London School of Economics and Political Science

Optimal Use of Communication Resources with Markovian Payoff Functions

Nicola Wittur

A thesis submitted to the Department of Mathematics of the London School of Economics for the degree of Doctor of Philosophy, London, September, 2018

Abstract

Our work is based on the model proposed in the paper "Optimal Use of Communication Resources" by Olivier Gossner, Penélope Hernández and Abraham Neyman, [6]. We propose two models that consider an alteration of the payoff function in [6]. The general setup is as follows. A repeated game is played between a team of two players, consisting of a forecaster and an agent, and nature. We assume that the forecaster and the agent share the same payoff function. The forecaster, contrary to the agent, is able to observe future states of nature that have an impact on the team's payoff. A given pair of strategies for the players induces a sequence of actions and thus implements an average distribution on the actions of interest, i.e., on those actions that determine the payoff. We let the team's stage payoff not only depend on actions played in one stage, but on actions played in two consecutive stages. We introduce two models that vary w.r.t. the specification of the payoff function and the actions played by nature, with the aim of characterizing the implementable average distributions. This characterization is achieved through an information inequality based on the entropy function, called the information constraint. It expresses a key feature of the strategies of the players, namely the fact that the information used by the agent cannot exceed the amount of information sent by the forecaster. In each model we develop an information constraint that characterizes the implementable distributions as follows. On the one hand, we show that every implementable distribution fulfills the information constraint. And on the other hand, we prove that a certain set of distributions that fulfill the designated information constraint is implementable.

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 29518 words.

Statement of Conjoint Work

I confirm that chapter 2 and chapter 3 both of of part I of the thesis was jointly co-authored with my second supervisor Dr Ron Peretz.

I confirm that chapter 2 of part II of the thesis was jointly co-authored with my first supervisor Prof Olivier Gossner.

Acknowledgements

Firstly, I would like to thank my first supervisor, Professor Olivier Gossner, for setting a challenging research question. Not only did this assignment and his expert advice teach me deep insights into the topic, but also the ability to persevere in harder times when no approaches seemed to work.
Furthermore, I would also like to express my sincere appreciation to my second supervisor, Dr Ron Peretz. Without his continued support and our many meetings and chats, this thesis would not have been possible.

I gratefully acknowledge the funding I received towards my PhD from the LSE PhD Studentship. Moreover, I would like to thank everyone in the LSE Mathematics Department who supported me in these past years. In particular, I would like to thank my Program Director, Professor Bernhard von Stengel, for his immense encouragement and for the very motivating meetings we had. A big thank you also goes to Rebecca Lumb, our previous Research Manager, who showed so much empathy and who cheered for me.

Finally, completing this work would have been all the more difficult were it not for my family. My gratitude goes to Marcel, my husband, who had so much patience with me and who boosted my moral. To my dear Samuel, my son, who always reminded me not to forget the most important things in life. I am much indebted to my mother in law, to my step dad and of course to my mother, who travelled to London so many times to give me time to work. And last but not least, I want to thank my grandparents, who always send me their good wishes and prayers.

Contents

Bibliography 116

Chapter 1

Introduction

In this work we consider a repeated game with an asymmetry of information among the players who share a markovian payoff function. The theory of repeated games took shape in the second half of the twentieth century. An excellent coverage of the area, in particular of the classic results, is provided by the work of Mertens, Sorin and Zamir, see [10]. One particular field in the theory of repeated games, that focuses on the information of the players, originated at the height of the Cold War, namely during the negotiations between the United States and the Soviet Union about mutual reductions in their nuclear arsenals. At that time, no party had concrete information about the other's arsenals. So the United States Arms Control and Disarmament Agency (ACDA) turned to the most well-known game theorists of their time to help with the strategic issues in these negotiations. It was during and after these negotiations that the first papers on repeated games with an asymmetry of information among the players were written. In 1986, Aumann and Maschler, [1], published their seminal work on this topic. In the model they present, one player lacks information about the state of nature, whereas the other player is informed about it. This paper has motivated many research directions. One direction focused on the problem of strategic information transmission between players with unequal information about states of nature and has been well studied since. Two notable examples are the papers by Crawford and Sobel, [3], and by Forge, [5]. Both papers present models in which an informed player needs to signal her (private) knowledge about states of nature to an uninformed second player, whose actions then influence both players' payoff. The informed player's signal to the uninformed player, however, does not directly influence the two players' payoff. In other words, the signal containing information about the state of nature is costless.

The first paper that proposed a model which takes into consideration that sending information can be costly in many circumstances (see, e.g., [13] in the case of organizations) was presented by Gossner, Hernández and Neyman in [6]. Here, the better informed player - called forecaster - can transmit her knowledge about the states of nature to the less informed player - called the agent - through her actions, which also affect the payoff of the players. Hence, sending information in this model is costly in the sense that the forecaster has to weigh up the pros and cons of the information transmission with respect to her payoff.

The research we present in this work is directly motivated by [6] and considers a particular alteration of the payoff function in the respective model. In order to formulate our research questions, let us briefly introduce the central features of the model in [6]. A repeated game between a team, consisting of a forecaster and an agent, and nature takes place. The sequence of the states of nature is assumed to be i.i.d. In each stage of the game, the team members are able to observe the past actions played and the past states of nature. In addition, unlike the agent, the forecaster is able to observe all future states of nature. The team receives a payoff in every stage, which depends on the current state of nature, as well as on the actions of the forecaster and the agent, which we call action triple in the following. The strategies of the team players induce an infinite sequence of random action triples and hence a limiting average distribution, Q, of an action triple. A distribution that is induced in such a way is also called an implementable distribution. The authors in [6] prove two important theorems that characterise the set of implementable distributions. This characterization involves an information theoretic inequality that applies the Shannon entropy function (see [15]), which is called the information constraint. This constraint can be interpreted as the fact that the amount of information used by the agent cannot be greater than the information she actually receives from the forecaster. The first result in [6] states that every implementable distribution fulfills the information constraint. For the second result the authors show that for every distribution Q that fulfills the information constraint, there exist strategies of the forecaster and of the agent that implement Q.

Inspired by these two results, we set out to investigate the following. Assume that the payoff of the team in one stage not only depends on the current actions of the players and on the state of nature, but also on the actions and states of nature of the preceding stage. That is, we consider two consecutive stages whose actions and states of nature influence the payoff. Note that depending on the precise payoff structure, we are interested in finding implementable distributions which depend on those actions and states of nature that influence the payoff of the team. The question that naturally arises can be formulated as follows: Can we develop similar (adapted) information constraints as in [6] to characterise the set of implementable distributions?

The reader should note that we only focus on the search for which distributions are implementable. This analysis is independent of the payoff maximization problem (which searches for the optimal implementable distribution w.r.t. maximizing the payoff). However, the results we present in this thesis present an optimal basis for research problems like these.

We present two different models that take this model setting with an altered payoff structure into account and analyze the implementable distributions. The following provides a short overview of these two models. Let us start with the second model. It is closer in structure to [6] than the first model we present in the sense that only the payoff function changes. In particular, we assume that the team payoff in one stage depends on the current state of nature and on the current action of the agent, as well as on the agent's action in the previous stage. Hence, the actions of the forecaster in this model do not influence the payoff (they are hence costless).

The first model we present not only differs from [6] in regards of the payoff function, but also regarding the assumption on the dynamics of the states of nature. More precisely, in this model we let nature react to past action triples, so that the sequence of the states of nature is not i.i.d. This assumption allows us to consider a stage payoff function that depends on the complete action triple of the current, as well as on the previous stage. The first model is therefore more complex and richer in details, and hence takes up the central part of this thesis.

For both models we are able to formulate an information constraint with a similar interpretation as in [6], that characterises the respective implementable distribution as follows. We present two main results for both models. In the first main result we show that every implementable distribution satisfies the information constraint. The second main result states that every distribution with certain properties that satisfies the information constraint can be implemented. Note, that the second main results are a little bit more restrictive than the second main result in [6], since we can only consider a certain set of distributions that are implementable if they satisfy the information constraint.

In order to prove the first main theorem for both models, we can apply a useful lemma on the concavity of the entropy function which is stated in [6]. The proof of the first main theorem for the first model requires an additional result we present on the entropy of an induced random sequence of action triples. We can prove this lemma with probabilistic tools such as the Hoeffding inequality.

In order to prove the second main theorems we need to construct strategies of the players that implement a given distribution. These constructions constitute most of the work of this thesis and hence consist of the majority of our contributions. Let us first briefly summarise the main ideas of the strategy construction for the first model. We provide a new conceptualization of a certain typical sequence which we term Locally Typical Sequence, and introduce the so-called Splicing Algorithm. With the help of these new concepts and their useful properties we develop, we establish the following mechanism. Let P denote the distribution we would like to implement. The strategies we construct first induce a set of locally typical sequences which in turn serves as an input set for the Splicing Algorithm. Then, this algorithm outputs a sequence of action triples that implements our desired probability distribution, P.
Let us now turn to the construction of the strategies of the second model. Even though this construction is less complex than the previous one, it applies the methods of [6] in an elegant way, using a new concept we develop, called Block Distributions. The idea is as follows. We first construct a markov chain with a transition matrix derived from the distribution we intend to implement, P. Next, we define a block distribution, Q, of a sequence of random variables drawn from the markov chain. We show that if Q satisfies the information constraint given in [6], then P is implementable.

The results and the new techniques we present provide novel and thorough insights into the complexity of strategic information transmission when the players are faced with a markovian payoff function. Note that that our results not only generalise the setting in [6], but also add a planning problem to the model: If present action choices also affect future rewards or payoffs, we have a non-trivial planning problem. Such planning problems appear, for instance, in dynamic pricing models (see, e.g., the seminal work by Rao and Bass, [14]). These models study, among others, the question of how setting a price today influences future sales or competitors' prices. Considering a markovian payoff function and thus adding a planning problem to the model of [6] hence increases the scope of its application. Furthermore, the results we provide can offer an ideal starting point for future research directions that directly build upon our results. In particular, it provides the basis for the analysis of the optimal payoff of the team, or for the characterization of the set of equilibrium payoffs.

Most of the tools we employ in this work are information theoretic techniques. In particular, we apply the Shannon entropy function and its properties, as well as the concept of typical sequences. These techniques are reviewed in the preliminaries. The remaining sections of the thesis are divided into two parts. The first part is dedicated to the first model. Chapter 1 in part I describes the model, defines the implementable distributions and introduces the two main theorems. In chapter 2 of part I, we prove the first main theorem. As indicated, this proof requires an additional lemma to estimate the entropy of an induced sequence of random action triples. Chapter 3 of part I contains the preparation and the proof of the second main theorem. It is the most elaborate chapter of the thesis. We begin with the introduction of the Splicing Algorithm and provide examples to demonstrate the functioning of this algorithm. We then continue with the definition of locally typical sequences and explain why this concept is needed for the construction of the strategies. In the remaining sections of chapter 3 we derive the strategies with the help of the concepts introduced earlier. We provide one section that explains the functioning of the strategies in an intuitive way, followed by a section that looks at the details of the strategy construction. In the last section we prove the second main theorem of the first model. The second part of the thesis is dedicated to the second model and consists of two chapters. Chapter 1 of part II introduces the model and defines the implementable distributions. In chapter 2 we prove the two main theorems. The proof of the second main theorem also requires some preparation, in particular the construction of a markov chain and the introduction of block distributions, which are also included in chapter 2.

Chapter 2

Preliminaries

In this section we present the key techniques and methods we will employ in part I and in part II.

Notation 1. Given a finite set A, we denote by $\Delta(A)$ the set of categorical distributions over A. That is, every $p \in \Delta(A)$ is a discrete probability distribution that describes the probability of observing one possible outcome in A. Throughout this paper, we will be dealing with discrete random variables, which are denoted by bold or capital letters. We write $\mathbf{x} \sim p$ to denote that the probability mass function (pmf) of the random variable $\mathbf{x}$ is p. $\diamond$

2.1 Entropy

From a very intuitive point of view, entropy is a measure of chaos, where chaos is considered to be a state of some system. Whenever the elements in the system are spread equally (think of the distance of particles in a container), the entropy of that system is maximal. Lower entropy occurs, if the elements are ordered in certain ways (e.g., by equal properties or parameter values). If we consider the entropy of a random variable, we often interpret it as a measure of uncertainty, or information - which is not so different from our intuitive notion. For instance, if we wanted to guess the outcome of a random experiment, then our uncertainty about this outcome is the average amount of information, or entropy, that we don't have, in order to be certain of the outcome. Now, whenever the outcomes of a random experiment have equal probability, it is of course much harder to guess the outcome. This is the case of maximal entropy. The more unequal the probabilities of the outcomes are, the lower the entropy becomes and the easier for us to guess the outcome. The concept of entropy as a measure of information was first properly defined by Claude Shannon in his groundbreaking work, “A Mathematical Theory of Information”, see [15], as the Shannon entropy function. It will play an important role in this work, especially for the formalization of the information constraints. Below, we review the most important definitions w.r.t. entropy and its elementary properties. A good introduction to the Shannon entropy function is given, for instance, in [2] or [9].

Definition (Entropy of Discrete Random Variables)**.** The entropy of a discrete random variable $\mathbf{x}$ with $\mathbf{x} \sim p$ is given by

$$H(\mathbf{x}) = -\sum_{x} p(x)\log(p(x)).$$

Equally, the entropy of two discrete random variables $\mathbf{x}, \mathbf{y}$ with $(\mathbf{x}, \mathbf{y}) \sim q$ is given by

$$H(\mathbf{x}, \mathbf{y}) = -\sum_{x,y} q(x,y)\log(q(x,y)),$$

and is called joint entropy. ◇

Definition (Conditional Entropy)**.** The conditional entropy of $\mathbf{y}$ given $\mathbf{x}$, with $(\mathbf{x}, \mathbf{y}) \sim q$, is given by

$$\begin{aligned} H(\mathbf{y}|\mathbf{x}) &= \sum_{x} q(x) H(\mathbf{y}|\mathbf{x} = x) \\ &= -\sum_{x} q(x) \sum_{y} q(y|x)\log(q(y|x)) \\ &= -\sum_{x,y} q(x,y)\log q(y|x). \end{aligned}$$

The conditional entropy can be thought of as the average uncertainty of $\mathbf{y}$ when we observe $\mathbf{x}$. ◇

Elementary Properties of Entropy

The above defined quantities have the following properties:

- **Non-negativity**: The entropy of a random variable is always non-negative; $H(\mathbf{x}) \geq 0$, and we have $H(\mathbf{x}) = 0$ only, iff $\mathbf{x}$ is deterministic.
- **Monotonicity**: Conditioning reduces entropy; $H(\mathbf{y}|\mathbf{x}) \leq H(\mathbf{y})$.
- **Maximum Entropy**: If $\mathbf{x}$ takes values in a finite set A, then $H(\mathbf{x}) \leq \log_2 |A|$. The maximum entropy is reached, if and only if $\mathbf{x}$ is uniformly distributed.

- **Non-increasing under functions**: For every (deterministic) function $f(\mathbf{x})$ of $\mathbf{x}$, it holds that $H(f(\mathbf{x})) \leq H(\mathbf{x})$.

- **Chain Rule**: Given a sequence of n random variables $(\mathbf{x}_1, \mathbf{x}_2, \ldots \mathbf{x}_n)$, the chain rule tells us that we can decompose the joint entropy of the n random variables as follows:

$$H(\mathbf{x}_1, \mathbf{x}_2, \ldots \mathbf{x}_n) = \sum_{i=1}^{n} H(\mathbf{x}_i | \mathbf{x}_{i-1}, \ldots \mathbf{x}_1).$$

  For two variables, $(\mathbf{x}, \mathbf{y})$ this becomes

$$H(\mathbf{x}, \mathbf{y}) = H(\mathbf{x}|\mathbf{y}) + H(\mathbf{y}) = H(\mathbf{y}|\mathbf{x}) + H(\mathbf{x}).$$

Notation 2. Occasionally, we write $H_p(\mathbf{x})$ ($H_q(\mathbf{y}|\mathbf{x})$), if we need to highlight that $\mathbf{x} \sim p$ ($(\mathbf{x}, \mathbf{y}) \sim q$). Moreover, instead of $H_p(\mathbf{x})$, we sometimes also write $H(p)$, if the context demands it. $\diamond$

Relative Entropy

Related to the concept of entropy is the relative entropy, also known as Kullback-Leibler distance. It is a measure of how much one probability distribution differs from another, where a lower relative entropy points to similar behavior.

Definition. Let P and Q denote two distributions in $\Delta(A)$. The Kullback-Leibler distance between P and Q (or short: KL-distance), is defined as follows:

$$D(P||Q) = \sum_{a \in A} P(a) \log \frac{P(a)}{Q(a)}.$$

$\diamond$

The KL-distance has the following basic properties:

- **Non-Negativity**: $D(P||Q) \geq 0$ and $D(P||Q) = 0$ iff $P = Q$.

- **Convexity**: $D(P||Q)$ is convex in (P, Q).

2.2 Types and Typical Sequences

Let us first give a small overview of how typical sequences have grown in importance in Information Theory. It was again Claude Shannon, who introduced typical sequences (in a more intuitive than technical sense) as a powerful tool to establish his Information Theory. One of his key discoveries (in a simplified version) was the following. Consider a distribution, p, over a finite alphabet A. Shannon connected typical sequences with the entropy of p, by observing that the negative logarithm of the probability of a *typical* sequence of length n, that is drawn i.i.d. from p, divided by n is close to $H(p)$. Furthermore, for n large enough, the (total) probability of all non-typical sequences is arbitrarily small. Consequently, Shannon concluded, for large n, typical sequences happen frequently, even though there are few of them in A^n (approximately $2^{nH(X)}$). These key discoveries are known today as the *AEP*, the Asymptotic Equipartition Theorem, and form the basis for Shannon's famous Coding Theorems (which won't be discussed here).

So far, we described important properties of typical sequences, without providing a proper definition. Typical sequences have been defined in various forms, from weakly typical sequences, to (ε-) strongly typical sequences. In the following, we first follow the method of types, first introduced in [4], to define typical sequences, and the results below are taken from [2], if not otherwise stated.

2.2.1 First Order Types

Definition (First Order Type)**.** Let A be a finite set. For a given integer $n \in \mathbb{N}$, let x^n be a sequence in A^n. For every $a \in A$, denote by $N(a|x^n)$ the number of occurrences of element a in x^n. The first order type of x^n is given by the empirical distribution of x^n, defined by

$$emp(x^n)[a] = \frac{1}{n}N(a|x^n) \text{ for all } a \in A. \tag{2.1}$$

◇

Remark. The set of all types with denominator n is denoted by $\mathscr{P}_n$, for which we have an upper bound:

$$|\mathscr{P}_n| \leq (n+1)^{|A|}. \tag{2.2}$$

◇

Definition (First Order Type Set)**.** Let $P \in \mathscr{P}_n$. The set of all sequences $x^n \in A^n$ with first order type P is given by

$$T_n(P) = \{x \in A^n : emp(x) = P\}, \tag{2.3}$$

and we call every sequence in $T_n(P)$ typical w.r.t. P. ◇

Remark (The Size of a First Order Type Set)**.** The lower and upper bounds of the size of $T_n(P)$ are well known:

$$\frac{1}{(n+1)^{|A|}} 2^{nH(P)} \leq |T_n(P)| \leq 2^{nH(P)}. \tag{2.4}$$

(See Theorem 11.1.3 on page 350 in [2].) ◇

2.2.2 Second Order Types

Second order types are related to the relative frequency of pairs in a sequence $x^n \in A^n$. Before we begin, we present some notation.

Notation 3. Let P_2 denote a distribution over a finite set $A \times A = A^2$. If P_2 has full support, then for every pair $(a,b) \in A^2$, $P_2(b|a)$ denotes the conditional probability of b given a. ◇

Definition (Second Order Type)**.** Consider a sequence $x^n \in A^n$. For every pair $(a,b) \in A^2$, $N(a,b|x^n)$ denotes the number of occurrences of the pair (a,b) in x^n. The empirical distribution over pairs in x^n, denoted by $emp_2(x^n)$, is given by the relative frequency of pairs in x^n:

$$emp_2(x^n)[a,b] = \tfrac{1}{n-1} N(a,b|x^n) \text{ for all } (a,b) \in A^2, \tag{2.5}$$

and we call $emp_2(x^n)$ the second order type of x^n. ◇

Definition (Second Order Type Set)**.** Let $P_2 \in \Delta(A^2)$. The set of all sequences $x^n \in A^n$ with second order type P_2 is called the second order type set and we write

$$T_n^2(P_2) = \{x \in A^n : emp_2(x^n) = P_2\}. \tag{2.6}$$

If $T_n^2(P_2) \neq \emptyset$, then a sequence $x^n \in T_n^2(P_2)$ is called 2-typical w.r.t. P_2. ◇

In the following remark we look at the marginal distribution of a second order type:

Remark 1. Let $x^n \in T_n^2(P_2) \neq \emptyset$ and denote by $x_{(1)} = (x_1, \ldots, x_{n-1})$ the sequence of the first $n-1$ elements of x^n and by $x_{(2)} = (x_2, \ldots, x_n)$ the sequence of the last $n-1$ elements of x^n. Let $P_{(1)}$ and $P_{(2)}$ denote the first order type of $x_{(1)}$ and $x_{(2)}$ respectively. Obviously, $P_{(1)}$ and $P_{(2)}$ only differ if $x_1 \neq x_n$. Moreover, observe that

$$\begin{aligned}\sum_{a \in A} P_2(a,b) &= \sum_{a \in A} emp(x^n)[a,b] \\ &= \tfrac{1}{n-1} \sum_{a \in A} N(a,b|x^n) \\ &= \tfrac{1}{n-1} N(b|x_{(2)}) \\ &= emp(x_{(2)})[b] \\ &= P_{(2)}(b),\end{aligned}$$

and similarly, $\sum_{b \in A} P_2(a,b) = P_{(1)}(a)$. For $n \to \infty$, we set $P_{(1)} \to P$ and $P_{(2)} \to P$, and we call P the (asymptotically) unique marginal distribution of P_2 on A. $\diamond$

We now develop a new set of sequences that we term conditional subsequences. These sequences will exclusively be of use in part 1, however, due to their relation to second types, we introduce them here.

Definition 1 (Conditional Subsequences)**.** Let $x^n \in A^n$ and let $x_{(1)}$ and $x_{(2)}$ be given as in Remark 1. For every $a \in A$, define the following subsequence of x^n:

$$x_a^n = (x_{a,1}^n, x_{a,2}^n, \ldots, x_{a,N(a|x_{(1)}^n)}^n),$$

where for every $i \in (1, \ldots, N(a|x_{(1)}^n))$, $x_{a,i}^n$ is the element in x^n that succeeds the ith occurrence of element a in x^n. We call x_a^n the conditional subsequence of x^n w.r.t a. $\diamond$

The following example demonstrates the intuitive concept behind a conditional subsequence:

Example 1. Let $A = \{0,1\}$ and let $x^n = (0,1,1,0,0,1,1,1)$. Then,

$$\begin{aligned} x_0^n &= (1,0,1) \\ x_1^n &= (1,0,1,1) \end{aligned}$$

◇

The following definition establishes the link between conditional subsequences and second order types.

Definition 2 (Typical Conditional Subsequences)**.** Let P_2 be a second order type with $T_n^2(P_2) \neq \emptyset$ and with full support. For every $x^n \in T_n^2(P_2)$ there are $|A|$ conditional subsequences of x^n, $\{x_a^n : a \in A\}$, which we call typical conditional subsequences. ◇

Let us now deduce important properties of typical conditional subsequences:

Remark 2. We continue with the notation of Definition 2. If n is large, then by Remark 1, P_2 has an (asymptotically) unique marginal distribution $P \in \Delta(A)$. Then, for every $x^n \in T_n^2(P_2)$, every typical conditional subsequence of x^n has length $(n-1)P(a)$ and has first order type $P_2(\cdot|a)$. This can be easily seen as follows. We know from Remark 1 that for every $a \in A$, $\frac{1}{(n-1)}N(a|x_{(1)}) = P(a)$. Therefore, the length of every typical conditional subsequence x_a^n of x^n is given by $N(a|x_{(1)}) = (n-1)P(a)$ (recall the notation of $x_{(1)}$ and $x_{(2)}$ as introduced in Remark 1 - the choice of $x_{(1)}$ and not $x_{(2)}$ is important here). Furthermore, the first order type of a typical conditional subsequence is derived below,

$$\begin{aligned} emp(x_a^n)[b] &= \tfrac{1}{(n-1)P(a)}N(b|x_a^n) \\ &= \tfrac{1}{(n-1)P(a)}N(a,b|x^n) \\ &= \tfrac{1}{(n-1)P(a)}(n-1)P_2(a,b) \\ &= P_2(b|a). \end{aligned}$$

◇

2.2.3 ε-Typical Sequences

In this section we present the concept of ε-typical sequences and sets. In the literature, this is also sometimes termed as *Strongly Typical Sequences and Sets*. ε-typical sequences prove to be a useful alternative to typical sequences if we can live with the fact that a first order type of a (sufficiently long) sequence is not exactly equal to a *true* distribution, but only close to it. Similar to typical conditional subsequence, ε-typical sequences will solely occur in part 1, but appear here due to the connection to typical sequences. In the following sections we adopt the notation from Chapter 14 in [16].

In the following, denote by $P \in \Delta(A)$ the true distribution over the elements of a finite set A and let $X_1, X_2, \ldots$ denote a sequence of i.i.d. random variables with $X_i \sim P$, $\forall i \geq 1$.

Definition (The ε-typical set $T_n^{\varepsilon}(P)$)**.** Let $\varepsilon > 0$. The ε-typical set w.r.t. a distribution $P \in \Delta(A)$, denoted by $T_n^{\varepsilon}(P)$, is given by the set of sequences whose first order type is ε-close to P in the following sense:

$$T_n^{\varepsilon}(P) = \left\{ x^n : \forall a \in A, \begin{cases} \left|\frac{1}{n}N(a|x^n) - P(a)\right| < \varepsilon, \text{ if } P(a) > 0 \\ \frac{1}{n}N(a|x^n) = 0, \text{ else.} \end{cases} \right\} \tag{2.7}$$

◇

ε-typical sequences and sets have many useful properties, the one of interest to us is stated below.

Remark (The Probability of an ε-Typical Sequence Occurring)**.** Let $X^n = (X_1, \ldots X_n)$ denote an i.i.d. sequence with law P and let P^n denote the product distribution over A^n derived from P. Furthermore, let $c = -\sum_{a \in A} \log P(a)$. Denote by $x^n \in I^n$ a realization of X^n. If $x^n \in T_n^{\varepsilon}(P(a))$, then

$$2^{-n(H(X)+c\varepsilon)} \leq P^n(x^n) \leq 2^{-n(H(X)-c\varepsilon)}. \tag{2.8}$$

(See Property 14.7.4 on page 424 in [16].) ◇

The concept of ε-typical sequences can be extended to conditional ε-typical sequences:

Definition (The Conditional ε-Typical Set)**.** Let A and B be two finite sets and denote by $P_{A\times B}$ a distribution over elements in $A \times B$ with marginal P_A on A and P_B on B. Let $x^n \in X^n$ be an ε-typical sequence in $T_n^{\varepsilon}(P_A)$. The conditional ε-typical set w.r.t. x^n and $P_{A\times B}$, denoted by $T_n^{\varepsilon}(P_{A\times B}|x^n)$, consists of all those sequences $y^n \in B^n$ such that for every pair $(a,b) \in A \times B$, $N(a,b|x^n,y^n)$ (the number of occurrences of the pair $(a,b) \in X \times Y$ in the sequence of pairs $(x^n,y^n) = ((x_1,y_1),(x_2,y_2),\ldots,(x_n,y_n)) \in (A \times B)^n$) is close to the product of $\frac{1}{n}N(a|x^n)$ with the true conditional distribution $P_{A\times B}(b|a)$ as follows:

$$\begin{aligned} T_n^{\varepsilon}(P_{A\times B}|x^n) = \{ & y^n \in B^n : \forall (a,b) \in A \times B, \\ & \left. \begin{cases} |N(a,b|x^n,y^n) - P_{A\times B}(b|a)N(a|x^n)| < n\varepsilon, \text{ if } P_{A\times B}(b|a) > 0 \\ N(a,b|x^n,y^n) = 0, \text{else.} \end{cases} \right\} \end{aligned} \tag{2.9}$$

◇

Similar to the probability of an ε-typical sequence occurring, there exists an equivalent on the probability of a conditional ε-typical sequence occurring:

Remark (The Probability of a Conditional ε-Typical Sequence Occurring)**.** We follow the notation in the definition of a conditional ε-typical set. Let $x^n \in T_n^{\varepsilon}(P_A)$ and let $Y^n = (Y_1, \ldots, Y_n)$ denote a sequence of n independent random variables in B^n distributed according to $\{P_{A\times B}(\cdot|a) : a \in A\}$ and x^n, i.e., $Pr(Y_j = b) = P_{A\times B}(b|x_j)$, for every $j \in (1, \ldots, n)$. Denote by y^n a realization of Y^n. If $y^n \in T_n^{\varepsilon}(P_{A\times B}|x^n)$, then

$$2^{-n(H(b|a)+2c'\varepsilon)} \leq P^n(y^n|x^n) \leq 2^{-n(H(b|a)-2c'\varepsilon)}, \tag{2.10}$$

where $P^n(y^n|x^n) = \prod_{j=1}^{n} P_{A\times B}(y_j|x_j)$ and $c' = \sum_{(a,b)\in A\times B}(\log P_{A\times B}(b|a) - P_{A\times B}(b|a)\log(P_{A\times B}(b|a)))$.

$\diamond$

2.2.4 The Random Empirical Distribution

Finally, let us consider the concept of a random empirical distribution.

Definition. Let X^n be a sequence of random variables $X^n = (X_1, \ldots, X_n)$ (not necessarily i.i.d), that takes values in A^n. Let $P^n \in \Delta(A^n)$ denote the pmf of X^n. Denote by P_i the marginal distribution of P^n on coordinate i, i.e., $P_i(a) = \sum_{x^n\in A^n, x_i=a} P^n(x^n)$. Furthermore, consider the indicator random variables $\mathbb{1}_{\{X_1=a\}}, \mathbb{1}_{\{X_2=a\}}, \ldots, \mathbb{1}_{\{X_n=a\}}$. The random empirical distribution of X^n at $a \in A$ is the sample mean of these indicator random variables,

$$\frac{1}{n}\sum_{i=1}^{n} \mathbb{1}_{\{X_i=a\}} = emp(X)[a] = \frac{1}{n}N(a|X^n). \tag{2.11}$$

$\diamond$

Remark. Note, that $\mathbb{E}[\mathbb{1}_{\{X_i=a\}}] = P_i(a)$, hence the expected empirical distribution is the average of the marginal distributions of P^n:

$$\mathbb{E}[emp(X)] = \frac{1}{n}\sum_{i=1}^{n} P_i. \tag{2.12}$$

Furthermore, if X^n is a sequence of i.i.d. random variables with law μ, then the above equality simplifies to

$$\mathbb{E}[emp(X)] = \mu. \tag{2.13}$$

$\diamond$

Part I

Model 1

Chapter 1

Properties of the Model and Main Results

1.1 Features of Model 1

Let us now introduce the features of the first model. Many assumptions we make below are similar to the model in [6], the major differences are our assumptions on the payoff function and on nature's play. We begin with the action sets.

Denote by I, J and K the action sets of nature, forecaster and agent respectively. Again, the forecaster and the agent form a team. We analyze a repeated game between the team and nature. In each stage $t \geq 0$, denote by x_t, y_t and z_t the actions played by nature, forecaster and agent, respectively. We call $a_t = (x_t, y_t, z_t)$ an action triple and assume that a_0 is chosen arbitrarily.

In each stage $t > 0$, the team is assigned a stage payoff, described by g_t. This function not only depends on an action triple played in stage t, but also on the action triple played in the previous stage, i.e.,

$$g_t : (I \times J \times K)^2 \to \mathbb{R}.$$

In a repeated game, we allow nature to react to action triples as follows. Let $\mu \in \Delta(I)$. Prior to the start of the game, for every $a \in I \times J \times K$, nature draws independent μ-distributed and I-valued random variables $\mathbf{u}_1^a, \mathbf{u}_2^a, \mathbf{u}_3^a, \ldots = (\mathbf{u}_t^a)_{t\geq 1}$. For every $a \in A$, we refer to $(\mathbf{u}_t^a)_{t\geq 1}$ as nature's conditional sequence w.r.t. a. Furthermore, denote by U the random matrix with rows $\{(\mathbf{u}_t^a)_{t\geq 1} : a \in A\}$, and we write U_r to denote a realization of U. We interpret a matrix element u_l^a in U_r to be nature's choice of action in the repeated game after the lth occurrence of action triple a in the past play. More precisely: at stage $t = 0$, we let nature choose an action arbitrarily. In every following stage $t \geq 1$ in the repeated game, nature first observes the previous action triple a_{t-1}, and then counts its number of occurrences, l, in the entire past play. In stage t, nature plays $x_t = u_l^{a_{t-1}}$.

