Communication with Unobservable Constraints

Ilya Khayutin *†

December 14, 2007

Abstract

We consider a repeated game with asymmetric information. The first player, called the forecaster, is fully informed in advance of the states of nature during the whole game. The second player, called the agent, can only deduce the future states of nature from the previous actions of the forecaster.

Gossner, Hernandez, and Neyman [6] studied similar games and solved the games with a signaling model in which at each stage of the repeated game the two players are informed of all the previous actions of nature, of the forecaster, and of the agent. They described the set of possible payoffs in the asymptotically long game using an information constraint characterization of the average distributions over the triple of the forecaster's action, the agent's action, and the state of nature (implementable distributions).

In this work we study a different signaling model in which at each stage the agent is informed of the previous actions of the forecaster but not of the past states of nature.

We show that the same information constraint completely characterizes the set of implementable distributions in this model. We conclude that the same average payoffs are achievable whether the states of nature are observable or not. To show this, we use the probabilistic method in hypergraph theory.

*Institute of Mathematics and Racah Institute of Physics, The Hebrew University of Jerusalem, Givat Ram, Jerusalem 91904, Israel (e-mail: ilya.khayutin@mail.huji.ac.il).

†This research was supported in part by Israeli Science Foundation grants 263/03 and 1123/06.

Acknowledgement

The author would like to thank Professor Abraham Neyman of the Hebrew University, under whose supervision this research was carried out.

1 Introduction

The field of repeated games with asymmetric incomplete information, although the subject of active current research, is still to a large extent unmapped territory. Communication is the key aspect of many open questions, since it is the device by which the gap of asymmetric information between several players and the resulting inefficiencies may be bridged.

Communication can be roughly divided into two qualitative categories: internal and external communication[1]. "external communication" means communication through some outside medium independent of the course of the game and its actions. Such communication may be limited either in capacity (the amount of information one player can reliably transmit to another) or by a price tag: the players must pay a price for transmitting a certain amount of information.

Internal communication models, on the other hand, do not assume some independent means of communication beyond the mechanism of the game itself. By internal communication we mean signaling from one player to another using the actions of the players in the game itself. For this to have any non-trivial meaning the game must be multi-stage and at least some players should be able somehow to observe the actions of the other players.

Internal communication is also limited in capacity: there is a limit to how much information one player can transmit to another using his actions alone. And, internal communication has a price, since transmitting information in high capacity usually results in inefficiencies in the game. On the other hand, the transmitted information can improve long-term efficiency, and thus many times the result is a delicate optimum between short-term inefficiencies due to internal information transmission and increased long-term efficiency.

In examining the general problem of games with internal communication, we limit ourselves to repeated games with two players and nature. Even though this might seem a rather simple class of games, it is nonetheless

[1] Called "online" and "offline" communication in [6].

extremely rich and gives rise to a complicated structure once information asymmetries between the two players are allowed.

Although different variations are possible on the theme of information asymmetries in repeated games we would like to focus on the case in which the first player, whom we shall call the forecaster, has in advance a complete knowledge of all the states of nature during the infinite repeated game. In each stage the state of nature defines the payoff functions for both players. The law by which a specific state of nature defines the payoff functions is known to both players, as is the distribution by which the states of nature are drawn. We shall discuss only the case where the states of nature are drawn independently according to the same law[2].

Before the forecaster learns about the infinite sequence of the states of nature, he and the second player, called the agent, get to coordinate a game strategy. This obviously allows communication through the game if at each stage the agent receives some sort of signal that depends on the actions of the players and the state of nature in the previous stage (perhaps by some probabilistic law). It is the variation on this signal that is the focus of this work.

Gossner, Hernandez and Neyman [6] gave a complete solution of games in which at each stage the agent and the forecaster receive complete information about each others action and the state of nature in the previous stages. To find all possible achievable payoffs in the asymptotically long game, they introduced the notion of implementable distributions, i.e., distributions over the triple of states of nature, the forecaster's actions and the agent's actions that are implementable as the average over all stages of the distribution of this triple in each stage in the infinite game. Achievable average payoffs can be calculated from the set of implementable distributions. This set is characterized using an information constraint and the proof uses information-theoretic tools.

The strategies constructed in [6] make direct use of the fact that the agent can at any stage observe both the previous forecaster's actions and the previous states of nature. This led the authors of [6] to ask what distributions can be implemented and what payoffs can be achieved if at any stage the agent can observe only the past actions of the forecaster but not the past states of nature. It is this question that we attempt to answer in this paper.

[2]Gossner, Hernandez and Neyman [6] show that similar techniques may be applicable to more general models.

One might at first assume that less can be achieved in the infinitely long game if the agent does not observe past states of nature, because at each stage he receives less information. But we shall show in the following sections that exactly the same average payoffs are achievable in the infinitely long game whether the agent observes past states of nature or not. To construct strategies that implement the required distributions in games in which the past states of nature are unobservable by the agent, we restate the problem using the notions of hypergraph theory and hypergraph colorings. We show that the required hypergraph colorings actually do exist using the probabilistic method.

2 The Model

Let I be the set of possible states of nature, J be the set of the forecaster's actions, and K be the set of the agent's actions. Assume I, J, K to be finite. Fix payoff functions for the forecaster and the agent: $g^f, g^a : I \times J \times K \to \mathbb{R}$. Fix also a distribution $\mu \in \Delta(I)$. The sequence of the states of nature is drawn as i.i.d. random variables with distribution μ. The game is defined by (I, J, K, μ, g^f, g^a).

The subject of this work is the USN model (unobservable states of nature). In the framework of the USN model a finite observable history of the game is an element of $(J \times K)^*$. In the USN model a pure strategy for the forecaster is a function $\sigma : I^{\mathbb{N}} \times (J \times K)^* \to J$ and a pure strategy for the agent is a function $\tau : (J \times K)^* \to K$. A pure strategy for the game is a pair (σ, τ) of such functions. Unless otherwise specified, we are discussing the USN model.

We shall also discuss the results of Gossner, Hernandez and Neyman [6] on the model in which a pure strategy for the agent is a function $\tau :$ $(I \times J \times K)^* \to K$ (the agent observes past states of nature). We shall refer to this model as the GHN model. Every sequence of actions implementable by some strategy (σ, τ) in the standard model is also implementable in the GHN model. The equivalent strategy to (σ, τ) in the GHN model is given by $\sigma_1 = \sigma$, $\tau_1(\mathbf{i}, \mathbf{j}, \mathbf{k}) = \tau(\mathbf{j}, \mathbf{k})$; in such a strategy the agent ignores the previous states of nature.

