{Crit}(\beta)\right\}\right\}$$

$$D_2 := \left\{(\beta, R) \in \mathbb{R}^2, R_{Crit}(\beta) < R \leq \min\left\{R^*(\beta), C(\beta)\right\}\right\}$$

$$D_3 := \left\{(\beta, R) \in \mathbb{R}^2, \max\left\{R_s(\beta), R^*(\beta)\right\} < R \leq C(\beta)\right\}$$

Finally, hypothesis 2) is satisfied if $R \leq R^*(\beta)$ *(see remark 34).*

To comment the previous results, we have that for $R \leq R^*(\beta)$ the error exponent $E_{err}(\beta, R)$ is equal to $I_{\beta,R}(0)$. If not, then $I_{\beta,R}(0)$.is only a lower bound. An analysis of these results, together with a comparison to the existing coding literature, will take place in the next chapter.

4.5 Appendix

4.5.1 Appendix A: binomial computations

In this appendix, given a positive integer N and a number $p \in [0,1]$, we shall denote by $B_{N,p}$ a binomial random variable with parameters N and p. Given three sequences $N_n \to \infty$ (of positive integers), $p_n \to 0$ (of numbers in $[0,1]$) and $x_n \geq 0$, we are interested in the exponential asymptotic behavior of $P(B_{N_n,p_n} > x_n)$ and $P(B_{N_n,p_n} < x_n)$. We assume that the following limits exist:

$$\alpha = \lim_{n\to\infty} \frac{1}{n} \log N_n$$

$$\beta = -\lim_{n\to\infty} \frac{1}{n} \log p_n$$

$$\gamma = \lim_{n\to\infty} \frac{1}{n} \log x_n$$

We write $a_n \sim b_n$ if $\frac{a_n}{b_n} \to 1$ as $n \to \infty$. Independently of β, we have

$$\gamma > \alpha \Rightarrow \lim_{n\to\infty} \frac{1}{n} \log P(B_{N_n,p_n} > x_n) = -\infty \tag{29}$$

Moreover,

Lemma 42 *i) If* $\gamma < 0$ *and* $\alpha < \beta$ *then*

$$\lim_{n\to\infty} \frac{1}{n} \log P(B_{N_n,p_n} > x_n) = \alpha - \beta \tag{30}$$

ii) If $\gamma < 0$ *and* $\alpha > \beta$ *then*

$$\lim_{n\to\infty} \frac{1}{n} \log P(B_{N_n,p_n} > x_n) = 0 \tag{31}$$

Proof. For $\gamma < 0$ and n so large that $x_n < 1$, we have

$$\begin{aligned} P\left(B_{N_n,p_n} > x_n\right) &= 1 - (1-p_n)^{N_n} = 1 - e^{N_n \log(1-p_n)} \\ &= 1 - e^{-N_n(p_n + o(p_n))} = 1 - e^{-N_n p_n + o(N_n p_n)} \end{aligned}$$

This easily implies the result. ■

Lemma 43 *If* $(\alpha - \beta)^+ < \gamma < \alpha$ *then*

$$\frac{1}{n} \log P\left(B_{N_n,p_n} > x_n\right) \sim (\alpha - \beta - \gamma)\, e^{n\gamma} \to -\infty \tag{32}$$

Proof. We give at least this proof in all the details, although elementary.
Step 1. We analyze

$$P\left(B_{N_n,p_n} = k_n\right) = \binom{N_n}{k_n} p_n^{k_n} (1-p_n)^{N_n - k_n}$$

where $k_n = \lceil x_n \rceil$ denotes the smallest integer strictly greater than x_n. Let us denote by ε_n every sequence such that $\varepsilon_n \to 0$ as $n \to \infty$ and apply to them Landau rules. Recall that

$$N! = \left(\frac{N}{e}\right)^N \sqrt{2\pi N}\,(1 + \varepsilon_N)$$

Since $k_n \to \infty$ and $N_n - k_n \to \infty$, we have

$$\begin{aligned} P\left(B_{N_n,p_n} = k_n\right) = \sqrt{\frac{N_n}{2\pi k_n \left(N_n - k_n\right)}} &\left(\frac{N_n p_n}{k_n}\right)^{k_n} \\ &\cdot \left(\frac{N_n (1-p_n)}{N_n - k_n}\right)^{N_n - k_n} (1 + \varepsilon_n) \end{aligned}$$