We assume that the forecaster observes a realization, U_r, of U before the game starts and that she is fully informed about the entire history of the play in every stage of the repeated game. Her (pure) strategy at stage t, denoted by σ_t, therefore depends on U_r and on the history of action triples

$$\sigma_t : (U_r \times I^{t-1} \times J^{t-1} \times K^{t-1}) \to J,$$

and her strategy for the repeated game, σ, is given by the sequence $(\sigma_t) = \sigma$. The agent has no further knowledge other than the history of action triples in the repeated game. Therefore, her strategy at stage t is given by a function

$$\tau_t : (I^{t-1} \times J^{t-1} \times K^{t-1}) \to K,$$

and her strategy of the entire game, τ, is given by the sequence $(\tau_t) = \tau$.

The random matrix U together with the strategies (σ, τ) induce a random sequence of action triples $\mathbf{a}_1, \mathbf{a}_2, \ldots$. We denote the corresponding probability distribution over $(I \times J \times K)^{\mathbb{N}}$ by $P_{U,\sigma,\tau}$.

Notation 4. Denote by $P^t_{U,\sigma,\tau}$ the marginal distribution of $P_{U,\sigma,\tau}$ over stage t's action triple, $\mathbf{a}_t$, and denote by $P^{t;2}_{U,\sigma,\tau}$ the marginal distribution over 2 consecutive action triples, $(\mathbf{a}_t, \mathbf{a}_{t+1})$. Let $Q^t_{U,\sigma,\tau} = \frac{1}{t}\sum_{t'=1}^{t} P^{t'}_{U,\sigma,\tau}$ be the average distribution up to stage t, and let $Q^{t;2}_{U,\sigma,\tau} = \frac{1}{t-1}\sum_{t'=1}^{t} P^{t';2}_{U,\sigma,\tau}$ be the average distribution over 2 consecutive action triples up to stage t (we also refer to the latter as the expected 2-step empirical distribution). $\diamond$

Definition 3 (Implementable Distribution)**.** Similar to [6], we call a distribution $Q \in \Delta((I \times J \times K)^2)$ *implementable* (*t-implementable*), if there is a strategy pair, (σ, τ), that implements (t-implements) the distribution Q, i.e., if $Q^{t;2}_{U,\sigma,\tau} \to Q$ as $t \to \infty$ ($Q^{t;2}_{U,\sigma,\tau} = Q$). We denote by $\mathscr{Q}$ (respectively, $\mathscr{Q}(t)$) the set of implementable (respectively, t-implementable) distributions. $\diamond$

Remark 3. It should be pointed out that an implementable distribution Q is the limit of an expected 2-step empirical distribution, $Q^{t;2}_{U,\sigma,\tau}$. In other words, if an expected 2-step empirical distribution, $Q^{t;2}_{U,\sigma,\tau}$, converges to Q, then we call Q implementable. $\diamond$

One important property of the sets $\mathcal{Q}$ and $\mathcal{Q}(t)$ which is used in the proof of Theorem 2 is stated below.

Remark 4. Every distribution $Q \in \mathcal{Q}(t)$ is implementable, i.e., $\mathcal{Q}(t)$ is contained in $\mathcal{Q}$. $\diamond$

(The proof of this remark is provided in the Appendix.)

1.2 Main Results

Our aim in this part is two-fold. On the one hand, we would like to describe the set of implementable distributions, $\mathcal{Q}$, in terms of an information constraint. On the other hand, we would like to specify distributions that are implementable - if possible - with the same information constraint. The results we produce in this paper achieve this goal - albeit under some restrictive assumptions. The information constraint we developed is given as follows.

Notation 5. Let $\mathbf{i}$, $\mathbf{i}'$ be I-, $\mathbf{j}$, $\mathbf{j}'$ be J- and $\mathbf{k}$, $\mathbf{k}'$ be K- valued random variables respectively. We write $(\mathbf{i}',\mathbf{j}',\mathbf{k}',\mathbf{i},\mathbf{j},\mathbf{k}) \sim P_2 \in \Delta((I \times J \times K)^2)$ to denote that the first random triple $(\mathbf{i}',\mathbf{j}',\mathbf{k}')$ precedes the second random triple, $(\mathbf{i},\mathbf{j},\mathbf{k})$. $\diamond$

Information Constraint

A distribution $P_2 \in \Delta((I \times J \times K)^2)$ is said to fulfill the information constraint if

$$H_{P_2}(\mathbf{i},\mathbf{j}|\mathbf{k},\mathbf{i}',\mathbf{j}',\mathbf{k}') \geq H_{P_2}(\mathbf{i}|\mathbf{i}',\mathbf{j}',\mathbf{k}'). \tag{1.1}$$

Equivalently, the information constraint can be stated as follows:

$$H_{P_2}(\mathbf{j}|\mathbf{k},\mathbf{i}',\mathbf{j}',\mathbf{k}') \geq H_{P_2}(\mathbf{i}|\mathbf{i}',\mathbf{j}',\mathbf{k}') - H_{P_2}(\mathbf{i}|\mathbf{k},\mathbf{i}',\mathbf{j}',\mathbf{k}').$$

In this form one can interpret the information constraint as the fact that the information used by the agent (right-hand side) cannot be greater than the information she receives from the forecaster (left-hand side), which is similar to the interpretation of the information constraint in [6].

In the first result we are able to show that every implementable distribution $Q \in \Delta((I \times J \times K)^2)$ satisfies the information constraint. In the second result we describe distributions, such that if they satisfy the information constraint, they are indeed implementable.

Lemma 1. *Every t- implementable distribution $P_2 \in \Delta((I \times J \times K)^2)$ fulfills the information constraint asymptotically, i,e., $H_{P_2}(\boldsymbol{i},\boldsymbol{j}|\boldsymbol{k},\boldsymbol{i}',\boldsymbol{j}',\boldsymbol{k}') \geq H_{P_2}(\boldsymbol{i}|\boldsymbol{i}',\boldsymbol{j}',\boldsymbol{k}') - \delta(t)$, where $\delta(t) \to 0$ as $t \to \infty$.*

Theorem 1. *Every implementable distribution $P_2 \in \Delta((I \times J \times K)^2)$ fulfills the information constraint.*

Theorem 2. *Let P_2 be a distribution over $(I \times J \times K)^2$ with the following properties:*

- *The marginal distribution of P_2 on the first and the second coordinates is identical, i.e., with $(\boldsymbol{i},\boldsymbol{i}',\boldsymbol{j},\boldsymbol{j}',\boldsymbol{k},\boldsymbol{k}') \sim P_2$, then $(\boldsymbol{i}',\boldsymbol{j}',\boldsymbol{k}') \sim (\boldsymbol{i},\boldsymbol{j},\boldsymbol{k}) \sim P$.*
- *$\boldsymbol{i}$ is independent of $(\boldsymbol{i}',\boldsymbol{j}',\boldsymbol{k}')$, i.e., $\boldsymbol{i} \sim \boldsymbol{i}|(\boldsymbol{i}',\boldsymbol{j}',\boldsymbol{k}') \sim \mu$.*
- *P_2 fulfills the information constraint.*

Then, P_2 is implementable.

Remark 5. The second property in Theorem 2 is a technical requirement in the proof of this Theorem, as we will see later on. Note that assuming $\mathbf{i}$ to be independent of $(\mathbf{i}',\mathbf{j}',\mathbf{k}')$ does not entirely reflect our assumption on nature's play in the game. As outlined in the model, the occurrance of an action triple $(\mathbf{i}',\mathbf{j}',\mathbf{k}')$ in one stage, t, of the game may affect the probability of nature's action, $\mathbf{i}$, in the next stage, $t+1$. This possible dependence can hence not be depicted in the implementable distributions of Theorem 2. $\diamond$

Chapter 2

Proof of Theorem 1

Theorem 1 follows directly from Lemma 1. The proof of Lemma 1 requires two additional Lemmas. In the first one, we show that the limiting average entropy of a random sequence of action triples is close to the entropy $H(\mu)$. The second result we need is the Concavity Lemma which is stated as Lemma 1 in [6]. Both lemmas are stated below.

Lemma 2. *Let $\boldsymbol{a} = (\boldsymbol{a}_1, \boldsymbol{a}_2, \dots \boldsymbol{a}_n)$ be a random sequence of action triples taking values in $A^n = (I \times J \times K)^n$, induced by strategies (σ, τ) and the random matrix U. Then,*

$$\tfrac{1}{n} H(\boldsymbol{a}) \geq H(\mu) - \delta(n),$$

where $\delta(n) \to 0$ as $n \to \infty$

Lemma (Concavity-Lemma)**.** *Let X and Y be finite sets. The function $Q \mapsto H_Q(\boldsymbol{y}|\boldsymbol{x})$ is concave on the set of probability measures on $X \times Y$.*

The result in Lemma 1 can now be derived as follows:

Proof of Lemma 1. Let $\mathbf{a} = (\mathbf{a}_1, \mathbf{a}_2, \dots \mathbf{a}_n) = (\mathbf{i}_1, \mathbf{j}_1, \mathbf{k}_1, \dots, \mathbf{i}_n, \mathbf{j}_n, \mathbf{k}_n)$ be a random sequence of action triples induced by strategies (σ, τ) and the random matrix U with $\mathbf{a} \sim P_{U,\sigma,\tau}$.
Let $Q^{n;2}_{U,\sigma,\tau} = \frac{1}{n} \sum_{t=1}^{n} P^{t;2}_{U,\sigma,\tau}$ denote the respective expected 2-step empirical distribution, as introduced in Notation 4. Applying the Concavity Lemma from GHN, we get

$$H_{Q^{n;2}_{U,\sigma,\tau}}(\mathbf{i}, \mathbf{j} | \mathbf{k}, \mathbf{i}', \mathbf{j}', \mathbf{k}') \geq \tfrac{1}{n} \sum_{t=1}^{n} H_{P^{t;2}_{U,\sigma,\tau}}(\mathbf{i}, \mathbf{j} | \mathbf{k}, \mathbf{i}', \mathbf{j}', \mathbf{k}').$$

Furthermore, we have

$$\begin{aligned}
\frac{1}{n}\sum_{t=1}^{n} H(\mathbf{i}_t,\mathbf{j}_t|\mathbf{k}_t,\mathbf{i}_{t-1},\mathbf{j}_{t-1},\mathbf{k}_{t-1}) &= \frac{1}{n}\sum_{t=1}^{n} H(\mathbf{i}_t,\mathbf{j}_t,\mathbf{k}_t|\mathbf{k}_t,\mathbf{i}_{t-1},\mathbf{j}_{t-1},\mathbf{k}_{t-1}) \\
&\geq \frac{1}{n}\sum_{t=1}^{n} H(\mathbf{i}_t,\mathbf{j}_t,\mathbf{k}_t|\mathbf{k}_t,\mathbf{i}_{t-1},\mathbf{j}_{t-1},\mathbf{k}_{t-1},\ldots,\mathbf{i}_1,\mathbf{j}_1,\mathbf{k}_1) \quad (2.1)\\
&\geq \frac{1}{n}\sum_{t=1}^{n} H(\mathbf{i}_t,\mathbf{j}_t,\mathbf{k}_t|\mathbf{i}_{t-1},\mathbf{j}_{t-1},\mathbf{k}_{t-1},\ldots,\mathbf{i}_1,\mathbf{j}_1,\mathbf{k}_1) \quad (2.2)\\
&= \frac{1}{n}H(\mathbf{i}_1,\mathbf{j}_1,\mathbf{k}_1,\ldots,\mathbf{i}_n,\mathbf{j}_n,\mathbf{k}_n) \quad (2.3)\\
&= \frac{1}{n}H(\mathbf{a}),
\end{aligned}$$

where inequality (2.1) follows since conditioning reduces entropy, inequality (2.2) is due to the fact that $\mathbf{k}_t$ is a function of $(\mathbf{i}_1,\mathbf{j}_1,\mathbf{k}_1,\ldots,\mathbf{i}_{t-1},\mathbf{j}_{t-1},\mathbf{k}_{t-1})$ and equality (2.3) is an application of the chain rule. Together with Lemma 2, we can conclude

$$H(\mu)-\delta(n) \leq \frac{1}{n}H(\mathbf{a}) \leq \frac{1}{n}\sum_{t=1}^{n} H_{P^{t;2}_{U,\sigma,\tau}}(\mathbf{i},\mathbf{j}|\mathbf{k},\mathbf{i}',\mathbf{j}',\mathbf{k}'),$$

and therefore

$$H_{Q^{n;2}_{U\sigma,\tau}}(\mathbf{i},\mathbf{j}|\mathbf{i}',\mathbf{j}',\mathbf{k},\mathbf{k}') \geq H(\mu)-\delta(n).$$

□

Let us now go back to Lemma 2. Before we look at the proof, we need to mention two important concepts that we are going to employ. The first one is Hoeffding's inequality, which provides a probability bound for the distance of a random variable to its expected value. From [7], we cite the following Theorem and Corollary:

Theorem. *(Hoeffding) Let $X_1,X_2,\ldots,X_n$ be i.i.d. random variables with $0 \leq X_i \leq 1$ ($i=1,\ldots,n$). Furthermore, let $\bar{X}=\frac{1}{n}\sum_{i=1}^{n} X_i$ and $\mu=\mathbb{E}\bar{X}$. Then, for every $t>0$:*

$$Pr(\bar{X}-\mu \geq t) \leq \exp(-2nt^2).$$

Corollary. *Let $X_1, X_2, \ldots, X_n$ be i.i.d. random variables as in Theorem (2). Then, for every $t \geq 0$:*

$$Pr(|\bar{X} - \mu| \geq t) \leq 2\exp(-2nt^2). \tag{2.4}$$

Second, we introduce the concept of a region:

Definition 4. [Region of U] Let $A = I \times J \times K$ and let $b \in \mathscr{P}_n(A)$ denote a first order type with denominator n. By definition, for every $a \in A$, $nb(a) \in \mathbb{N}$. A region in the random matrix U specified by b, defines, for every $a \in A$, the segment of the first $nb(a)$ elements of the row sequence $(\mathbf{u}_k^a)_k$ in U. We write $\mathbf{x}_b^n = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n)$ to denote the joined sequence of nature's actions in a region. The order of the elements in $\mathbf{x}_b^n$ follows the order of the rows in U, i.e., we first add the elements of of the first row of the specified region in U to $\mathbf{x}_b^n$, then the elements of the second row, and so on. By construction, the length of $\mathbf{x}_b^n$ is n. ◇

Remark 6. Since a region is solely specified by a type in $\mathscr{P}_n(A)$, the sequence of nature in such a region is i.i.d. with law μ. ◇

We now have all the tools at hand to prove Lemma 2:

Proof of Lemma 2. Let $\bar{\mathbf{a}}$ denote the (random) empirical distribution of $\mathbf{a}$. Further, let $\mathbf{u_a}$ be the random sequence of nature's actions in $\mathbf{a}$. Note that $\mathbf{u_a}$ is a function of $\mathbf{a}$, hence we have $H(\mathbf{a}) \geq H(\mathbf{u_a})$. Therefore, it is sufficient to prove

$$\tfrac{1}{n}H(\mathbf{u_a}) \geq H(\mu) - \delta(n).$$

Let $\mathbf{u_a} \sim \rho^n \in \Delta(I^n)$. Recall, that $\mathscr{P}_n(A)$ denotes the set of empirical distributions over A. Let $\mathbf{u}_b$ be the (random) action sequence of nature given $\bar{\mathbf{a}} = b$ with $b \in \mathscr{P}_n(A)$ and we set $\mathbf{u}_b \sim \rho_b^n$. Further, let $\mu^n = \mu \times \mu \times \cdots \mu$ denote the product measure on I^n.

We will make use of the following operator, introduced in [12]. For a distribution q that is absolutely continuous[1] w.r.t a distribution p define the following linear operator:

$$L_p(q) := H(q) + D(q||p) = -\sum_i \log(p_i)q_i.$$

[1] q is absolutely continuous w.r.t p if $p(x) = 0 \to q(x) = 0$ for all x.

Now,

$$
\begin{aligned}
H(\mathbf{u_a}) &\geq H(\mathbf{u_a}|\bar{\mathbf{a}}) \\
&= \sum_{b\in\mathscr{P}_n(A)} Pr(\bar{\mathbf{a}}=b)H(\mathbf{u_a}|\bar{\mathbf{a}}=b) \\
&= \sum_{b\in\mathscr{P}_n(A)} Pr(\bar{\mathbf{a}}=b)H(\mathbf{u}_b) \\
&= \sum_{b\in\mathscr{P}_n(A)} Pr(\bar{\mathbf{a}}=b)H(\rho_b^n) \\
&= \sum_{b\in\mathscr{P}_n(A)} Pr(\bar{\mathbf{a}}=b)(L_{\mu^n}(\rho_b^n)-D(\rho_b^n||\mu^n)) \\
&= \sum_{b\in\mathscr{P}_n(A)} Pr(\bar{\mathbf{a}}=b)L_{\mu^n}(\rho_b^n)-\sum_{b\in\mathscr{P}_n(A)} Pr(\bar{\mathbf{a}}=b)D(\rho_b^n||\mu^n)) \qquad (2.5)
\end{aligned}
$$

The remaining proof is split into two claims that focus on the sums in equation (2.5):

Claim 1. *For all $\delta > 0$ there is N_δ s.t. for all $n \geq N_\delta$:*

$$\sum_{b\in\mathscr{P}_n(A)} Pr(\bar{\boldsymbol{a}}=b)L_{\mu^n}(\rho_b^n) \geq n(H(\mu)-\delta).$$

Claim 2.

$$\sum_{b\in\mathscr{P}_n(A)} Pr(\bar{\boldsymbol{a}}=b)D(\rho_b^n||\mu^n)) \leq \mathscr{O}(\log(n))$$

We prove Claim 1 as follows: Since $L_{\mu^n}(\rho_b^n)$ is a linear operator in ρ_b^n, we get:

$$\sum_{b\in\mathscr{P}_n(A)} Pr(\bar{\mathbf{a}}=b)L_{\mu^n}(\rho_b^n) = L_{\mu^n}\left(\sum_{b\in\mathscr{P}_n(A)} Pr(\bar{\mathbf{a}}=b)\rho_b^n\right) = L_{\mu^n}(\rho^n).$$

Now, we show

$$L_{\mu^n}(\rho^n) \geq \sum_{j=1}^{n} L_\mu(\rho^j), \qquad (2.6)$$

where ρ^j denotes the marginal distribution of ρ^n on coordinate j, i.e., for $a \in I$

$$\rho^j(a) = \sum_{\substack{x^n\in I^n \\ x_j=a}} \rho^n(x).$$

Recall, that for $x^n \in I^n$ we have $\mu^n(x^n) = \prod_{j=1}^n \mu(x_j)$, hence

$$
\begin{aligned}
L_{\mu^n}(\rho^n) &= -\sum_{x^n \in I^n} \log(\mu^n(x^n))\rho^n(x^n) \\
&= -\sum_{x^n \in I^n} \sum_{j=1}^{n} \log(\mu(x_j))\rho^n(x^n) \\
&= -\sum_{j=1}^{n} \sum_{x^n \in I^n} \log(\mu(x_j))\rho^n(x^n) \\
&= -\sum_{j=1}^{n} \sum_{a \in I} \sum_{\substack{x^n \in I^n \\ x_j = a}} \log(\mu(x_j))\rho^n(x^n) \\
&= -\sum_{j=1}^{n} \sum_{a \in I} \sum_{\substack{x^n \in I^n \\ x_j = a}} \log(\mu(a))\rho^n(x^n) \\
&= -\sum_{j=1}^{n} \sum_{a \in I} \log(\mu(a)) \sum_{\substack{x^n \in I^n \\ x_j = a}} \rho^n(x^n) \\
&= -\sum_{j=1}^{n} \sum_{a \in I} \log(\mu(a))\rho^j(a) \\
&= \sum_{j=1}^{n} L_\mu(\rho^j), \qquad (2.7)
\end{aligned}
$$

which proves (2.6). Now, recall that $L_\mu(\rho^j)$ is linear in ρ^j, therefore, we can simplify the expression in (2.7) to

$$
\sum_{j=1}^{n} L_\mu(\rho^j) = n\left(\frac{1}{n}\sum_{j=1}^{n} L_\mu(\rho^j)\right) = n\left(L_\mu\left(\frac{1}{n}\sum_{j=1}^{n}\rho^j\right)\right),
$$

where $\frac{1}{n}\sum_{j=1}^{n}\rho^j$ is the expected empirical distribution of nature's actions $\mathbf{u_a}$. Note, that if $\bar{\mathbf{u}}_\mathbf{a}$ denotes the (random) empirical distribution of $\mathbf{u_a}$, we have $\frac{1}{n}\sum_{j=1}^{n}\rho^j = \mathbb{E}[\bar{\mathbf{u}}_\mathbf{a}]$ (recall, that $\bar{\mathbf{u}}_\mathbf{a}$ takes values in $\mathscr{P}_n(I)$). The result in (2.6) is therefore equivalent to

$$
L_{\mu^n}(\rho^n) \geq nL_\mu(\mathbb{E}[\bar{\mathbf{u}}_\mathbf{a}]). \qquad (2.8)
$$

Now, we show

$$
\|\mathbb{E}[\bar{\mathbf{u}}_\mathbf{a}] - \mu\| \to 0 \text{ as } n \to \infty. \qquad (2.9)
$$

First, recall that by Definition 4 and by Remark 6 every fixed type $b \in \mathscr{P}_n(A)$ corresponds to a region in the random matrix U, in which the sequence of nature, $\mathbf{x}_b^n = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n)$ is i.i.d. with law μ. Let $\bar{\mathbf{x}}_b^n$ denote the (random) empirical distribution of $\mathbf{x}_b^n$ which takes values in $\mathscr{P}_n(I)$. Now, let $\varepsilon > 0$ and fix $b \in \mathscr{P}_n(A)$. Denote by $E_{b,\varepsilon}$ the event that the L_1-distance between $\bar{\mathbf{x}}_b^n$ and μ is larger than ε, i.e., $\left\|\bar{\mathbf{x}}_b^n - \mu\right\| > \varepsilon$. Hence, the union of these events, $\bigcup_{b \in \mathscr{P}_n(A)} E_{b,\varepsilon}$, depicts the event that there exists at least one type $b \in \mathscr{P}_n(A)$, with $\left\|\bar{\mathbf{x}}_b^n - \mu\right\| > \varepsilon$. Applying the Hoeffding inequality, we get

$$Pr(E_{b,\varepsilon}) = Pr(|\bar{\mathbf{x}}^n - \mu| > \varepsilon) \leq 2\exp(-2n\varepsilon^2),$$

and

$$\begin{aligned} Pr\Big(\bigcup_{b \in \mathscr{P}_n(A)} E_{b,\varepsilon}\Big) &\leq \sum_{b \in \mathscr{P}_n(A)} Pr(E_{b,\varepsilon}) \\ &\leq (n+1)^{|A|} 2\exp(-2n\varepsilon^2). \end{aligned}$$

Therefore,

$$Pr\Big(\bigcup_{b \in \mathscr{P}_n(A)} E_{b,\varepsilon}\Big) \to 0, \text{ as } n \to \infty. \tag{2.10}$$

Now, by the law of total expectation, letting $E_\varepsilon = \bigcup_{b \in \mathscr{P}_n(A)} E_{b,\varepsilon}$, and applying (2.10), it holds that

$$\begin{aligned} \|\mathbb{E}[\bar{\mathbf{u}}_\mathbf{a}] - \mu\| &= \|\mathbb{E}[\bar{\mathbf{u}}_\mathbf{a}|E_\varepsilon] Pr(E_\varepsilon) + \mathbb{E}[\bar{\mathbf{u}}_\mathbf{a}|E_\varepsilon^c](1 - Pr(E_\varepsilon)) - (\mu Pr(E_\varepsilon) + \mu(1 - Pr(E_\varepsilon)))\| \\ &= \|(\mathbb{E}[\bar{\mathbf{u}}_\mathbf{a}|E_\varepsilon] - \mu) Pr(E_\varepsilon) + (\mathbb{E}[\bar{\mathbf{u}}_\mathbf{a}|E_\varepsilon^c] - \mu)(1 - Pr(E_\varepsilon))\| \\ &\leq Pr(E_\varepsilon) \|\mathbb{E}[\bar{\mathbf{u}}_\mathbf{a}|E_\varepsilon] - \mu\| + (1 - Pr(E_\varepsilon)) \|\mathbb{E}[\bar{\mathbf{u}}_\mathbf{a}|E_\varepsilon^c] - \mu\| \\ &\leq \varepsilon \text{ for } n \text{ large enough,} \end{aligned}$$

since $\|\mathbb{E}[\bar{\mathbf{u}}_\mathbf{a}|E_\varepsilon^c] - \mu\| \leq \varepsilon$ and $\|\mathbb{E}[\bar{\mathbf{u}}_\mathbf{a}|E_\varepsilon] - \mu\|$ is bounded by 2. Hence, since this holds for all $\varepsilon > 0$, the result in (2.9) follows:

$$\|\mathbb{E}[\bar{\mathbf{u}}_\mathbf{a}] - \mu\| \to 0 \text{ as } n \to \infty.$$

We can now finalise the proof of Claim 1. Note, that the linear operator L_μ is continuous and therefore Lipschitz continuous (continuity and Lipschitz continuity are equivalent for linear operators). Hence, there exists $\mu^* > 0$ (in fact, $\mu^* = \min_{i \in I} \mu_i$) such that

$$\left|L_\mu(\mathbb{E}[\bar{\mathbf{u}}_\mathbf{a}]) - L_\mu(\mu)\right| \leq \mu^* \|\mathbb{E}[\bar{\mathbf{u}}_\mathbf{a}] - \mu\| \to 0 \text{ as } n \to \infty.$$

Therefore, for all $\delta > 0$ there is N_δ, s.t. for all $n > N_\delta$:

$$L_\mu(\mathbb{E}[\bar{\mathbf{u}}_{\mathbf{a}}]) \geq L_\mu(\mu) - \delta.$$

This completes the proof of Claim 1, since $L_\mu(\mu) = H(\mu)$.

Let us now turn to the proof of Claim 2. Let $p_b = (Pr(\bar{\mathbf{a}} = b), Pr(\bar{\mathbf{a}} \neq b))$ denote a binary distribution. Further, define $\rho^n_{b^c} = \frac{\mu^n - \rho^n_b Pr(\bar{\mathbf{a}}=b)}{Pr(\bar{\mathbf{a}} \neq b)}$ s.t.

$$\mu^n = Pr(\bar{\mathbf{a}} = b)\rho^n_b + Pr(\bar{\mathbf{a}} \neq b)\rho^n_{b^c}.$$

Applying Proposition 3 in the Appendix, we get

$$Pr(\bar{\mathbf{a}} = b)D(\rho^n_b||\mu^n) + Pr(\bar{\mathbf{a}} \neq b)D(\rho^n_{b^c}||\mu^n) \leq H(p_b),$$

hence,

$$\sum_{b \in \mathscr{P}_n(A)} Pr(\bar{\mathbf{a}} = b)D(\rho^n_b||\mu^n) \leq \sum_{b \in \mathscr{P}_n(A)} H(p_b).$$

Now, letting $k = |\mathscr{P}_n(A)|$ we show

$$\tfrac{1}{k} \sum_{b \in \mathscr{P}_n(A)} H(p_b) \leq H(\tfrac{1}{k}). \tag{2.11}$$

Since the entropy is a concave function in p_b, we have

$$\tfrac{1}{k} \sum_{b \in \mathscr{P}_n(A)} H(Y^b) \leq H(\tfrac{1}{k} \sum_{b \in \mathscr{P}_n(A)} p_b).$$

With $p = \frac{1}{k}\sum_{b \in \mathscr{P}_n(A)} p_b$, p is a Bernoulli distribution with support in $\{0,1\}$ and with

$$p = \tfrac{1}{k} \sum_{b \in \mathscr{P}_n(A)} Pr(\bar{\mathbf{a}} = b) = \tfrac{1}{k}.$$

Hence, the inequality in (2.11) follows.

Finally, we can deduce

$$\begin{aligned}\sum_{b\in\mathscr{P}_n(A)} H(Y^b) &\leq kH(\tfrac{1}{k}) \\ &= -k(\tfrac{1}{k}\log(\tfrac{1}{k}) + (1-\tfrac{1}{k})\log(1-\tfrac{1}{k})) \\ &= -(-\log(k) + \log(1-\tfrac{1}{k})^{k-1}) \\ &= \mathscr{O}(\log(n)),\end{aligned}$$

where the last equality follows since $\log(1-\frac{1}{k})^{k-1}$ is bounded by $-\frac{1}{\ln 2}$ as $k \to \infty$ and $k \leq (n+1)^{|A|}$. This concludes the proof of Claim 2.

Claim 1 and Claim 2 now imply the result in Lemma 2.

□

Chapter 3

Proof of Theorem 2

The proof of Theorem 2 is more complex than the proof of Theorem 1 and requires substantial preparation. We begin with the motivation of our proof idea, followed by a general outline of our approach. These introductions will facilitate the understanding of the subsequent sections.

3.1 Motivation for the Proof of Theorem 2

Some of the key ideas that we are going to apply in the proof of Theorem 2 originate from the proof of Theorem 2 in [6], which is stated below for completeness.

Theorem (Theorem 2 in [6])**.** *Any distribution $Q \in \Delta(I \times J \times K)$ that satisfies the information constraint ($H_Q(i, j|k) > H_Q(i)$) and has marginal μ on I is implementable.*

We now explain how this result influenced our proof idea. To this end, let us compare our model with the model in [6]. One major difference between the two models is the assumption on the sequence of nature's actions. While in [6] the sequence of nature is assumed to be i.i.d., in our model nature's actions played in one stage depend on the actions played of all players in the previous stages. Only the conditional sequences of nature are assumed to be i.i.d. The general approach in [6] to implement a distribution Q was to construct strategies that induce a sequence of action triples with a first order type close to Q. This approach led us to the following idea. Let $P_2' \in \Delta((I \times J \times K)^2)$ be a distribution with the same properties as in Theorem 2, and for $a \in A$, let $P_2'(\cdot|a)$ denote a conditional distribution of P_2'. With the help of the results in [6], we know how to implement $P_2'(\cdot|a)$. I.e., we need to construct strategies that induce a sequence with first order type $P_2'(\cdot|a)$. However, we not only want to implement one conditional distribution $P_2'(\cdot|a)$, the aim is to construct strategies such that for every $a \in A$ a sequence with first order type $P_2'(\cdot|a)$ is induced. Then, in a next step, we need to construct a mechanism to join these sequences together appropriately, so that they yield a sequence with second order type close to P_2'.

Note, however, that for some $a \in A$, $P_2'(\cdot|a)$ may not be defined, which may cause problems inducing sequences with such a type. However, by Lemma 16 in the Appendix, we can find a distribution $P_2 \in \Delta((I \times J \times K)^2)$ that is close to P_2', satisfies the conditions stated in Theorem 2, and, in addition, has full support, so that for every $a \in A$ the conditional distribution $P_2(\cdot|a)$ is well-defined.

3.2 The General Approach

Let us now first outline the basic set up for the proof, before we specify the proof idea outlined above a little bit more. The aim is to construct strategies for the forecaster and for the agent that together with nature's actions induce a certain game with the following properties. Given a specified integer $n \in \mathbb{N}$, we first divide the stages of the game into blocks of length n. The strategies will be designed in such a way, that in each block $k > 1$ and dependent on the actions of nature, they induce action sequences for the forecaster and for the agent, that will be combined into a sequence of action triples $\alpha^n \in (I \times J \times K)^n$. We will show that in almost every block $k > 1$, α^n has a second order type close to the distribution we intend to implement.

If P_2' is the distribution we want to implement, let us now choose a distribution P_2 that is close to P_2' with the same properties as stated in Theorem 2 and with full support. Following the idea in the previous section, our approach to prove Theorem 2 is to induce a sequence $\alpha^n \in (I \times J \times K)^n$ in every block $k > 1$, that has a second order type close to P_2, via first inducing $|A|$ sequences with first order type $P_2(\cdot|a)$, for every $a \in A$. An important question that arises is concerned with the length of these $|A|$ sequences, if we have a fixed block length, n. The following observation provides some insights on this question.