At the beginning of a game a random i.i.d. sequence $\mathbf{i} = (i_t)_{t=1}^{t=\infty}$ is drawn. At stage t the forecaster and the agent choose their actions according to

$$j_t = \sigma(\mathbf{i}, ((j_1, k_1), (j_2, k_2), \ldots, (j_{t-1}, k_{t-1})))$$

$$k_t = \tau((j_1, k_1), (j_2, k_2), \ldots, (j_{t-1}, k_{t-1}))$$

The payoff at stage t is given by $g^f(i_t, j_t, k_t)$ for the forecaster and $g^a(i_t, j_t, k_t)$ for the agent. Note that in the USN model the agent is not informed of the payoff at the end of the stage.

A strategy (σ, τ) together with the distribution μ defines a probability measure for each stage t on the triples of the state of nature, the forecaster's action, and the agent's action $I \times J \times K$. It is the marginal on stage t of the probability measure over all game histories. We denote the probability according to this distribution by $\mathrm{Pr}^t_{\mu,\sigma,\tau}$. The average payoff at stage t is $\mathrm{E}^t_{\mu,\sigma,\tau}\, g^f(i, j, k)$ for the forecaster and $\mathrm{E}^t_{\mu,\sigma,\tau}\, g^a(i, j, k)$ for the agent. The average payoff in a t-stage game is $g^{f,t}_{\sigma,\tau,\mu} = \frac{1}{t}\sum_{s=1}^{t} \mathrm{E}^s_{\mu,\sigma,\tau}\, g^f(i, j, k)$ for the forecaster and $g^{a,t}_{\sigma,\tau,\mu} = \frac{1}{t}\sum_{s=1}^{t} \mathrm{E}^s_{\mu,\sigma,\tau}\, g^f(i, j, k)$ for the agent. We shall focus on the payoffs in the asymptotically long game $g^f_{\sigma,\tau,\mu} = \lim_{t\to\infty} g^{f,t}_{\sigma,\tau,\mu}$,$g^a_{\sigma,\tau,\mu} = \lim_{t\to\infty} g^{a,t}_{\sigma,\tau,\mu}$.

Because every strategy in the USN model gives rise to an equivalent strategy in the GHN model, every payoff achievable either in the t-staged or in the asymptotically long game in the USN model is also achievable in the GHN model. Yet because our USN model is weaker than the GHN model, a reasonable assumption would be that some payoffs achievable in the GHN model are not achievable in the USN model.

Let a distribution $Q \in \Delta(I \times J \times K)$ be t-implementable if there exists a pure strategy (σ, τ) so that $Q = \frac{1}{t}\sum_{s=1}^{t} \mathrm{Pr}^s_{\mu,\sigma,\tau}(i_t, j_t, k_t)$. Also let a distribution $Q \in \Delta(I \times J \times K)$ be implementable if there exists a pure strategy (σ, τ) such that $Q = \lim_{t\to\infty} \frac{1}{t}\sum_{s=1}^{t} \mathrm{Pr}^s_{\mu,\sigma,\tau}(i_t, j_t, k_t)$. We shall denote the set of t-implementable distributions by $\mathcal{Q}(t)$ and the set of implementable distributions by $\mathcal{Q}$. Note that if the strategy (σ, τ) t-implements (implements) the distribution $Q \in \Delta(I \times J \times K)$, then $g^{f,t}_{\sigma,\tau,\mu} = \mathrm{E}_Q(g^f(i, j, k))$, $g^{a,t}_{\sigma,\tau,\mu} = \mathrm{E}_Q(g^a(i, j, k))$, and so the t-staged (asymptotic) payoff is a linear function of the t-implementable (implementable) distribution.

The sets of t-implementable and implementable distributions in the USN model have the same basic properties as those sets in the GHN model.

Remark 1

1. *$\mathcal{Q}(t)$ is closed;*

2. *$\mathcal{Q}(t) \subseteq \mathcal{Q}$;*

3. *$\frac{s}{s+t}\mathcal{Q}(s) + \frac{t}{s+t}\mathcal{Q}(t) \subseteq \mathcal{Q}(s+t)$;*

4. the Hausdorff distance between $\mathcal{Q}(s)$ and $\mathcal{Q}(s+t)$ is bounded by $\frac{2t}{s+t}$.

From this basic properties we can conclude:

Remark 2

1. *The limit $\lim_{t\to\infty} \mathcal{Q}(t)$ exists and equals the closed convex hull of $\bigcup_{t\geq 1} \mathcal{Q}(t)$;*
2. *$\mathcal{Q}(t) \to_{t\to\infty} \mathcal{Q}$;*
3. *$\mathcal{Q}$ is closed and convex.*

2.1 Example: Coordination with Nature

We take this example from [6] and [5]. In this example there is a common payoff function for both the agent and the forecaster. We have $I = J = K = \{0, 1\}$ and the distribution μ, according to which the states of nature are drawn, is uniform. The common payoff function is given by

$$g(i, j, k) = \begin{cases} 1 & \text{if } i = j = k \\ 0 & \text{otherwise} \end{cases}$$

This can be represented by the following payoff matrices

$i = 0$:

	0	1
0	1	0
1	0	0

$i = 1$:

	0	1
0	0	0
1	0	1

where nature chooses the matrix, the forecaster chooses the row and the agent chooses the column. The goal of this game is to achieve maximum coordination between the agent, the forecaster and nature.

The first trivial strategy is when the forecaster always plays the same bit as nature while the agent plays randomly with uniform distribution (or plays a series of interchanging 0s and 1s). It is clear that the long-term payoff for each player will be $\frac{1}{2}$. Note that in this strategy no information is transferred to the agent about future states of nature, and in particular there is no need for the agent to observe past states of nature.

A second strategy is the one in which at even stages the forecaster plays the same bit as nature but at odd stages the forecaster plays the bit that nature is going to play at the next stage. The agent, on the other hand, at even stages plays the bit that the forecaster played at the previous stage and at odd stages plays a random bit. A quick calculation shows that the expected long-term payoff for each player in this strategy is $\frac{5}{8}$. In this strategy information is transmitted from the forecaster to the agent, although the agent needs only to observe the previous forecaster's action and may be oblivious of any other past actions. In particular, he does not need to observe any past states of nature.

We shall now consider a strategy played in blocks of three stages. The strategy is such that at each block, except perhaps the first, at least two-thirds of the bits are correlated (nature, forecaster and agent). At each three-stage block the forecaster signals what is the majority bit of the states of nature in the next three-stage block. The agent plays at each block the same bit three times, according to the signal he received from the forecaster in the previous block.

To transmit the majority bit of the next block, the forecaster at each block plays two bits in coordination with both nature and the agent and one signaling bit which is equal to the majority bit of the states of nature in the next block. If at the current block all three states of nature (and all the three actions of the agent) are equal, the signaling bit for the forecaster will be the last bit in the block. Otherwise it will be the bit where the nature bit is different from the agent's bit.