Moreover,

$$\lim_{n\to\infty} \frac{1}{n} \log \sqrt{\frac{N_n}{2\pi k_n \left(N_n - k_n\right)}} = -\frac{\gamma}{2}, \tag{33}$$

$$\begin{aligned} &\log \left(\frac{N_n (1-p_n)}{N_n - k_n}\right)^{N_n - k_n} \\ &= \left(N_n - k_n\right) \left(\frac{N_n (1-p_n)}{N_n - k_n} - 1\right) \frac{\log\left(1 + \left(\frac{N_n(1-p_n)}{N_n - k_n} - 1\right)\right)}{\left(\frac{N_n(1-p_n)}{N_n - k_n} - 1\right)} \\ &= \left(-N_n p_n + k_n\right)(1 + \varepsilon_n) \end{aligned} \tag{34}$$

hence, noting that $N_n p_n = o(k_n)$,

$$\frac{1}{n} \log \left(\frac{N_n (1-p_n)}{N_n - k_n}\right)^{N_n - k_n} = \frac{1}{n} k_n (1 + \varepsilon_n)$$

and finally

$$\begin{aligned}\frac{1}{n}\log\left(\frac{N_n p_n}{k_n}\right)^{k_n} &= k_n \frac{1}{n}\log\left(\frac{N_n p_n}{k_n}\right)\\ &= k_n\left(\alpha-\beta-\gamma+\varepsilon_n\right)\end{aligned} \tag{35}$$

Therefore we have

$$\frac{1}{n}\log P\left(B_{N_n,p_n}=k_n\right)=\frac{\gamma}{2}+\varepsilon_n+k_n\left(\alpha-\beta-\gamma+\varepsilon_n\right),$$

so

$$\frac{1}{n}\log P\left(B_{N_n,p_n}=k_n\right)\sim\left(\alpha-\beta-\gamma\right)e^{n\gamma}\to-\infty$$

Step 2. We now show that $P\left(B_{N_n,p_n}=k_n\right)$ gives us the correct asymptotic of $P\left(B_{N_n,p_n}>x_n\right)$. Given a positive integer k, let $C_{k,n}=\frac{N_n-k}{k+1}\cdot\frac{p_n}{1-p_n}$, so that $P\left(B_{N_n,p_n}=k+1\right)=C_{k,n}P(B_{N_n,p_n}=k)$ and notice that $C_{k,n}$ is decreasing in k. Hence

$$k\geq k_n\Rightarrow P(B_{N_n,p_n}=k)\leq\left(C_{k_n,n}\right)^{k-k_n}P\left(B_{N_n,p_n}=k_n\right)$$

$$\begin{aligned}P(B_{N_n,p_n}>x_n) &\leq P\left(B_{N_n,p_n}=k_n\right)\sum_{k=k_n}^{N_n}\left(C_{k_n,n}\right)^{k-k_n}\\ &\leq P\left(B_{N_n,p_n}=k_n\right)\frac{1}{1-C_{k_n,n}}\end{aligned}$$

On the other hand of course $P(B_{N_n,p_n}>x_n)\geq P\left(B_{N_n,p_n}=k_n\right)$. The final result follows from the fact that

$$C_{k_n,n}=\frac{N_n-k_n}{k_n+1}\cdot\frac{p_n}{1-p_n}\to 0$$

The proof is complete. ■

When $\alpha>\beta$ and $0<\gamma<\alpha-\beta$, both $E\left[B_{N_n,p_n}\right]=N_n p_n$ and x_n diverge to $+\infty$, but x_n is much smaller than $E\left[B_{N_n,p_n}\right]$, so that for n large enough $P\left(B_{N_n,p_n}>x_n\right)\geq\frac{1}{2}$. This proves the following:

Lemma 44 *If $\alpha>\beta$ and $0<\gamma<\alpha-\beta$ then*

$$\lim_{n\to\infty}\frac{1}{n}\log P\left(B_{N_n,p_n}>x_n\right)=0 \tag{36}$$

Finally, arguing as above, we have (notice that $C\cdot N_n$ is much larger than $E\left[B_{N_n,p_n}\right]$):

Lemma 45 *For every $C\in(0,1)$*

$$\lim_{n\to\infty}P\left(B_{N_n,p_n}>C\cdot N_n\right)=0 \tag{37}$$

Now we turn our attention to $P\left(B_{N_n,p_n} < x_n\right)$. Note that

$$\lim_{n\to\infty} \frac{1}{n} \log P\left(B_{N_n,p_n} > x_n\right) \neq 0 \quad \Longrightarrow \quad \lim_{n\to\infty} \frac{1}{n} \log P\left(B_{N_n,p_n} < x_n\right) = 0$$

There are only two cases left. In both it will turn out that the rate is $-\infty$.