Observation 1. Let $A = I \times J \times K$ and let $P_2 \in \Delta(A^2)$ be a distribution that has full support. Fix $n \in \mathbb{N}$, such that P_2 has a non-empty second order type set w.r.t. length $n+1$, i.e., $T^2_{n+1}(P_2) \neq \emptyset$. Let $\alpha^{n+1} \in T^2_{n+1}(P_2)$. Then, by Remark 2, for every $a \in A$, the typical conditional subsequence α_a^{n+1} of α^{n+1} has length $nP(a)$ and has first order type $P_2(\cdot|a)$. $\diamond$

Remark 7. The strategies we are going to construct in the proof of Theorem 2 actually produce a sequence $\alpha^{n+1} = (\alpha_0, \alpha_1, \ldots, \alpha_n)$ of length $n+1$ per block $k > 0$, with the property that α_0 is the last element played in the previous block. We will assume n to be large in our proof, therefore, the difference in the second order type of the sequences α^{n+1} and α^n is negligible. $\diamond$

With the help of Observation 1, we now present the general approach which summarises our first (intuitive) idea of how to induce a sequence $\alpha^n \in (I \times J \times K)^n$ that has a second order type close to P_2. This approach should be considered as a guideline that will undergo a process of amendments until we eventually arrive at the actual approach that we take in the proof of Theorem 2.

The General Approach

Step 1: (**Sequence Induction**) At the beginning of block $k > 1$, we induce a set of sequences $\left\{\alpha^{nP(a)} \in T_{nP(a)}(P_2(\cdot|a)) : a \in A\right\}$ (note that every sequence in this set has the same properties as the typical conditional subsequences of a sequence $\alpha^{n+1} \in T^2_{n+1}(P_2)$).

Step 2: (**Sequence Splicing**) We combine the sequences from Step 1 into a sequence α^n, s.t. its second order type is close to P_2.

In the section that follows we will focus on Step 2. We will test whether the splicing can be accomplished with the given set of sequences given in Step 1. For this, we provide an algorithm that uses these sequences as input and combines them into one sequence with the aim of maintaining the correct transition frequencies stipulated by $\{P_2(\cdot|a) : a \in A\}$. We will then come across several problems concerning the length and the second order type of the output sequence. These issues will force us to make changes to the set of (input-) sequences in Step 1, which will be highlighted at the end of the relevant paragraphs.

3.3 Sequence Splicing

3.3.1 The Splicing Algorithm

Definition 5 (Input Set)**.** Fix $n \in \mathbb{N}$ and let A be a finite set ($|A| \geq 2$). Furthermore, let P_2 be a distribution over A^2 with full support, unique marginal distribution $P \in \Delta(A)$, and non-empty second order type set $T^2_{n+1}(P_2) \neq \emptyset$. For every $a \in A$, denote the marginal distributions by $\rho^a = P_2(\cdot|a)$. Since P_2 has full support, then by Remark 2 there are $|A|$ first order type sets of the form $\left\{T_{nP(a)}(\rho^a) : a \in A\right\}$. Denote by

S a choice of $|A|$ sequences, one from each set $T_{nP(a)}(\rho^a)$, $a \in A$:

$$S = \left\{ \alpha^{nP(a)} \in T_{nP(a)}(\rho^a) : a \in A \right\}.$$

We call S an input set for the Splicing Algorithm w.r.t. n and P_2. ◇

Notation 6. In order to ease notation we sometimes abbreviate the sequence $\alpha^{nP(a)} \in S$ to α^a for every $a \in A$ if the associated distribution P_2 is known. ◇

We now introduce the so-called Splicing Algorithm (in Pseudo Code) that takes the set S as input and combines its sequences into one sequence, α, while maintaining, as best as possible, the correct transition frequencies specified by $\{\rho^a : a \in A\}$:

Algorithm 1 The Splicing Algorithm

Require: An input set $S = \left\{ \alpha^a \in T_{nP(a)} : a \in A \right\}$ w.r.t n and P_2 (given as a set of lists); an initial value $\alpha^0 \in A$
Ensure: A finite sequence $\alpha = (\alpha_0, \alpha_1, \ldots)$ (unspecified length)

```
start a list α with α = list(α^0).
for j ≥ 1 do
    l = last element added to α (l ∈ A)
    if α^l not empty then
        α_j = α^l[1]
        α.append(α_j)
        del(α^l[1]) (delete the element just added to α from α^l)
    else {algorithm terminates}
    end if
end for
```

The following example demonstrates the mechanism of the Splicing Algorithm.

Example 2. Let $A = \{a, b\}$ and let $P_2 = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$ so that $\rho^a = \rho^b = (\frac{1}{2}, \frac{1}{2})$. Fix $n = 8$. Let $\alpha_0 = a$ and $S = \left\{ \alpha^a \in T_4(\rho^a), \alpha^b \in T_4(\rho^b) \right\}$, with

$$\alpha^a = (a, b, b, a)$$

$$\alpha^b = (b, a, b, a).$$

The Splicing Algorithm applied to S and α_0 builds up a sequence α as follows: initiated by $\alpha_0 = a$, in the first step we identify a as the last element added to α, add the first element of the sequence α^a, a, to α and delete a (i.e. the first element) from α^a. After the first step we have $\alpha = (a, a)$ and a reduced sequence $\alpha^a = (b,b,a)$. In the second step, we identify a as the last element added to α, add the first element from the reduced sequence α^a, b, to α and delete b from α^a. Hence, we have $\alpha = (a,a,b)$ and a further reduced sequence $\alpha^a = (b,a)$, etc. We summarise the steps of the Algorithm in the table below.

Step	α^a	α^b	α
start	(a,b,b,a)	(b,a,b,a)	(a)
1	(b,b,a)	(b,a,b,a)	(a,a)
2	(b,a)	(b,a,b,a)	(a,a,b)
3	(b,a)	(a,b,a)	(a,a,b,b)
4	(b,a)	(b,a)	(a,a,b,b,a)
5	(a)	(b,a)	(a,a,b,b,a,b)
6	(a)	(a)	(a,a,b,b,a,b,b)
7	(a)	$\emptyset$	(a,a,b,b,a,b,b,a)
8	$\emptyset$	$\emptyset$	(a,a,b,b,a,b,b,a,a)

Table 3.1: Application of the Algorithm to Example 2

Hence, the Algorithm produces a sequence $\alpha = (a,a,b,b,a,b,b,a,a)$ of length 9 with second order type P_2. ◇

Definition 6. The sequence α produced by the Splicing Algorithm is called the output sequence. If $\alpha \in A^{n+1}$, i.e., if α has full length, then the algorithm exhausted all elements from the input sequences in S, and we call α an optimal output sequence. ◇

Note, that the output sequence in Example 2 is optimal. Every input sequence has been exhausted completely by the end of the algorithm. However, this is not always the case, as shown in the next example.

Example 3. As before, let $A = \{a,b\}$, $P_2 = (\frac{1}{4},\frac{1}{4},\frac{1}{4},\frac{1}{4})$ and $\rho^a = \rho^b = (\frac{1}{2},\frac{1}{2})$. Again, fix $n = 8$ and $\alpha_0 = a$. Consider the set $S = \left\{\alpha^a \in T_4(\rho^a), \alpha^b \in T_4(\rho^b)\right\}$ with

$$\alpha^a = (a,a,b,b)$$

$$\alpha^b = (a,a,b,b)$$

Running the Algorithm on this Example yields $\alpha = (a,a,a,b,a,b,a) \in A^7$ as shown in Table 3.2. One can see that the sequence α^b is not exhausted when the Algorithm ends. This happens since the sequence α^a is exhausted too early. ◇

We can observe from the two examples above that for a given n, a set of sequences S, and an initial value α_0, if the algorithm does not run until all sequences in S are exhausted, it only produces a sequence α of

Step	α^a	α^b	α
start	(a,a,b,b)	(a,a,b,b)	(a)
1	(a,b,b)	(a,a,b,b)	(a,a)
2	(b,b)	(a,a,b,b)	(a,a,a)
3	(b)	(a,a,b,b)	(a,a,a,b)
4	(b)	(a,b,b)	(a,a,a,b,a)
5	$\emptyset$	(a,b,b)	(a,a,a,b,a,b)
6	$\emptyset$	(b,b)	(a,a,a,b,a,b,a)

Table 3.2: Application of the Algorithm to Example 3

length $< n+1$.

To resolve this problem, we are now going to use longer input sequences with certain properties, so that the output sequence has a guaranteed length $n+1$. In order to do so, we now introduce another set of sequences, a so-called set of tails, T, that we add to a given input set S. We then show that running the algorithm on the appended sets S and T will produce a sequence α of length at least $n+1$.

Definition 7 (Set of Tails)**.** Let S be an input set of the Splicing Algorithm w.r.t n and P_2 as introduced in Definition 5. Let $m \in \mathbb{N}$ s.t. for all $a \in A$, $T_{mP(a)}(\rho^a) \neq \emptyset$. We call

$$T = \left\{ t^{mP(a)} \in T_{mP(a)}(\rho^a) : a \in A \right\}$$

a set of tails for the input set S w.r.t. m and P_2 ◇

Notation 7. Given an input set S and a set of tails T w.r.t. n, m and P_2, we denote by $S \oplus T$ the set of the coupled sequences in S and T, i.e.,

$$S \oplus T = \left\{ \alpha^{nP(a)} \oplus t^{mP(a)} \in T_{(n+m)P(a)}(\rho^a) : a \in A \right\},$$

with $\alpha^{nP(a)} \oplus t^{mP(a)} = (\alpha_1, \ldots, \alpha_{nP(a)}, t_1, \ldots, t_{mP(a)})$. Furthermore, similar as above, if clear from the context, then for every $a \in A$ we abbreviate the sequence $\alpha^{nP(a)} \oplus t^{mP(a)} \in S \oplus T$ to $\alpha^a \oplus t^a$. ◇

Let us now run the Splicing Algorithm on an input set S that is coupled with a set of tails, T:

Example 4. We continue with Example 3 by coupling the following tails to the input set S: for $m = 4$, let $T = \{t^a = (a,b), t^b = (a,b)\}$. Then, $S \oplus T$ consists of the two sequences

$$\alpha^a \oplus t^a = (a,a,b,b,a,b)$$
$$\alpha^b \oplus t^b = (a,a,b,b,a,b).$$

Running the Algorithm on $S \oplus T$ with initial value $\alpha_0 = a$ yields

$$\alpha = (a,a,a,b,a,b,a,a,b,b,b,a) \in A^{12},$$

as shown in Table 3.3.

Step	$\alpha^a \oplus t^a$	$\alpha^b \oplus t^b$	α
start	(a,a,b,b,a,b)	(a,a,b,b,a,b)	(a)
1	(a,b,b,a,b)	(a,a,b,b,a,b)	(a,a)
2	(b,b,a,b)	(a,a,b,b,a,b)	(a,a,a)
3	(b,a,b)	(a,a,b,b,a,b)	(a,a,a,b)
4	(b,a,b)	(a,b,b,a,b)	(a,a,a,b,a)
5	(a,b)	(a,b,b,a,b)	(a,a,a,b,a,b)
6	(a,b)	(b,b,a,b)	(a,a,a,b,a,b,a)
7	(b)	(b,b,a,b)	(a,a,a,b,a,b,a,a)
8	$\emptyset$	(b,b,a,b)	(a,a,a,b,a,b,a,a,b)
9	$\emptyset$	(b,a,b)	(a,a,a,b,a,b,a,a,b,b)
10	$\emptyset$	(a,b)	(a,a,a,b,a,b,a,a,b,b,b)
11	$\emptyset$	(b)	$(a,a,a,b,a,b,a,a,b,b,b,a)$

Table 3.3: Application of the Splicing Algorithm to $S \oplus T$

This example shows that adding a set of tail sequences T to a given input set S prompts the Splicing Algorithm to run for 12 stages and does not lead to a break up before stage 8 is reached. $\diamond$

The above example motivates the following Lemma:

Lemma 3. *Given an input set S w.r.t. n and P_2 and a set of tails T w.r.t. m and P_2, then for every initial value α_0, applying the Splicing Algorithm to $S \oplus T$ produces a sequence α of length at least $n+1$ (i.e., the algorithm does not break up before reaching stage n).*

Remark 8. It is not necessary to provide the proof of this Lemma here, since we are going to prove a variant of this Lemma at a later stage (which will also be essential for the proof of Theorem 2, whereas the result above won't). $\diamond$

First Amendment to General Approach

Let us now compare the results so far with Step 2 in the general approach. For a given integer $n \in \mathbb{N}$, a distribution P_2 with full support and an input $S = \left\{\alpha^{nP(a)} \in T_{nP(a)}(P_2(\cdot|a)) : a \in A\right\}$, our aim was to produce a sequence of length $n+1$, that has a second order type equal or close to P_2. The Splicing Algorithm indeed produces a sequence α^{n+1}, but only if we add a set of tails, T, to the input set S. Even though this addition of T preserves the first order type of the input sequences (observe, that for every $a \in A$, the input sequence $\alpha^a \oplus t^a$ has - by construction - first order type ρ^a), we have to extend the length of the input sequences from $nP(a)$ to $(n+m)P(a)$, for every $a \in A$. Therefore, the first amendment we have to make to our general approach applies to step 1:

Step 1: - First Amendment: Let $n, m \in \mathbb{N}$, then at the beginning of each block $k > 1$, we need to induce a set of sequences of the form

$$S \oplus T = \left\{\alpha^{nP(a)} \oplus t^{mP(a)} \in T_{(n+m)P(a)}(\rho^a) : a \in A\right\}$$

where S is an input set w.r.t. n and P_2 and T is a set of tails w.r.t. m and P_2.

So far, we have not yet analyzed the second order type of an output sequence of the Splicing Algorithm applied to an extended input set $S \oplus T$. We will see that, unfortunately, we cannot ensure that such an output sequence has a second order type equal or close to P_2.

Definition 8 (n-stage Output Sequence)**.** Let α be the output sequence of the Splicing Algorithm applied to $S \oplus T$ and an initial value α_0. Let α^{n+1} denote the first $n+1$ elements of α (which is the sequence produced after the nth stage in the algorithm). We call α^{n+1} the n-stage output sequence of the algorithm. Furthermore, the conditional subsequences of α^{n+1} (see Definition 1) are denoted by α_a^{n+1}, for every $a \in A$. ◇

Observation 2. The reader should observe, that the n-stage output sequence may have subsequences with first order types that differ widely from the desired conditional distributions $\{\rho^a : a \in A\}$. This can happen, since it is not guaranteed that all elements of the sequences in $S \oplus T$ are added to the output sequence, α^{n+1}. In particular, it may be the case that some input sequences $\alpha^a \oplus t^a$ in $S \oplus T$ are skewed in such a way, that the subsequence of $\alpha^a \oplus t^a$ that is eventually added to α^{n+1}, is not typical w.r.t.

ρ^a. ◇

The above Observation shows, that we need to impose more assumptions on the input sequences in $S \oplus T$ in order to produce an n-stage output sequence α^{n+1} with second order type P_2. In particular, we would like our input sequences in $S \oplus T$ to have a certain "locally typical" structure, s.t. for every $a \in A$, the conditional subsequence α_a^{n+1} of α^{n+1} - no matter its length - has a first order type close to ρ^a. The property we are looking for in the input sequences is called *Local Typicality* which will be the topic of the next section.

3.3.2 Local Typicality

Notation 8. As before, we set $A = I \times J \times K$ and denote by P_2 a distribution on A^2 with full support and with (identical) marginal distribution $P \in \Delta(A)$. Furthermore, for all $a \in A$, let $\{\rho^a : a \in A\}$ denote the conditional distributions of P_2. Note, that we don't require P_2 to have a non-empty second order type set anymore. ◇

Given a distribution P over a finite set A, locally typical sequences w.r.t. P not only have a first order type close to P, but also possess contiguous subsequences with first order type close to P. Locally typical sequences hence have the structure our input sequences in the previous section lacked so far. We will introduce a new input set of sequences that are locally typical w.r.t. the conditional distributions $\{\rho^a : a \in A\}$. We will then show that the Splicing Algorithm applied to this new input set produces an n-stage output sequence α^{n+1} that has a second order type close to P_2.

Recall that sequences with first order type close to a distribution P are called ε-typical sequences. Let us now properly define locally typical sequences:

Definition 9 (Local Typicality)**.** Let P be a distribution over a finite set A and let $N, l \in \mathbb{N}$. Given a sequence $x^N \in A^N$, denote by $x_{t,l}^N$ a subsequence of l successive elements of x^N starting at index t. For every $\varepsilon > 0$ we call x^N l-locally typical w.r.t. P, if for every $t \in \{1, 2, \ldots, N-l\}$, $x_{t,l}^N \in T_l^\varepsilon(P)$, where $T_l^\varepsilon(P)$ denotes the ε-typical set of length l w.r.t. P. We denote the set of all l-locally typical sequences by $T_{N,l}^\varepsilon(P)$. ◇

The following useful result states that if N is large enough, then the empirical distribution of an l-locally typical sequence x^N w.r.t. P is also close to P:

Lemma 4. *If $N > \frac{2l}{\varepsilon}$, then every l-locally typical sequence $x^N \in T^{\varepsilon}_{N,l}(P)$ is 2ε-typical w.r.t P, i.e.,*

$$|\tfrac{1}{N}N(a|x^N) - P(a)| < 2\varepsilon \quad \forall a \in A.$$

Proof. See Appendix C.2. □

Another important concept we will need in later sections is the concept of locally conditional typical sequences, defined below.

Definition 10 (Local Conditional Typicality)**.** Let $A = I \times K$ be a finite set and let $P \in \Delta(A)$ be a distribution with full support and with marginals P_I on I and P_K on K. Furthermore, for every $i \in I$, denote by ρ^i a conditional distribution on K derived from P, i.e., for every $k \in K$, $\rho^i(k) = P(k|i)$. Let $n \in \mathbb{N}$ and fix a locally typical sequence, $x^n \in T^{\varepsilon}_{n,l}(P_I)$. A sequence $y^n \in K^n$ is called locally conditional typical, if for every $t \in \{1, 2, \ldots, n-l\}$: $y^n_{t,l} \in T^{\varepsilon}_l(P|x^n)$, where again $y^n_{t,l}$ is a subsequence of y^n of l consecutive elements starting at index t and $T^{\varepsilon}_l(P|x^n)$ is the conditional ε-typical set w.r.t. P and x^n. The set of all locally conditional typical sequences is denoted by $T^{\varepsilon}_{n,l}(P|x^n)$. ◇

We now define a (new) set of input sequences for the Splicing Algorithm. The reader may observe that the idea of adding tail sequences, as in the previous paragraphs, is incorporated into this new input set.

Definition 11 (Locally Typical Input Set $\mathscr{S}$)**.** Fix $\varepsilon > 0$ and let $P_2 \in \Delta(A^2)$ be a distribution as in Notation 8. Let n, m, l be integers that satisfy the following properties:

$$nP(a), mP(a) \in \mathbb{N}, \forall a \in A \tag{3.1}$$

$$\min_{a \in A} nP(a) > \tfrac{2l}{\varepsilon} \tag{3.2}$$

$$\min_{a \in A} mP(a) > \tfrac{2l}{\varepsilon} \tag{3.3}$$

Setting $n_a = nP(A)$, $m_a = mP(a)$ and $r_a = n_a + m_a$, we call every set of l-locally typical sequences of the form

$$\mathscr{S} = \left\{ x^{r_a} \in T^{\varepsilon}_{r_a,l}(\rho^a) : a \in A \right\}$$

a locally typical input set w.r.t. n, m, l and P_2. ◇

Remark 9. Note, that the properties imposed on the parameters n, m and l in the above definition ensure that the input sequences in $\mathscr{S}$ are 2ε-typical as stated in Lemma 4. ◇

Similar to the previous section, we now want to show that the Splicing Algorithm applied to locally typical sequences produces a sequences of length at least $n+1$:

Lemma 5 (The Splicing Algorithm applied to Locally Typical Sequences)**.** *Let $\mathscr{S}$ be a locally typical input set w.r.t. n,m,l and P_2 and let $\alpha_0 \in A$. Fix $\varepsilon > 0$ and let $p_2 = \min_{(a,b)\in A^2} P_2(a,b)$. Then, if $\frac{n}{m} < \frac{p}{4\varepsilon} - 2$, then Splicing Algorithm applied to α_0 and $\mathscr{S}$ yields an output sequence of length at least $n+1$.*

The proof of this Lemma makes use of the following notation:

Notation 9. In order to ease readability, we call for every $a \in A$ an input sequence $x^{r_a} \in \mathscr{S}$ the row w.r.t a, or simply $row(a)$. Every row in S can be written as $row(a) = x^{n_a} \oplus x^{m_a}$. We call x^{n_a} the *body* and x^{m_a} the *tail* of $row(a)$. We summarise the bodies and the tails of the rows of $\mathscr{S}$ in two different sets:
the set of bodies is denoted by

$$B = \left\{x^{n_a} \in T^{\varepsilon}_{n_a,l}(\rho^a) : a \in A\right\},$$

and the set of tails by

$$T = \left\{x^{m_a} \in T^{\varepsilon}_{m_a,l}(\rho^a) : a \in A\right\},$$

so that

$$\mathscr{S} = B \oplus T.$$

Note, that by assumption, each body in B and each tail in T is also 2ε-typical.
Furthermore, for every pair $(a,b) \in A^2$, we say that an element a in $row(b)$ is a *reference* to $row(a)$ from $row(b)$. If $\alpha_0 = a$, then α_0 is also called a reference to $row(a)$. Note that the total number of references to $row(a)$ from $row(b)$ is given by $N(a|x^{r_b})$, and since x^{r_b} is 2ε-typical, we have

$$N(a|x^{r_b}) \leq r_b(\rho^b(a) + 2\varepsilon). \tag{3.4}$$

Similarly, an upper bound of the number of references to $row(a)$ from the tail of $row(b)$ is given by

$$N(a|x^{m_a}) \leq m_b(\rho^b(a) + 2\varepsilon). \tag{3.5}$$

$\diamond$

Proof of Lemma 5. Let α denote the output sequence of the algorithm. Towards a contradiction, assume that the algorithm stops before it has reached n stages, i.e., before it has added at least n elements to α. This means that for some $a \in A$, the algorithm has exhausted all elements of $row(a)$ (i.e., it has added the last element of $row(a)$ to the output sequence) at some stage k, and at a later stage $k' \geq k$ the algorithm terminated by adding element a to α and $k' < n$ (since $row(a)$ is exhausted, the algorithm cannot add anymore elements from $row(a)$ to α).

Now, let us count the maximal number of references to $row(a)$ in α. First, note that by assumption it is not possible that the Algorithm exhausted all elements of the bodies from each row of $\mathscr{S}$; since for every $a \in A$, the body of $row(a)$ has n_a elements, the algorithm would have otherwise added $\sum_{a \in A} n_a = n$ elements to the output sequence.

Therefore, we assume that there is at least one element $a^* \in A$, such that the body of the corresponding row has not been completely exhausted by the algorithm. This means that from $row(a^*)$ we can maximally count $N(a|x^{n_{a^*}})$ references to $row(a)$. From every other $row(b)$, $b \neq a^*$, there are maximally $N(a|x^{r_b})$ references to $row(a)$ present in the output sequence α.

Finally, recall that possibly α_0 is also a reference to $row(a)$. Therefore, the maximal number of references to $row(a)$ that can be counted in α (denoted by $\max ref(a)$), is given by

$$\begin{aligned}
\max ref(a) &\leq 1 + \sum_{a^* \neq b \in A} N(a|x^{r_b}) + N(a|x^{n_{a^*}}) \\
&= 1 + \sum_{b \in A} N(a|x^{r_b}) - N(a|x^{m_{a^*}}) \\
&\leq 1 + \sum_{b \in A} r_b(\rho^b(a) + 2\varepsilon) - m_{a^*}(\rho^{a^*}(a) - 2\varepsilon) \\
&= 1 + \sum_{b \in A} (n+m)P(b)(\rho^b(a) + 2\varepsilon) - mP(a^*)(\rho^{a^*}(a) - 2\varepsilon) \\
&= 1 + (n+m)\left(\sum_{b \in A} P_2(a,b) + \sum_{b \in A} P(b)2\varepsilon\right) - m(P_2(a,a^*) - P(a^*)2\varepsilon) \\
&= 1 + (n+m)(P(a) + 2\varepsilon) - m(P_2(a,a^*) - P(a^*)2\varepsilon),
\end{aligned}$$

where the second inequality is due to the inequalities in (3.4) and (3.5). By assumption, the maximal number of references to $row(a)$ must be larger than the number of elements of $row(a)$ (otherwise, the

Algorithm would not terminate as assumed). I.e., it must hold that

$$\begin{aligned}(n+m)P(a) &< \max ref(a)\\ &\leq 1+(n+m)(P(a)+2\varepsilon)-m(P_2(a,a^*)-P(a^*)2\varepsilon),\end{aligned}$$

which is equivalent to

$$mP_2(a,a^*)-1 < ((n+m)+mP(a^*))\,2\varepsilon$$

or

$$\frac{mP_2(a,a^*)-1}{2((n+m)+mP(a^*))} < \varepsilon. \tag{3.6}$$

However, inequality (3.6) cannot be verified if $\frac{n}{m} < \frac{p}{4\varepsilon} - 2$ holds, which we assume in the Lemma. We can therefore conclude, that the Splicing Algorithm applied to α_0 and $\mathscr{S}$ produces a sequence of length at least $n+1$. □

In the following section we can finally show that the Splicing Algorithm applied to locally typical sequences produces an output sequences (of length $n+1$) with a second order type close to P_2.

3.3.3 The Second Order Type of a Typical Output Sequence

Notation 10 (n-stage Typical Output Sequence)**.** Let $\mathscr{S}$ be a locally typical input set w.r.t. the integers n,m,l and P_2, that satisfy the conditions in Lemma 5. Let α denote the output sequence of the Splicing Algorithm applied to $\mathscr{S}$ and an initial value α_0 and denote by α^{n+1} the first $n+1$ elements of α. We call α^{n+1} the n-stage typical output sequence and for every $a \in A$ we denote by α_a^{n+1} the conditional subsequences of α^{n+1}. ◇

Remark 10. [The Difference of Rows and Conditional Subsequences] Given an n-stage typical output sequences, α^{n+1}, a conditional subsequence α_a^{n+1} of α^{n+1} is a prefix of $row(a)$ (an input sequence of the locally typical input set $\mathscr{S}$). More precisely, if the algorithm added *all* elements from $row(a)$ to α^{n+1}, then $row(a) = \alpha_a^{n+1}$. Otherwise, α_a^{n+1} is equal to the first elements in $row(a)$ that were added to α^{n+1}. ◇

The last Remark raises the question of how many elements from each row of a locally typical input set $\mathscr{S}$ are at least added to an n-stage typical output sequences α^{n+1}, which we discuss in the following.

Lemma 6. *Let $\mathscr{S}$ be a locally typical input set w.r.t. n, m, l and P_2. For every $a \in A$, the length of a conditional subsequence α_a^{n+1} of an n-stage typical output sequence α^{n+1} is given by $N(a|\alpha^{n+1})$. If $\frac{m}{n} > \min_a P(a)$, then*

$$N(a|\alpha^{n+1}) \geq (n+m)P(a) - m.$$

Proof. Given an element $a^* \in A$, let's assume that all but one $row(a^*)$ are exhausted in the n-stage typical output sequence α^{n+1}, i.e., for every $a \neq a^*$, $row(a)$ is the conditional subsequence of α^{n+1}, $row(a) = \alpha_a^{n+1}$. Note that this situation describes the most extreme case in which the number of elements from $row(a^*)$ that are added to the output sequence is minimal. In any other situation, we would have at least two rows ($row(a^*)$ inclusive) that are not entirely exhausted in the output sequence. In these cases the number of elements added from $row(a^*)$ to the output sequence would only but increase, compared to the situation first described. Hence, assuming that all but one $row(a^*)$ are exhausted in the output sequence is sufficient for our analysis. Every $row(a)$ has length $(n+m)P(a)$, therefore, the following must hold for the length of α^{n+1}

$$n+1 = \sum_{a \in A, a \neq a^*} (n+m)P(a) + N(a^*|\alpha^{n+1}) + 1,$$

equivalently,

$$N(a^*|\alpha^{n+1}) = (n+m)P(a) - m.$$

Since $\frac{m}{n} > \min_a P(a)$, we have $N(a^*|\alpha^{n+1}) > 0$. □

Observation 3. [Conditional Subsequences are Typical] Observe, that for an n-stage typical output sequence α^{n+1} that corresponds to a locally typical input set $\mathscr{S}$ w.r.t. n, m, l and P_2, if $\frac{m}{n} > \min_a P(a)$, then it holds that for every $a \in A$, $(n+m)P(a) - m > mP(a)$. Furthermore, since $mP(a) > \frac{2l}{\varepsilon}$ by Definition 11, it follows from Lemma 6 and Lemma 4 that every conditional subsequence α_a^{n+1} of α^{n+1} has a first order type close to ρ^a, $\forall a \in A$. This is an important property of the conditional subsequences of an n-stage typical output sequence. Recall, that this property was not inherent in the conditional subsequences of an n-stage output sequence in the previous section, which is why we introduced locally typical sequences. ◇

With the above observation, we can now finally conclude this section and show that n-stage typical output sequences have a second order type that is close to P_2.

Theorem 3. *Let α^{n+1} be an n-stage typical output sequence of the Splicing Algorithm applied to a locally typical input set $\mathscr{S}$ w.r.t. n,m,l and P_2. Then, for all $\varepsilon > 0$, it holds that*

$$\left\|emp_2(\alpha^{n+1}) - P_2\right\|_{TV} < \delta(\varepsilon),$$

where $\|\cdot\|_{TV}$ denotes the total variation distance[1] *and $\delta(\varepsilon)$ is a function with $\delta(\varepsilon) \to 0$ as $\varepsilon \to 0$.*

In order to prove the above Theorem, we first use the fact that the function that maps a strictly positive stochastic matrix to its (unique) stationary distribution is continuous (see Lemma below). This function helps bounding the distance of the the stationary distribution of P_2 and $emp_2(\alpha^{n+1})$. Furthermore, as noted in Observation 3, the empirical distributions of the conditional subsequences of α^{n+1} are 2ε- close to the conditional probabilities of P_2. These two relations will then be used to show that the distance of $emp_2(\alpha^n)$ and P_2 is close.

Lemma 7. *Let A be a finite set and denote by $\mathscr{M}$ the set of all strictly positive stochastic matrices over A. The mapping $f : \mathscr{M} \to \Delta(A)$, with $f(\mathbb{P}) = \pi$, where $\mathbb{P} \in \mathscr{M}$ with (unique) stationary distribution π, is continuous.*

Remark 11. Note that the function f is well-defined since f maps every strictly positive stochastic matrix to its unique stationary distribution. $\diamond$

The proof of the above Lemma is added in the Appendix.

Proof of Theorem (3). Let $\mathbb{P}$ be the stochastic matrix derived from P_2 where the rows display the conditional distributions $\{\rho^a : a \in A\}$. We denote the entries of $\mathbb{P}$ by $p(a,b) = \rho^a(b)$, for every $a,b \in A^2$. Furthermore, denote by $\tilde{\mathbb{P}}$ another stochastic matrix that we derive from the n-stage typical output sequence α^{n+1}: the rows of $\tilde{\mathbb{P}}$ represent the empirical distributions of the conditional subsequences $\left\{\alpha_a^{n+1} : a \in A\right\}$, hence the entries of $\tilde{\mathbb{P}}$ are denoted by $\tilde{p}(a,b) = emp(\alpha_a^{n+1})[b]$ for every pair $(a,b) \in A^2$. As seen in Observation 3, the conditional subsequences of α^{n+1} are 2ε close to the conditional distributions of P_2, i.e.,

$$|emp(\alpha_a^{n+1}) - \rho^a| < 2\varepsilon,$$

or, in matrix notation:

$$|p(a,b) - \tilde{p}(a,b)| < 2\varepsilon.$$

[1]Given two probability distributions $p,q \in \Delta(A)$, the **total variation distance** between p and q, denoted $\|p-q\|_{TV}$, is given by $\max_{a\in A} |p(a) - q(a)|$

Let us now derive the stationary distributions of $\mathbb{P}$ and $\tilde{\mathbb{P}}$:

First, the probability vector π with $\pi(a) = P(a) = \sum_{b \in A} P_2(a,b)$, $\forall a \in A$ is obviously the unique stationary distribution of $\mathbb{P}$.