The agent is always capable of interpreting the signal without observing past states of nature. When he looks at the previous block, if all three bits of the forecaster's actions are equal then this bit is the majority bit of the states of nature in the next block. Otherwise the minority bit of the forecaster's three actions in the previous block is the signaling bit.

A careful calculation shows that the expected distribution over the triple $I \times J \times K$ implemented by this strategy in the long-run game is given by

	0	1
0	17/48	1/16
1	1/48	1/16

$i = 0$

	0	1
0	1/16	1/48
1	1/16	17/48

$i = 1$

and the expected average payoff for each player is $\frac{17}{24}$, which is larger than $\frac{5}{8}$ (the average payoff in the previous two-stage strategy). Note that in this strategy, too, the agent need not observe the past states of nature, yet he is required to observe the forecaster's last three actions.

Nonetheless, if the agent can observe the previous states of nature, sometimes the forecaster can transmit additional information to the agent using a similar strategy. In such blocks that all the agent's three bits are the same as the three nature bits (perfect blocks), the forecaster can actually use any bit in the block to send the signal. If the majority bit of the states of nature in the next block is different from the majority bit in the current block, then the choice which bit to use as the signal can be used to transmit additional information to the agent. This is possible only in the GHN model, because in the USN model the agent does not know what previous blocks where perfect. Even though only little information can be sent using this technique, in the long run it can be used to increase the average payoff. This example demonstrates how observing past states of nature might be useful.

3 The Information Constraint

3.1 Entropy and Information

Information-theoretic tools are very effective in the study of repeated games with asymmetric information. The notions of Shannon entropy and mutual information [7] of random variables on a finite alphabet are necessary for our characterization of $\mathcal{Q}$. We include here the definitions of those basic information-theoretic notions used in our analysis.

Definition 1 *The entropy of a discrete random variable* $\mathbf{x}$ *with law* p *and values in a finite set* X *is given by*

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

Entropy is a measure of the randomness of a random variable (a random variable with an atomic law has the minimal possible entropy $H(\mathbf{x}) = 0$, while a random variable with a uniform law has the maximal possible entropy $H(\mathbf{x}) = \log(|X|)$). Entropy can also be viewed as the optimal rate at which an information source with law p can transmit information.

Definition 2 *The conditional entropy of a discrete random variable* $\mathbf{x}$ *given a random variable* $\mathbf{y}$ *with a joint law* p *of* $(\mathbf{x}, \mathbf{y})$ *and values in a finite set*

$X \times Y$, is

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

The conditional entropy is a measure of the randomness of a random variable $\mathbf{x}$ as seen by an observer knowing the value of a random variable $\mathbf{y}$. As one would expect, if the two random variables are independent, the knowledge of $\mathbf{y}$ has no effect on the entropy of $\mathbf{x}$: $H(\mathbf{x} \mid \mathbf{y}) = H(\mathbf{x})$. At the opposite extreme, if the value of $\mathbf{x}$ is a deterministic function of $\mathbf{y}$, an observer knowing $\mathbf{y}$ will see no randomness in $\mathbf{x}$: $H(\mathbf{x} \mid \mathbf{y}) = 0$. As with standard entropy the conditional entropy can be interpreted as the optimal rate at which an information source with a given law can transmit information to a person knowing the value of some random variable.

Definition 3 *The mutual information of two discrete random variable* $\mathbf{x}$,$\mathbf{y}$ *with values in a finite set $X \times Y$ is*

$$I(\mathbf{x};\mathbf{y}) = H(\mathbf{x}) - H(\mathbf{x} \mid \mathbf{y}) = H(\mathbf{y}) - H(\mathbf{y} \mid \mathbf{x})$$
$$= H(\mathbf{x}) + H(\mathbf{y}) - H(\mathbf{x},\mathbf{y})$$

The mutual information is the measure of the information stored in one random variable about another one. Note that the mutual information is symmetric in both random variables. As with entropy one can define the mutual information of two random variables conditioned on a third one.

3.2 Characterizing Implementable Distributions

About the GHN model we know from [6] that a distribution $Q \in \Delta(I\times J\times K)$ is in $\mathcal{Q}_{GHN}$ (the set of implementable distributions in the GHN model) if and only if it has marginal μ on I and it fulfills the information constraint:

$$H_Q(j \mid i,k) \geq I_Q(i;k) \tag{1}$$

This formulation of the information constraint gives rise to a simple interpretation. The right-hand part of the inequality is a measure of the information the agent has on the state of nature, while the left-hand part is a measure of the randomness of the variable describing the forecaster's action conditioned on the state of nature and the agent's action, such that it can be seen as the capacity of the forecaster to transmit information while adhering to the constraints of choosing actions that implement the distribution Q given the state of nature and the agent's action. In this interpretation the information constraint simply states in the correct terms that all the information the agent has about the state of nature comes from the information transmitted to him by the forecaster.

We have already noticed that because every strategy in the USN model gives rise to an equivalent strategy in the GHN model, every payoff achievable in the USN model is also achievable in the GHN model. The same remark applies to implementable distributions. Every distribution t-implementable (implementable) in the USN model is also implementable (t-implementable) in the GHN model. We have $\mathcal{Q}(t) \subseteq \mathcal{Q}_{GHN}(t)$, $\mathcal{Q} \subseteq \mathcal{Q}_{GHN}$. In particular, we have the following proposition:

Proposition 1 *If a distribution $Q \in \Delta(I \times J \times K)$ is implementable in the USN model, then it fulfills the information constraint* (1) *and has marginal μ on I:*

$$\begin{aligned} &Q_I = \mu \\ &H_Q(j \mid i, k) \geq I_Q(i; k) \end{aligned}$$

□

A priori there is no clear reason to believe that the converse is true. Since at each stage in the USN model the agent is less informed about the history of the game than in the GHN model, it may be so that the forecaster has less capacity to transmit to the agent the necessary information to implement some distributions that are implementable in the GHN model. Specially, it may be impossible to achieve asymptotic optimal payoffs in those games (like the internal matching pennies game studied by Gossner, Hernandez and Neyman in [5]) in which achieving those payoffs requires full use of an effective communication channel between the forecaster and the agent (in the asymptotically long game no more information can be transmitted between the forecaster and the agent).

Nevertheless, we show that the converse is true.

4 The Main Result

As stated, in the asymptotically long game the observation of the state of nature by the agent (as in the GHN model) contributes nothing to the amount of information the agent can receive from the forecaster. This is stated explicitly in the following theorem which is the main result of this work:

Theorem 1 *A distribution $Q \in \Delta(I \times J \times K)$ is implementable in the USN model if and only if it fulfills the information constraint* (1) *and has marginal μ on I:*

$$\begin{aligned} &Q_I = \mu \\ &H_Q(j \mid i, k) \geq I_Q(i; k) \end{aligned}$$

5 The Combinatorial Properties of Information

Our construction of implementing strategies is based upon playing in blocks of large size. We need to approximate the size of the sets of forecaster's and agent's possible actions, that will induce the required distribution in such large blocks. To do so, we turn to the mathematical device of typical sequences.