Lemma 46 *If $\gamma < 0$ and $\alpha > \beta$ then*

$$\lim_{n\to\infty} \frac{1}{n} \log P\left(B_{N_n,p_n} < x_n\right) = -\infty \tag{38}$$

Proof. If n is so large that $x_n < 1$, we have

$$\frac{1}{n} \log P\left(B_{N_n,p_n} < x_n\right) = \frac{1}{n} \log\left(1-p_n\right)^{N_n} = \frac{N_n}{n} \log\left(1-p_n\right) \leq -\frac{N_n p_n}{n}$$

This easily implies the result. ■

Lemma 47 *If $\alpha > \beta$ and $0 < \gamma < \alpha - \beta$ then*

$$\frac{1}{n} \log P\left(B_{N_n,p_n} < x_n\right) \sim -\frac{1}{n} N_n p_n \to -\infty \tag{39}$$

Proof. The proof is completely analogous to that of (32).
Step 1. We analyze

$$P\left(B_{N_n,p_n} = k_n\right) = \binom{N_n}{k_n} p_n^{k_n} \left(1-p_n\right)^{N_n-k_n}$$

where $k_n = \lfloor x_n \rfloor$ denotes the greatest integer strictly smaller than x_n. Equations (33), (34) and (35) are still valid, but here $k_n = o(N_n p_n)$, so that

$$\frac{1}{n} \log\left(\frac{N_n\left(1-p_n\right)}{N_n - k_n}\right)^{N_n-k_n} = -\frac{1}{n} N_n p_n \left(1+\varepsilon_n\right),$$

and therefore

$$\frac{1}{n} \log P\left(B_{N_n,p_n} = k_n\right) = \frac{\gamma}{2} + \varepsilon_n - \frac{1}{n} N_n p_n \left(1 + \varepsilon_n - n\frac{k_n}{N_n p_n}\left(\alpha - \beta - \gamma + \varepsilon_n\right)\right),$$

so

$$\frac{1}{n} \log P\left(B_{N_n,p_n} = k_n\right) \sim -\frac{1}{n} N_n p_n \to -\infty$$

Step 2. Given a positive integer k, let $C_{k,n} = \frac{k}{N_n - k + 1} \cdot \frac{1-p_n}{p_n}$, so that $P\left(B_{N_n,p_n} = k-1\right) = C_{k,n} P(B_{N_n,p_n} = k)$; then $C_{k,n}$ is increasing in k. Hence

$$k \leq k_n \Rightarrow P(B_{N_n,p_n} = k) \leq \left(C_{k_n,n}\right)^{k-k_n} P\left(B_{N_n,p_n} = k_n\right)$$

$$\begin{aligned} P(B_{N_n,p_n} < x_n) &\leq P\left(B_{N_n,p_n} = k_n\right) \sum_{k=0}^{k_n} \left(C_{k_n,n}\right)^{k-k_n} \\ &\leq P\left(B_{N_n,p_n} = k_n\right) \cdot \frac{1}{1 - C_{k_n,n}} \end{aligned}$$

Again $P(B_{N_n,p_n} < x_n) \geq P(B_{N_n,p_n} = k_n)$, so that the thesis follows from

$$C_{k_n,n} = \frac{k_n}{N_n - k_n + 1} \cdot \frac{1 - p_n}{p_n} \to 0$$

■

4.5.2 Appendix B: the LDP-sum lemma

For sake of completeness and with the aim of showing the easiness of rate function manipulations to the unexperienced but interested reader, we prove here a simple lemma which derives the rate functional of a sum of two independent sequences of random variables from their individual rate functionals. A reference to general large deviation properties could be [10].