Denote by $\tilde{\pi}$ the stationary distribution of $\tilde{\mathbb{P}}$. Let $\alpha^n_{(1)}$ denote the sequence of α^{n+1} reduced by the last element, α_n, and let $\alpha^n_{(2)}$ be the sequence of α^{n+1} reduced by the first element α_0. Assuming n large, we set $N(a|\alpha^n_{(1)}) = N(a|\alpha^n_{(2)})$ for every $a \in A$. Setting $\tilde{\pi}(b) = \frac{1}{n} N(b|\alpha^n_{(2)})$ for every $b \in A$, yields $\tilde{\pi} = \tilde{\pi}\tilde{\mathbb{P}}$ as shown below:

$$\begin{aligned}
\tilde{\pi}\tilde{\mathbb{P}}(b) &= \sum_{a \in A} \frac{1}{n} N(a|\alpha^n_{(2)}) \tilde{p}(a,b) \\
&= \frac{1}{n} \sum_{a \in A} N(a|\alpha^n_{(2)}) emp(\alpha^{n+1}_a)[b] \\
&= \frac{1}{n} \sum_{a \in A} N(a|\alpha^n_{(2)}) \frac{N(b|\alpha^{n+1}_a)}{N(a|\alpha^n_{(1)})} \\
&= \frac{1}{n} \sum_{a \in A} N(b|\alpha^{n+1}_a) \\
&= \frac{1}{n} N(b|\alpha^n_{(2)}) \\
&= \tilde{\pi}(b).
\end{aligned}$$

Therefore, for all $(a,b) \in A^2$, we have

$$\begin{aligned}
P_2(a,b) &= \pi(a)p(a,b), \text{ and} \\
emp_2(\alpha^{n+1})[a,b] &= \tilde{\pi}(a)\tilde{p}(a,b).
\end{aligned}$$

Now, by Lemma 21, with sufficiently small ε, we have for all $\delta > 0$ and for all $(a,b) \in A^2$,

$$|p(a,b) - \tilde{p}(a,b)| < 2\varepsilon \Rightarrow |\pi(a) - \tilde{\pi}(a)| < \delta,$$

and thus

$$\begin{aligned}
emp_2(\alpha^{n+1})[a,b] - P_2(a,b) &= \tilde{\pi}(a)\tilde{p}(a,b) - \pi(a)p(a,b) && (3.7) \\
&< (\pi(a) + \delta)(p(a,b) + 2\varepsilon) - \pi(a)p(a,b) && (3.8) \\
&< 2\varepsilon + \delta + 2\delta\varepsilon && (3.9)
\end{aligned}$$

similarly,

$$P_2(a,b) - emp_2(\alpha^{n+1})[a,b] < 2\varepsilon + \delta + 2\delta\varepsilon.$$

Letting $\delta = \varepsilon$ completes the proof. □

Second Amendment to General Approach

With the application of the Splicing Algorithm on locally typical sequences we are able to achieve an output sequence of length at least $n+1$ that also has a second order type close to P_2, as desired. Step 2 in the general approach can therefore be achieved, if we make the necessary changes in Step 1 of the general approach:

Step 1 - Second Amendment Fix $\varepsilon > 0$ and let n, m, l be integers with $\frac{1}{p} < \frac{n}{m} < \frac{p_2}{4\varepsilon} - 2$, where $p = \min_a P(a)$ and $p_2 = min_{a,b} P_2(a,b)$. At the beginning of each block $k > 1$ we need to induce a locally typical input set $\mathscr{S}$ w.r.t. n, m, l and P_2.

In the following sections, we will therefore focus on this amended Step 1.

3.4 Sequence Induction

In the previous sections, we developed *which* sequences need to be induced, so that they can be spliced to produce a sequence with a second order type that is close to P_2. In these upcoming sections, we show *how* these sequences can be induced, i.e., how the strategies of the players need to be constructed to induce these sequences.

In the following, we will refer to a distribution P_2 with properties stated below. Even though this notation has been used in the previous sections, we only add additional properties to (the former) P_2, and hence the results in the previous sections also hold for this distribution.

Notation 11. Let $\varepsilon > 0$. Denote by P_2 a distribution over A^2, $A = I \times J \times K$, with full support that satisfies the properties stated in Theorem 2. That is,

- The marginal of P_2 on the first and the second coordinates is identical and given by $P \in \Delta(A)$.

- For every $a \in A$, the marginal of the conditional distribution $P_2(\cdot|a)$ on I is μ.
- P_2 fulfills the following information constraint: $H_{P_2}(\mathbf{i},\mathbf{j}|\mathbf{k},\mathbf{i}',\mathbf{j}',\mathbf{k}') - H_{P_2}(\mathbf{i}) \geq \varepsilon$, where $(\mathbf{i},\mathbf{j},\mathbf{k},\mathbf{i}',\mathbf{j}',\mathbf{k}') \sim P_2$.

Furthermore, we denote by $\{\rho^a : a \in A\}$ the conditional distributions of P_2, i.e., for all $a \in A$, $P_2(\cdot|a) = \rho^a$. The marginal distribution of ρ^a on I, K, $I \times K$, etc., are denoted by ρ^a_I, ρ^a_K, $\rho^a_{I\times K}$, respectively. ◇

The strategies we are going to construct will implement a distribution P_2 as denoted above.

Remark 12. Recall, that by Lemma 16 in the Appendix, for a given distribution $P_2' \in \Delta((I \times J \times K)^2)$ which satisfies the properties stated in Theorem 2, for every $\varepsilon > 0$ there exists a distribution $P_2 \in \Delta((I \times J \times K)^2)$ with full support and equal properties as P_2', such that $\|P_2' - P_2\| < 2\varepsilon$. If we can show that such a distribution P_2 is implementable, then P_2' is also implementable, since the set of implementable distributions is closed by Remark 25. ◇

3.4.1 Strategy Outline

In the following paragraphs we only provide an outline of the strategies that will implement P_2. The details of the strategy construction will be discussed in the subsequent section.

First of all, we divide the game that is induced by the strategies of the players and by the actions of nature, into blocks of a given length n. We associate to every such block k an input set, denoted by $\mathscr{S}[k]$. Such an input set consists of an input set for nature, $\mathscr{S}_I[k]$, forecaster, $\mathscr{S}_J[k]$ and agent, $\mathscr{S}_I[k]$, such that

$$\mathscr{S}[k] = \mathscr{S}_I[k] \times \mathscr{S}_J[k] \times \mathscr{S}_K[k].$$

Denote by $\alpha^n[k] = (x[k], y[k], z[k]) \in (I \times J \times K)^n$ the actual elements played by the players in block k. We want this sequence of action triples to have a second order type close to P_2. We can accomplish this with the following approach:

We construct strategies, such that for given integers $n, m, l \in \mathbb{N}$ and $\varepsilon > 0$, in asymptotically almost every block $k > 1$, an input set $\mathscr{S}[k]$ is induced that is locally typical w.r.t. n, m, l and P_2 according to Definition 11. Once this is accomplished, we run the Splicing Algorithm on $\mathscr{S}[k]$. By Lemma 5, this produces a typical n-stage output sequence for block k of length $n+1$, which we denote by

$$\alpha^{n+1}[k] = (\alpha_0[k], \alpha_1[k], \ldots, \alpha_n[k]).$$

We set $\alpha_0[k] = \alpha_n[k-1]$, i.e., $\alpha_0[k]$ is the last element of the sequence in block $k-1$. For every $j \in [0, \ldots n]$, we write $\alpha_j[k] = (x_j[k], y_j[k], z_j[k]) \in I \times J \times K$ to denote the actions of the individual players in stage j in block k.

Recall, that by Theorem 3, the second order type of $\alpha^{n+1}[k]$ gets arbitrarily close to P_2 in the total variation norm as $n \to \infty$.

For the above approach to work, the sequences in $\mathscr{S}[k]$ need to satisfy certain requirements. First of all, the parameters $n, m, l \in \mathbb{N}$ and ε need to be chosen such that they satisfy the properties in Definition 11 and in Lemma 5. Moreover, the sequences in the input set $\mathscr{S}[k]$ must be ε-typical w.r.t. the conditional distributions of P_2, $\{\rho^a : a \in A\}$. In order to achieve this, we impose the following conditions on the input sets of nature, forecaster and agent:

Conditions on the Input Sets

Let $\varepsilon > 0$. For every $a \in A$, setting $r_a = (n+m)P(a)$, the following conditions must hold on the input set of nature, of the forecaster and of the agent:

$$\mathscr{S}_I[k] \ni x^{r_a} \in T^{\varepsilon}_{r_a,l}(\rho^a_I) \tag{3.10}$$

$$\mathscr{S}_K[k] \ni z^{r_a} \in T^{\varepsilon}_{r_a,l}(\rho^a_{I\times K}|x^{r_a}) \text{ for } x^{r_a} \in \mathscr{S}_I[k] \tag{3.11}$$

$$\mathscr{S}_J[k] \ni y^{r_a} \in T^{\varepsilon}_{r_a,l}(\rho^a|x^{r_a}, z^{r_a}) \text{ for } x^{r_a} \in \mathscr{S}_I[k], z^{r_a} \in \mathscr{S}_K[k]. \tag{3.12}$$

Note that the input sequences of the forecaster are assumed to be locally conditional typical sequences according to Definition (10), and so it must hold that for every sequence $(x^{r_a}, y^{r_a}, z^{r_a}) \in (\mathscr{S}_I[k] \times \mathscr{S}_J[k] \times \mathscr{S}_K[k])$,

$$(x^{r_a}, y^{r_a}, z^{r_a}) \in T^{\varepsilon}_{r_a,l}(\rho^a).$$

The reader should observe that the concept of local conditional typicality keeps ε fix, i.e., there is no need for any change on ε. Therefore, the input set $\mathscr{S}[k] = (\mathscr{S}_I[k] \times \mathscr{S}_J[k] \times \mathscr{S}_K[k])$ that satisfies the conditions in (3.10), (3.11) and in (3.12) indeed satisfies all properties of a locally typical input set according to Definition 11. Moreover, by Lemma 5, there is an n-stage output sequence $\alpha^{n+1}[k]$ of the Splicing Algorithm applied to $\mathscr{S}[k]$, that has a second order type close to P_2 according to Theorem 3.

The strategies we will construct are designed such that the input set of the agent and of the forecaster, $\mathscr{S}_K[k]$ and $\mathscr{S}_J[k]$, have the properties stated in (3.11) and (3.12). Note, that the input set of nature may not always have the desired property in (3.10).

Remark 13. [Nature's Input Set] The sequences in $\mathscr{S}_I[k]$ are segments from nature's conditional sequences which are fixed from the beginning of the play. Obviously, we cannot guarantee that nature will play $x^{r_a} \in T^{\varepsilon}_{r_a,l}(\rho^a_I)$ for every $a \in A$, as in (3.10). However, we know that this happens with high frequency throughout the game by Lemma 18 in the Appendix. We will later discuss what we do in case nature does not play as in (3.10). $\diamond$

So far, we have outlined what the strategies of the players need to achieve, namely the induction of an input set $\mathscr{S}[k]$ w.r.t. n,m,l and P_2. As announced, we will provide the detailed construction of the strategies in the upcoming section. In the remaining part of this section, we take a step ahead and present a simplified overview of how the players act under the strategies we are about to construct. The purpose of this proceeding is to introduce the key features of the strategies in an informal way. Thereby, the reader is given an idea of the strategies, before we immerse into the details of their construction.

How the Players Act

In the following, let us assume that a play is initiated by the strategies of the players, that can induce a locally typical input set $\mathscr{S}[k]$ with properties stated in (3.10),(3.11) and (3.12) in every block (note, that this is an idealised assumption, since by Remark 13, property (3.10) may not always be satisfied).

First, recall that at the beginning of the play, the agent has no knowledge about future states of nature, nor of the future play of the forecaster, whereas the forecaster has complete knowledge about all future states of nature's conditional sequences. Note, however, that the agent (as well as the forecaster) can observe the complete history of the play at any stage in the game.

In order to know her input set $\mathscr{S}_K[k]$ at the beginning of a block $k > 1$, the agent has to rely on her observation of the past play. Since we want $\mathscr{S}_K[k]$ to be chosen in dependence of $\mathscr{S}_I[k]$ as in (3.11), it is the task of the forecaster to send the necessary information in block $k-1$ to the agent. More precisely, the forecaster conditions her actions in block $k-1$ on nature's actions in block k in such a way, that it will transmit sufficient information to the agent to know $\mathscr{S}_K[k]$ at the beginning of the block. At the same time, the forecaster also conditions her actions in $\mathscr{S}_J[k-1]$ on $\mathscr{S}_I[k-1]$ and on $\mathscr{S}_K[k-1]$ as indicated in (3.12), such that the input set $\mathscr{S}[k-1]$ is locally typical as desired.

We refer to *information transmission* as the process of sending and receiving information about future states of nature. More precisely, we say that information transmission from block $k-1$ to block k from the forecaster to the agent is possible, if the forecaster is able to send enough information to the agent in block $k-1$, so that the agent knows $\mathscr{S}_K[k]$ at the beginning of block k, while at the same time the forecaster can match nature's and agent's actions in block $k-1$, so that the input set $\mathscr{S}[k-1]$ is induced. Hence, information transmission is the key feature of the strategies of the players. Let us now shortly demonstrate that in order for the information transmission to work, the information constraint plays a vital role.
First, we will see that the information transmission process is formalised with the introduction of a *Message Set* of the forecaster and a *Set of Action Plans* of the agent, which we will properly define in the upcoming sections. Next, we show that the information transmission process can only work, if the message set is larger in size than the set of action plans. Only then it is possible for the forecaster to send enough information about future states of nature to the agent, so that the agent can deduce the desired actions. Finally, we will see that the information constraint can guarantee that the relative sizes of the message set and the set of action plans are as required.
In the next paragraphs, we intend to provide a first idea of the functioning of a message and of an action plan.

A message of the forecaster in block $k-1$ is a subset of the input set of the forecaster, $\mathscr{S}_J[k-1]$. It has the property that every element in the message appears in the output sequence, $\alpha^{n+1}[k-1]$, so that at the beginning of the next block, k, the agent is able to observe the message. In every block $k-1$, $k > 1$, the forecaster sends a message to the agent, that entails information about a so-called *Hypothetical Input Set* *of Nature* of block k, which we denote by $\mathscr{I}[k]$. This set consists of those segments of the conditional sequences of nature, that contain the sequences of the actual input set of nature, $\mathscr{S}_I[k]$. More precisely, every sequence in $\mathscr{S}_I[k]$ is a middle segment of a sequence in $\mathscr{I}[k]$. Hence, the forecaster sends more information in her message to the agent, than what will actually be needed. Let us illustrate this idea with the two figures below:

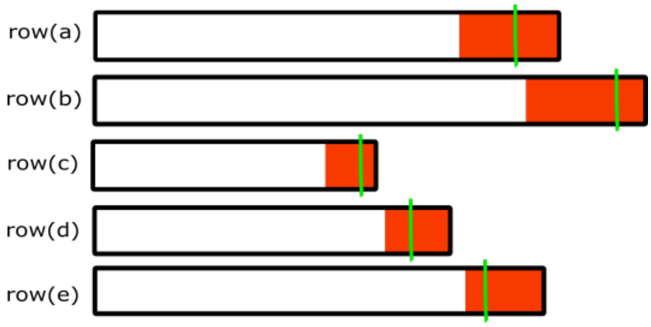

Figure 3.1: Input Sequences

Every row in Figure 3.1 represents an input sequence of the locally typical input set $\mathscr{S}[k-1]$. In every such row we add an orange suffix and a green pointer. To understand these additions, recall that when applying the Splicing Algorithm to a locally typical input set, not necessarily every element of an input sequence is eventually added to the output sequence. However, Lemma 6 informs us of the minimal number of elements from an input sequence that are added to the output sequence ($(n+m)P(a)-m$). Now, the orange suffix marks those elements that are potentially not added to the output sequence, i.e., there are exactly $(n+m)P(a)-m$ elements to the left of the orange suffix. Now, the the green pointer can only move in the orange suffix and indicates exactly the number of elements that are added to the output sequence, i.e., all elements to the left of the green pointer. The elements to the right of the pointer do not appear in the output sequence of block $k-1$. The green pointers also marks the beginning of the next block, k. I.e., it marks the starting points for the input sequences of $\mathscr{S}[k]$.

Figure 3.2 focuses on one single row from Figure 3.1 (with two additional dashed boxes added) and demonstrates the procedure of the Information Transmission. As before, the row represents an input sequence from the input set $\mathscr{S}[k-1]$, this time it is divided into the three sequences of nature, forecaster and agent (in that order). Here, we can see, that the message (i.e., an element of the message) of the forecaster is placed before the orange suffix. It includes information about the hypothetical input set of nature (a sequence of the same), which begins right after the minimal length of the sequence of the input set of nature in $\mathscr{S}[k-1]$ (marked as a blue dashed box). Then, at the beginning of block k (indicated

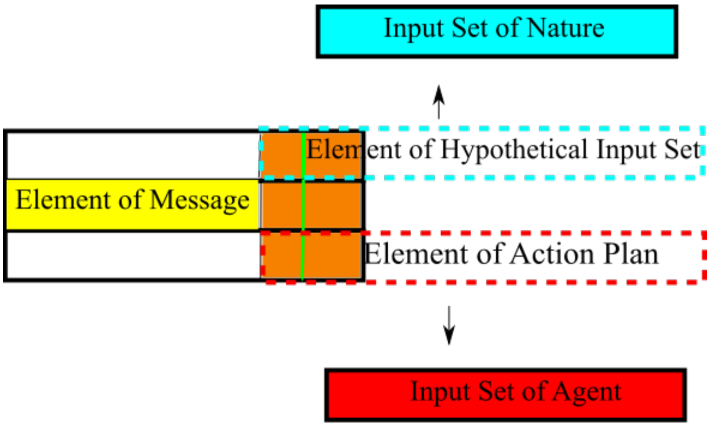

Figure 3.2: Information Transmission

by the green pointers), the agent observes the message of the forecaster in block $k-1$, as well as the complete history of the play (not included in the figures above). From that, she deduces a so-called action plan, which is a set of sequences that contains her input set $\mathscr{S}_K[k]$ for block k. In Figure 3.2 we marked one sequence of the action plan (the red dashed box). The action plan has the property that it matches the hypothetical input set of nature in the same way as the agent's input set $\mathscr{S}_K[k]$ matches nature's set $\mathscr{S}_I[k]$ as stated in (3.11). With the help of her observation of the past play, the agent is then able to extract her input set from her action plan.

This outline of how the players act is certainly not exhaustive, as indicated before. The following section will now add the details to this mechanism.

3.4.2 Information Transmission

In this section we look at the features of the strategies introduced in the previous passages in more depths. In particular, we define the hypothetical input set of nature, as well as the the message set of the forecaster and the set of action plans of the agent. We begin with the definition of a set of parameters that we are going to employ.

Definition 12. [Fit Parameters] Let $A = I \times J \times K$ and let $\varepsilon > 0$. Let P_2 be a distribution over A^2 as given in Notation 11. We call the set $\{P_2, \varepsilon, n, m, l\}$ a set of fit parameters, if the following holds:

- $\min_{a \in A} nP(a) > \frac{2l}{\varepsilon}$
- $\min_{a \in A} mP(a) > \frac{2l}{\varepsilon}$

- $l > \frac{\ln(2n|A|)}{2\varepsilon^2}$
- $\frac{1}{p} < \frac{n}{m} < \frac{p_2}{4\varepsilon} - 2$
  where $p = \min_{a \in A} P(a)$ and $p_2 = \min_{a,b \in A^2} P_2(a,b)$.

$\diamond$

Definition 13. We say that a set of fit parameters $\{P_2, \varepsilon, n, m, l\}$ behaves asymptotically appropriately, if there exists a function $v(\varepsilon)$ with $v(\varepsilon) \to 0$ as $\varepsilon \to 0$ and

$$\tfrac{m}{n} < v(\varepsilon).$$

$\diamond$

Remark 14. Note, that a set of fit parameters as given in Definition 12 obviously exists whenever n, m are assumed to be sufficiently large. Furthermore, one can find a set of fit parameters, $\{P_2, \varepsilon, n, m, l\}$, that behaves asymptotically appropriately with $v(\varepsilon) = \gamma\varepsilon$, where $\gamma \in \mathbb{R}$ is sufficiently large ($\gamma > \frac{4}{p_2 - 8\varepsilon}$). $\diamond$

Next, we need to introduce some notation regarding the structure of a play:

Notation 12. We say that a game induced by the strategies of the players and nature is an induced play. In an induced play, denote by α^{k-1} the entire past sequence of actions played up until the beginning of (and not including) block k. Furthermore, denote by $N(a|\alpha^{k-1})$ the number of occurrences of a in α^{k-1}. $\diamond$

In the following sections, we will often refer to an induced play, without having specified the strategies of the players yet. The reader should observe, however, that in doing so, we do not presuppose certain features of the strategies and only need to describe the correct setting, that is independent of the strategies.

The Hypothetical Input Set of Nature

Nature's Conditional Sequences: First, recall that at the beginning of a game, nature draws for every $a \in A$ and according to μ, an infinite conditional sequence $(x_1^a, x_2^a, \ldots)$ of states, which the forecaster observes. Remember, that if at stage t in the course of a game, element a is played, i.e., $(x_t, y_t, z_t) = a$, then nature plays $x_{t+1} = x_i^a$, if a has occurred i times in the history of the game up until stage t.

Remark 15. At the beginning of each block $k \geq 1$ in an induced play, the forecaster can observe the entire past sequence of actions played by nature, forecaster and agent, α^{k-1}. For every $a \in A$, she can also work out $N(a|\alpha^{k-1})$. Therefore, the forecaster knows nature's input set in block k, $\mathscr{S}_I[k]$, since she can observe all the elements nature will play after the upcoming occurrences of elements $a \in A$, i.e., forecaster observes

$$\mathscr{S}_I[k] \ni x^{r_a}[k] = (x_1[k], \ldots, x_{r_a}[k]) = (x^a_{N(a|\alpha^{k-1})}, x^a_{N(a|\alpha^{k-1})+1}, \ldots, x^a_{N(a|\alpha^{k-1})+r_a-1}).$$

◇

To be precise, the forecaster not only knows $\mathscr{S}_I[k]$, but the entire future conditional sequences of nature. However, this is not what she will transmit to the agent (this is obviously too much information). The forecaster restricts her knowledge that she transmits to the so-called hypothetical input set of nature:

Definition 14 (The Hypothetical Input Set of Nature)**.** Let $\{P_2, \varepsilon, n, m, l\}$ be a set of fit parameters. Let k be a block in an induced play, such that $\mathscr{S}[k]$ is a locally typical input set w.r.t. n, m, l and P_2. The hypothetical input set of nature of block $k+1$, denoted by $\mathscr{I}[k+1]$, is a set of $|A|$ segments of nature's conditional sequences $\{x^a : a \in A\}$, with the property that for every $a \in A$, a segment has length $r_a + m$ and is locally typical w.r.t. μ:

$$\mathscr{S}[k+1] = \left\{\iota^{r_a+m}[k+1] \in T^{\varepsilon}_{r_a+m,l}(\mu) : a \in A\right\},$$

where for every $a \in A$, the segment $\iota^{r_a+m}[k+1] \in \mathscr{I}[k+1]$ has the following start and end points:

$$\iota^{r_a+m}[k+1] = \left(x^a_{\min N(a|\alpha^k)}, \ldots, x^a_{\max N(a|\alpha^k)+r_a}\right),$$

where $\min N(a|\alpha^k) = N(a|\alpha^{k-1}) + r_a - m$ and $\max N(a|\alpha^k) = N(a|\alpha^{k-1}) + r_a$. ◇

It should be observed that the hypothetical input set of nature contains locally typical sequences w.r.t. μ. Again, it is not guaranteed that nature's conditional sequences are of these types, but if they are, it will be this set which the forecaster informs the agent about. The following remark provides some motivation and further explanation for the introduction of a hypothetical input set.

Remark 16. Let $\alpha^{n+1}[k]$ be the n-stage output sequence of the Splicing Algorithm applied to a locally typical input set $\mathscr{S}[k]$. For every $a \in A$, the length of a subsequence $\alpha^{n+1}_a[k]$ of $\alpha^{n+1}[k]$, denoted by

$L(\alpha_a^{n+1}[k])$, is given by

$$r_a - m \leq L(\alpha_a^{n+1}[k]) \leq r_a, \text{ with } r_a = (n+m)P(a), \tag{3.13}$$

by Lemma 6 (recall, that $L(\alpha_a^{n+1}[k]) = N(a|\alpha^{n+1}[k])$ for all $a \in A$). As noted in Remark 10, we can see here again that a conditional subsequence $\alpha_a^{n+1}[k]$ of α^{n+1} is not necessarily equal to the respective input sequence $\alpha^{r_a}[k] \in \mathscr{S}[k]$. $\alpha_a^{n+1}[k]$ is only a prefix of $\alpha^{r_a}[k]$. Hence, since the lengths of the subsequences $\left\{\alpha_a^{n+1}[k] : a \in A\right\}$ can vary in between the interval (3.13), and since these lengths (equivalently, the numbers $N(a|\alpha^{n+1}[k])$, for $a \in A$) determine nature's actions in block $k+1$, we introduced the hypothetical input set of nature, that takes into account the shortest possible lengths of the conditional subsequences of the output sequence in block k. Finally, note, that obviously the input set of nature is a subset of the hypothetical input set of nature in block $k+1$, i.e., $\mathscr{S}_I[k+1] \subset \mathscr{I}[k+1]$. ◇

Observation 4 (Elements Added and Elements Discarded)**.** Since not all elements from each sequence in $\mathscr{S}[k]$ are added to the output sequence $\alpha^{n+1}[k]$, the question arises of what to do with unused elements, i.e., elements of sequences in $\mathscr{S}[k]$ that are not added to $\alpha^{n+1}[k]$. Given the bounds of a conditional subsequence (3.13), there may be up to m unused, or unadded elements of every locally typical input sequence $\alpha^{r_a}[k] \in \mathscr{S}[k]$. While we can simply discard the unused elements of forecaster and agent, and then start anew in a locally typical input set in the next block, we cannot simply discard the unused elements of nature. In other words, unused elements of nature in the input sequences in $\mathscr{S}[k]$ have to reappear in the input set in the next block, $\mathscr{S}[k+1]$, while unused elements of forecaster and agent can be discarded. ◇

The Set of Action Plans of the Agent

Let us now look at the concept of an action plan of the agent in more detail. As indicated in the outline of the strategies, the action plan is a set of $|A|$ sequences of length $r_a + m$ for every $a \in A$, that matches a given hypothetical input set of nature, $\mathscr{I}[k]$ in block k (that is only observed by the forecaster in the previous block) in the following sense. If κ^{r_a+m} denotes a sequence in the action plan, then for every $a \in A$, it holds that

$$(\iota^{r_a+m}[k], \kappa^{r_a+m}) \in T^{\varepsilon}_{r_a+m,l}(\rho^a_{I\times K}), \text{ for every } a \in A,$$

where $\iota^{r_a+m}[k] \in \mathscr{I}[k]$. The action plan can hence be regarded as an extension of the input set of the agent for block k. In fact, an important feature of the action plan is that upon observing the past play until block k, the agent is able to extract her input set $\mathscr{S}_K[k]$ from the action plan. Let us now properly define the set of action plans for the agent. The reader may observe the similarity to the Definition of the set of action plans in [6].

Definition 15 (The Set of Action Plans of the Agent)**.** Let $\{n,m,l,\varepsilon,P_2\}$ be a set of fit parameters. Given $a \in A$, we say that the set of action plans w.r.t. a, denoted by AP^a, is a subset of $T^{\varepsilon}_{r_a+m,l}(\rho^a_K)$ of minimal size, s.t. for all $x \in T^{\varepsilon}_{r_a+m,l}(\rho^a_I)$, there exists an element $z \in AP^a$, with $(x,z) \in T^{\varepsilon}_{r_a,l}(\rho^a_{I\times K})$. The (total) set of action plans is given by AP, with

$$AP = \bigtimes_{a\in A} AP^a.$$

Elements of AP will be denoted by $\mathscr{K}$. ◇

The size of the set of action plans for the agent plays an important role in the construction of the strategies, which is stated below.

Lemma 8. *[The Size of the Set of Action Plans] Let $\{n,m,l,\varepsilon,P_2\}$ be a set of asymptotically fit parameters. Let $\boldsymbol{a}$ denote a random triple in $A = I\times J\times K$, and let $(\boldsymbol{i},\boldsymbol{j},\boldsymbol{k},\boldsymbol{a})$ denote two random triples in A^2 with $(\boldsymbol{i},\boldsymbol{j},\boldsymbol{k},\boldsymbol{a}) \sim P_2$. Further, let $\lambda(\varepsilon)$ be a function of ε with $\lambda(\varepsilon)\to 0$ as $\varepsilon \to 0$. Then,*

$$|AP| \le 2^{(1+\lambda(\varepsilon))n(H(\boldsymbol{k}|\boldsymbol{a})-H(\boldsymbol{k}|\boldsymbol{i},\boldsymbol{a})+\xi(\varepsilon))},$$

for $\xi(\varepsilon) = 3\tilde{c}\varepsilon$, with $\tilde{c} = -\sum_{k\in K}\log P_K(k)$.

Proof. Let $(\mathbf{i}^a,\mathbf{k}^a)$ denote a pair of random variables distributed according to $\rho^a_{I\times K}$. By Theorem 7 in the Appendix, for every $a\in A$ and for every $\xi > 2\tilde{c}\varepsilon$, there exists a set of action plans w.r.t element a of size

$$|AP^a| = 2^{n_a(1+\lambda_a)(H(\mathbf{k}^a)-H(\mathbf{k}^a|\mathbf{i}^a)+\xi)},$$

with $\lambda_a = \frac{m}{n}(1+\frac{1}{P(a)})$. Since $\frac{m}{n} < \nu(\varepsilon)$, we have $\lambda_a < \nu(\varepsilon)(1+\frac{1}{p})$, with $p = \min_{a\in A} P(a)$. Setting $\xi = \xi(\varepsilon) = 3\tilde{c}\varepsilon$,

$$\begin{aligned}|AP^a| &= 2^{n_a(1+\lambda_a)(H(\mathbf{k}^a)-H(\mathbf{k}^a|\mathbf{i}^a)+\xi(\varepsilon))}\\ &< 2^{n_a(1+\lambda(\varepsilon))(H(\mathbf{k}^a)-H(\mathbf{k}^a|\mathbf{i}^a)+\xi(\varepsilon))}.\end{aligned}$$

Since $AP = \times_{a \in A} AP^a$, we have $|AP| = \prod_{a \in A} |AP^a|$ and

$$\prod_{a \in A} |AP^a| < \prod_{a \in A} 2^{n_a(1+\lambda(\varepsilon))(H(\mathbf{k}^a)-H(\mathbf{k}^a|\mathbf{i}^a)+\xi(\varepsilon))} \tag{3.14}$$

$$= 2^{n(1+\lambda(\varepsilon))(H(\mathbf{k}|\mathbf{a})-H(\mathbf{k}|\mathbf{i},\mathbf{a})+\xi(\varepsilon))}. \tag{3.15}$$

□

We now show that at the beginning of every block $k > 1$, the agent is able to deduce an input set $\mathscr{S}_K[k]$ from her observation of the past play, and, in particular, from her observation of an action plan.

Remark 17. In the later construction of the strategies we will see that if the agent is able to deduce an action plan $\mathscr{K} \in AP$ from her observation of the past play at the beginning of some block k, then this must have been possible due to the forecaster's ability to observe a hypothetical input set in block k, $\mathscr{I}[k]$ (and her transmission of information of the same to the agent). This remark is vital for the following lemma. ◇

Lemma 9. *Assume, that in block* $k-1$, $k > 1$, *the forecaster and the agent are able to induce a locally typical input set* $\mathscr{S}[k-1]$. *Moreover, assume that the agent can deduce an action plan* $\mathscr{K} \in AP$ *from her observation of the past play. Then, the agent is able to construct an input set* $\mathscr{S}[k]$ *from* $\mathscr{K}$ *for block k, that satisfies property* (3.11).

Proof. Recall, that given the assumptions stated in the current lemma, the action plan $\mathscr{K}$, received by the agent at the beginning of block k, corresponds to a hypothetical input set of nature, that was observed by the forecaster in block $k-1$. Observe, that since every $\kappa^{r_a+m} \in \mathscr{K}$ is locally conditional typical, i.e., $\kappa^{r_a+m} \in T^{\varepsilon}_{r_a+m,l}(\rho^a_{I\times K}|\iota^{r_a+m}[k])$ for every $a \in A$, then by Definition 10, every subsequence of r^a consecutive elements of κ^{r_a+m} is also locally conditional typical. Hence, the only thing the agent needs to do, is to determine the correct segment in κ^{r_a+m} of length r_a for every $a \in A$, which will constitute the set $\mathscr{S}_K[k]$. For this, the agent needs to observe the ending points of each conditional subsequence $\alpha_a^{n+1}[k-1]$ of the output sequence $\alpha^{n+1}[k-1]$. These ending points determine the starting points for the sequences in the input set $\mathscr{S}[k]$ of block k (these ending points are equivalent to the green pointers in Figures 3.1 and 3.2).