5.1 Combinatorial Entropy and Typical Sequences

We start with a notion of combinatorial entropy that is parallel to the notion of the Shannon entropy of a random variable.

Definition 4 *Let A be a finite set. The combinatorial entropy of A, denoted by $H(A)$, is the base* 2 *logarithm of $|A|$: $H(A) = \log |A|$.*

The connection between the two notions of entropy becomes clear when we introduce the notions of empirical distribution and typical sequences.

Definition 5

1. *For a finite sequence $\mathbf{l} = (l_1, l_2, \ldots, l_n)$ over a finite alphabet L let $\rho(\mathbf{l}) \in \Delta(L)$ denote the empirical distribution induced over L by the sequence $\mathbf{l}$: $\rho(\mathbf{l})[l] = \frac{1}{n} \sum_{t=1}^{n} \mathbb{I}_{l_t = l}$.*

2. *A sequence* $\mathbf{l} = (l_1, l_2, \ldots, l_n)$ *is said to be a typical sequence of the distribution* $\nu \in \Delta(L)$ *if* $\rho(\mathbf{l}) = \nu$.

3. *Define* $T_n(\nu) = \{\mathbf{l} \in L^n \mid \rho(\mathbf{l}) = \nu\}$ *to be the* n*-type set of* ν.

4. *We denote by* $\mathbb{T}_n(L)$ *the set of all distributions over* L *having a typical sequence of length* n: $\{\nu \in \Delta(L) \mid T_n(\nu) \neq \emptyset\}$.

The entropy of a random variable with law ν is a good approximation to the combinatorial entropy of $T_n(\nu)$ if this set is not empty. We cite the following proposition (see [3, Theorem 12.1.3, p. 282]):

Proposition 2 *For* $\nu \in \Delta(L)$ *(*L *is a finite set), if* $T_n(\nu) \neq \emptyset$*, then*

$$\frac{2^{nH(\nu)}}{(n+1)^{|L|}} \leq |T_n(\nu)| \leq 2^{nH(\nu)}$$

Equivalently, the combinatorial entropy of the n*-type set fulfills*

$$nH(\nu) - |L| \log(n+1) \leq H(T_n(\nu)) \leq nH(\nu)$$

The notions of a typical sequence and the n-type set of some distribution can be extended to the notions of a conditional typical sequence and the conditional n-type set.

Definition 6

1. *Let* $\nu \in \Delta(A \times B)$ *have marginal distributions* $\nu_A \in \Delta(A)$ *and* $\nu_B \in \Delta(B)$. *Let* $\mathbf{b} = (b_1, b_2, \ldots, b_n) \in T_n(\nu_B)$.
   A sequence $\mathbf{x} = ((a_1, c_1), (a_2, c_2), \ldots, (a_n, c_n) \in (A \times B)^n$ *is a typical sequence of* ν *conditional on* $\mathbf{b}$ *(denoted* $\nu|\mathbf{b}$*) if* $\mathbf{x} \in T_n(\nu)$ *and* $c_t = b_t$ *for all* $1 \leq t \leq n$. *Equivalently, a sequence* $\mathbf{a} = (a_1, a_2, \ldots, a_n) \in A^n$ *is* n*-typical of* $\nu|\mathbf{b}$ *if* $(\mathbf{a}, \mathbf{b}) = ((a_1, b_1), (a_2, b_2), \ldots, (a_n, b_n)) \in T_n(\nu)$.

2. *Denote by* $T_n(\nu|\mathbf{b})$ *the set of all* n*-typical sequences of* $\nu|\mathbf{b}$ *in* $(A \times B)^n$. *This set can be identified with the subset of* A^n *of the equivalent* n*-typical sequences of* $\nu|\mathbf{b}$.

Note that the size of $T_n(\nu|\mathbf{b})$ is independent of the specific choice of $\mathbf{b} \in T_n(\nu_B)$ and $|T_n(\nu|\mathbf{b})| \cdot |T_n(\nu_B)| = |T_n(\nu)|$. This implies a law of additivity of combinatorial entropies parallel to the law of additivity of entropies: $H(T_n(\nu|\mathbf{b})) + H(T_n(\nu_B)) = H(T_n(\nu))$.

As with standard combinatorial entropy, the conditional Shannon entropy of a random variable with law ν is a good approximation to the combinatorial entropy of the set of conditional n-typical sequences. We cite the following proposition (see [6, Lemma 2, p. 1633]):

Proposition 3 *For $\nu \in \Delta(A \times B)$ ($A \times B$ is a finite set), if, $T_n(\nu) \neq \emptyset$, then for any $\mathbf{b} \in T_n(\nu_B)$*

$$\frac{2^{n(H(\nu)-H(\nu_B))}}{(n/|B|+1)^{|A\times B|}} \leq |T_n(\nu|\mathbf{b})| \leq 2^{n(H(\nu)-H(\nu_B))}$$

Equivalently, the combinatorial entropy of the n-type set fulfills

$$n(H(\nu)-H(\nu_B))-|A\times B|\log(n/|B|+1) \leq H(T_n(\nu|\mathbf{b})) \leq n(H(\nu)-H(\nu_B))$$

Also, $H(\nu)-H(\nu_B) = H_\nu(a|b)$ for a,b random variables of the projections of $A \times B$ over A,B respectively with the probability measure ν over $A \times B$.

We now turn briefly to discuss the asymptotic behavior of the n-type sets for large n. If a distribution ν is in $\mathbb{T}_k(A)$, then it is in $\mathbb{T}_{kn}(A)$ for every integer n. Proposition 2 implies

$$\lim_{n\to\infty} \frac{H(T_{kn}(\nu))}{n} = kH(\nu) \tag{2}$$

A similar statement can be made with regard to the conditional combinatorial entropy. Let ν be a distribution in $\mathbb{T}_k(A \times B)$; then we have $\nu \in \mathbb{T}_{kn}(A \times B)$ for all integers n. Let $\mathbf{b}_{kn} \in T_{kn}(\nu_B)$ for $n = 1$ to ∞ be a sequence of typical sequences. Remember that the size of $T_{kn}(\nu|\mathbf{b}_{kn})$ does not depend on the specific choice of $\mathbf{b}_{kn} \in T_{kn}(\nu_B)$. Proposition 3 implies

$$\lim_{n\to\infty} \frac{H(T_{kn}(\nu|\mathbf{b}_{kn}))}{n} = k(H(\nu)-H(\nu_B)) = kH_\nu(a|b) \tag{3}$$

5.2 Action Sets

One of the vital parts of constructing an implementing strategy in our method is the set of action plans for the agent. For this we make the following broad definition:

Definition 7

1. *Let $Q \in \mathbb{T}_n(A \times B)$. A set $S \subseteq T_n(Q_B)$ is an n-action set if*

$$(\forall \mathbf{a} \in T_n(Q_A)) \quad T_n(Q|\mathbf{a}) \cap S \neq \emptyset$$

2. *Let $Q \in \mathbb{T}_n(A \times B)$ and let f be an integer. A set $S \subseteq T_n(Q_B)$ is an f-fold n-action set if*

$$(\forall \mathbf{a} \in T_n(Q_A)) \quad |T_n(Q|\mathbf{a}) \cap S| \geq f$$

Note that the largest f for which an n-action set exists is $|T_n(Q|\mathbf{a})| = |T_n(Q)|/|T_n(Q_A)|$ and there is only one n-action set with this maximal f: $T_n(Q_B)$. This is also the maximal possible action set.