Lemma 48 *Let $a_1 = \{a_{1,n}\}_{n\in\mathbb{N}}$ and $a_2 = \{a_{2,n}\}_{n\in\mathbb{N}}$ be two mutually independent sequences of real random variables satisfying large deviation principles with speed n and rate functionals I_1 and I_2 respectively. Let a_3 be defined as $a_{3,n} := a_{1,n} + a_{2,n}$ for all n. Then a_3 satisfies a LDP with speed n and rate functional $I_3(x) = \inf_{r\in\mathbb{R}} \{I_1(x-r) + I_2(r)\}$.*

Proof. In this proof we use only the definition of LDP. Let us first build the product random sequence $a_p = \{a_{p,n}\}_{n\in\mathbb{N}}$ as follows:

$$a_{p,n} = (a_{1,n}, a_{2,n}) \in \mathbb{R}^2$$
$$P(a_{p,n} \in A_1 \times A_2) = P(a_{1,n} \in A_1)\, P(a_{2,n} \in A_2)$$

Where A_1 and A_2 are Borel sets. This definition implies the following facts:

$$-\frac{1}{n}\log P(a_{p,n} \in A_1 \times A_2) = -\frac{1}{n}\log P(a_{1,n} \in A_1) - \frac{1}{n}\log P(a_{2,n} \in A_2)$$

$$\begin{aligned}\liminf_{n\to\infty} \frac{1}{n}\log P(a_{p,n} \in A_1 \times A_2) &\geq \liminf_{n\to\infty} \frac{1}{n}\log P(a_{1,n} \in A_1) \\ &\quad + \liminf_{n\to\infty} \frac{1}{n}\log P(a_{2,n} \in A_2) \\ &\geq -\inf_{t\in\overset{\circ}{A}_1} I_1(t) - \inf_{s\in\overset{\circ}{A}_2} I_2(s) \\ &= -\inf_{\{t,s\}\in(A_1\times A_2)^\circ} \{I_1(t) + I_2(s)\}\end{aligned}$$

and

$$\begin{aligned}\limsup_{n\to\infty} \frac{1}{n}\log P(a_{p,n} \in A_1 \times A_2) &\leq \limsup_{n\to\infty} \frac{1}{n}\log P(a_{1,n} \in A_1) \\ &\quad + \limsup_{n\to\infty} \frac{1}{n}\log P(a_{2,n} \in A_2) \\ &\leq -\inf_{t\in\overline{A}_1} I_1(t) - \inf_{s\in\overline{A}_2} I_2(s) \\ &= -\inf_{\{t,s\}\in\overline{(A_1\times A_2)}} \{I_1(t) + I_2(s)\}\end{aligned}$$

showing that a_p satisfies a LDP with speed n and rate functional $I_p(x_1, x_2) = I_1(x_1) + I_2(x_2)$ in $\mathbb{R}^2$. Now, for every couple of reals (s, u) with $s \leq u$ let us consider the set $A_{s,u} := \{(x_1, x_2) \in \mathbb{R}^2 : s \leq x_1 + x_2 \leq u\}$. Notice that

$$-\frac{1}{n}\log P(a_{p,n} \in A_{s,u}) = -\frac{1}{n}\log P(s \leq a_{3,n} \leq u)$$

Now by the LDP of a_p we have

$$\begin{aligned}\limsup_{n\to\infty}\frac{1}{n}\log P(a_{p,n} \in A_{s,u}) &\leq -\inf_{(s,t)\in\overline{A_{s,u}}} I_p(s,t) \\ &= -\inf_{r\in\mathbb{R}, s\leq v\leq u} I_p(v-r, r) \\ &= -\inf_{s\leq v\leq u}\inf_{r\in\mathbb{R}}\{I_1(v-r) + I_2(r)\}\end{aligned}$$

In the same way we also have

$$\begin{aligned}\liminf_{n\to\infty}\frac{1}{n}\log P(a_{p,n} \in A_{s,u}) &\geq -\inf_{(s,t)\in\overset{\circ}{A}_{s,u}} I_p(s,t) \\ &= -\inf_{r\in\mathbb{R}, s<v<u} I_p(v-r, r) \\ &= -\inf_{s<v<u}\inf_{r\in\mathbb{R}}\{I_1(v-r) + I_2(r)\}\end{aligned}$$