In order to compute the exact ending points of each conditional subsequence, she needs to observe, for every $a \in A$, the number of occurrences of a in the past play up to block $k-1$, $N(a|\alpha^{k-2})$, as well as the number of occurrences of element a in block $k-1$, given by $L(\alpha_a^{n+1}[k-1])$. Inequality (3.13) in Remark 16 tells us that

$$r_a - m \leq L(\alpha_a^{n+1}[k-1]) \leq r_a.$$

The sequence in the action plan and in the hypothetical input set start right after the minimal length of $L(\alpha_a^{n+1}[k-1])$, as was demonstrated in Figures 3.1 and 3.2. Therefore, if $L(\alpha_a^{n+1}[k-1]) > r_a - m$, the agent needs to cut the first $L(\alpha_a^{n+1}[k-1]) - (r_a - m)$ elements of the prefix of κ^{r_a+m}, for every $a \in A$. Since for every $a \in A$ a sequence in $\mathscr{S}_K[k]$ has r_a elements, we need to make sure that redundant elements at the end of κ^{r_a+m} are also cut off. This means that the agent needs to cut $r_a + m - \big(L(\alpha_a^{n+1}[k-1]) - (r_a - m)\big)$ elements in the suffix of κ^{r_a+m}.

Note, that nature proceeds in exactly the same way. The resulting sequences, denoted by $(\iota^{r_a}[k], \kappa^{r_a})$, have the desired length r_a. Hence, $\{\iota^{r_a}[k] : a \in A\}$ are the sequences of nature's input set, $\mathscr{S}_I[k]$, and $\{\kappa^{r_a} : a \in A\}$, are the desired sequences of agent's input set, $\mathscr{S}_K[k]$. □

The Message Set of the Forecaster

Before we define the message set of the forecaster, let us give an overview of the main properties of a message:

A message of the forecaster in block k consists of a set of sequences, one for every $a \in A$. Every such sequence is of length $r_a - m$, and is therefore shorter than a sequence in the forecaster's input set, $\mathscr{S}_J[k]$. However, one important feature of a message is its extendability, i.e., a message can be extended into an input set $\mathscr{S}_J[k]$, that satisfies the property stated in (3.12). Another important feature of a message is its ability to transmit the forecaster's knowledge of the hypothetical input set $\mathscr{I}[k+1]$ to the agent. The following paragraphs will look at these properties now in more detail.

Let us first introduce the concept of extendability:

Definition 16 (Extendable Locally Conditional Typical Sequence)**.** Let $P \in \Delta(I \times J)$ be a distribution with marginal distribution $P_I \in \Delta(I)$. Furthermore, let $n, m, l \in \mathbb{N}$ be integers with $n, m > l$ and let $x^{n+m} \in T^{\varepsilon}_{n+m,l}(P_I)$ be a locally typical sequence. Note, that the sequences of the first n and of the last

m elements of x^{n+m}, denoted by x^n and x^m, respectively, are also locally typical w.r.t. P_I. We say that a locally conditional typical sequence $y^n \in T^{\varepsilon}_{n,l}(P|x^n)$ is m-extendable, if there exists a sequence $y^m \in T^{\varepsilon}_{m,l}(P|x^m)$, s.t. the concatenation of y^n and y^m, $y^n \oplus y^m$, is an element of $T^{\varepsilon}_{n+m,l}(P|x^{n+m})$. We denote the set of m-extendable locally conditional typical sequences by $T^{\varepsilon}_{n,l}(P|x^n, ext(m))$. ◇

The message set of the forecaster is an extendable locally conditional typical set in the following sense:

Definition 17 (The Message Set of the Forecaster)**.** Let $\{P_2, \varepsilon, n, m, l\}$ be a set of fit parameters that behaves asymptotically appropriately. Furthermore, let $\mathscr{S}_I[k]$ and $\mathscr{S}_K[k]$ denote the input sets of nature and of the agent in block k, that satisfy the properties stated in (3.10) and in (3.11), respectively. For every $a \in A$, set $r_a = (n+m)P(a)$ and denote by $(x^{r_a - m}, z^{r_a - m})$ the prefix of the first $r_a - m$ elements of $(x^{r_a}, z^{r_a}) \in \mathscr{S}_I[k] \times \mathscr{S}_K[k]$. The message set of the forecaster w.r.t. an element $a \in A$, denoted by M^a_k, is then given by the m-extendable locally conditional typical set

$$M^a_k = T^{\varepsilon}_{r_a - m, l}(\rho^a | (x^{r_a - m}, z^{r_a - m}), ext(m)).$$

The (total) message set, M_k, is given by

$$M_k = \bigtimes_{a \in A} M^a_k.$$

◇

Remark 18 (Length of a Message)**.** Note, that for every $a \in A$, a sequence in a message is of length $r_a - m$. This is no coincidence. Recall, the minimal length of a conditional subsequence α^{n+1}_a of the output sequence α^{n+1} has the same length, as stated in (3.13). A message of length $r_a - m$ hence ensures that it will appear in the output sequence (any message of a longer length might not fully appear in the output sequence). ◇

Remark 19. [Extending a Message to a Locally Typical Input Set] By Definition 16, for every $a \in A$ and for every pair $(x^{r_a}, z^{r_a}) \in \mathscr{S}_I[k] \times \mathscr{S}_K[k]$, a sequence $y^{r_a - m} \in M^a_k$ can be m-extended to a sequence $y^{r_a} \in T^{\varepsilon}_{r_a, l}(\rho^a | (x^{r_a}, z^{r_a}))$. The set of these m-extended sequences is exactly the conditional locally typical input set $\mathscr{S}_J[k]$ stated in (3.12). ◇

With the help of Corollary 10 in the Appendix, we can now derive the size of the message set.

Lemma 10 (The Size of the Message Set)**.** *Let $\varepsilon > 0$ and let $\delta(\varepsilon)$ be a function of ε such that $\delta(\varepsilon) \to 0$ as $\varepsilon \to 0$. Furthermore, with $\boldsymbol{a} = (\boldsymbol{i}', \boldsymbol{j}', \boldsymbol{k}')$, denote by $(\boldsymbol{a}, \boldsymbol{i}, \boldsymbol{j}, \boldsymbol{k})$ two triples of random variables distributed* *according to P_2 Then, for every $\delta > 0$ and for n sufficiently large,*

$$|M_k| \geq 2^{(1-\delta(\varepsilon))n(H(\boldsymbol{j}|\boldsymbol{i},\boldsymbol{k},\boldsymbol{a})-2\varepsilon d-\delta)},$$

with $d = \sum_{a\in A} H(\boldsymbol{j}|\boldsymbol{i},\boldsymbol{k},\boldsymbol{a}=a)$

Proof. By Corollary 10 (in the Appendix), with $d = \sum_{a\in A} H(\mathbf{j}|\mathbf{i},\mathbf{k},\mathbf{a}=a)$, we can bound the message set of the forecaster with respect to an element $a \in A$ as follows:

$$|M_k^a| > 2^{(r_a-m)(H(\mathbf{j}|\mathbf{i},\mathbf{k},a)-2\varepsilon d-\delta)}.$$

Now, since we assumed that the set of fit parameters $\{P_2, \varepsilon, n, m, l\}$ behaves asymptotically appropriately, there exists a function $v(\varepsilon)$ with $v(\varepsilon) \to 0$ for $\varepsilon \to 0$. Setting $\delta(\varepsilon) = \frac{v(\varepsilon)}{\min_{a\in A} P(a)}$ and $n_a = nP(a)$, note that we can deduce the following bound for $r_a - m$ for every $a \in A$:

$$\begin{aligned}
r_a - m &= (n+m)P(a) - m \\
&= n_a(1 + \tfrac{m}{n} - \tfrac{m}{nP(a)}) \\
&> n_a(1 - \tfrac{m}{nP(a)}) \\
&> n_a(1 - \tfrac{v(\varepsilon)}{\min_{a\in A} P(a)}) \\
&= n_a(1 - \delta(\varepsilon)).
\end{aligned}$$

Therefore,

$$\begin{aligned}
|M_k| &= \prod_{a\in A} |M_k^a| \\
&> \prod_{a\in A} 2^{(r_a-m)(H(\mathbf{j}|\mathbf{i},\mathbf{k},a)-2\varepsilon d-\delta)} \\
&> 2^{\sum_{a\in A} n_a(1-\delta(\varepsilon))(H(\mathbf{j}|\mathbf{i},\mathbf{k},a)-2\varepsilon d-\delta)} \\
&= 2^{n(1-\delta(\varepsilon))\sum_{a\in A} P(a)(H(\mathbf{j}|\mathbf{i},\mathbf{k},a)-2\varepsilon d-\delta)} \\
&= 2^{n(1-\delta(\varepsilon))(H(\mathbf{j}|\mathbf{i},\mathbf{k},\mathbf{a})-2\varepsilon d-\delta)}.
\end{aligned}$$

□

The Role of the Information Constraint

The information constraint is a vital element in the construction of the strategies of the players. Only with the help of the information constraint we can show that information transmission from the forecaster to the agent is possible. We elaborate on this in the following paragraphs.

In the following Corollary we introduce the so-called transmission function:

Corollary 1. *Let $\{P_2, \varepsilon, n, m, l\}$ be a set of fit parameters that behaves asymptotically appropriately. Given a message set for block k, M_k, and given the set of action plans, AP, there exists a surjective mapping $f_k : M_k \to AP$, as $\varepsilon \to 0$. We call f_k the transmission function w.r.t. block k.*

Proof. We show that $|M_k| \geq |AP|$ as $\varepsilon \to 0$.
Let $\mathbf{a} = (\mathbf{i}', \mathbf{j}', \mathbf{k}')$ and let $(\mathbf{i}, \mathbf{j}, \mathbf{k}, \mathbf{a}) \sim P_2$.. As stated in Notation 11, P_2 satisfies the following information constraint:

$$H(\mathbf{i},\mathbf{j}|\mathbf{k},\mathbf{a}) \geq H(\mathbf{i}) + \varepsilon. \tag{3.16}$$

Now, by Lemma 8, we have

$$|AP| \leq 2^{(1+\lambda(\varepsilon))n(H(\mathbf{k}|\mathbf{a})-H(\mathbf{k}|\mathbf{i},\mathbf{a})+\xi(\varepsilon))},$$

with $\lambda(\varepsilon) \to 0$ and $\xi(\varepsilon) \to 0$ as $\varepsilon \to 0$.
Further, setting $\delta = \varepsilon$ in Lemma 10, we get

$$|M_k| \geq 2^{(1-\delta(\varepsilon))n(H(\mathbf{j}|\mathbf{i},\mathbf{k},\mathbf{a})-2\varepsilon d-\varepsilon)}.$$

We now show that

$$H(\mathbf{j}|\mathbf{i},\mathbf{k},\mathbf{a}) - \varepsilon \geq H(\mathbf{k}|\mathbf{a}) - H(\mathbf{k}|\mathbf{i},\mathbf{a}). \tag{3.17}$$

First, applying the chain rule, we get

$$H(\mathbf{j}|\mathbf{i},\mathbf{k},\mathbf{a}) - \varepsilon = H(\mathbf{i},\mathbf{j}|\mathbf{k},\mathbf{a}) - H(\mathbf{i}|\mathbf{k},\mathbf{a}) - \varepsilon \tag{3.18}$$

Next, since by assumption, $H(\mathbf{i}|\mathbf{a}) = H(\mu) = H(\mathbf{i})$, we have

$$\begin{aligned} H(\mathbf{k}|\mathbf{a}) - H(\mathbf{k}|\mathbf{i},\mathbf{a}) &= H(\mathbf{k}|\mathbf{a}) - (H(\mathbf{k},\mathbf{i}|\mathbf{a}) - H(\mathbf{i}|\mathbf{a})) \\ &= -(H(\mathbf{k}|\mathbf{a}) + H(\mathbf{i}|\mathbf{k},\mathbf{a}) - H(\mathbf{k}|\mathbf{a})) + H(\mathbf{i}) \\ &= H(\mathbf{i}) - H(\mathbf{i}|\mathbf{k},\mathbf{a}). \end{aligned} \tag{3.19}$$

Therefore, by equations (3.18) and (3.19), the inequality in (3.17) is equivalent to

$$H(\mathbf{i},\mathbf{j}|\mathbf{k},\mathbf{a}) - H(\mathbf{i}|\mathbf{k},\mathbf{a}) - \varepsilon \geq H(\mathbf{i}) - H(\mathbf{i}|\mathbf{k},\mathbf{a}),$$

which is equivalent to the information constraint stated in (3.16).

The result follows by setting $\varepsilon \to 0$. □

Our aim is to describe the process of the the information transmission in detail. In order to do so, we need the transmission function and the action plan function, which we denote in the following:

Notation 13. By the construction of the set of action plans, AP, we can define the following surjective function,

$$f_{AP} : HI \to AP,$$

where $HI = \times_{a \in A} T^{\varepsilon}_{r_a+m,l}(\mu)$, i.e., HI includes all possible hypothetical input sets of nature. ◇

Definition 18 (Information Transmission-Process)**.** Let $\{P_2, \varepsilon, n, m, l\}$ denote a set of parameters that behaves asymptotically appropriately. Let us assume that there exist strategies, that together with nature's actions, induce a play. Denote by k and $k+1$ two consecutive blocks in the play, s.t. in block k the players are able to induce a locally typical input set $\mathscr{S}[k]$ w.r.t. n, m, l and P_2. We call the sending and the receiving of information, as outlined below, the information transmission-process:

Sending Information:

1. At the beginning of block k, the forecaster observes the hypothetical input set of nature of block $k+1$, $\mathscr{I}[k+1]$.

2. The forecaster then computes the corresponding action plan for the agent using the action plan function, i.e., $f_{AP}(\mathscr{I}[k+1]) = \kappa \in AP$.

3. This information is sent to the agent in block k via the transmission function, i.e., the forecaster chooses a message $\mathscr{M}_k \in f_k^{-1}(\kappa) \subset M_k$.

<u>Receiving Information:</u>

1. At the beginning of block $k+1$, the agent observes the past play up to and including block k, α^k. In particular, she observes the message $\mathscr{M}_k$ in the output sequence $\alpha^{n+1}[k]$.

2. The agent receives her action plan by applying the transmission function to the message, i.e., $f_k(\mathscr{M}_k) = \mathscr{K}$.

$\diamond$

Corollary 2 (Inducing P_2)**.** *Given the assumptions in Definition 18, the information transmission process over two consecutive blocks k and $k+1$ defines strategies for the players such that in block $k+1$ distribution P_2 is induced via a locally typical input set $\mathscr{S}[k+1]$.*

Proof. First of all, the hypothetical input set of nature in block $k+1$, $\mathscr{I}[k+1]$, satisfies property (3.10), since the sequences in $\mathscr{I}[k+1]$ are locally typical w.r.t. μ. Moreover, as described in Definition 18, at the beginning of block $k+1$, the agent receives an action plan $f_k(\mathscr{M}_k) = \mathscr{K}$ that was sent to her by a message from the forecaster, $\mathscr{M}_k$, in block k. By Lemma 9, the agent can extract an input set $\mathscr{S}_K[k+1]$ from her action plan $\mathscr{K}$ that satisfies property (3.11). The actions of the forecaster depend on her observation of nature's actions: If the forecaster observes a hypothetical input set for the upcoming block, $\mathscr{I}[k+2]$, she first computes the corresponding action plan for the agent, $f_{AP}(\mathscr{I}[k+2]) = \mathscr{K}'$ and then chooses a message $\mathscr{M}_{k+1} \in f_{k+1}^{-1}(\mathscr{K}') \subset M_{k+1}$. By Remark 19, the forecaster then extends $\mathscr{K}_{k+1}$ into an input set $\mathscr{S}_J[k+1]$, that satisfies property (3.12). Otherwise, if the the forecaster does not observe a hypothetical input set for the upcoming block, she can directly choose an input set $\mathscr{S}_J[k+1]$ as in (3.12) (which does not contain a message).
Hence, by construction, the set

$$\mathscr{S}_I[k+1] \times \mathscr{S}_J[k+1] \times \mathscr{S}_K[k+1] = \mathscr{S}[k+1]$$

constitutes a locally typical input set w.r.t. n, m, l and P_2. $\square$

Observation 5. The reader should observe that for the information transmission process to work, it is necessary that the forecaster indeed observes a hypothetical input set of nature. In particular, this means that the sequences in $\mathscr{I}[k+1]$ are locally typical w.r.t. μ. Otherwise, the forecaster is unable to find an action plan for the agent and the players cannot induce a locally typical input set $\mathscr{I}[k+1]$ in block $k+1$. ◇

Notation 14. If, as described in the above observation, the forecaster cannot observe a hypothetical input set of nature in block $k+1$, then we say that the information transmission failed and block $k+1$ is then called a lost block. ◇

Restarting after a Lost Block

If the players have encountered a lost block, it is not the end of the world. First, we know that by Lemma 18 in the Appendix, this does not happen frequently, and second, it is possible to resume the information transmission process. Obviously, this only depends on nature's conditional sequences. Once they are locally typical again, the players can restart. However, since the information transmission has failed before, the players cannot immediately induce the desired distribution P_2 (since the agent does not have the necessary information). More precisely, we need a recovery process of at least two blocks: if a block k is lost, then the earliest block in which the players can induce a distribution P_2 again is block $k+2$ (assuming nature plays typical). To see that, it is important to note that in a lost block we don't have a mechanism (such as the Splicing Algorithm) that helps us to control the length of the conditional subsequences in that block. Therefore, it is impossible for the forecaster with the tools at hand to send any information of future states of nature in the lost block to the agent (so that they are able to induce a distribution P_2 in the block following a lost block). Instead, if nature plays locally typical again after a lost block, we show that a different distribution, Q_2, can be induced in the block after a lost block (for which no previous information transmission is needed) which enables the information transmission process to work again (to induce P_2) thereafter.

The following Lemma focuses on the first block after a lost block in which the players can induce such a distribution Q_2.

Lemma 11. *Let $Q = \mu \times U_J$, where U_J denotes the uniform distribution over J. Fix $z \in K$, set $\tilde{A} = I \times J \times z$ and denote by $Q_2 = Q \times Q$ a distribution over $\tilde{A}^2$. Furthermore, let $\{P_2, \varepsilon, n, m, l\}$ be a set of fit pa-*

rameters, such that $nQ(\tilde{a}), mQ(\tilde{a}) \in \mathbb{N}$, *for all* $\tilde{a} \in \tilde{A}$. *Set* $r_{\tilde{a}} = (n+m)Q(\tilde{a})$. *Finally, let* k^* *denote the first block after a lost block in which* $\mathscr{S}_I[k^*]$ *satisfies property* (3.10), *i.e.,* $\mathscr{S}_I[k^*] = \left\{x^{r_{\tilde{a}}} \in T^{\varepsilon}_{r_{\tilde{a}},l}(\mu) : \tilde{a} \in \tilde{A}\right\}$. *Then, the players are able to induce* Q_2 *in block* k^*.

Proof. We show that the players can induce a locally typical input set $\mathscr{S}[k^*]$ w.r.t. n, m, l and Q_2, such that the corresponding output sequence of the Splicing Algorithm applied to $\mathscr{S}[k^*]$ has a second order type close to Q_2. To this end, we need $\{Q_2, \varepsilon, n, m, l\}$ to be a set of fit parameters (note, that we only replaced P_2 by Q_2 in the above set of fit parameters), which we demonstrate below:

Denote by $\left\{\lambda^{\tilde{a}} : \tilde{a} \in \tilde{A}\right\}$ the conditional distributions of Q_2. Note, that by assumption, Q_2 has full support and for every $\tilde{a} \in \tilde{A}$, it holds that the marginal of $\lambda^{\tilde{a}}$ on I, $\lambda^{\tilde{a}}_I$, is equal to μ. Furthermore, if $\min_{\tilde{a}\in\tilde{A}} Q(\tilde{a}) > \min_{a\in A} P(a)$ and $\min_{\tilde{a},\tilde{b}\in\tilde{A}^2} Q(\tilde{a},\tilde{b}) > \min_{a,b\in A^2} P(a,b)$, then $\{Q_2, \varepsilon, n, m, l\}$ is a set of fit parameters.

Indeed, for all $a = i, j, k \in A$, we have

$$\begin{aligned}\min_{(i,j,k)\in A} P(i,j,k) &= \min_{(i,j,k)\in A} P(j|k,i)P(k|i)\mu(i)\\ &< \tfrac{1}{|J|}\min_{i\in I}\mu(i)\\ &= \min_{\tilde{a}\in\tilde{A}} Q(\tilde{a}).\end{aligned}$$

In the same way, one can show $\min_{\tilde{a},\tilde{b}\in\tilde{A}^2} Q(\tilde{a},\tilde{b}) > \min_{a,b\in A^2} P(a,b)$, hence $\{Q_2, \varepsilon, n, m, l\}$ satisfies the properties of a set of fit parameters.

It remains to specify the actions the players need to play in order to induce a locally typical input set $\mathscr{S}[k^*]$: The input set of nature is assumed to be

$$\mathscr{S}_I[k^*] = \left\{x^{r_{\tilde{a}}} \in T^{\varepsilon}_{r_{\tilde{a}},l}(\mu) : \tilde{a} \in \tilde{A}\right\}. \tag{3.20}$$

If the agent plays a sequence of the fixed element z, i.e.,

$$\mathscr{S}_K[k^*] = \left\{z^{r_{\tilde{a}}} = (z, z, \dots, z) \in K^{r_{\tilde{a}}} : \tilde{a} \in \tilde{A}\right\}, \tag{3.21}$$

and the agent matches nature and agent, i.e.,

$$\mathscr{S}_J[k^*] = \left\{ y^{r_{\tilde{a}}} \in T^{\varepsilon}_{r_{\tilde{a}},l}(\lambda^{\tilde{a}} | x^{r_{\tilde{a}}}, z^{r_{\tilde{a}}}) : \tilde{a} \in \tilde{A} \right\}, \tag{3.22}$$

then for every $(x^{r_{\tilde{a}}}, y^{r_{\tilde{a}}}, z^{r_{\tilde{a}}}) \in \mathscr{S}_I[k^*] \times \mathscr{S}_J[k^*] \times \mathscr{S}_{\mathscr{K}}[k^*] = \mathscr{S}[k^*]$ it holds that

$$(x^{r_{\tilde{a}}}, y^{r_{\tilde{a}}}, z^{r_{\tilde{a}}}) \in T^{\varepsilon}_{r_{\tilde{a}},l}(\lambda^{\tilde{a}}),$$

which completes the proof. □

Now, if the players are able to induce a distribution Q_2 in a block k^*, then the information transmission process can be resumed if the forecaster observes a hypothetical input set of nature in the following block, k^*+1. However, we need to be precise here. We have to adapt the hypothetical input set of nature to the fact that the players induced a distribution Q_2 in block k^* instead of the distribution P_2 (compare with Definition 14). The sequences in the adapted hypothetical input set need to have slightly different start and end indices, but the length of each sequence remains $r_a + m$, for all $a \in A$. This stems from the fact that $\tilde{A}$, the set on which Q_2 is defined, is only a subset of A. Hence, for some elements $a \in A$, the conditional subsequences α_a^{n+1} of the output sequence in block k^* are empty. The following Definition clarifies these statements.

Definition 19 (Q-adapted Hypothetical Input Set of Nature)**.** Let k^* denote a block in the play in which the distribution Q, as introduced in Lemma 11, is induced. Furthermore, denote by $\{P_2, \varepsilon, n, m, l\}$ a set of fit parameters. The Q-adapted hypothetical input set of nature, denoted by $\mathscr{I}_Q[k^*+1]$ is a set of A segments of nature's conditional sequences $\{x^a : a \in A\}$, with the property that a segment $\iota^{r_a+m}[k^*+1]$ has length $r_a + m$ and is locally typical w.r.t. μ:

$$\mathscr{I}_Q[k^*+1] = \left\{ \iota^{r_a+m}[k^*+1] \in T^{\varepsilon}_{r_a+m,l}(\mu) : a \in A \right\},$$

and for every $a \in A$,

$$\iota^{r_a+m}[k^*+1] = (x^a_{\min N(a|\alpha^{k^*})}, \ldots, x^a_{\max N(a|\alpha^{k^*})+r_a}),$$

where $\min N(a|\alpha^{k^*}) = N(a|\alpha^{k^*-1}) + r_{\tilde{a}} - m$ and $N(a|\alpha^{k^*}) = N(a|\alpha^{k^*-1}) + r_{\tilde{a}}$.
Note, that for a given element $a = (i, j, k) \in A$ in the above equations for the indices, $\tilde{a}$ is only defined if $k = z$. If $k \neq z$, then we set $r_{\tilde{a}} = 0$. ◇

Note, that since the length of the sequences in the Q-adapted hypothetical input sets of nature hasn't changed, the set of action plans for the agent does also not change. However, we need to adapt the message set of the forecaster to a block k^* in which distribution Q is induced:

Definition 20 (The Q-adapted Message Set of the Forecaster)**.** Let Q be a distribution as introduced in Lemma 11 and let $\{Q_2, \varepsilon, n, m, l\}$ be a set of fit parameters that behaves asymptotically appropriately. As before, denote by $\left\{\lambda^{\tilde{a}} :\in \Delta(\tilde{A}) : \tilde{a} \in \tilde{A}\right\}$ the set of conditional distributions of Q_2. Furthermore, let $\mathscr{S}_I[k^*]$ and $\mathscr{S}_K[k^*]$ denote the input sets of nature and of the agent in block k^*, that satisfy the properties stated in (3.20) and in (3.21), respectively. For every $\tilde{a} \in \tilde{A}$, set $r_{\tilde{a}} = (n+m)Q(\tilde{a})$ and denote by $(x^{r_{\tilde{a}}-m}[k^*], z^{r_{\tilde{a}}-m}[k^*])$ the prefix of the first $r_{\tilde{a}} - m$ elements of $(x^{r_{\tilde{a}}}[k^*], z^{r_{\tilde{a}}}[k^*]) \in \mathscr{S}_I[k^*] \times \mathscr{S}_K[k^*]$. The message set of the forecaster w.r.t. an element $\tilde{a} \in \tilde{A}$, denoted by $M_{k^*}^{\tilde{a}}$, is then given by the m-extendable locally conditional typical set

$$M_{k^*}^{\tilde{a}} = T_{r_{\tilde{a}}-m,l}^{\varepsilon}(\lambda^{\tilde{a}} | (x^{r_{\tilde{a}}-m}[k^*], z^{r_{\tilde{a}}-m}[k^*]), ext(m)).$$

The (total) message set, M_{k^*}, is given by

$$M_{k^*} = \underset{\tilde{a} \in \tilde{A}}{\times} M_{k^*}^{\tilde{a}}.$$

$\diamond$

The size of the Q-adapted message set of the forecaster can now be directly deduced from Lemma 10:

Corollary 3 (The Size of the Q-adapted Message Set)**.** *Let $\delta(\varepsilon)$ be a function of ε such that $\delta(\varepsilon) \to 0$ as $\varepsilon \to 0$. Then, for n sufficiently large,*

$$|M_{k^*}| \geq 2^{(1-\delta(\varepsilon))n(\log_2 |J| - 2\varepsilon\tilde{d} - \varepsilon)},$$

with $\tilde{d} = |\tilde{A}| \log_2 |J|$.

Proof. The proof follows directly from Lemma 10. $\square$

Observation 6. Observe, that $\log_2 |J| > H(\mathbf{j}|\mathbf{i}, \mathbf{k}, \mathbf{a})$, and hence $|M_{k^*}| > |M_k| > |AP|$. Therefore, Corollary 1 can be applied to a Q-adapted message set of the forecaster, i.e., there exists a surjective mapping $f_{k^*} : M_{k^*} \to AP$, which we also refer to as the transmission function w.r.t. block k^*. Moreover, there still exists the action plan function $f_{AP}: HI \to AP$, so that the information transmission process as outlined in Definition 18 can be directly adapted to blocks k^* and k^*+1, where distribution Q_2 is induced in block k^*. $\diamond$

3.5 Proof of Theorem 2

We now have all the tools at hand to finally proof Theorem 2:

Proof of Theorem 2. Let $A = (I \times J \times K)$ and let $P_2' \in \Delta(A^2)$ be a distribution that satisfies the assumptions in Theorem 2. For all $\varepsilon > 0$, we know by Lemma 16 that there exists a distribution $P_2 \in \Delta(A^2)$ with properties stated in Notation 11. In particular, it holds that

$$\left\|P_2 - P_2'\right\|_1 \leq 2\varepsilon. \tag{3.23}$$

Furthermore, let $Q_2 \in \Delta(\tilde{A}^2)$, with $\tilde{A} = I \times J \times z$ and $z \in K$, denote the distribution introduced in Lemma 11.

Fix $\varepsilon > 0$ and let $l, n, m \in \mathbb{N}$ be integers where n, m are chosen sufficiently large, so that the following holds:

- $nQ(\tilde{a}), mQ\tilde{a} \in \mathbb{N}$ for all $\tilde{a} \in \tilde{A}$
- $\{P_2, \varepsilon, n, m, l\}$ is a set of fit parameters that behaves asymptotically appropriately.

We show that there exist strategies (σ, τ) that together with nature induce a sequence of random variables, $(\mathbf{x}_t, \mathbf{y}_t, \mathbf{z}_t)$, $t \geq 1$, s.t. the expected empirical distribution over two action triples after $t = nw$ stages (for $w \in \mathbb{N}$ sufficiently large), denoted by P_2^{nw}, is close to P_2, i.e.,

$$\|P_2^{nw} - P_2\| < 2\varepsilon + \delta(\varepsilon), \tag{3.24}$$

where $\delta(\varepsilon)$ is a function of ε with $\delta(\varepsilon) \to 0$ as $\varepsilon \to 0$.

Note, that P^{nw} is an nw-implementable distribution, and by Remark 25, P^{nw} is also implementable. By equations (3.23) and (3.24) one can then conclude that

$$\left\|P_2^{nw} - P_2'\right\| \leq 4\varepsilon + \delta(\varepsilon),$$

and since the set of implementable distributions is closed (also by Remark 25), Theorem 2 follows. We will henceforth focus on the proof of inequality (3.24).

A pair of strategies and nature's actions induce a play, i.e., a sequence of action triples, which we divide into blocks of length n. We first show that there exist strategies, (σ, τ), such that in asymptotically almost every block $k > 1$, distribution P_2 is induced via a locally typical input set $\mathscr{S}[k]$. To be precise, P_2 is not directly induced by the strategies, but a distribution close to P_2. I.e., in every block, the strategies induce a sequence of action triples, denoted by $\alpha^{n+1}[k]$ (we apply the same notation as introduced in section 3.4.1 (Strategy Outline)), with the property that

$$\left\|emp_2(\alpha^{n+1}[k]) - P_2\right\| < \delta(\varepsilon).$$

In the following, when stating that P_2 is induced, we think of the induction in the sense of the above inequality.

Recall that in Corollary 2 and in Lemma 11 we have already developed strategies for singled-out scenarios in an induced play. Observe that the above assumptions on the parameters ε, n, m, l and on the distributions P_2 and Q_2 are the same as in those scenarios. Therefore, we can now assemble these strategies for an entire play.

In block $k = 1$ we have $n + 1$ stages. We assume that in the first stage of block 1, the players play arbitrarily. Now, if nature's input set, $\mathscr{S}_I[1]$, is typical, i.e., satisfies property (3.10), the players can induce the distribution Q_2 as outlined in Lemma 11. If, in addition, the forecaster observes a Q-adapted hypothetical input set of nature, $\mathscr{I}_Q[2]$ for block 2, then the players can induce P_2 in block 2 (this follows directly from Observation 6 and from Corollary 2).

Whenever the players are able to induce a distribution P_2 in a block $k \geq 2$ and if the forecaster observes a hypothetical input set $\mathscr{I}[k+1]$ at the beginning of block k, they are able to induce P_2 in block $k+1$. Again, this is outlined in Corollary 2.

If, in block $k=1$, nature's input set, $\mathscr{S}_I[1]$, is not typical, i.e., does not satisfy property (3.10), then we consider the first block as a lost block and the forecaster cannot transmit any information about future states of nature. If this happens, the agent still plays a fixed element z and the forecaster plays a random sequence. In this case, the strategies of the players for block 2 are the same as in block 1: if nature's input set, $\mathscr{S}_I[2]$, satisfies property (3.10), the players can induce Q_2 in block 2, etc.