We are interested in the case of $Q_{I\times K}$ for some $Q \in \Delta(I \times J \times K)$. In this case an n-action set is essentially a set of possible action for the agent, such that for every possible sequence of n states of nature there is a sequence of actions for the agent in the action set that together with the sequence of the states of nature induces the required empirical distribution $Q_{I\times K}$.

The construction in [6] of implementing strategies in the GHN model uses the maximal action set $T_n(Q_K)$ (called a set of action plans there). The possibility of using smaller action sets is discussed in [6, section 7.1] and such sets are used in the construction of the strategies in [5].

We can also use the maximal action set $T_n(Q_K)$ to construct implementing strategies in our USN model (more on this in Section 7.3.2), but we prefer to use the smallest action set possible. This has several advantages over using the maximal action set.

We are interested in optimal action sets, i.e., those of minimal size. Let us define the following:

Definition 8 *Let $Q \in \mathbb{T}_n(A \times B)$. We denote by $\mathcal{A}_{n,f}(Q; A, B)$ the size of a minimal f-fold n-action set $S \subseteq T_n(Q_B)$. In most cases when it is clear what the relevant sets A,B are we abuse the notation and write $\mathcal{A}_{n,f}(Q)$. For the case $f = 1$ we use the notation $\mathcal{A}_n(Q)$.*

Of course, the minimal of all f-fold n-actions set for all f is the minimal 1-fold n-action set. That is why the value of $\mathcal{A}_n(Q)$ is the most interesting to us.

A lower bound for the size of an action set is easy to state and prove.

Lemma 1 *Let $Q \in \mathbb{T}_n(A \times B)$ and let f be an integer. If $S \subseteq T_n(Q_B)$ is an f-fold n-action set,*

$$|S| \geq \frac{|T_n(Q_A)||T_n(Q_B)|}{|T_n(Q)|} f \geq \frac{2^{nI_Q(a;b)}}{(n+1)^{|A|+|B|}} f$$

Specifically, for an optimal action set: $\mathcal{A}_{n,f}(Q) \geq f 2^{nI_Q(a;b)}/(n+1)^{|A|+|B|}$.

Proof. Note that $T_n(Q_A) = \bigcup_{\mathbf{b} \in S} T_n(Q|\mathbf{b})$, so $|T_n(Q_A)| \leq |T_n(Q|\mathbf{b})||S|$ (remember that the size of $T_n(Q|\mathbf{b})$ does not depend on the specific choice of $\mathbf{b} \in T_n(Q_B)$).

If we look more carefully at the disjoint union $\bigsqcup_{\mathbf{b} \in S} T_n(Q|\mathbf{b})$ we see that each element of $T_n(Q_A)$ appears in at least f distinct sets of the disjoint union. We can conclude that $|T_n(Q_A)| f \leq |T_n(Q|\mathbf{b})||S|$. As $|T_n(Q|\mathbf{b})||T_n(Q_B)| = |T_n(Q)|$ we have $|S| \geq f|T_n(Q_A)||T_n(Q_B)|/|T_n(Q)|$.

The second inequality of the lemma is a straightforward application of proposition 2 to $|T_n(Q_A)||T_n(Q_B)|/|T_n(Q)|$. □

This result implies a relation between the size of minimal action sets and mutual information. A restatement of the above result in term of combinatorial entropy is $H(S) \geq nI_Q(a;b) + \log f - (|A|+|B|)\log(n+1)$ (if S fulfills the conditions in the lemma).

What about an upper bound? We cite without proof the following result from [6] regarding the existence of action sets.

Proposition 4 *Let $Q \in \mathbb{T}_n(A \times B)$. For all $1 \leq f \leq |T_n(Q)|/|T_n(Q_A)|$ and $g > 1$ such that $\frac{f(g-1)^2}{2g} > n \ln 2 H(Q_A)$, there exists an f-fold n-action set $S \subseteq T_n(Q_B)$ with*

$$\frac{|T_n(Q_A)| \cdot |T_n(Q_B)|}{|T_n(Q)|} fg \leq |S| < 1 + \frac{|T_n(Q_A)| \cdot |T_n(Q_B)|}{|T_n(Q)|} fg$$

The following corollary gives an upper bound for the size of an optimal action set.

Corollary 1 *Let $Q \in \mathbb{T}_k(A \times B)$. There exists an f-fold n-action set $S \subseteq T_n(Q_B)$ such that*

$$\begin{aligned} |S| &< 1 + 2(n \ln 2H(Q_A) + f)\frac{|T_n(Q_A)||T_n(Q_B)|}{|T_n(Q)|} \\ &\leq 4(1 + \ln 2H(Q_A)n)(1+n)^{|A \times B|} 2^{nI_Q(a;b)} f \end{aligned}$$

Proof. Let $g = 2(n \ln 2H(Q_A)/f + 1)$. Then

$$f\frac{(g-1)^2}{2g} = f(g/2 - 1 + 1/(2g)) > f(g/2-1) = n \ln 2H(Q_A)$$

By the above-stated proposition there exists an f-fold n-action set satisfying

$$|S| < 1 + \frac{|T_n(Q_A)||T_n(Q_B)|}{|T_n(Q)|} fg = 1 + 2(n \ln 2H(Q_A) + f)\frac{|T_n(Q_A)||T_n(Q_B)|}{|T_n(Q)|}$$

□

Now we can state our bounds in terms of the combinatorial entropy of an optimal action set as

$$\begin{aligned}-(|A| + |B|)\log(n+1) &\leq \log(\mathcal{A}_{n,f}(Q)/f) - nI_Q(a;b) \\ &< |A \times B| \log(n+1) + \log(n \ln 2H(Q_A) + 1) + 2\end{aligned}$$

For a distribution $Q \in \mathbb{T}_k(A \times B)$ we have $Q \in \mathbb{T}_{kn}(A \times B)$ for all integers n. The inequality above implies the following limit for such a distribution:

$$\lim_{n\to\infty} \frac{\log(\mathcal{A}_{kn,f}(Q)/f)}{n} = kI_Q(a;b) \tag{4}$$