Since the intervals $[u, s]$ with $s \leq u$ are a base of the Borel σ-algebra, we obtain the thesis. ■

4.5.3 Appendix C: the separation bound

Proposition 49 *Let X, Y and Z be real independent random variables on a space $(\Omega, \mathcal{F}, \mathbb{P})$ with expectation $\mathbb{E}$. Then*

$$\mathbb{P}(X > Y, X > Z) \geq \mathbb{P}(X > Y)\mathbb{P}(X > Z)$$

Proof. By the independence of $\mathbb{E}[1_{X>Y}|Y]$ and $\mathbb{E}[1_{X>Z}|Z]$, we get

$$\begin{aligned}\mathbb{P}(X > Y)\mathbb{P}(X > Z) &= \mathbb{E}[\mathbb{E}[1_{X>Y}|Y] \cdot \mathbb{E}[1_{X>Z}|Z]] \\ &\leq \mathbb{E}[\mathbb{E}[1_{X>Y}|Y]; Y > Z] + \mathbb{E}[\mathbb{E}[1_{X>Z}|Z]; Y \leq Z]\end{aligned}$$

moreover, $\mathbb{P}$-a.s. on $\{Y > Z\}$,

$$\mathbb{E}[1_{X>Y}|Y] = \mathbb{E}[1_{X>Y}|Y, Z] = \mathbb{E}[1_{X>Y}1_{X>Z}|Y, Z]$$

and similarly, $\mathbb{P}$-a.s. on $\{Y \leq Z\}$,

$$\mathbb{E}[1_{X>Z}|Z] = \mathbb{E}[1_{X>Y}|Y, Z] = \mathbb{E}[1_{X>Y}1_{X>Z}|Y, Z]$$

The proof is completed by substituting these identities in the above inequality. ■

Proof. (alternative) Let us observe that

$$\mathbb{P}(X > Y, X > Z) = \mathbb{E}\left\{\mathbb{P}(\,X > Y|\,X)\mathbb{P}(\,X > Z|\,X)\right\}$$

because of conditional independence of $\{X > Y\}$ and $\{X > Z\}$ given X. Now the two functions $\mathbb{P}(\,X > Y|\,X)$ and $\mathbb{P}(\,X > Z|\,X)$ are both monotone non-increasing in X, so they are positively correlated. In such a case

$$\begin{aligned}\mathbb{E}\left\{\mathbb{P}(\,X > Y|\,X)\mathbb{P}(\,X > Z|\,X)\right\} &\geq \mathbb{E}\left\{\mathbb{P}(\,X > Y|\,X)\right\}\mathbb{E}\left\{\mathbb{P}(\,X > Z|\,X)\right\} \\ &= \mathbb{P}(X > Y)\mathbb{P}(X > Z)\end{aligned}$$

■

5 Conclusions

In this thesis we have studied the Shannon problem, which we defined as a channel coding problem with a memoriless binary symmetric channel of parameter p and a random code picked from the Shannon random ensemble $SRE\left(R,n\right)$, in the limit $n\to\infty$.

We adopted an approach from the point of view of statistical mechanics pioneered by Nicolas Sourlas, and we applied large deviations techniques in order to compute the relevant asymptotic quantities of the problem, namely the capacity function $C\left(\beta\right)$ and the error exponent $E_{err}\left(\beta,R\right)$.

The Sourlas approach introduces an extra parameter β, the inverse temperature of the system, which parametrises a general decoding strategy. When $\beta=\beta_{Trans}:=\log\frac{1-p}{p}$ we recover the actual bitwise MAP decoding case, while the limit $\beta\to\infty$ gives the wordwise MAP decoding. In order to compare with the coding literature, let us observe that for "$\beta=\infty$" the results are known through various classic techniques (see for instance [5]). The case $\beta=\beta_{Trans}$ is known, too, but with less detail and usually through indirect techniques. The intermediate range $\beta_{Trans}<\beta<\infty$ is new, although heuristically it interpolates two known cases. The outer range $\beta<\beta_{Trans}$ is completely new.

The values for capacity and error exponent we recover are tight in the subset of the (R,β) phase diagram defined by $R\leq R^{*}(\beta)$ (we will resume the definition of $R^{*}(\beta)$ a little further into this conclusion), while they are only lower bounds for $1\geq R>R^{*}(\beta)$. We conjecture the results to be tight also in this case, but to the moment we have no rigorous proof.