Finally, if in a block $k \geq 1$ the players have either induced Q_2 or P_2 and the forecaster is not able to observe a (Q-adapted) hypothetical input set of nature for the upcoming block $k+1$, then P_2 cannot be induced in block $k+1$ and no information can be transmitted from the forecaster to the agent. Hence, block $k+1$ is a lost block. Also, observe that in this case no information is sent to the agent in block k, i.e., the forecaster does not send a message to the agent (nevertheless, the forecaster's input set $\mathscr{S}_J[k]$ still has property (3.12)). In the lost block $k+1$, the agent again plays a fixed sequence z and the forecaster plays a random sequence. In block $k+2$ after the lost block $k+1$, the strategies of the players are the same as in block 1, i.e., they start anew.

Denote by $(\mathbf{x}_t, \mathbf{y}_t, \mathbf{z}_t)$, $1 \leq t \leq n \cdot w$, $w \in \mathbb{N}$, a sequence of random triples induced by the strategies (σ, τ) and by nature's actions. Furthermore, denote by $(\mathbf{x}[k], \mathbf{y}[k], \mathbf{z}[k]) = \boldsymbol{\alpha}[k]$ the kth block of n random triples, with $1 \leq k \leq w$. Now, recall, that we call a block $k > 1$ a lost block, if the players cannot induce P_2. In particular, in a lost block it holds that $\|emp_2(\boldsymbol{\alpha}[k]) - P_2\| > \delta(\varepsilon)$, where $\delta(\varepsilon)$ is a function of ε with $\delta(\varepsilon) \to 0$ as $\varepsilon \to 0$. Denote by LB the event of such a lost block. In a lost block, at least one of nature's conditional sequences is not locally typical w.r.t. μ. Now, recall, that it depends on the observation of the (possibly Q-adapted) hypothetical input set of nature in block k, whether that block is lost or not. More precisely, if, for every $a \in A$, the sequences in the hypothetical input set are locally typical w.r.t. μ, block k is not lost. Now, by Lemma 18, it holds that for every $a \in A$, and for every segment of $r_a + m$ elements in nature's (random) conditional sequence (note, that $r_a + m$ is the length of a sequence in the hypothetical input set),

$$Pr(\mathbf{x}^{r_a+m} \in T^{\varepsilon}_{r_a+m,l}(\mu)) \to 1 \text{ as } n \to \infty.$$

It directly follows that for the complement of LB, which we call a good block, or GB,

$$PR(GB) = PR(\|emp_2(\boldsymbol{\alpha}[k]) - P_2\| \leq \delta(\varepsilon)) \to 1 \text{ as } n \to \infty.$$

Therefore, it is indeed the case that in asymptotically almost every block $k > 1$ the strategies induce a good block.

We conclude the proof by showing that the expected empirical distribution over pairs of the random sequence $(\mathbf{x}_t, \mathbf{y}_t, \mathbf{z}_t)$, denoted by P_2^{nw}, is close to P_2, i.e., we prove

$$\|P_2^{nw} - P_2\| < \delta(\varepsilon) + 2\varepsilon. \tag{3.25}$$

First, by the law of total expectation, and for every $k > 1$, we have

$$\mathbb{E}[emp_2(\boldsymbol{\alpha}[k])] = \mathbb{E}[emp_2(\boldsymbol{\alpha}[k])|GB]PR(GB) + \mathbb{E}[emp_2(\boldsymbol{\alpha}[k])|LB]PR(LB),$$

therefore, for n sufficiently large,

$$\begin{aligned}
\|\mathbb{E}[emp_2(\boldsymbol{\alpha}[k])] - P_2\| &< \|\mathbb{E}[emp_2(\boldsymbol{\alpha}[k])|GB] - P_2\| + \varepsilon \\
&< \delta(\varepsilon) + \varepsilon.
\end{aligned}$$

It follows, that if w is sufficiently large,

$$\begin{aligned}
\|P_2^{nw} - P_2\| &= \|\mathbb{E}[emp_2(\mathbf{x}_t, \mathbf{y}_t, \mathbf{z}_t)] - P_2\| \\
&= \frac{1}{w}\left\|\mathbb{E}\left[\sum_{k=1}^{w} emp_2(\boldsymbol{\alpha}[k])\right] - wP_2\right\| \\
&= \frac{1}{w}\left\|\sum_{k=1}^{w} \mathbb{E}\left[emp_2(\boldsymbol{\alpha}[k])\right] - wP_2\right\| \\
&\leq \frac{1}{w}\|\mathbb{E}\left[emp_2(\boldsymbol{\alpha}[1])\right] - P_2\| + \frac{1}{w}\sum_{k=2}^{w}\|\mathbb{E}[emp_2(\boldsymbol{\alpha}[k])] - P_2\| \\
&\leq \delta(\varepsilon) + 2\varepsilon,
\end{aligned}$$

which completes the proof. □

Part II

Model 2

Chapter 1

The Properties of Model 2

We now introduce a second model that is more closely related to the model in [6]. The reader will notice that other than in the first model, the forecaster's actions do not influence the players' payoff. The forecaster's only task is to submit her knowledge about future states of nature. The structure of the payoff function hence reduces the complexity of the proofs of the main theorems to an extent. Nevertheless, we cannot adapt the methods from [6] to this model one to one. The key idea here is to introduce another concept, called block distributions, that prove to be useful in applying the results from [6], especially in the proof of the second main theorem.

1.1 The Description of the Model

Let us first present the features of the new model. As before, we denote the action sets of nature, forecaster and agent by I, J and K respectively. We will consider a repeated game, and in each stage $t \geq 1$, the actions played by nature, forecaster and agent are denoted by $x_t \in I$, $y_t \in J$ and $z_t \in K$, respectively.

Just like in Model 2, the team, that again consists of the forecaster and the agent, is assigned a payoff per stage. This time, it does not depend on the forecaster's actions, but only on agent's and on nature's actions in the current stage, as well as on the agent's actions in the previous stage. Hence, the stage-payoff function is given as follows:

$$g_t : K \times I \times K \to \mathbb{R}.$$

The assumptions on the knowledge of the players are similar to the model in [6]. At the beginning of a repeated game, the forecaster can observe all future states of nature. She can also observe the past play during the game. The agent, on the other hand, does not have any knowledge at the beginning of the game and in every stage $t > 1$, she can only observe the past play. More precisely, the forecaster's strategy, denoted by $\sigma = (\sigma_t)_t$, is expressed in dependence of her observation of nature's sequence of actions, and of her observation of the past play. That is, in every stage $t > 1$,

$$\sigma_t : I^{\mathbb{N}} \times J^{t-1} \times K^{t-1} \to J$$

describes the action of the forecaster, y_t, in stage t. The agent's strategy, given by $\tau = (\tau_t)_t$, is expressed in dependence of the past play, i.e., for every stage $t > 1$,

$$\tau_t : I^{t-1} \times J^{t-1} \times K^{t-1} \to K$$

describes the action of the agent, z_t in stage t.

The assumptions on nature are also identical to the model in [6]: at the beginning of the play, nature draws a realisation $x = x_1, x_2, \ldots, \in I^{\mathbb{N}}$ of an i.i.d. sequence $\mathbf{i} = \mathbf{i}_1, \mathbf{i}_2, \ldots$ with law μ.

A pair of strategies (σ, τ) together with μ induce a sequence of random action triples $\mathbf{a}_1, \mathbf{a}_2, \ldots$ with $\mathbf{a}_t = (\mathbf{i}_t, \mathbf{j}_t, \mathbf{k}_t)$, for $t \geq 1$. The corresponding induced probability distribution over $(I \times J \times K)^{\mathbb{N}}$ is denoted by $Q_{\mu,\sigma,\tau}$.

1.2 Implementable Distributions

In this section, we introduce 2 types of implementable distributions that will be of interest in the upcoming sections. We will first denote two important distribution that can be deduced from an induced probability distribution $Q_{\mu,\sigma,\tau}$.

Notation 15. Denote by $Q^{(t-1,t)}_{K\times(I\times K)}$ the marginal distribution of $Q_{\mu,\sigma,\tau}$ on $K \times (I \times K)$ in stages $((t-1), t)$. The average distribution up to stage $T \geq 1$ of these distributions is written as

$$Q^{T}_{K\times(I\times K)} = \frac{1}{T} \sum_{t=2}^{T} Q^{(t-1,t)}_{K\times(I\times K)},$$

and we refer to this distribution as the expected empirical 2-stage distribution. Let us now divide the stages into consecutive blocks $l = 1,2,\ldots$, each of size T (i.e., we have T stages in each block). Denote by $Q^{(l;T)}_{\mu,\sigma,\tau}$ the marginal of $Q_{\mu,\sigma,\tau}$ on block l and denote by

$$\bar{Q}^{(r;T)}_{\mu,\sigma,\tau} = \frac{1}{r}\sum_{l=1}^{r} Q^{(l;T)}_{\mu,\sigma,\tau}$$

the average of these marginal distributions up to block r. We call this distribution the expected empirical T-block distribution. $\diamond$

Given the above notation, we can now introduce two types of implementable distributions:

Definition 21. A distribution $P \in \Delta(K \times (I \times K))$ is implementable, if there exists a pair of strategies (σ,τ), s.t. $Q^{T}_{K\times(I\times K)} \to P$, for $T \to \infty$. P is called T- implementable, if $Q^{T}_{K\times(I\times K)} = P$. The set of this type of implementable (respectively, T-implementable) distributions is denoted by $\mathscr{P}$ (respectively, $\mathscr{P}(T)$). $\diamond$

Definition 22. A distribution $Q \in \Delta(I^T \times J^T \times K^T)$ is implementable, if there exists a pair of strategies (σ,τ), s.t. $\bar{Q}^{(r;T)}_{\mu,\sigma,\tau} \to Q$ as $r \to \infty$. We call Q r- implementable if $\bar{Q}^{(r;T)}_{\mu,\sigma,\tau} = Q$. The set of this type of implementable (respectively, r-implementable) distributions is denoted by $\mathscr{Q}$ (respectively, $\mathscr{Q}(r)$). $\diamond$

Below, we state properties of the sets $\mathscr{P}$ and $\mathscr{P}(t)$ which directly follow from Remark 1 and Remark 2 in [6]:

Remark 20. [Properties of Sets of Implemebtable Distributions]

- The set of implementable distribution $\mathscr{P}$ as denoted in Definition 21 is closed.
- The set of T-implementable distributions, $\mathscr{P}(T)$, is contained in $\mathscr{P}$.

$\diamond$

1.3 The Information Constraint

Let P be a distribution over $K \times (I \times K)$, and denote by $(\mathbf{k}',\mathbf{i},\mathbf{k})$ a triple of random variables distributed according to P. We say that P fulfills the information constraint, if

$$\log|J| \geq H(\mathbf{k}|\mathbf{k}') - H(\mathbf{k}|\mathbf{k}',\mathbf{i}).$$

Equivalently, one can write

$$\log|J| \geq H(\mathbf{i}|\mathbf{k}') - H(\mathbf{i}|\mathbf{k},\mathbf{k}').$$

We can interpret the left-hand side as the information sent from the forecaster to the agent and the right-hand side as the information used by the agent. The latter is the reduction of uncertainty that $\mathbf{k}$ gives on the conditional random variables $\mathbf{i}|\mathbf{k}'$. Note, that the right-hand side is also known as the conditional mutual information of the random variabes $(\mathbf{k},\mathbf{i},\mathbf{k}')$.

Chapter 2

Results

We will prove two major results. In the first one, the information constraint defines the set of implementable distributions:

Theorem 4. *Let $P \in \Delta(K \times (I \times K))$ be an implementable distribution with $(\boldsymbol{k}', \boldsymbol{i}, \boldsymbol{k}) \sim P$ and $\boldsymbol{i} \sim \mu$. Then, P fulfills the information constraint*

$$\log|J| \geq H(\boldsymbol{i}|\boldsymbol{k}') - H(\boldsymbol{i}|\boldsymbol{k}, \boldsymbol{k}').$$

In the second result, we characterise distributions $P \in \Delta(K \times (I \times K))$ that are implementable:

Theorem 5. *Let $P \in \Delta(K \times (I \times K))$ be a distribution with full support and with marginal distribution $P_I = \mu$ on I and with unique marginal distribution P_K on K. Denote by $(\boldsymbol{k}', \boldsymbol{i}, \boldsymbol{k})$ a random triple with $(\boldsymbol{k}', \boldsymbol{i}, \boldsymbol{k}) \sim P$ and let $\mathbf{i}$ be independent of $\boldsymbol{k}'$. If P fulfills the the information constraint, P is implementable.*

Remark 21. The assumption in 5 that $\mathbf{i}$ is independent of $\mathbf{k}'$ is a necessary technical assumption as we will see later on in the proof of the Theorem. It reduces the information constraint to

$$\log|J| \geq H(\mathbf{i}) - H(\mathbf{i}|\mathbf{k}).$$

Note, that assuming J to be larger in size than I would directly imply this information constraint (since then $\log|J| \geq \log|I| \geq H(\mathbf{i})$). The independence assumption, however, does not immediately lead to the information constraint to be fulfilled (at least not without further assumptions). Hence, however restrictive this independence assumption may appear regarding the set of distributions we want to describe as implementable, it does not restrict the sizes of the action sets and hence does not imply a trivial (immediate) fulfilment of the information constraint. ◇

2.1 Proof of Theorem 4

Proof. First, we prove the Theorem for every T-implementable distribution $P \in \Delta(K \times (I \times K))$:

By Lemma 1 in [6] the function $f : \Delta(K \times (I \times K)) \to \mathbb{R}$, $f(Q) = H_Q(\mathbf{i}|\mathbf{k},\mathbf{k}')$ is concave. Recall, that every T-implementable distribution $P \in \Delta(K \times (I \times K))$ can be expressed as the expected empirical distribution of 2-step triples $(\mathbf{k}_{t-1}, \mathbf{i}_t, \mathbf{k}_t)$ in an induced game up to stage $T \geq 1$, $P = \frac{1}{T}\sum_{t=1}^{T} Q^{t-1,t}_{K\times(I\times K)}$. Therefore, we get

$$H_P(\mathbf{i}|\mathbf{k}',\mathbf{k}) \geq \frac{1}{T}\sum_{t=1}^{T} H_{Q^{t-1,t}_{K\times(I\times K)}}(\mathbf{i}|\mathbf{k}',\mathbf{k}).$$

Furthermore,

$$\begin{aligned}
\sum_{t=1}^{T} H_{Q^{t-1,t}_{K\times(I\times K)}}(\mathbf{i}|\mathbf{k}',\mathbf{k}) &= \sum_{t=1}^{T} H(\mathbf{i}_t|\mathbf{k}_t,\mathbf{k}_{t-1}) \\
&= \sum_{t=1}^{T} H(\mathbf{i}_t,\mathbf{k}_t|\mathbf{k}_t,\mathbf{k}_{t-1}) \\
&\geq \sum_{t=1}^{T} H(\mathbf{i}_t,\mathbf{k}_t|\mathbf{i}_1,\ldots,\mathbf{i}_{t-1},\mathbf{j}_1,\ldots\mathbf{j}_T,\mathbf{k}_1,\ldots,\mathbf{k}_{t-1},\mathbf{k}_t) &(2.1)\\
&= \sum_{t=1}^{T} H(\mathbf{i}_t,\mathbf{k}_t|\mathbf{i}_1,\ldots,\mathbf{i}_{t-1},\mathbf{j}_1,\ldots,\mathbf{j}_T,\mathbf{k}_1,\ldots,\mathbf{k}_{t-1}) &(2.2)\\
&= H(\mathbf{i}_1,\ldots,\mathbf{i}_T,\mathbf{k}_1,\ldots,\mathbf{k}_T|\mathbf{j}_1,\ldots,\mathbf{j}_T) \\
&= H(\mathbf{i}_1,\mathbf{j}_1,\mathbf{k}_1,\ldots,\mathbf{i}_T,\mathbf{j}_T,\mathbf{k}_T) - H(\mathbf{j}_1,\ldots,\mathbf{j}_T) \\
&\geq H(\mathbf{i}_1,\ldots,\mathbf{i}_T) - T\log|J| \\
&= T(H(\mathbf{i}) - \log|J|),
\end{aligned}$$

where inequality (2.1) follows from the monotonicity property of the entropy (conditioning reduces entropy) and equality (2.2) is due to the fact that $\mathbf{k}_t$ is a deterministic function of the past. In the remaining rearrangements the chain rule is applied, as well as the property of maximum entropy. Finally, note that the last equality results from the assumption that nature's sequence $\mathbf{i}_1,\ldots,\mathbf{i}_T$ is i.i.d.

Hence, we have

$$TH_P(\mathbf{i}|\mathbf{k}',\mathbf{k}) \geq T(H(\mathbf{i}) - \log|J|) \tag{2.3}$$

$$\geq T(H_P(\mathbf{i}|\mathbf{k}') - \log|J|). \tag{2.4}$$

We have thus shown that every T-implementable distribution fulfills the information constraint.

Now, recall that every implementable distribution $P' \in \mathscr{P}$ is the limiting distribution of a T-implementable distribution. Since the mappings $P \mapsto H_P(\mathbf{i}|\mathbf{k},\mathbf{k}')$, and $P \mapsto H_P(\mathbf{i}|\mathbf{k}')$ are uniformly continuous, we can conclude that the Information Constraint must also hold for all implementable distributions. □

2.2 Proof of Theorem 5

The proof consists of several steps. The rough outline is as follows. We first define a so-called block distribution $\tilde{Q}$ of a sequence of random variables with transition probabilities derived from P. We can then use the results in [6] to show that $\tilde{Q}$ is implementable if P (as given in Theorem 5) satisfies the information constraint. In a final step we will then show that the strategies that implement $\tilde{Q}$ also implement P.

Notation 16. Let P denote a distribution over $K \times (I \times K)$ with identical properties as in Theorem 5. Furthermore, for every $(z',x,z) \in K \times (I \times K)$, we write

$$\mathbb{P}(\mathbf{k}',\mathbf{i},\mathbf{k} = z',x,z) = P(z',x,z),$$

$$\mathbb{P}(\mathbf{i},\mathbf{k} = x,z) = P_{I\times K}(x,z),$$

$$\mathbb{P}(\mathbf{k} = z) = P_K(z),$$

$$\mathbb{P}(\mathbf{i} = x) = \mu(x), \text{ and}$$

$$\mathbb{P}(\mathbf{i},\mathbf{k} = x,z|\mathbf{k}' = z') = P(x,z|z') = \frac{P(z',x,z)}{P_K(z')}.$$

◇

We will now define a markov chain with a transition matrix derived from P:

Definition 23. Denote by $((\mathbf{i}_t,\mathbf{k}_t);t \in \mathbb{N})$ a time-homogeneous markov chain with state space $I \times K$, initial distribution $P_{I\times K}$ and with the following property: for every $n \geq 1$ and for every sequence $((x_1,z_1),\ldots,(x_{n-2},z_{n-2}),(x',z'),(x,z)) \in (I\times K)^n$, let

$$\begin{aligned}
&\mathbb{P}(\mathbf{i}_n,\mathbf{k}_n = x,z|(\mathbf{i}_1,\mathbf{k}_1 = x_1,z_1),\ldots,(\mathbf{i}_{n-2},\mathbf{k}_{n-2} = x_{n-2},z_{n-2}),(\mathbf{i}_{n-1},\mathbf{k}_{n-1} = x',z'))\\
&= \mathbb{P}(\mathbf{i}_n,\mathbf{k}_n = x,z|\mathbf{i}_{n-1},\mathbf{k}_{n-1} = x',z')\\
&= \mathbb{P}(\mathbf{i}_n,\mathbf{k}_n = x,z|\mathbf{k}_{n-1} = z')\\
&= P(x,z|z')
\end{aligned}$$

The transition matrix of this markov chain is given by

$$\mathbf{P}_t : (I\times K)\times(I\times K) \to [0,1],$$

with $\mathbf{P}_t(x',z';x,z) = P(x,z|z'), \quad \forall x',\in I.$

$\diamond$

Properties of the Markov Chain

Proposition 1. *The markov chain $((\boldsymbol{i}_t,\boldsymbol{k}_t);t\in\mathbb{N})$ with transition matrix $\boldsymbol{P}_t$ is stationary with stationary distribution $P_{I\times K}$.*

Proof. First, since P has full support, the transition matrix $\mathbf{P}_t$ has full support and thus the markov chain has a unique stationary distribution. Further, consider $P_{I\times K}$ as a stochastic row vector of length $|I\times K|$. We show

$$P_{I\times K}\mathbf{P}_t = P_{I\times K}.$$

For every $(x,z) \in I \times K$, we have

$$\begin{aligned}
\sum_{(i,k)\in I\times K} P_{I\times K}(i,k)\mathbf{P}_t(i,k;x,z) &= \sum_{(i,k)\in I\times K} P_{I\times K}(i,k)P(x,z|k) \\
&= \sum_{k\in K} P(x,z|k) \sum_{i\in I} P_{I\times K}(i,k) \\
&= \sum_{k\in K} P(x,z|k)P_K(k) \\
&= \sum_{k\in K} P(k,x,z) \\
&= P_{I\times K}(x,z)
\end{aligned}$$

□

Proposition 2. *For $n \in \mathbb{N}_+$ let $\boldsymbol{i}^n = \boldsymbol{i}_1, \ldots, \boldsymbol{i}_n$ be a sequence of random variables drawn from the markov chain $((\boldsymbol{i}_t, \boldsymbol{k}_t); t \in \mathbb{N})$ with transition matrix $\boldsymbol{P}_t$. Then, $\boldsymbol{i}^n$ is an i.i.d. sequence with $\mathbf{i}_t \sim \mu$, $t \in (1, \ldots, n)$.*

Proof. Let $(x_1, \ldots, x_n) \in I^n$. To simplify the notation, we replace $\mathbb{P}(\mathbf{i}_1 = x_1, \ldots, \mathbf{i}_n = x_n)$ by $\mathbb{P}(x_1, \ldots, x_n)$. We will show $\mathbb{P}(x_1, \ldots, x_n) = \prod_{t=1}^n \mu(x_t)$.

For $t \in (1, \ldots, n)$ and for any integer j, $1 \leq j < t$, we have

$$\begin{aligned}
\mathbb{P}(x_t|x_{t-1}, \ldots, x_{t-j}) &= \sum_{z'\in K} \mathbb{P}(x_t|x_{t-1}, \ldots, x_{t-j}, \mathbf{k}_{t-1} = z')\mathbb{P}(\mathbf{k}_{t-1} = z'|x_{t-1}, \ldots, x_{t-j}) \\
&= \sum_{z'\in K} P(x_t|z')\mathbb{P}(\mathbf{k}_{t-1} = z'|x_{t-1}, \ldots, x_{t-j}) \\
&= \sum_{z'\in K} P(x_t)\mathbb{P}(\mathbf{k}_{t-1} = z'|x_{t-1}, \ldots, x_{t-j}) \\
&= P(x_t) = \mu(x_t),
\end{aligned}$$

where the third equality is due to the fact that $\mathbf{i}$ is independent of $\mathbf{k}'$. Hence,

$$\begin{aligned}
\mathbb{P}(x_1, \ldots, x_n) &= P(x_1) \prod_{t=2}^n \mathbb{P}(x_t|x_{t-1}, \ldots, x_1) \\
&= \prod_{t=1}^n P(x_t) \\
&= \prod_{t=1}^n \mu(x_t).
\end{aligned}$$

□

Information Constraint of a block distribution $\tilde{Q}$

We now define a block distribution $\tilde{Q} \in \Delta(I^T \times J^T \times K^T)$ that is derived from the above markov chain:

Definition 24. Let T be a fixed integer and let $(\mathbf{i}^T, \mathbf{k}^T) = (\mathbf{i}_1, \mathbf{k}_1, \ldots, \mathbf{i}_T, \mathbf{k}_T)$ be a sequence of random variables drawn from the markov chain $((\mathbf{i}_t, \mathbf{k}_t); t \in \mathbb{N})$ with transition matrix $\mathbf{P}_t$. Set $(\mathbf{i}^T, \mathbf{k}^T) \sim \tilde{P} \in \Delta(I^T \times K^T)$. Furthermore, let $\mathbf{j}^T = (\mathbf{j}_1, \ldots, \mathbf{j}_T)$ be an i.i.d. sequence with $\mathbf{j}_t \sim U_J$, where U_J denotes the uniform distribution over J. We write $\mathbf{j}^T \sim U_J^{\otimes T}$ and we let $\mathbf{j}^T$ to be independent of $\mathbf{i}^T$ and $\mathbf{k}^T$. The block distribution $\tilde{Q} \in \Delta(I^T \times J^T \times K^T)$ derived from the markov chain $((\mathbf{i}_t, \mathbf{k}_t); t \in \mathbb{N})$ is then defined as the product distribution of $\tilde{P}$ and $U_J^{\otimes T}$, i.e.

$$\tilde{Q} = \tilde{P} \times U_J^{\otimes T}.$$

◇

Remark 22. As an immediate consequence from Proposition 2, the block distribution $\tilde{Q}$ defined above has the marginal distribution $\tilde{Q}_I = \mu^{\otimes T}$ on I^T. ◇

The following result introduces the information constraint on block distributions and can be directly inferred from [6].

Lemma 12. *For every $T \in \mathbb{N}$, let Q denote a distribution over $I^T \times J^T \times K^T$, with marginal $Q_I = \mu^{\otimes T}$ on I^T. If Q satisfies the following inequality*

$$H_Q(\boldsymbol{i}^T, \boldsymbol{j}^T | \boldsymbol{k}^T) \geq H_Q(\boldsymbol{i}^T), \tag{2.5}$$

then Q is implementable. We will refer to inequality 2.5 as the GHN- information constraint, or simply the GHN IC.

The next Lemma establishes the link between the information constraint defined in this paper and the GHN IC:

Lemma 13. *Let $\tilde{Q} \in \Delta(I^T \times J^T \times K^T)$ be defined as in Definition 24. If P satisfies the information constraint*

$$\log |J| \geq H(\boldsymbol{k}|\boldsymbol{k}') - H(\boldsymbol{k}|\boldsymbol{k}', \boldsymbol{i}),$$

then $\tilde{Q}$ satisfies the GHN IC

$$H_{\tilde{Q}}(\boldsymbol{i}^T,\boldsymbol{j}^T|\boldsymbol{k}^T) \geq H_{\tilde{Q}}(\boldsymbol{i}^T).$$

Proof. Since $\mathbf{j}^T$ is assumed to be independent of $\mathbf{i}^T$ and $\mathbf{k}^T$, we have

$$\begin{aligned} H_{\tilde{Q}}(\mathbf{i}^T,\mathbf{j}^T|\mathbf{k}^T) &= H_{\tilde{Q}}(\mathbf{i}^T|\mathbf{j}^T,\mathbf{k}^T) + H_{\tilde{Q}}(\mathbf{j}^T|\mathbf{k}^T) \\ &= H_{\tilde{Q}}(\mathbf{i}^T|\mathbf{k}^T) + TH(U_J). \end{aligned}$$

Now,

$$\begin{aligned} H_{\tilde{Q}}(\mathbf{i}^T|\mathbf{k}^T) = H_{\tilde{P}}(\mathbf{i}^T|\mathbf{k}^T) &= H_{\tilde{P}}(\mathbf{i}^T,\mathbf{k}^T) - H_{\tilde{P}}(\mathbf{k}^T) \\ &= \sum_{t=1}^{T} H_{\tilde{P}}(\mathbf{i}_t,\mathbf{k}_t|\mathbf{i}_{t-1},\mathbf{k}_{t-1},\ldots,\mathbf{i}_1,\mathbf{k}_1) - \sum_{t=1}^{T} H_{\tilde{P}}(\mathbf{k}_t|\mathbf{k}_{t-1},\ldots,\mathbf{k}_1) \\ &= \sum_{t=1}^{T} H_{\tilde{P}}(\mathbf{i}_t,\mathbf{k}_t|\mathbf{i}_{t-1},\mathbf{k}_{t-1}) - \sum_{t=1}^{T} H_{\tilde{P}}(\mathbf{k}_t|\mathbf{k}_{t-1}) \\ &= \sum_{t=1}^{T} H_{P}(\mathbf{i}_t,\mathbf{k}_t|\mathbf{k}_{t-1}) - \sum_{t=1}^{T} H_{P}(\mathbf{k}_t|\mathbf{k}_{t-1}) \\ &= T(H_P(\mathbf{i},\mathbf{k}|\mathbf{k}') - H_P(\mathbf{k}|\mathbf{k}')) \\ &= T(H_P(\mathbf{i}|\mathbf{k},\mathbf{k}') + H_P(\mathbf{k}|\mathbf{k}') - H_P(\mathbf{k}|\mathbf{k}')) \\ &= TH_P(\mathbf{i}|\mathbf{k}). \end{aligned}$$

Furthermore, note, that since we assumed $\mathbf{i}$ to be independent of $\mathbf{k}'$, the information constraint of P reduces to

$$\log|J| \geq H_P(\mathbf{i}) - H_P(\mathbf{i}|\mathbf{k}).$$

Also, note that $H(U_J) = \log|J|$, hence

$$TH(U_J) \geq T(H_P(\mathbf{i}) - H_P(\mathbf{i}|\mathbf{k})).$$

Therefore, it follows that

$$\begin{aligned} H_{\tilde{Q}}(\mathbf{i}^T|\mathbf{k}^T) + TH(U_J) &\geq TH_P(\mathbf{i}|\mathbf{k}') + T(H_P(\mathbf{i}) - H_P(\mathbf{i}|\mathbf{k})) \\ &\geq TH_P(\mathbf{i}) \\ &= H_{\tilde{Q}}(\mathbf{i}^T), \end{aligned}$$

where the last equality follows due to Proposition 2. □

Lemma 12 and Lemma 13 can be combined into the following Corollary:

Corollary 4. *Let $\tilde{Q}$ be the distribution as constructed in Definition (24). If P satisfies the information constraint*

$$\log|J| \geq H(\boldsymbol{k}|\boldsymbol{k}') - H(\boldsymbol{k}|\boldsymbol{k}', \boldsymbol{i}),$$

then $\tilde{Q}$ is implementable.

Even though the Corollary follows immediately from Lemma 12 and from Lemma 13, we nevertheless sketch the proof of this Corollary (which is not intended to be exhaustive). The approach is exactly the same as in [6], but since $\tilde{Q}$ is a block distribution, we first have to adjust some of the applied notation and concepts. We then provide an outline of the construction of the strategies (σ, τ) that implement $\tilde{Q}$. Since we will (later) deduce that the strategies that implement $\tilde{Q}$ also implement P, many of the concepts and notations that we establish in the following will be applied in the proof of the implementation of P.