In the case of $Q \in \Delta(I \times J \times K)$ we see that the size of an optimal action set for the agent over a block of n stages for large n is approximately $2^{nI_Q(i;k)}$. Let us look again at our interpretation of the information constraint $H_Q(j \mid i,k) \geq I_Q(i;k)$. We interpret the right-hand part of this inequality the amount of information the agent has about the state of nature. We see now that this gives an approximate size of an optimal set of action plans for the agent. That is, in a block of n stages the agent has to receive enough information from the forecaster to pinpoint an element from a set of approximate size $2^{nI_Q(i;k)}$ for large n. The left-hand part of the information constraint inequality in our interpretation is the capacity of the forecaster to transmit information to the agent. And, indeed, if in a block of n stages the forecaster has to choose a sequence of actions that sends a message to the agent, yet he's restricted by the dictated states of nature and the agent's actions in this block, then the set of possible messages for him to send is $T_n(Q|\mathbf{i},\mathbf{k})$. Its size by equation (3) is approximately $2^{nH_Q(j|i,k)}$ for large n. So a necessary condition for the forecaster to be able to pinpoint the correct action plan for the agent is that the set of possible messages be greater than or equal to the set of required action plans. For large n this gives approximately $H_Q(j|i,k) \geq I_Q(i;k)$, which is exactly the information constraint.

In the case where the agent knows the past states of nature, as in the GHN model, it is quite straightforward using the device of typical sets to show that this is sufficient and this has been done in [6]. Yet in the USN model the agent does not know from which set $T_n(Q|\mathbf{i},\mathbf{k})$ the forecaster has had to choose his message since the agent does not know the past states of nature. To prove that the information constraint is sufficient in the USN model as well, let us now turn to the results developed in the next section.

6 Hypergraphs and Communication

In the construction of the implementing strategies for some distribution that fulfills the information constraint (1) we need a function from a set of allowed actions for the forecaster to a set of action plans for the agent. Using such a function the agent will be able to choose an appropriate action plan according to the past actions of the forecaster.

In the signaling model we examine, because the agent does not observe the past states of nature, he has incomplete knowledge of the constraints, that were imposed on the forecaster in his past actions. So while the forecaster has to choose his actions (which are the effective messages to the agent) from a set of allowed actions limited by the states of nature, the agent has little knowledge about what is the set of allowed actions for the forecaster.

This situation from the agent's point of view can be described as a hypergraph (see Definition 9) in which the vertices are all the sequences of allowed actions of the forecaster for some sequence of states of nature drawn i.i.d. according to μ and the edges are the sets of allowed actions for the forecaster conditioned on some specific sequence of states of nature. While the forecaster is limited in his choice of actions to some specific edge the agent does not know which edge it is.

Those hypergraphs that we shall construct when proving the existence of implementing strategies have some simple combinatorial properties. In particular, they are uniform (all the edges have the same cardinality) and the sizes of the vertex sets and the edges behave asymptotically due to combinatorial information approximations.

What we require in the construction is a communication function from the vertex set to a given set of action plans for the agent so that for every action plan in this set there will be a source in each edge. Using such a function the forecaster will be able to pinpoint every action plan for the agent no matter which edge he is limited to by nature.

6.1 Full-Colorings of Uniform Hypergraphs

We prove the existence of such a communication function using the probabilistic method [1]. For the sake of generality we state our results in terms of general hypergraph theory.

Definition 9

1. *A hypergraph is a pair $H = (V, E)$ where V is a finite set and $E \subseteq 2^V$. The elements of V are called vertices and the elements of E are called edges. A hypergraph $H = (V, E)$ is a graph if all the edges are of size 2.*

2. *A hypergraph $H = (V, E)$ is said to be r-uniform if every edge $X \in E$ is of size r. r is called the rank of the hypergraph. The rank of a non-uniform hypergraph is the size of the largest edge.*

We are chiefly interested in uniform hypergraphs. We now present the notion of a full coloring of a hypergraph. This notion corresponds to the communication function discussed above.

Definition 10 *Let c be an integer. A full c-coloring (called a good c-coloring by Berge [2]) of a hypergraph $H = (V, E)$ is a function $f : V \to \mathbb{N}_c = \{1, 2, \ldots, c\}$ such that $(\forall k \in \mathbb{N}_c)(\forall X \in E) f^{-1}(\{k\}) \cap X \neq \emptyset$. That is a full c-coloring of a hypergraph is a coloring of its vertices such that there is a vertex of each color in every edge.*

Note that for a full c-coloring to exist the size of each edge of the hypergraph must be greater than or equal to c.

Following Erdős and Hajnal, who in [4] studied full 2-colorings of hypergraphs (in the case of 2-colorings the notion of a full coloring coincides we the usual notion of hypergraph coloring[3]), we denote by $m_f(r, c)$ the smallest integer such that there exists an r-uniform hypergraph with $m_f(r, c)$ edges with no full c-coloring. We prove the following bound on $m_f(r, c)$ (which is a generalization of the result by Erdős and Hajnal from [4]).

[3]The usual notion of hypergraph coloring is a vertex coloring such that each edge is colored by at least two different colors.

Proposition 5 *Let $H = (V, E)$ be a general hypergraph with m edges of size $r_1, r_2, \ldots, r_m$ that fulfills the following inequality*

$$c \cdot \sum_{i=1}^{m} (1 - \frac{1}{c})^{r_i} < 1$$

for some integer $c \leq \min(r_1, r_2, \ldots, r_m)$. In such a case H has a full c-coloring.

For uniform hypergraphs we deduce that for every pair of integers $r \geq c \geq 1$ we have

$$m_f(r, c) \geq \frac{1}{c}(1 - \frac{1}{c})^{-r}$$

Proof. Let $H = (V, E)$ be a hypergraph. We denote $E = \{X_1, X_2, \ldots, X_m\}$ such that $|X_i| = r_i$ for each $1 \leq i \leq m$. We need to prove for every $c \leq \min(r_1, r_2, \ldots, r_m)$ that if $c \cdot \sum_{i=1}^{m} (1 - \frac{1}{c})^{r_i} < 1$, then there exists a full c-coloring of H. We use the probabilistic method.

First, we color the vertices of H randomly in c colors (each vertex is colored independently according to a uniform law over $\mathbb{N}_c$). Let A denote the event that the coloring is not a full one. Let $A_{i,k}$ be the event that the edge X_i is not colored by the color k. Note that $\Pr(A_{i,k}) = (\frac{c-1}{c})^{r_i} = (1 - \frac{1}{c})^{r_i}$. We now have

$$\Pr(A) = \Pr(\bigcup_{1 \leq i \leq m, 1 \leq k \leq x} A_{i,k}) \leq \sum_{i=1}^{m} \sum_{k=1}^{c} \Pr(A_{i,k}) = c \cdot \sum_{i=1}^{m} (1 - \frac{1}{c})^{r_i} < 1$$

We conclude that there exist at least one full c-coloring of H.