Our large deviation analysis followed a general-to-particular approach: we developed a rather general theorem for the LDP of sequences of "partition function-like" quantities. We then reduced its scope of application to "empirical measures" induced by independent and identically distributed random variables, obtaining a much lighter formulation which we consequently applied to the Shannon problem. Although a problem-adapted theorem might have been much simpler to prove, our strategy enabled us to perceive which conditions are truly necessary to the interesting large deviation behaviour of the a posteriori measure induced by the code and the noise. Such a choice aligns with our intention of developing a bottom-up approach to coding theory, and not only to a particular problem or a classic setup. Nevertheless, we picked one of such problems to develop our approach.

5.1 Capacity

Let us consider a concept of capacity of the full system encoder/channel/decoder, defined as the supremum of those rates for which our system is error-free in the limit as $n\to\infty$. The capacity function we obtain for the Shannon problem is:

$$C\left(\beta\right):=\begin{cases}\dfrac{\left(1-2p\right)\frac{\beta}{2}-\log\cosh\frac{\beta}{2}}{\log 2} & \text{if}\quad\beta<\beta_{Trans}\\ D_{2}\left(p||\frac{1}{2}\right) & \text{if}\quad\beta\geq\beta_{Trans}\end{cases}$$

where:

$$\beta_{Trans} := \log \frac{1-p}{p}$$

It coincides with the well-known channel capacity obtained by Shannon for all $\beta \geq \beta_{Trans}$, where it depends only on the channel (the value of p, here). On the contrary it is β-dependent for $\beta < \beta_{Trans}$. In this range we have $C(\beta) \leq C(\beta_{Trans})$, as shown in Figure 2. Knowledge of the capacity in this range may be important in the case of a wrong estimate of p, see below.

5.2 Error exponent

Our candidate for the error exponent

$$E_{err}(\beta, R) := \lim_{n\to\infty} -\frac{1}{n} \log p_n^{err}$$

is $I_{\beta,R}(0)$, whose (complicate) definition is resumed below. We actually proved equality for the most interesting values of the parameters,

$$\lim_{n\to\infty} -\frac{1}{n} \log p_n^{err} = I_{\beta,R}(0) \quad \text{if } R \leq R^*(\beta)$$

and we have an upper bound on the error probability otherwise, which translates into a lower bound for the error exponent

$$\underline{\lim}_{n\to\infty} -\frac{1}{n} \log p_n^{err} \geq I_{\beta,R}(0) \quad \text{if } R > R^*(\beta)$$

The complete description of $I_{\beta,R}(0)$ follows:

$$I_{\beta,R}(0) = \begin{cases} \log\left(\dfrac{2}{1+2\sqrt{p(1-p)}}\right) - R\log 2 & \text{if} \quad (\beta,R) \in D_1 \\ D_e\left(\delta_{GV}(R)\,||p\right) & \text{if} \quad (\beta,R) \in D_2 \\ D_e\left(\frac{1}{\beta}\left[\frac{\beta}{2} - \log\cosh\frac{\beta}{2} - R\log 2\right] \middle\| p\right) & \text{if} \quad (\beta,R) \in D_3 \end{cases}$$

where we defined

$$D_1 := \left\{(\beta,R) \in \mathbb{R}^2, 0 \leq R \leq \min\left\{R_s(\beta), R_{Crit}(\beta)\right\}\right\}$$

$$D_2 := \left\{(\beta,R) \in \mathbb{R}^2, R_{Crit}(\beta) < R \leq \min\left\{R^*(\beta), C(\beta)\right\}\right\}$$

$$D_3 := \left\{(\beta,R) \in \mathbb{R}^2, \max\left\{R_s(\beta), R^*(\beta)\right\} < R \leq C(\beta)\right\}$$

and

$$R_{Crit}(\beta) := \begin{cases} \dfrac{(1-2\delta_{Crit})\frac{\beta}{2} - \log\cosh\frac{\beta}{2}}{\log 2} & \text{if} \quad \beta < \beta_{Crit} \\ D_2\left(\delta_{Crit}||\frac{1}{2}\right) & \text{if} \quad \beta \geq \beta_{Crit} \end{cases}$$