Let us first introduce empirical T-block distributions:

Definition 25 (Empirical T-Block Distribution)**.** Let A be a finite set and let $n, T \in \mathbb{N}$, such that n is a multiple of T, i.e., we write $n = pT$, for $p > 0$. Let $a^n = a_1, a_2, \ldots, a_n$ denote a sequence in A^n, which can also be written as blocks of length T:

$$a^n = a^1, a^2, \ldots, a^p \text{ with } a^l = a_{(l-1)T+1}, \ldots, a_{lT}, l \in (1, \ldots, p).$$

Given a (block-) element $\alpha^T \in A^T$, we define the block-wise relative frequency of α^T in a sequence a^n as follows:

$$emp^T(a^n)[\alpha^T] = \tfrac{1}{p}\left|\left\{l \geq 1 : a^l = \alpha^T\right\}\right|.$$

We call $emp^T(a^n)$ the empirical T-block distribution of a^n. $\diamond$

Remark 23. Note, that by block-wise relative frequency of a block-element $\alpha^T \in A^T$ in a^n we mean the relative number of times the block-element α^T appears in the fixed blocks $a^1, a^2, \dots, a^p$. This definition needs to be distinguished from the relative frequency of a block-element in a^n, which accounts for the relative number of appearances of a block-element throughout the whole sequence a^n, irrespective of fixed blocks. The block-wise empirical frequency can hence be treated as the block-version of the relative frequency of a single element $\alpha \in A$ in a given sequence a^n. $\diamond$

Notation 17. [Actions played in Long Blocks] For given integers $r, m, T \in \mathbb{N}$ with $m = rT$, we call a block of length m that consists of r blocks of T stages a long block. The sequences played by nature, forecaster and agent in a long block l are denoted by $x[l], y[l]$ and $z[l]$ respectively. For every long block l, the sequence in l can either be given as a sequence of (stage) action triples, denoted by

$$(x[l], y[l], z[l]) = (x_1[l], y_1[l], z_1[l], \dots, x_m[l], y_m[l], z_m[l]),$$

with $(x_t[l], y_t[l], z_t[l]) \in I \times J \times K$, $\forall t \in (1, \dots, m)$, or as a sequence of r (small) blocks of action triples of length T, denoted by

$$(x[l], y[l], z[l]) = (x_1^T[l], y_1^T[l], z_1^T[l], \dots, x_r^T[l], y_r^T[l], z_r^T[l]),$$

with $(x_1^T[l], y_1^T[l], z_1^T[l]) \in X^T \times Y^T \times Z^T$, $\forall b \in (1, \dots, r)$. $\diamond$

Definition (Stage-and Block-Hamming Distance with Empirical T-Block Distributions)**.** Let m denote the length of a long block as in Notation 17, i.e., set $m = rT$. Let $\tilde{Q}_I \in \Delta(I^n)$ and let $x^m \in I^n$ and $\tilde{x}^m \in T_m^T(\tilde{Q}_I) \neq \emptyset$. Recall that both sequences can be written as sequences of blocks of length T as in Definition 25. We define the block hamming distance of $(x^m, \tilde{x}^m)$ as the number of blocks $1 \leq b \leq r$ with $x_b^T \neq \tilde{x}_b^T$, where x_b^T $(\tilde{x}_b^T)$ denotes block b in x^m $(\tilde{x}^m)$. The block hamming distance of $(x^m, \tilde{x}^m)$ needs to be distinguished from the usual (stage-) hamming distance of $(x^m, \tilde{x}^m)$, which is given by the number of stages $1 \leq t \leq rT$ with $x_t^m \neq \tilde{x}_t^m$. $\diamond$

We are now able to provide an outline of the construction of strategies (σ, τ) that implement a distribution $\tilde{Q}$ as stated in Corollary 4:

Implementing $\tilde{Q}$

We divide the stages of the game into long blocks of length m with $m = rT$ as in Notation 17, s.t. every long block consists of r blocks of T stages. W.l.o.g., we assume that $T_m^T(\tilde{Q}) \neq \emptyset$.

The Strategies: The strategies (σ, τ) are based on the following mechanism. Before the start of each long block $l > 0$, the forecaster observes the sequence played by nature in the next long block $l+1$, $x[l+1]$. She then chooses a sequence $\tilde{x}[l+1] \in T_m^T(\tilde{Q}_I)$ (note that since $T_m^T(\tilde{Q}) \neq \emptyset$, we have $T_m^T(\tilde{Q}_I) \neq \emptyset$) in such a way, so that it minimises the block- hamming distance to $x[l+1]$, as well as the (stage-) hamming distance to $x[l+1]$ (the minimization of the (stage-) hamming distance is not a requirement for the implementation of $\tilde{Q}$, however, this additional assumption simplifies the proof of the implementation of P). Since $\tilde{Q}$ satisfies the GHN IC, in each block $l \geq 1$, the forecaster is able to play a sequence $y[l]$, that entails a message to the agent about what to play in the subsequent long block $[l+1]$. In block $l+1$ the forecaster and agent are then able to play sequences that together with the slightly changed sequence of nature, $\tilde{x}[l+1]$, induce $\tilde{Q}$, i.e., $(\tilde{x}[l], y[l], z[l]) \in T_m^T(\tilde{Q})$, $\forall l > 1$.

The Expected Empirical Distribution: Let us now consider a long game with n long blocks, i.e., let $N = nrT$ denote the length of the long game. Following the proof of Theorem 2 in [6], the strategies (σ, τ) implement a sequence of random variables $(\mathbf{x}_t, \mathbf{y}_t, \mathbf{z}_t)_t$, where $(\mathbf{x}_t)_t$ is the i.i.d. sequence of nature with $\mathbf{x}_t \sim \mu$, as well as the sequence $(\tilde{\mathbf{x}}_t, \mathbf{y}_t, \mathbf{z}_t)_t$, for $1 \leq t \leq N$ (note, that a realization of $(\tilde{\mathbf{x}}_t)_t$ in a long block is an element in $T_m^T(\mu^{\otimes T})$. We can also denote these sequences as sequences of block sequences, i.e.,

$$(\mathbf{x}_t, \mathbf{y}_t, \mathbf{z}_t)_t = (\mathbf{x}_b^T, \mathbf{y}_b^T, \mathbf{z}_b^T)_b, \tag{2.6}$$

and

$$(\tilde{\mathbf{x}}_t, \mathbf{y}_t, \mathbf{z}_t)_t = (\tilde{\mathbf{x}}_b^T, \mathbf{y}_b^T, \mathbf{z}_b^T)_b, \tag{2.7}$$

for $1 \leq b \leq n$. Observe, that by construction, every realization of $(\tilde{\mathbf{x}}_b^T, \mathbf{y}_b^T, \mathbf{z}_b^T)_b$ has the property that in every long block after the first, it is T-block typical w.r.t. $\tilde{Q}$. Hence, the expected empirical distribution of a long block in $(\tilde{\mathbf{x}}_b^T, \mathbf{y}_b^T, \mathbf{z}_b^T)_b$ is $\tilde{Q}$:

$$\mathbb{E}[emp^T(\tilde{\mathbf{x}}[l], \mathbf{y}[l], \mathbf{z}[l])] = \tilde{Q} \quad \forall l > 1.$$

It will be useful to observe the following (equivalent) notation. If $(\tilde{\mathbf{x}}_b^T, \mathbf{y}_b^T, \mathbf{z}_b^T) \sim Q_b^T$, we have

$$\frac{1}{r}\sum_{b=1}^{r} Q_b^T = \tilde{Q}. \tag{2.8}$$

To finalise the sketch of the proof of Corollary 4, one can follow the exact same steps as in the proof of Theorem 2 in [6], to come to the conclusion that the strategies (σ, τ) implement $\tilde{Q}$.

Observe, that we assumed that $T_m^T(\tilde{Q}) \neq \emptyset$. In the following remark, one can see that this assumptions is not restrictive, since even if $T_m^T(\tilde{Q}) = \emptyset$, we can approximate $\tilde{Q}$ by a distribution $\bar{Q}$, such that $T_m^T(\bar{Q}) \neq \emptyset$.

Remark 24. If $T_m^T(\tilde{Q}) = \emptyset$, then by Lemma 4 in [6], it holds that for all $\varepsilon > 0$ there exists a distribution $\bar{Q} \in \Delta(I^T \times J^T \times K^T)$, s.t. $T_m^T(\tilde{Q}) \neq \emptyset$ with the following properties:

- $\|\bar{Q} - \tilde{Q}\| < 7\varepsilon$
- $\|\bar{Q}_I - \mu^{\otimes T}\| < 7\varepsilon$
- $\|\bar{P} - \tilde{P}\| < 7\varepsilon$
- $\|\bar{P}^{t-1;t} - P\| < 7\varepsilon, \forall t \in (1, \ldots, T-1)$

and

- $H_{\bar{Q}}(\mathbf{i}^T, \mathbf{j}^T | \mathbf{k}^T) - H_{\bar{Q}}(\mathbf{i}^T) \geq \varepsilon,$

where $\bar{P}$ denotes the marginal of $\bar{Q}$ on $I^T \times K^T$ and $\bar{P}^{t-1;t}$ denotes the marginal of $\bar{P}$ on $(K \times (I \times K))$ at stages $(t-1; t)$. ◇

Implementation of P

Let us quickly recap our approach so far. We started off with a distribution $P \in \Delta(K \times (I \times K))$ that satisfies the properties in Theorem 5. In the previous paragraphs we constructed a (block-) distribution $\tilde{Q} \in \Delta(I^T \times J^T \times K^T)$ which we derived from P in Definition 24 and we were able to show that $\tilde{Q}$ is implementable if P satisfies the information constraint. In the following, we will show that the strategies that implement $\tilde{Q}$ in Corollary 4 also implement P. An important step in this proof is the next Theorem.

Theorem 6. *Fix $\varepsilon > 0$. Let $P \in \Delta(K \times (I \times K))$ be a distribution that satisfies the properties in Theorem 5. Let (σ, τ) be the strategies that implement $\tilde{Q}$ from Definition 24 and set $n, r, t \in \mathbb{N}$ with $N = nrT$.*

Furthermore, let $(\boldsymbol{x}_t, \boldsymbol{y}_t, \boldsymbol{z}_t)_t$, $1 \leq t \leq N$ *denote the sequence of random actions induced by* (σ, τ) *and by nature's i.i.d. sequence* $(\boldsymbol{x}_t)_t$*. If* $P^N_{\mu,\sigma,\tau}$ *denotes the expected empirical 2-stage distribution over* $(K \times (I \times K))$ *of* $(\boldsymbol{x}_t, \boldsymbol{y}_t, \boldsymbol{z}_t)_t$*, then* $P^N_{\mu,\sigma,\tau}$ *is close to* P *in the* L_1*-norm i.e.,*

$$\left\| P^N_{\mu,\sigma,\tau} - P \right\| < 19\varepsilon.$$

We will split the proof of Theorem 6 into several Lemmas. First, consider the following.

Lemma 14. *Let* $\tilde{P}^m_{\mu,\sigma,\tau}$ *denote the expected empirical 2-stage distribution over* $(K \times (I \times K))$ *in a long block* $l > 1$ *of length* $m = rT$ *of the sequence* $(\tilde{\boldsymbol{x}}_t, \boldsymbol{y}_t, \boldsymbol{z}_t)_t$*, which is induced by the strategies* (σ, τ) *of the players that implement* $\tilde{Q}$*. Then,*

$$\tilde{P}^m_{\mu,\sigma,\tau} \to P \text{ as } T \to \infty.$$

Proof. Fix a long block $l > 1$, and for $1 \leq t \leq m = rT$ denote by $P^{t-1;t} \in \Delta(K \times (I \times K))$ the distribution of the random triple $(\mathbf{z}_{t-1}, \tilde{\mathbf{x}}_t, \mathbf{z}_t)$. Furthermore, for a (small) block b in l with $1 \leq b \leq r$ denote by $P^{t-1;t}_b$ the distribution of the random 2-stage triple $(\mathbf{z}^{t-1}_b, \tilde{\mathbf{x}}^t_b, \mathbf{z}^t_b)$, where $1 \leq t \leq T$. (the reader should observe that these distributions are not different from each other, they just differ w.r.t. the notation of a random triple in a (small) block, b, or in a long block. Furthermore, recall that the marginal of $\tilde{Q}$ on $I^T \times K^T$ is given by $\tilde{P}$, and by Definition 24, the marginal of $\tilde{P}$ on $K \times (I \times K)$ at stages $(t-1, t)$ is P for every $t \in (2, \ldots, T)$. Let $\tilde{P}^{t-1;t}$ denote the marginal of $\tilde{P}$ on $(K \times (I \times K))$ at stages $(t-1;t)$.
Now, if $T^T_m(\tilde{Q}) \neq \emptyset$, then by equation (2.8), it holds that

$$\frac{1}{r} \sum_{b=1}^{r} P^{t-1;t}_b = \tilde{P}^{t-1;t} = P. \tag{2.9}$$

and

$$\begin{aligned}
\tilde{P}^{rT}_{\mu,\sigma,\tau} &= \frac{1}{rT} \sum_{t=2}^{rT} P^{t-1;t} \\
&= \frac{1}{rT} \left(\sum_{t=2}^{T-1} P^{t-1;t} + \sum_{t=T+1}^{2T-1} P^{t-1;t} + \ldots + \sum_{t=(r-1)T+1}^{rT-1} P^{t-1;t} + \sum_{b=1}^{r} P^{T-1;T}_b \right) \\
&= \frac{1}{rT} \left(\sum_{t=2}^{T-1} P^{t-1;t}_1 + \sum_{t=2}^{T-1} P^{t-1;t}_2 + \ldots + \sum_{t=2}^{T-1} P^{t-1;t}_r + \sum_{b=1}^{r} P^{T;T-1}_b \right) \\
&= \frac{1}{rT} \left(\sum_{t=2}^{T-1} \sum_{b=1}^{r} P^{t-1;t}_b + \sum_{b=1}^{r} P^{T-1;T}_b \right).
\end{aligned}$$

Therefore, by substituting equation (2.9) into the last equation above, we can conclude that

$$\tilde{P}^{rT}_{\mu,\sigma,\tau} = \frac{1}{rT}\left((T-2)rP + \sum_{b=1}^{r} P_b^{T-1;T}\right) \to P \text{ for } T \to \infty.$$

The computation of $\tilde{P}^{rT}_{\mu,\sigma,\tau}$ changes slightly if $T_m^T(\tilde{Q}) = \emptyset$, but the result is not affected. By Remark 24, we know that for all $\varepsilon > 0$ there exists a distribution $\bar{Q} \in \Delta(I^T \times J^T \times K^T)$ with $T_m^T(\bar{Q}) \neq \emptyset$ that is close to $\tilde{Q}$ and has marginal $\bar{P}$ on $I^T \times K^T$. Hence, the result in (2.8) changes to

$$\frac{1}{r}\sum_{b=1}^{r} Q_b^T = \bar{Q},$$

and we have to replace equation (2.9) by

$$\frac{1}{r}\sum_{b=1}^{r} P_b^{t-1;t} = \bar{P}^{t-1;t},$$

where $\bar{P}^{t-1;t}$ denotes the marginal of $\bar{P}$ on stages $(t-1,t)$. Then, since $\left\|\bar{P}^{t-1;t} - P\right\| < 7\varepsilon, \forall t \in (2,\ldots,T)$, we can deduce that

$$\frac{1}{rT}\sum_{t=2}^{T-1} \bar{P}^{t-1;t} \to P \text{ as } T \to \infty.$$

□

So far, we have shown that the expected empirical distribution over triples $(\mathbf{z}_{t-1}, \tilde{\mathbf{x}}_t, \mathbf{z}_t)$ in $(\tilde{\mathbf{x}}_t, \mathbf{y}_t, \mathbf{z}_t)_t$ in every long block l after the first is arbitrarily close to P. One can summarise the result in the following Corollary:

Corollary 5. *For $N = nrT$, if $\tilde{P}^N_{\mu,\sigma,\tau}$ denotes the expected empirical 2-stage distribution over $(K \times (I \times K))$ of the sequence $(\tilde{\mathbf{x}}_t, \mathbf{y}_t, \mathbf{z}_t)_t$ for $1 \le t \le N$, then, assuming T large enough,*

$$\left\|\tilde{P}^N_{\mu,\sigma,\tau} - P\right\| < \varepsilon.$$

In the final Lemma, we connect the expected empirical 2-stage distribution over $(K \times (I \times K))$ of the sequence $(\tilde{\mathbf{x}}_t, \mathbf{y}_t, \mathbf{z}_t)_t$ with the expected empirical 2-stage distribution over $(K \times (I \times K))$ of the sequence $(\mathbf{x}_t, \mathbf{y}_t, \mathbf{z}_t)_t$:

Lemma 15. *Fix $\varepsilon > 0$. Let $N = rnT$ and let $r > \frac{|I|^2}{\varepsilon^3}$. Let $P^N_{\mu,\sigma,\tau}$ be given as in Theorem 6 and let $\tilde{P}^N_{\mu,\sigma,\tau}$*

be given as in Lemma 14. Then,

$$\left\|\tilde{P}^N_{\mu,\sigma,\tau} - P^N_{\mu,\sigma,\tau}\right\| < 18\varepsilon.$$

Proof. The proof of this Lemma can be deduced from the proof of Theorem 2 and from Corollary 2, both in [6]. The ideas from [6] only need to be adapted to 2− stage empirical distributions. In particular, we make use of the following notation.

For $(x^n, y^m) \in (I \times K)^m$, we write $emp_2^{K\times(I\times K)}(x^n, y^n)$ to denote the following 2- stage empirical distribution:

$$emp_2(x^n,y^n)[k',i,k] = \tfrac{1}{n-1}N(k',i,k|x^n,y^n), \quad \forall (k',i,k) \in K \times (I \times K).$$

Now, in a first step, we show that for every large block $l > 0$:

$$\left\|emp_2^{K\times(I\times K)}(\tilde{x}[l],z[l]) - emp_2^{K\times(I\times K)}(x[l],z[l])\right\| \leq \tfrac{2}{rT}\sum_{t=1}^{rT}\mathbb{1}_{\tilde{x}_t[l]\neq x_t[l]}.$$

To simplify the notation, we write $(\tilde{x},z)$ and (x,z) instead of $(\tilde{x}[l],z[l])$ and $(x[l],z[l])$, with $(\tilde{x},z),(x,z) \in (I\times K)^{rT}$. Furthermore, let $B^* \subset K\times(I\times K)$, s.t. $emp_2(\tilde{x},z)[k',i,k] > emp_2(x,z)[k',i,k], \forall (k',i,k) \in B^*$, and let B_I^* denote the respective subset on I. Then,

$$\begin{aligned}
\left\|emp_2^{K\times(I\times K)}(\tilde{x}[l],z[l]) - emp_2^{K\times(I\times K)}(x[l],z[l])\right\| &= \sum_{\substack{(k',i,k)\\ \in K\times(I\times K)}} \left|emp_2(\tilde{x},z)[k',i,k] - emp_2(x,z)[k',i,k]\right| \\
&\leq 2\sum_{(k',i,k)\in B^*} \left|emp_2(\tilde{x},z)[k',i,k] - emp_2(x,z)[k',i,k]\right| \\
&\leq \tfrac{2}{rT}\sum_{i\in B_I^*}\left(\left|\{t\geq 1 : \tilde{x} = i\}\right| - \left|\{t \geq 1 : x_t = i\}\right|\right) \\
&= \tfrac{2}{rT}\sum_{t=1}^{rT}\mathbb{1}_{\tilde{x}_t\neq x_t}
\end{aligned}$$

Now, applying Corollary 2 from [6], we have

$$\mathbb{P}\Big(\sum_{t=1}^{rT}\mathbb{1}_{\tilde{x}_t[l]\neq x_t[l]} \geq 8rT\varepsilon\Big) \leq \frac{|I|^2}{\varepsilon^2 rT},$$

and hence, due to the choice of r, $\mathbb{P}(\sum_{t=1}^{rT}\mathbb{1}_{\tilde{x}_t[l]\neq x_t[l]} \geq 8rT\varepsilon) \leq \varepsilon$. With this, we can then conclude that for $1 \leq t \leq N$,

$$\begin{aligned}
\left\|\tilde{P}^N_{\mu,\sigma,\tau} - P^N_{\mu,\sigma,\tau}\right\| &= \left\|\mathbb{E}[emp_2^{K\times(I\times K)}(\tilde{\mathbf{x}}_t,\mathbf{z}_t)_t] - \mathbb{E}[emp_2^{K\times(I\times K)}(\mathbf{x}_t,\mathbf{z}_t)_t]\right\| \\
&\leq \frac{2}{rT}\mathbb{E}\sum_{t=1}^{rT}\mathbb{1}_{\tilde{\mathbf{x}}_t\neq\mathbf{x}_t} \\
&= \frac{2}{rT}\sum_{j=1}^{rT}\mathbb{P}(\sum_{t=1}^{rT}\mathbb{1}_{\tilde{\mathbf{x}}_t\neq\mathbf{x}_t} \geq j) \\
&\leq \frac{2}{rT}(8\varepsilon rT + \varepsilon rT) \\
&= 18\varepsilon,
\end{aligned}$$

where we use the fact that the sum $\sum_{t=1}^{rT}\mathbb{1}_{\tilde{x}_t\neq x_t}$ is minimised.

□

One can now immediately observe that Theorem 6 follows from the results in Corollary 5 and in Lemma 15 by applying the triangle inequality.

Proof of Theorem 5

Let us now conclude this section with the proof of Theorem 5, which is a Corollary of Theorem 6.

Proof of Theorem 5. Let $P \in (K \times (I \times K))$ denote a distribution that satisfies the properties in Theorem 5. By Theorem 6, we know that there exist strategies (σ,τ) that N-implement a distribution $P^N_{\mu,\sigma,\tau}$ (the expected empirical 2-stage distribution over $(K \times (I \times K))$), where $P^N_{\mu,\sigma,\tau}$ is close to P in the L_1-norm. The result now follows, since by Remark 20, the set of implementable distributions is closed and contains the set of N-implementable distributions. □

Appendix A

Complementary Results for Chapter 1 in Part 1

Remark 25. Every distribution $Q \in \mathscr{Q}(t)$ is implementable, i.e., $\mathscr{Q}(t)$ is contained in $\mathscr{Q}$. $\diamond$

Proof. Let Q be a t-implementable distribution, i.e., following Notation 4, there exist strategies (σ,τ), s.t.

$$Q = \tfrac{1}{t-1}\sum_{t'=1}^{t-1} P^{t':2}_{U,\sigma,\tau}.$$

We will show that there are strategies (σ',τ'), s.t.

$$\tfrac{1}{t-1}\sum_{t'=1}^{t-1} P^{t':2}_{U,\sigma',\tau'} \to Q \text{ as } t \to \infty.$$

Consider the following game. Assume that the first t stages are induced by the strategy pair (σ,τ), so that in those first t stages Q is implemented. After the first t stages we repeat the same strategy, i.e., we induce the stages $t+1, t+2, \dots, 2t$ with the same strategy pair, (σ,τ). We repeat this process n times, i.e., we construct blocks of t stages, such that in each block the random sequences of action triples are induced by the strategy pair (σ,τ). Hence, if we consider each block seperately, we t-implement distribution Q in each such block.

Now, denote by (σ',τ') the strategy pair that implements the first tn stages of the game we just described. We can think of (σ',τ') as the n-times repeated strategy pair (σ,τ). The average distribution, Q', which is implemented after the first tn stages of the game and which is induced by (σ',τ') can be written as follows:

$$Q' = \tfrac{1}{nt-1} \sum_{t'=1}^{nt} P^{t':2}_{U,\sigma',\tau'}. \tag{A.1}$$

The sum in the above equation can be rewritten as follows:

$$\sum_{t'=1}^{nt} P^{t':2}_{U,\sigma',\tau'} = \sum_{t'=1}^{t-1} P^{t':2}_{U,\sigma,\tau} + \sum_{t'=1}^{2t-1} P^{t':2}_{U,\sigma,\tau} + \ldots + \sum_{t'=1}^{nt-1} P^{t':2}_{U,\sigma,\tau} \tag{A.2}$$

$$+ P^{t:2}_{U,\sigma',\tau'} + P^{2t:2}_{U,\sigma',\tau'} + \ldots + P^{(n-1)t:2}_{U,\sigma',\tau'} \tag{A.3}$$

$$= n(t-1)Q + \sum_{b=1}^{n-1} P^{bt:2}_{U,\sigma',\tau'}. \tag{A.4}$$

Note that every term in $\sum_{b=1}^{n-1} P^{bt:2}_{U,\sigma',\tau'}$ denotes the distribution of two consecutive action triples at a transition point from one block of t stages to the next.

We can now show for every $n \in \mathbb{N}$ that the distance between Q and Q' in the L_1-norm converges to zero as $t \to \infty$:

Since

$$Q' = \tfrac{n(t-1)}{nt-1} Q + \tfrac{1}{nt-1} \sum_{b=1}^{n-1} P^{bt:2}_{U,\sigma',\tau'},$$

we have

$$\|Q - Q'\| \leq \left(1 - \frac{n - \frac{n}{t}}{n - \frac{1}{t}}\right) \|Q\| + \frac{1}{nt+1} \left\| \sum_{b=1}^{n-1} P^{bt:2}_{U,\sigma',\tau'} \right\|.$$

Hence, since both terms on the right-hand side in the expression above converge to 0 as $t \to \infty$, the result follows.

Appendix B

Complementary Results for Chapter 2 in Part 1

Proposition 3. *Let* $\mu \in \Delta(I)$ *be a convex combinations of two distributions* $\mu_0, \mu_1 \in \Delta(I)$*, i.e.,* $\mu = \beta_0\mu_0 + \beta_1\mu_1$ *and let* $X \sim \mu$*. Furthermore, denote by* Y *a binary indicator variable with the following property: let* $X \sim \mu_0$*, if* $Y = 0$ *and let* $X \sim \mu_1$*, if* $Y = 1$*. Then,*

$$\beta_0 D(\mu_0||\mu) + \beta_1 D(\mu_1||\mu) \leq H(Y)$$

with $D(\cdot||\cdot)$ *denoting the Kullback-Leibler distance.*

Proof.

$$
\begin{aligned}
\beta_0 D(\mu_0||\mu) + \beta_1 D(\mu_1||\mu) &= \beta_0 \sum_{x\in I} \mu_0(x)(\log(\mu_0(x)) - \log(\mu(x))) \\
&+ \beta_1 \sum_{x\in I} \mu_1(x)(\log(\mu_1(x)) - \log(\mu(x))) \\
&= -\sum_{x\in I} (\beta_0\mu_0 + \beta_1\mu_1)(x)\log(\mu(x)) \\
&+ \beta_0 \sum_{x\in I} \mu_0(x)\log(\mu_0(x)) + \beta_1 \sum_{x\in I} \mu_1(x)\log(\mu_1(x)) \\
&= -\sum_{x\in I} \mu(x)\log(\mu(x)) - \beta_0 H(\mu_0) - \beta_1 H(\mu_1) \\
&= H(X) - H(X|Y) \\
&= H(Y) - H(Y|X) \\
&\leq H(Y)
\end{aligned}
$$

□

Appendix C

Complementary Results for Chapter 3 in Part 1

C.1 Approximation of Probabilities

Lemma 16. *Let $\mu \in \Delta(I)$ be a distribution with full support. For all $\varepsilon > 0$ and for all $P_2' \in \Delta(A^2)$ ($A = I \times J \times K$) with $P_2'(i|a) = \mu(i)$, for all $i \in I$ and $a \in A$, identical marginals P' on A, and with $H_{P_2'}(\boldsymbol{i},\boldsymbol{j}|\boldsymbol{k},\boldsymbol{i}',\boldsymbol{j}',\boldsymbol{k}') \geq H_{P_2'}(\mu)$, $\exists P_2 \in \delta(A^2)$ with full support, s.t.*

$$H_{P_2}(\boldsymbol{i},\boldsymbol{j}|\boldsymbol{k},\boldsymbol{i}',\boldsymbol{j}',\boldsymbol{k}') - H(\mu) \geq \varepsilon \tag{C.1}$$

and

$$\left\|P_2 - P_2'\right\|_1 \leq 2\varepsilon \tag{C.2}$$

Proof. Fix $\varepsilon > 0$. Denote by P_I the marginal of P on I. Observe, that since $P_2'(i|a) = \mu(i)$, it also holds that the marginal of P' on I, P_I, is equal to μ.

Let $R = (P_I' \times U_J \times U_K)$ and $R_2 = R \times R$, where U_J and U_K denote the uniform distribution on J and K respectively, define $P_2 = \varepsilon R_2 + (1-\varepsilon)P_2'$. Observe, that P_2 has full support (since R_2 has full support) and it holds that $P_2(i|a) = \mu(i)$, for all $i \in I$ and $a \in A$.

Now, by the concavity of the entropy function $P_2 \mapsto H_{P_2}(\mathbf{i},\mathbf{j}|\mathbf{k},\mathbf{i}',\mathbf{j}',\mathbf{k}')$, it holds that

$$\begin{aligned} H_{P_2}(\mathbf{i},\mathbf{j}|\mathbf{k},\mathbf{i}',\mathbf{j}',\mathbf{k}') &\geq \varepsilon H_{R_2}(\mathbf{i},\mathbf{j}|\mathbf{k},\mathbf{i}',\mathbf{j}',\mathbf{k}') + (1-\varepsilon) H_{P_2'}(\mathbf{i},\mathbf{j}|\mathbf{k},\mathbf{i}',\mathbf{j}',\mathbf{k}') \\ &\geq \varepsilon H_{R_2}(\mathbf{i}) + \varepsilon \log_2 |J| + (1-\varepsilon) H(\mu) \\ &= \varepsilon + H(\mu). \end{aligned}$$

Hence, it follows that

$$H_{P_2}(\mathbf{i},\mathbf{j}|\mathbf{k},\mathbf{i}',\mathbf{j}',\mathbf{k}') - H(\mu) \geq \varepsilon.$$

Furthermore,

$$\begin{aligned} \left\|P_2 - P_2'\right\|_1 &= \left\|\varepsilon R_2 + (1-\varepsilon)P_2' - P_2'\right\| && \text{(C.3)} \\ &= \varepsilon \left\|R_2 - P_2'\right\| && \text{(C.4)} \\ &\leq 2\varepsilon. && \text{(C.5)} \end{aligned}$$

□

C.2 Locally Typical Sequences

Lemma 17. *If $N > \frac{2l}{\varepsilon}$, then every l-locally typical sequence $x^N \in T^{\varepsilon}_{N,l}(P)$ is 2ε-typical, i.e.,*

$$\left|\tfrac{1}{N} N(a|x^N) - P(a)\right| < 2\varepsilon \quad \forall a \in A.$$

Proof. Let $x^N \in T^{\varepsilon}_{N,l}(P)$. First, let $N = ql$, $1 \leq q \in \mathbb{N}$. Then, for every $a \in A$:

$$\begin{aligned} |N(a|x^N) - NP(a)| &= \left|\sum_{t=0}^{q-1} (N(a|x^N_{tl+1,l}) - lP(a))\right| \\ &\leq \sum_{t=0}^{q-1} |N(a|x^N_{tl+1,l}) - lP(a)| \\ &< ql\varepsilon, \end{aligned}$$

where the last inequality follows directly from the definition of l-locally typical sequences.

Now, let $N = ql + l_d$, with $0 < l_d < l$. Denote by x^{ql} the first ql coordinates and by x^{l_d} the last l_d coordinates of x^N. Then,

$$
\begin{aligned}
|N(a|x^N) - nP(a)| &= |N(a|x^{ql}) - qlP(a) + N(a|x^{l_d}) - l_dP(a)| \\
&\leq |N(a|x^{ql}) - qlP(a)| + |N(a|x^{l_d}) - l_dP(a)|. \qquad \text{(C.6)}
\end{aligned}
$$

From above, we know that $|N(a|x^{ql}) - qlP(a)| < ql\varepsilon$. Furthermore, we have

$$|N(a|x^{l_d}) - l_dP(a)| < l(1 + \tfrac{1}{\varepsilon})\varepsilon,$$

since for every $t \in (1, \ldots, N - l)$,

$$
\begin{aligned}
N(a|x^{l_d}) - l_dP(a) &\leq N(a|x^N_{t,l}) \\
&< l(P(a) + \varepsilon) \\
&< l(1 + \varepsilon) \\
&= l(1 + \tfrac{1}{\varepsilon})\varepsilon,
\end{aligned}
$$

and

$$
\begin{aligned}
N(a|x^{l_d}) - l_dP(a) &> -l_dP(a) \\
&> -l \\
&> -l(1 + \tfrac{1}{\varepsilon})\varepsilon.
\end{aligned}
$$

Returning to inequality (C.6), we conclude

$$|N(a|x^N) - NP(a)| < ql\varepsilon + l(1 + \tfrac{1}{\varepsilon})\varepsilon,$$

and for N large enough, i.e., $N > l(1 + \frac{1}{\varepsilon})$, or, more generously, for $N > \frac{2l}{\varepsilon}$, we have

$$|N(a|x^n) - NP(a)| < 2N\varepsilon.$$

□

The following Corollary follows directly from the above Lemma C.2 and from the property of ε-typical sequences (stated in the Preliminaries):

Corollary 6. *Let $X^n = (X_1, X_2, \ldots, X_n)$ be an i.i.d. sequence with $X_i \sim P$. Denote by x^n a realization of X^n. If $x^n \in T^{\varepsilon}_{n,l}(P)$, and if $n > \frac{2l}{\varepsilon}$, then*

$$2^{-n(H(X)+2c\varepsilon)} \leq P^n(x^n) \leq 2^{-n(H(X)-2c\varepsilon)},$$

where P^n denotes the product distribution derived from P and $c = -\sum_{a\in A} \log P(a)$.