For an r-uniform hypergraph we have

$$\Pr(A) \leq c \cdot \sum_{i=1}^{m} (1 - \frac{1}{c})^{r_i} = mc(1 - \frac{1}{c})^{r}$$

and so every r-uniform hypergraph with $m < \frac{1}{c}(1 - \frac{1}{c})^{-r}$ has a full c-coloring. □

The next corollary supplies an asymptotic bound for $m_f(r, c)$ for very large $r \geq c$. Essentially, we see that in the case that r/c grows to infinity the growth rate of $m_f(r, c)$ is at least exponential in the growth rate of r/c, yet if r/c is bounded while c grows to infinity our bound tends to zero and is useless for large values of c.

Corollary 2 *As $r \geq c \to \infty$ we have*

$$m_f(r,c) \geq \frac{1}{c} e^{r/c(1+o(1))}$$

Proof.

$$\begin{aligned} m_f(r,c) &\geq \frac{1}{c}(1-\frac{1}{c})^{-r} = \frac{1}{c}\exp[-\ln(1-\frac{1}{c})\cdot r] \\ &= \frac{1}{c}\exp[-(-\frac{1}{c}-o(1/c))\cdot r] \\ &= \frac{1}{c}\exp[r/c(1+c\cdot o(1/c))] \\ &= \frac{1}{c}e^{r/c(1+o(1))} \end{aligned}$$

□

Finally, we prove the following second corollary which implies that in the case that r and c grow exponentially and r grows faster than c, if the growth rate of the amount of edges is also exponentially bounded, then for r,c large enough we always have a full c-coloring for an r-uniform hypergraph with an amount of edges bounded by the above-mentioned exponential growth rate.

Corollary 3 *For two series of integers $(r_n)_{n=1}^{\infty}$, $(c_n)_{n=1}^{\infty}$ such that $\lim_{n\to\infty}\frac{\log r_n}{n} = \rho$, $\lim_{n\to\infty}\frac{\log c_n}{n} = \gamma$ for $\rho,\gamma \in \mathbb{R}$ and $\rho > \gamma$, we have $\lim_{n\to\infty}\frac{\log m_f(r_n,c_n)}{n} = \infty$.*

In particular, for a series of integers $(m_n)_{n=1}^{\infty}$ such that $m_n \leq a_n$ for all n and $\lim_{n\to\infty}\frac{\log a_n}{n} = \alpha$ for some $\alpha \in \mathbb{R}$ for n large enough, we always have $m_n < m_f(r_n,c_n)$.

Proof. Note that $r_n/c_n = 2^{n(\rho-\gamma+o(1))} \to_{n\to\infty} \infty$, and so for n large enough r_n is necessarily $\geq c_n$ and $m_f(r_n,c_n)$ exists.

By Corollary 2 we know that

$$\begin{aligned} \frac{\log m_f(r_n,c_n)}{n} &\geq \frac{1}{n\ln(2)}\frac{r_n}{c_n}(1+o(1)) - \log c_n/n \\ &= \frac{1}{n\ln(2)}2^{n(\rho-\gamma+o(1))}(1+o(1)) - \gamma + o(1) \to_{n\to\infty} \infty \end{aligned}$$

For a series $(m_n)_{n=1}^{\infty}$ (as above) it is sufficient to prove that for n large enough we always have $a_n < m_f(r_n,c_n)$:

$$\log\frac{m_f(r_n,c_n)}{a_n} \geq \frac{1}{n}\log\frac{m_f(r_n,c_n)}{a_n} = \frac{1}{n}\log m_f(r_n,c_n) - \alpha + o(1) \to_{n\to\infty} \infty$$

Let n be large enough so that $\log \frac{m_f(r_n,c_n)}{a_n} > 0$; then necessarily $m_f(r_n, c_n) > a_n$. □

6.2 Strategic Hypergraphs

We now turn to those hypergraphs and full colorings that will be used in the construction of implementing strategies. We construct those hypergraphs for a distribution $Q \in \Delta(I \times J \times K)$ that fulfills some properties not common to all the distributions that fulfill the information constraint (1). In the proof of Theorem 1 we will use approximations of distributions to overcome this difficulty.

Definition 11 *Let $Q \in \Delta(I \times J \times K)$ such that*

$$(\exists m \in \mathbb{N}) \quad Q \in \mathbb{T}_m(I \times J \times K) \tag{5}$$

$$(\exists \varepsilon > 0) \quad H_Q(j|i,k) - I_Q(i;k) \geq \varepsilon \tag{6}$$

We call such a distribution a **strategic distribution**.

1. *The mn-strategic hypergraph for an integer n and $\mathbf{k} \in T_{mn}(Q_K)$ is a hypergraph $SH_{mn}(Q) = (V, E)$ with a vertex set $V = T_{mn}(Q_{J\times K}|\mathbf{k})$ and a set of edges $E = \{T_{mn}(Q|\mathbf{i}, \mathbf{k}) \mid \mathbf{i} \in T_{mn}(Q_I)\}$.*

2. *An mn-strategic action set $A_{mn}(Q)$ for an integer n is a minimal 1-fold mn-action set $A_{mn}(Q) \subseteq T_{mn}(Q_K)$ of the distribution $Q_{I\times K} \in \mathbb{T}_{mn}(I \times K)$.*

The choice of the specific $\mathbf{k} \in T_{mn}(Q_K)$ will be irrelevant to the properties we prove in this section. The next lemma is central to the construction of implementing strategies to be presented in Section 7.

Lemma 2 *Let $Q \in \Delta(I \times J \times K)$ fulfill (5) and (6). For each integer n large enough and $\mathbf{k} \in T_{mn}(Q_K)$ there exists an mn-communication function $\mathrm{Com}_{mn} : T_{mn}(Q_{J\times K}|\mathbf{k}) \to A_{mn}(Q)$ such that*

$$(\forall \mathbf{a} \in A_{mn}(Q))(\forall \mathbf{i} \in T_{mn}(Q_I)) \quad T_{mn}(Q|\mathbf{i}, \mathbf{k}) \cap \mathrm{Com}_{mn}^{-1}(\{\mathbf{a}\}) \neq \emptyset \tag{7}$$

Proof. We need to prove that the strategic hypergraph $SH_{mn}(Q)$ has a full c_n-coloring for $c_n = |A_{mn}(Q)|$. $SH_{mn}(Q)$ is a uniform hypergraph of rank