$$\beta_{Crit} := \frac{1}{2}\log\frac{1-p}{p}$$

$$\delta_{Crit} := \frac{\sqrt{p}}{\sqrt{p} + \sqrt{1-p}}$$

$$R^*(\beta) := D_2\left(\frac{1}{e^\beta + 1}\middle\|\frac{1}{2}\right) = \frac{\frac{\beta}{2}\tanh\frac{\beta}{2} - \log\cosh\frac{\beta}{2}}{\log 2}$$

while $R_s(\beta)$ is defined for $0 \leq \beta \leq \beta_{Crit}$ as the unique root of the equation

$$\log\left(\frac{2}{1 + 2\sqrt{p(1-p)}}\right) - R\log 2 - D_e\left(\frac{1}{\beta}\left[\frac{\beta}{2} - \log\cosh\frac{\beta}{2} - R\log 2\right]\middle\|p\right) = 0$$

within the domain $R_{Crit}(\beta) \geq R \geq R^*(\beta)$. For $\beta > \beta_{Crit}$ instead, $R_s(\beta)$ is arbitrarily set to 1.

As for the capacity, for $\beta \geq \beta_{Trans}$ the error exponent does not depend on β and is equal to the classic Shannon results (but without the refinement known as "expurgation", see [11]). Notice that $\beta \geq \beta_{Trans}$ corresponds to sections of the regions D_1 nd D_2, excluding D_3.

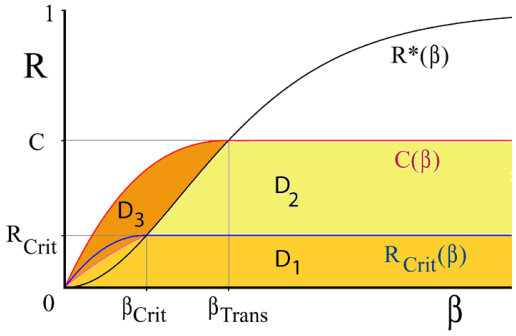

Figure 2: the (R, β) phase diagram

5.3 Towards a choice of β

The approach developed in this thesis is aimed at offering a more structured base for the analysis of random coding problems. Here is an instance of an operative suggestion coming from the above results.

As Prof. H. Loeliger (ETH Zurich) recently pointed out to us, in case we want to perform bitwise MAP decoding and we do not now precisely the value of p (the "flipping probability" of the binary simmetric channel), we can underestimate it safely when applying the β-parametrised decoding technique, according to our previous results. Indeed, the only way p enters into the computation of the a posteriori probability is through β_{Trans}, and for all $\beta > \beta_{Trans}$ the capacity and the error exponent are equal to the maximum value. Since β_{Trans} is monotone decreasing in p, we got the thesis. Operatively, this translates into

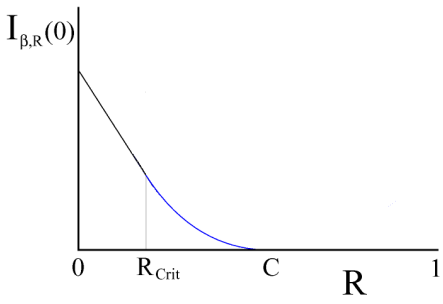

(a) $\beta \geq \beta_{Trans}$

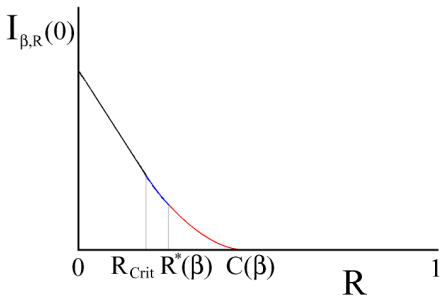

(b) $\beta_{Crit} < \beta = 1.4 < \beta_{Trans}$

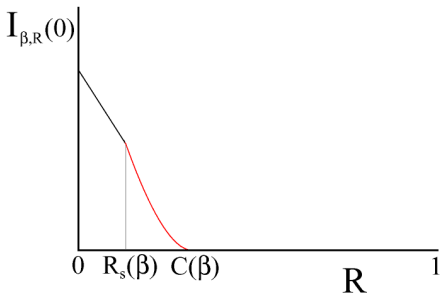

(c) $\beta = 0.6 \leq \beta_{Crit}$

Figure 3: set of plots of $I_{\beta,R}(0)$ with $p = 0.1$: in this case $\beta_{Trans} \simeq 2.2$, $\beta_{Crit} \simeq 1.1$ and for β high $R_{Crit} \simeq 0.189$ and $C \simeq 0.531$

raising to a positive exponent the a posteriori probabilities. Renormalisation of $e^{-\beta H}$ by the common denominator Z, although formally necessary, is of no consequence for the decoding algorithm.