Lemma 18. *Let $X_1, X_2, \ldots$ be drawn i.i.d according to $P \in \Delta(A)$ and let $X^n = (X_1, \ldots, X_n)$. Then, for every $l > \frac{\ln(2n|A|)}{2\varepsilon^2}$, we have*

$$\lim_{n\to\infty} Pr(X^n \in T^{\varepsilon}_{n,l}(P)) = 1.$$

Proof. First, note that

$$\begin{aligned}
Pr(X^n \in T^{\varepsilon}_{n,l}(P)) &= Pr(\forall t \in (1, \ldots, n-l) : X^n_{t,l} \in T^{\varepsilon}_l(P)) \\
&= 1 - Pr(\exists t \in (1, \ldots, n-l) : X^n_{t,l} \notin T^{\varepsilon}_l(P)) \\
&= 1 - Pr(\bigcup_{t=1}^{n-l} X^n_{t,l} \notin T^{\varepsilon}_l(P)) \\
&\geq 1 - \sum_{t=1}^{n-l} Pr(X^n_{t,l} \notin T^{\varepsilon}_l(P))
\end{aligned}$$

Furthermore, for every $t \in (1, \ldots, n-l)$,

$$\begin{aligned}
Pr(X^n_{t,l} \notin T^{\varepsilon}_l(P)) &= Pr(\exists a \in A : |\tfrac{1}{l} N(a|x^n_{t,l}) - P(a)| \geq \varepsilon) \\
&= Pr(\bigcup_{a\in A} |\tfrac{1}{l} N(a|x^n_{t,l}) - P(a)| \geq \varepsilon) \\
&\leq \sum_{a\in A} Pr(|\tfrac{1}{l} N(a|x^n_{t,l}) - P(a)| \geq \varepsilon) \\
&\leq |A| 2\exp(-2l\varepsilon^2),
\end{aligned}$$

where the last inequality is a direct application of Hoeffding's inequality.

Therefore,

$$\sum_{t=1}^{n-l} Pr(X_{t,l}^n \notin T_{n,l}^{\varepsilon}(P)) \leq \sum_{t=1}^{n-l} |A| 2\exp(-2l\varepsilon^2)$$
$$= (n-l)|A| 2\exp(-2l\varepsilon^2)$$
$$\to 0 \text{ as } n \to \infty,$$

and the last expression goes to 0 as $n \to \infty$ for every $l > \sqrt{n}$, or, in dependence of ε, for every $l > \frac{\ln(2n|A|)}{2\varepsilon^2}$. This completes the proof. □

Corollary 7. *Let $X \sim P$ and set $c = -\sum_{a\in A} \log P(a)$. Let n, l be integers with $n > \frac{2l}{\varepsilon}$ and $l > \frac{\ln(2n|A|)}{2\varepsilon^2}$. If n is sufficiently large, then*

$$(1-\delta)2^{n(H(X)-2c\varepsilon)} \leq |T_{n,l}^{\varepsilon}(P)| \leq 2^{n(H(X)+2c\varepsilon)},$$

for every $\delta > 0$.

Proof. Since

$$Pr(X^n \in T_{n,l}^{\varepsilon}(P)) = \sum_{x^n \in T_{n,l}^{\varepsilon}(P)} P^n(x^n),$$

with

$$2^{-n(H(X)+2c\varepsilon)} \leq P^n(x^n) \leq 2^{-n(H(X)-2c\varepsilon)}$$

(by Corollary 6), we get

$$\sum_{x^n \in T_{n,l}^{\varepsilon}(P)} 2^{-n(H(X)+2c\varepsilon)} \leq Pr(X^n \in T_{n,l}^{\varepsilon}(P)) \leq \sum_{x^n \in T_{n,l}^{\varepsilon}(P)} 2^{-n(H(X)-2c\varepsilon)},$$

equivalently,

$$2^{-n(H(X)+2c\varepsilon)}|T_{n,l}^{\varepsilon}(P)| \leq Pr(X^n \in T_{n,l}^{\varepsilon}(P)) \leq 2^{-n(H(X)-2c\varepsilon)}|T_{n,l}^{\varepsilon}(P)|.$$

By Lemma 18, we know that for n sufficiently large the following relation holds for every $\delta > 0$:

$$1-\delta \leq Pr(X^n \in T_{n,l}^{\varepsilon}(P)) \leq 1,$$

and so the result follows. □

C.3 Results on Locally Conditional Typical Sequences

We derive a couple of very helpful properties of locally conditional typical sequences.

Lemma 19. *Let $x^n \in T^{\varepsilon}_{n,l}(P_I)$. Then, if $n > \frac{2l}{\varepsilon}$, every locally conditional typical sequence $y^n \in T^{\varepsilon}_{n,l}(P|x^n)$ is conditionally 2ε-typical, i.e., for every $(i,k) \in I \times K$ the following holds:*

$$|N(i,k|x^n,y^n) - \rho^i(k)N(i|x^n)| < 2n\varepsilon.$$

Remark 26. Observe that in Definition 10 we assumed P to have full support. This assumptions simplifies the Definition of conditional ε-typical sequences slightly, since for all $(i,k) \in I \times K$, it holds that $\rho^i(k) > 0$. ◇

Proof of Lemma 19. The proof of Lemma 19 is very similar in structure to the proof of Lemma 4. Therefore, we only provide a sketch of the proof and highlight the parts that slightly differ.
Let $x^n \in T^{P_I}_{n,l}$. Recall, that this means that for every $t \in (1,\dots,(n-l))$ and for every $i \in I$ the following inequality holds:

$$N(i|x^n_{t,l}) < l(\varepsilon + P_I(i)). \tag{C.7}$$

Similar to the proof of Lemma 4, we first assume $n = ql$, with $q \in \mathbb{N}$ and $q \geq 1$. It then trivially holds, that for all $(i,k) \in I \times K$,

$$|N(i,k|x^n,y^n) - \rho^i(k)N(i|x^n)| < n\varepsilon. \tag{C.8}$$

We now set $n = ql + l_d$ with $0 < l_d < l$ and denote by (x^{ql}, y^{ql}) the first ql coordinates and by (x^{l_d}, y^{l_d}) the last l_d coordinates of (x^n, y^n). By the triangle inequality, we have for all $(i,k) \in I \times K$

$$\begin{aligned}|N(i,k|x^n,y^n) - \rho^i(k)N(i|x^n)| \leq\ & |N(i,k|x^{ql},y^{ql}) - \rho^i(k)N(i|x^{ql})| \\ & + |N(i,k|x^{l_d},y^{l_d}) - \rho^i(k)N(i|x^{l_d})|.\end{aligned}$$

Since by inequality C.8 we have

$$|N(i,k|x^{ql},y^{ql}) - \rho^i(k)N(i|x^{ql})| < ql\varepsilon,$$

it remains to bound $|N(i,k|x^{l_d},y^{l_d}) - \rho^i(k)N(i|x^{l_d})|$.

In the same manner as in the proof of Lemma 4 and taking into account inequality C.7, we can show

$$|N(i,k|x^{l_d},y^{l_d}) - \rho^i(k)N(i|x^{l_d})| < l(2+\tfrac{1}{\varepsilon})\varepsilon.$$

From here, we can directly deduce the desired lower bound for $|N(i,k|x^n,y^n) - \rho^i(k)N(i|x^n)|$ for $n > \frac{2l}{\varepsilon}$. □

Similar to Corollary 6, the following corollary can now be directly deduced from Lemma 19 and from the remark about conditional ε-typical sequences occurring, stated in the Preliminaries:

Corollary 8. *Let $x^n \in T^{\varepsilon}_{n,l}(P_I)$ and let $Y_1, Y_2, \ldots$ be random variables in K^n distributed according to $\{\rho^i : i \in I\}$ and x^n, i.e., $Pr(Y_j = k) = \rho^{x_j}(k)$, for every $j \geq 1$. Denote by y^n a realization of $Y^n = (Y_1, \ldots, Y_n)$. If $y^n \in T^{\varepsilon}_{n,l}(P|x^n)$ and $c' = \sum_{(i,k)\in I\times K}(\log\rho^i(k) - \rho^i(k)\log\rho^i(k))$, and if $n > \frac{2l}{\varepsilon}$, then*

$$2^{-n(H(k|i)+2c'\varepsilon)} \leq P^n(y^n|x^n) \leq 2^{-n(H(k|i)-2c'\varepsilon)},$$

where $P^n(y^n|x^n) = \prod_{j=1}^n \rho^{x_j}(y_j)$.

Again, similar to Lemma 18, we can make the following statement about a random sequence Y^n that is distributed according to $\{\rho^i : i \in I\}$ and x^n:

Lemma 20. *Let $x^n \in T^{\varepsilon}_{n,l}(P_I)$ and let $Y_1, Y_2, \ldots$ be independent random variables distributed according to $\{\rho^i : i \in I\}$ and x^n. Let $Y^n = (Y_1, \ldots, Y_n)$, then, for every $l > \frac{\ln(2n|A|)}{2\varepsilon^2}$,*

$$\lim_{n\to\infty} Pr(Y^n \in T^{\varepsilon}_{n,l}(P|x^n)) = 1.$$

Proof. The proof of this lemma is similar in structure to the proof of Lemma 18. It remains to show the following: let $Y^l = Y_1, \ldots, Y_l$ denote a subsequence of Y^n of length l, and let $x^l \in T^{\varepsilon}_l(P_I)$. Then,

$$Pr(Y^l \notin T^{\varepsilon}_l(P|x^l)) \leq |I||J|2\exp(-2l\varepsilon^2). \tag{C.9}$$

To prove the above inequality, we first order the sequence x^l according to the set I, i.e., we write

$$x^l = \underbrace{x_1,\ldots,x_1}_{N(x_1|x^l)},\underbrace{x_2,\ldots,x_2}_{N(x_2|x^l)},\ldots\underbrace{x_{|I|},\ldots,x_{|I|}}_{N(x_{|I|}|x^l)},.$$

Now, according to the order of x^l, we group together those random variables in Y^l, that have the same conditioning event, i.e.,

$$Y^l = Y^{N(x_1|x^l)}, Y^{N(x_2|x^l)}, \ldots, Y^{N(x_{|I|}|x^l)},$$

where for every $x \in I$, $Y^{N(x|x^l)} = Y_1^x, \ldots, Y_{N(x|x^l)}^x$ is now an i.i.d. sequence with $Y_k^x \sim P(\cdot|x)$, for all $k \in (1, \ldots, N(x|x^l))$. Now, the event $Y^l \in T_l^\varepsilon(P|x^l)$ is equivalent to the event that for all $x \in I$, $Y^{N(x|x^l)} \in T_{N(x|x^l)}^\varepsilon(P(\cdot|x))$. Therefore, the following holds:

$$\begin{aligned} Pr(Y^l \notin T_l^\varepsilon(P|x^l)) &= Pr\left(\exists (x,y) \in I \times J : \left|\frac{N(y|Y^{N(x|x^l)})}{N(x|x^l)} - P(y|x)\right| \geq \varepsilon\right) \\ &\leq \sum_{x,y \in I \times J} Pr\left(\left|\frac{N(y|Y^{N(x|x^l)})}{N(x|x^l)} - P(y|x)\right| \geq \varepsilon\right) \\ &\leq |I||J| 2\exp(-2l\varepsilon^2), \end{aligned}$$

where the last inequality is again a direct application of the Hoeffding inequality. This can be easily seen as follows. Since

$$\frac{N(y|Y^{N(x|x^l)})}{N(x|x^l)} = \frac{1}{N(x|x^l)} \sum_{k=1}^{N(x|x^l)} \mathbb{1}_{\{Y_k^x = y\}},$$

and since $Y_1^x, \ldots, Y_{N(x|x^l)}^x$ is an i.i.d. sequence with law $P(\cdot|x)$,

$$\mathbb{E}\left[\frac{1}{N(x|x^l)} \sum_{k=1}^{N(x|x^l)} \mathbb{1}_{\{Y_k^x = y\}}\right] = P(y|x),$$

the result follows from the Hoeffding inequality. □

Finally, we can deduce the following result w.r.t. a conditional ε-typical set:

Corollary 9. *Let $x^n \in T_{n,l}^\varepsilon(P_I)$ and let $n > \frac{2l}{\varepsilon}$. For n sufficiently large,*

$$(1-\delta')2^{n(H(k|i)-2c'\varepsilon)} \leq |T_{l,n}^\varepsilon(P|x^n)| \leq 2^{n(H(k|i)+2c'\varepsilon)},$$

for every $\delta' > 0$ and with $c' = \sum_{(i,k) \in I \times K}(\log \rho^i(k) - \rho^i(k) \log \rho^i(k))$.

Proof. The proof follows the exact same structure as the proof of Corollary 7. □

C.4 Mapping a Strictly Positive Stochastic Matrix to its unique Stationary Distribution

In this section we add the proof of Lemma 21, which is stated below for completeness.

Lemma 21. *Let A be a finite set and denote by $\mathscr{M}$ the set of all strictly positive stochastic matrices over A. The mapping $f : \mathscr{M} \to \Delta(A)$, with $f(\mathbb{P}) = \pi$, where $\mathbb{P} \in \mathscr{M}$ with (unique) stationary distribution π, is continuous.*

Proof. Let $\mathbb{P}$ denote a strictly positive stochastic matrix and let $\{X_k\} = X_0, X_2, \ldots$ be a markov chain with transitions according to $\mathbb{P}$ and with initial distribution $\nu \in \Delta(A)$ (given as a column probability vector), i.e., $X_0 \sim \nu$. For every $k > 0$, let g_k^ν be a function from $\mathscr{M}$, the set of all strictly positive stochastic matrices, to $\Delta(A)$, with $g_k^\nu(\mathbb{P}) = \nu^T \mathbb{P}^k$ (where ν^T denotes the transpose of ν). Note, that $g_k^\nu(\mathbb{P})$ describes the distribution of the kth coordinate of the markov chain $\{X_k\}$, since $X_k \sim \nu^T \mathbb{P}^k$. Furthermore, it is well known that for every distribution $\nu \in \Delta(A)$ and for every $\mathbb{P} \in \mathscr{M}$ it holds that

$$\lim_{k\to\infty} g_k^\nu(\mathbb{P}) = f(\mathbb{P})$$

(see, for instance, Chapter 1.8 in [11]). With f being the limit of the functions $\left\{g_k^\nu\right\}_k$, it is sufficient to show that the functions $\left\{g_k^\nu\right\}_k$ converge uniformly to f in a neighborhood of $\mathbb{P}$. Since each function g_k^ν is continuous, it then follows from the uniform limit theorem that f is continuous. Hence, it is sufficient to show the following claim:

Claim 3. *For every distribution $\nu \in \Delta(A)$ and for every $\mathbb{P} \in \mathscr{M}$ with stationary distribution π, we have for every $k > 0$,*

$$\|g_k^\nu(\mathbb{P}) - f(\mathbb{P})\|_{TV} \leq (1-p)^k,$$

where $p = \min_{i,j\in A} p(i,j)$.

We prove the claim via a coupling argument (coupling methods are introduced e.g., in Chapter 4 in [8]): in addition to $\{X_k\}$, let $\{Y_k\}$ denote another markov chains with transition matrix $\mathbb{P}$. We have $X_0 \sim \nu$ and we set $Y_0 \sim \pi$, where π is the stationary distribution of $\mathbb{P}$. Define a coupling $\{X_k, Y_k\}_k$ of these chains as follows. Let $t \geq 0$ denote the first time the two markov chains meet. Then, for all $k \geq t$, let $X_k = Y_k$ and for all $k \leq t$, assume X_k and Y_k to be independent. Obviously, the bivariate chain $\{X_k, Y_k\}_k$ is a Markov

Chain over state space $A \times A$ that has a transition matrix $\mathbb{Q}$ with entries $q((a',b'),(a,b))$ that satisfy

$$q((a',b'),(a,b)) = \begin{cases} p(a',a)p(b',b) & \text{if } a' \neq b' \\ p(a',a) & \text{if } a' = b' \text{ and } a = b \\ 0 & \text{if } a' = b' \text{ and } a \neq b. \end{cases}$$

Recall, that $X_k \sim \nu^T \mathbb{P}^k$ and we write $\nu^T \mathbb{P}^k(a) = Pr(X_k = a)$, for all $a \in A$. Furthermore, since π is the stationary distribution of $\mathbb{P}$, $Y_k \sim \pi$ for every $k \geq 0$. Therefore, for all $a \in A$,

$$\begin{aligned} \nu^T \mathbb{P}^k(a) - \pi(a) &= Pr(X_k = a) - Pr(Y_k = a) \\ &= Pr(X_k = a, t \leq k) + Pr(X_k = a, t > k) \\ &- (Pr(Y_k = a, t \leq k) + Pr(Y_k = a, t > k)) \\ &= Pr(X_k = a, t > k) - Pr(Y_k = a, t > k) \\ &\leq Pr(t > k) \\ &= Pr(X_k \neq Y_k), \end{aligned}$$

where the third equality follows since $X_k = Y_k$ for all $t \leq k$.

Since $\nu^T \mathbb{P}^k = g_k^\nu(\mathbb{P})$, our results so far summarize to

$$\|g_k^\nu(\mathbb{P}) - \pi\|_{TV} \leq Pr(X_k \neq Y_k). \tag{C.10}$$

Now, by construction of the coupling, $Pr(X_k \neq Y_k | X_{k-1} = Y_{k-1}) = 0$, hence,

$$Pr(X_k \neq Y_k) = Pr(X_{k-1} \neq Y_{k-1}) Pr(X_k \neq Y_k | X_{k-1} \neq Y_{k-1}).$$

Furthermore, for every $(b,c) \in A^2$, $b \neq c$,

$$\begin{aligned} Pr(X_k = Y_k | X_{k-1} = b, Y_{k-1} = c) &= \sum_{a \in A} q((b,c),(a,a)) \\ &= \sum_{a \in A} p(b,a)p(c,a) \\ &> p, \end{aligned}$$

and thus $Pr(X_k = Y_k | X_{k-1} \neq Y_{k-1}) > p$.

The claim now follows from

$$\begin{aligned} Pr(X_k \neq Y_k) &< Pr(X_{k-1} \neq Y_{k-1})(1-p) \\ &< (1-p)^k \end{aligned}$$

and from inequality (C.10). □

C.5 The Existence and the Size of the Set of Action Plans - with Locally Typical Sequences

The next Theorem provides the key to prove the size of the set of action plans of the agent in Lemma 8:

Theorem 7. *Let $P \in \Delta(I \times K)$ be a distribution with marginals P_I on I and P_K on K. Let $(\boldsymbol{i},\boldsymbol{k})$ denote a pair of random variables distributed according to P. For any fixed $\xi > 0$, let $M = 2^{N(H(\boldsymbol{k})-H(\boldsymbol{k}|\boldsymbol{i})+\xi)}$ and let $S = \{s_1, s_2, \ldots s_M\}$ be a set of i.i.d. random variables uniformly distributed in $T^{\varepsilon}_{n,l}(P_K)$. Now, the probability that for every locally typical sequence $x^n \in T^{\varepsilon}_{N,l}(P_I)$, there is at least one element $s \in S$, s.t. $(x,s) \in T^{\varepsilon}_{N,l}(P)$ goes to 1 as $n \to \infty$, i.e.,*

$$\lim_{n\to\infty} Pr(\forall x^n \in T^{\varepsilon}_{n,l}(P_I) \exists s \in S : (x,s) \in T^{\varepsilon}_{l,n}(P)) = 1$$

Proof. We show

$$\lim_{n\to\infty} Pr(\exists x^n \in T^{\varepsilon}_{n,l}(P_I) : \forall s \in S : (x,s) \notin T^{\varepsilon}_{n,l}(P)) = 0.$$

First, applying the union bound, we have

$$
\begin{aligned}
Pr(\exists x^n \in T^{\varepsilon}_{n,l}(P_I) : \forall s \in S : (x^n, s) \notin T^{\varepsilon}_{n,l}(P)) &= Pr(\bigcup_{x^n \in T^{\varepsilon}_{n,l}(P_I)} \forall s \in S : (x^n, s) \notin T^{\varepsilon}_{n,l}(P)) \\
&\leq \sum_{x^n \in T^{\varepsilon}_{n,l}(P_I)} Pr(\forall s \in S : (x^n, s) \notin T^{\varepsilon}_{n,l}(P)) \\
&= \sum_{x^n \in T^{\varepsilon}_{n,l}(P_I)} Pr(\bigcap_{i=1}^{M} s_i \notin T^{\varepsilon}_{n,l}(P|x^n)) \\
&= \sum_{x^n \in T^{\varepsilon}_{n,l}(P_I)} \prod_{i=1}^{M} Pr(s_i \notin T^{\varepsilon}_{n,l}(P|x^n)) \\
&= \sum_{x^n \in T^{\varepsilon}_{n,l}(P_I)} \prod_{i=1}^{M} \big(1 - Pr(s_i \in T^{\varepsilon}_{n,l}(P|x^n))\big) \\
&< |T^{\varepsilon}_{n,l}(P_I)| \big(1 - Pr(s_1 \in T^{\varepsilon}_{n,l}(P|\bar{x}^n))\big)^M,
\end{aligned}
$$

where the last inequality holds since $Pr(s_i \in T^{\varepsilon}_{n,l}(P|x^n))$ is identical for all $i \in (1, \ldots, M)$ and $\bar{x}^n = \arg\max_{x^n \in T^{\varepsilon}_{n,l}(P_I)} \prod_{i=1}^{M} \Big(1 - Pr(s_i \in T^{\varepsilon}_{n,l}(P|x^n))\Big)$. Now, observe that for any $0 < \alpha < 1$, we have $(1 - \alpha)^{\frac{1}{\alpha}} < e^{-1}$. Furthermore, recall that from Corollary 7, $|T^{\varepsilon}_{n,l}(P_I)| \leq 2^{n(H(i)+2c\varepsilon)}$ with $c = -\sum_{i \in I} \log P_I(i)$ hence with $\alpha = Pr(s_1 \in T^{\varepsilon}_{n,l}(P|\bar{x}^n))$ we have

$$
\begin{aligned}
|T^{\varepsilon}_{n,l}(P_I)| \big(1 - Pr(s_1 \in T^{\varepsilon}_{n,l}(P|\bar{x}^n))\big)^M &= |T^{\varepsilon}_{n,l}(P_I)|(1-\alpha)^M \\
&= |T^{\varepsilon}_{n,l}(P_I)| \left((1-\alpha)^{\frac{1}{\alpha}}\right)^{\alpha M} \\
&< |T^{\varepsilon}_{n,l}(P_I)| \exp(-\alpha M) \\
&< 2^{n(H(i)+2c\varepsilon)} \exp(-\alpha M).
\end{aligned}
$$

Since

$$
\alpha = Pr(s_1 \in T^{\varepsilon}_{n,l}(P|\bar{x}^n)) = \frac{|T^{\varepsilon}_{n,l}(P|\bar{x}^n)|}{|T^{\varepsilon}_{n,l}(P_K)|},
$$

we can apply Corollaries 7 and 9 with $c' = \sum_{i,k \in I \times K} \log \rho^i(k)(1 - \rho^i(k))$ and $\tilde{c} = -\sum_{k \in K} \log P_K(k)$ to get

$$
\begin{aligned}
\alpha = \frac{|T^{\varepsilon}_{n,l}(P|\bar{x}^n)|}{|T^{\varepsilon}_{n,l}(P_K)|} &\geq \frac{(1-\delta')2^{n(H(k|i)-2c'\varepsilon)}}{2^{n(H(k)+2\tilde{c}\varepsilon)}} \\
&= (1-\delta')2^{n(H(k|i)-H(i)-2\varepsilon(c'+\tilde{c}))} \\
&> (1-\delta')2^{n(H(k|i)-H(i)-2\varepsilon\tilde{c})}
\end{aligned}
$$

(note that $c' < 0$).Hence, with $M = 2^{n(H(k)-H(k|i)+\xi)}$ for $\xi > 2\varepsilon\tilde{c}$,

$$\begin{aligned}\alpha M &\geq (1-\delta')2^{n(H(k|i)-H(i)-2\varepsilon(c'+\tilde{c}))}2^{n(H(k)-H(k|i)+\xi)}\\ &= (1-\delta')2^{n(\xi-2\varepsilon\tilde{c})},\end{aligned}$$

therefore

$$2^{n(H(i)+2c\varepsilon)}\exp(-\alpha M) \to 0 \text{ as } n \to \infty,$$

and hence

$$\lim_{n\to\infty} Pr(\exists x \in T^{\varepsilon}_{n,l}(P_I) : \forall s \in S : (x,s) \notin T^{\varepsilon}_{n,l}(P)) = 0.$$

This completes the proof. □

C.6 *m*-Extendable Locally Conditional Typical Sequences

In the following, we prove important properties of *m*-extendable locally typical conditional sequences which play an important role for the size of the message set of the forecaster.

First, consider the following observation that applies Lemma 20 to *m*-extendable locally typical conditional sequences:

Observation 7. Let $n,m,l \in \mathbb{N}$ be integers with $n,m > l$ and let $x^{n+m} \in T^{\varepsilon}_{n+m,l}(P_I)$. Assume $Y^{n+m} = (Y_1,\dots,Y_{n+m})$ is a sequence of independent random variables with $Y_i \sim P(\cdot|x_i)$ for every $i \in (1,\dots,n+m)$. We refer to $Y^n = (Y_1,\dots,Y_n)$ as the prefix of Y^{n+m}. Then, for every $l \in \mathbb{N}$ with $l > \frac{\ln(2n|A|)}{2\varepsilon^2}$ and $l > \frac{\ln(2m|A|)}{2\varepsilon^2}$, the probability that the prefix of Y^{n+m} is an *m*-extendable locally conditional typical sequence converges to 1, i.e., it holds that

$$\lim_{n\to\infty} Pr(Y^n \in T^{\varepsilon}_{n,l}(P|x^n, ext(m))) = 1.$$

◇

The following Theorem applies Lemma 20 and Observation 7 and provides a vital result to deduce the size of an *m*-extendable locally typical set.

Theorem 8. *Let $(X,Y) \sim P$, fix $\varepsilon > 0$ and $x^n \in T^{\varepsilon}_{n,l}(P_I)$. Let the integers n,m,l and the random sequence Y^{n+m} have the same properties as in Observation 7. Denote by EXT the event that the prefix of Y^{n+m} is* *extendable, i.e.,* $EXT = Y^n \in T^{\varepsilon}_{n,l}(P|x^n, ext(m))$*. Then,*

$$\tfrac{1}{n}H(Y^n|EXT) \to H(Y|X) - 2\varepsilon d, \quad \text{as } n \to \infty,$$

with $d = \sum_{a \in I} H(Y|X = a)$*.*

Proof. First, we show

$$H(Y^n) > n(H(Y|X) - 2\varepsilon d).$$

Since Y^n is a sequence of independent random variables, it holds that

$$\begin{aligned} H(Y^n) &= \sum_{i=1}^{n} H(Y_i) \\ &= \sum_{i=1}^{n} (-\sum_{y \in J} P(y|x_i) \log P(y|x_i)) \\ &= \sum_{i=1}^{n} H(Y|X = x_i), \end{aligned} \tag{C.11}$$

where $x_i \in I$ is the i-th element in the locally typical sequence $x^n \in T^{\varepsilon}_{n,l}(P_I)$. Since P has full support, it also holds that for every element $a \in I$, we have $N(a|x^n) > n(P(a) - 2\varepsilon)$. Now, we group the last sum (C.11) into sums of entropies with the same conditioning event:

$$\begin{aligned} \sum_{i=1}^{n} H(Y|X = x_i) &= \sum_{a \in I} \sum_{i: x_i = a} H(Y|X = x_i) \\ &= \sum_{a \in I} N(a|x^n) H(Y|X = a) \\ &> \sum_{a \in I} n(P(a) - 2\varepsilon) H(Y|X = a) \\ &= n\left(\sum_{a \in I} P(a) H(Y|X = a) - 2\varepsilon \sum_{a \in I} H(Y|X = a)\right) \\ &= n(H(Y|X) - 2\varepsilon \sum_{a \in I} H(Y|X = a)), \end{aligned}$$

which completes the first part of the proof.

Next, let $\mathbb{1}_{\{EXT\}}$ denote a random indicator function with $Pr(\mathbb{1}_{\{EXT\}} = 1) = Pr(EXT)$ and $Pr(\mathbb{1}_{\{EXT\}} =$

$0) = Pr(EXT^c)$, where EXT^c denotes the complement of EXT, i.e., $Y^n \notin T^{\varepsilon}_{n,l}(P|x^n, ext(m))$. Then,

$$\begin{aligned} H(Y^n|\mathbb{1}_{\{EXT\}}) &= Pr(EXT)H(Y^n|\mathbb{1}_{\{EXT\}} = 1) + Pr(EXT^c)H(Y^n|\mathbb{1}_{\{EXT\}} = 0) \\ &= Pr(EXT)H(Y^n|EXT) + Pr(EXT^c)H(Y^n|EXT^c). \end{aligned}$$

Furthermore, by the chain rule of conditional entropy, we have

$$H(Y^n|\mathbb{1}_{\{EXT\}}) = H(Y^n, \mathbb{1}_{\{EXT\}}) - H(\mathbb{1}_{\{EXT\}}),$$

therefore,

$$Pr(EXT)H(Y^n|EXT) = H(Y^n, \mathbb{1}_{\{EXT\}}) - H(\mathbb{1}_{\{EXT\}}) - Pr(EXT^c)H(Y^n|EXT^c).$$

Now, since $H(\mathbb{1}_{\{EXT\}}) \leq \log_2 2 = 1$, $H(Y^n|EXT^c) \leq n \log |J|$ and
$H(Y^n, \mathbb{1}_{\{EXT\}}) \geq H(Y^n)$, we get

$$\begin{aligned} \tfrac{1}{n}H(Y^n|EXT) &\geq \frac{H(Y^n)}{nPr(EXT)} - \frac{(1 + Pr(EXT^c)n \log |J|)}{nPr(EXT)} \\ &> \frac{(H(Y|X) - 2\varepsilon d)}{Pr(EXT)} - \frac{(1 + Pr(EXT^c)n \log |J|)}{nPr(EXT)}. \end{aligned}$$

By Observation 7, $Pr(EXT) \to 1$ as $n \to \infty$, hence we can conclude that

$$\tfrac{1}{n}H(Y^n|EXT) \to H(Y|X) - 2\varepsilon d \quad \text{as } n \to \infty.$$

□

Corollary 10 (The Size of $T^{\varepsilon}_{n,l}(P|x^n, ext(m))$)**.** *Let $\varepsilon > 0$. Then, for every $\delta > 0$ and n sufficiently large, we have*

$$|T^{\varepsilon}_{n,l}(P|x^n, ext(m))| > 2^{n(H(Y|X) - 2\varepsilon d - \delta)}$$

Proof. The statement follows immediately from Theorem 8 and from the fact that

$$H(Y^n|EXT) \leq \log_2 |T^{\varepsilon}_{n,l}(P|x^n, ext(m))|.$$

□

Bibliography

[1] Robert J. Aumann and Michael Maschler. “Repeated Games of Incomplete Information: The Zero-Sum Extensive Case”. Mathematica, Inc Report ST-143, 37-116, Princeton, 1968.

[2] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory.* Wiley, New Jersey, 2006.

[3] Vincent P. Crawford and Joel Sobel. “Strategic Information Transmission.” Econometrica, Vol. 50, No.6, 1431-1451, 1982.

[4] Imre Csiszár. “The Method of Types.” IEEE Transactions on Information Theory, Vol. 44, No. 6, 1998.

[5] Françoise Forges. “Non-Zero-Sum Repeated Games and Information Transmission.” Essays in Game Theory in Honor of Michael Maschler, Springer Verlag, No. 6, 65-95, 1994.

[6] Olivier Gossner, Penélope Hernández, Abraham Neyman. “Optimal Use of Communications Resources.” Econometrica, Vol.74, No.6, 1603-1636, 2006.

[7] Wassily Hoeffding. “Probability Inequalities for Sums of Bounded Random Variables.” Journal of the American Statistical Association, Vol. 58, No. 301, 13-30, 1963

[8] David A. Levin, Yuval Peres, Elizabeth L. Wilmer. *Markov Chains and Mixing Times.* American Mathematical Society, 2009.

[9] David J.C. MacKay. *Information Theory, Inference, and Learning Algorithms.* Cambridge University Press, Cambridge, 2005.

[10] Jean-François Mertens, Sylvain Sorin, Shmuel Zamir. *Repeated Games.* Cambridge University Press, Cambridge, 2015

[11] James R. Norris. *Markov Chains.* Cambridge University Press, Cambridge, 2012

[12] Ron Peretz. “The Strategic Value of Recall.” Games and Economic Behavior, No. 74, 332-351, 2012.

[13] Roy Radner. “The Organization of Decentralized Information Processing.” Econometrica, Vol. 61, No. 5, 1109-1146, 1993.

[14] Ram C. Rao and Frank M. Bass. “Competition, Strategy and Price Dynamics: A Theoretical and Empirical Investigation.” Journal of Marketing Research, Vol. 22, No. 3, 283-296, 1985.

[15] Claude E. Shannon. “A Mathematical Theory of Communication.” The Bell System Technical Journal, Vol. 27, No. 3, 379-423, 623-656, 1948.

[16] Mark M. Wilde. *Quantum Information Theory.* Cambridge University Press, Cambridge, 2017.