5.4 Open problems

Close lines of development include first of all a tight result for capacity and error exponent in the region $R > R^*(\beta)$, intendendly according to our aforementioned conjecture. This would add necessity to the robustness argument of the underestimation of p.

Moreover, we would be interested into adding a linear constraint to our random code model, as most of technologically implemented codes are so, and expanding the channel model to a more general linear additive channel.

One of our far goal would instead be to decouple the large deviation behaviours of the code and the channel, so to be able to compute the error exponent not for the ensemble of codes, but for an optimal subgroup.

References

[1] A. Barg, D. Forney, *Random codes: Minimal Distances and Error Exponents*, IEEE Trans. on Information Theory 48 (2002), n. 9, 2568-2573.

[2] C. Berrou, A. Glavieux, and P. Thitimajshima, *Near Shannon Limit Error-Correcting Coding and Decoding: Turbo Codes*, Proc.Int. Conf. Comm. 1993, 1064-1070

[3] A. Bovier, I. Kurkova, *Rigorous results on some simple spin glass models*, Markov Proc. Rel. Fields 9 (2003), 209-242.

[4] T. M. Cover, J. A. Thomas, *Elements of Information Theory*, Wiley, New York 1991.

[5] I. Csiszar, J. Korner, *Information Theory: Coding Theorems for Discrete Memoryless Systems*, Akademiai Kiado 1997

[6] A. Dembo, O. Zeitouni, *Large Deviations Techniques and Applications*, Springer-Verlag, New York, 1998.

[7] B. Derrida, *Random-energy model: an exactly solvable model of disordered systems*, Phys. Review B 24 (1981), n. 5, 2613-2626.

[8] J.-D. Deuschel, D. W. Stroock, *Large Deviations*, Academic Press, Inc., Boston, MA, 1989.

[9] T.C. Dorlas and J.R. Wedagedera, *Large Deviations and the Random Energy Model*, Internat. J. Modern Phys. B 15 (2001), n. 1, 1-15.

[10] R. S. Ellis, Entropy, *Large Deviations and Statistical Mechanics*, Springer-Verlag, New York, 1985.

[11] R. G. Gallager, *Information Theory and Reliable Communication*, Wiley, New York, 1968.

[12] R. C. Gallager, *Low Density Parity Check Codes*, Reasearch monograph series n. 21, MIT Press, Cambridge Mass., 1963.

[13] R. G. Gallager, *A Simple Derivation of the Coding Theorem and Some Applications*, IEEE Trans. Information Theory 11 (1965), 3-18.

[14] R. G. Gallager, *The Random Coding Bound is Tight for the Average Code*, IEEE Trans. Information Theory 19 (1973), n. 2, 244-246.

[15] A. N. Kolmogorov and V. A. Uspensky, *Algorithms and randomnes*s, - SIAM J. Theory of Probability and Its Applications, vol. 32 (1987), pp. 389-412.

[16] A. Montanari, *The glassy phase of Gallager codes*, Eur. Phys. J. B 23 (2001), 121-136.

[17] A. Montanari, N. Sourlas, *The statistical mechanics of turbo codes*, Eur. Phys. J. B 18 (2000), 107-119.

[18] C. E. Shannon, *A Mathematical theory of Communication*, Bell Sys. Tech. J. 27 (1948), 379-423 and 623-656.

[19] N. Sourlas, *Spin-glass models as error-correcting codes*, Nature 339 (1989), 693-695.

[20] N. Sourlas, *Statistical mechanics and capacity-approaching error-correcting codes*, cond-mat/0106570.

[21] M. Talagrand, *A first course on spin glasses*, Lectures on probability theory and statistics, Saint-Flour 2000, Lecture Notes in Math. 1816 (2003), 181-285.

[22] M. Talagrand, *Spin Glasses, a Challenge to Mathematicians*, Springer, New York, 2003.