$r_n = |T_{mn}(Q|\mathbf{i},\mathbf{k})|$ (the size of $T_{mn}(Q|\mathbf{i},\mathbf{k})$ does not depend on the specific choice of $\mathbf{i} \in T_{mn}(Q_I)$). Denote by $\rho = mH_Q(j|i,k)$. By (3) we have

$$\lim_{n\to\infty} \frac{\log r_n}{n} = \rho$$

The number of edges in $SH_{mn}(Q)$ ($=m_n$) is bounded by the size of $T_{mn}(Q_{I\times K}|\mathbf{k})$. Let $a_n = |T_{mn}(Q_{I\times K}|\mathbf{k})|$. Let $\alpha = mH_Q(i|k)$. The limit (3) implies that

$$\lim_{n\to\infty} \frac{\log a_n}{n} = \alpha$$

The number of required colors is $c_n = |A_{mn}(Q)| = \mathcal{A}_n(Q)$. From the limit (4) we can deduce for $\gamma = mI_Q(i;k)$,

$$\lim_{n\to\infty} \frac{\log c_n}{n} = \gamma$$

Because Q fulfills the sharp information constraint (6): $H_Q(j|i,k) - I_Q(i;k) \geq \varepsilon$ for some $\varepsilon > 0$, we have $\rho - \gamma \geq m\varepsilon > 0$. Now all the requirements of Corollary 3 (of Proposition 5) are fulfilled. Corollary 3 implies that for all n large enough $m_n < m_f(r_n, c_n)$ and the hypergraph $SH_{mn}(Q)$ has a full c_n-coloring as required. □

7 Construction of Implementing Strategies

In this section we present the proof of the main theorem, Theorem 1. The proof relies on the ability to approximate any distribution having marginal μ on I and fulfilling the information constraint (1) by a distribution that can be implemented by the strategies we construct. The closeness of $\mathcal{Q}$ (see Remark 2) will thus imply that any distribution fulfilling the information constraint is implementable.

7.1 Approximating Distributions

The are two steps in the approximation technique as in [6]. The first step is to approximate a distribution $Q \in \Delta(I \times J \times K)$ having marginal μ on I and fulfilling the information constraint (1) by a strategic distribution having a strategic hypergraph (such a distribution fulfills (5) and (6)).

The following lemma from [6] allows us to make this first approximation:

Lemma 3 *For every $\varepsilon > 0$ there exists an integer N such that for all $Q \in \Delta(I \times J \times K)$ with $H_Q(i, j \mid k) - H_Q(i) \geq 0$ and $n \geq N$ there exists $\bar{Q} \in \Delta(I \times J \times K)$ such that*

$$\begin{aligned} &\bar{Q} \in \mathbb{T}_n(I \times J \times K) \\ &H_{\bar{Q}}(i, j \mid k) - H_{\bar{Q}}(i) \geq \varepsilon \\ &\|Q - \bar{Q}\|_1 < 7\varepsilon \end{aligned}$$

In particular, $\bar{Q}$ is a strategic distribution.

An approximating strategic distribution $\bar{Q}$ does not necessarily fulfill $\bar{Q}_I = \mu$ (sometimes this is actually impossible, as in cases where μ assigns irrational probabilities to some elements of I). Because every implementable distribution clearly must induce μ on I we do not have a direct way to implement any strategic distribution.

Yet because $\bar{Q}$ approximates Q the distribution $\bar{Q}_I$ approximates μ. By the second step of approximation, we can design a way to implement a distribution that is sufficiently close to $\bar{Q}$. This is carried out using a translation function that gives for any sequence of length n of states of nature distributed i.i.d. according to μ a different sequence of states of nature distributed according to some other distribution ν (in our use this distribution ν will be the marginal of the strategic distribution $\bar{Q}$ on I). The following lemma from [6] states that if we use such a translation function that minimizes the amount of states of nature in the sequence changed by the translation function, then for a sequence of states of natures distributed according to μ the probability for a large amount of errors to occur because of the translation is small for large n.

Lemma 4 *Let $\nu \in \Delta(I)$ and $n \in \mathbb{N}$ such that $\nu \in \mathbb{T}_n(I)$. There exists a translation map $f : I^n \to T_n(\nu)$ such that for $\mu \in \Delta(I)$ we have*

$$\Pr_{\mu^{\otimes n}}\left(\sum_{t=1}^{n} \mathbb{I}_{x_t \neq f_t(\mathbf{x})} > \|\nu - \mu\|_1 n + \varepsilon n\right) \leq \frac{|I|^2}{\varepsilon^2 n}$$

In particular, if $\mathrm{Ts}^n : I^n \to T_n(\nu)$ is a map that minimizes the sum $\sum_{t=1}^{n} \mathbb{I}_{x_t \neq f_t(x)}$ and $\|\nu - \mu\|_1 \leq 7\varepsilon$ then

$$\Pr_{\mu^{\otimes n}}\left(\sum_{t=1}^{n} \mathbb{I}_{x_t \neq (\mathrm{Ts}^n(\mathbf{x}))_t} > 8\varepsilon n\right) \leq \frac{|I|^2}{\varepsilon^2 n} \tag{8}$$

□

7.2 Proof of Theorem 1

We already know that any implementable distribution fulfills the information constraint. Now we shall prove the converse.

Let $Q \in \Delta(I \times J \times K)$ fulfill the information constraint (1) and have marginal μ on I

$$Q_I = \mu$$
$$H_Q(j \mid i, k) \geq I_Q(i; k)$$

We prove that for every $\varepsilon > 0$ there exists a distribution $Q' \in \mathcal{Q}$ such that $\|Q - Q'\|_1 \leq 26\varepsilon$. Because $\mathcal{Q} = \bar{\mathcal{Q}}$ this implies that $Q \in \mathcal{Q}$.

By Lemma 3 there exists an integer $m > |I|^2/\varepsilon^3$ and an approximating strategic distribution $\bar{Q} \in \mathbb{T}_m(I \times J \times K)$ such that $H_{\bar{Q}}(i, j \mid k) - H_{\bar{Q}}(i) \geq \varepsilon$ and $\|Q - \bar{Q}\|_1 < 7\varepsilon$. In particular, $\|\mu - \bar{Q}_I\|_1 < 7\varepsilon$.

By Lemma 2 there exists an integer n large enough and a communication function $\text{Com} = \text{Com}_{mn} : T_{mn}(\bar{Q}_{J\times K}|\mathbf{k}) \to A_{mn}(\bar{Q}_{I\times K})$ fulfilling (7), where $A_{mn} = A_{mn}(\bar{Q}_{I\times K})$ is an optimal mn-action set for the distribution $\bar{Q}_{I\times K}$.

Denote $l = mn$. Because $\|\mu - \bar{Q}_I\|_1 < 7\varepsilon$ by Lemma 4, there exists a translation function $\text{Ts} = \text{Ts}^l : I^l \to T_l(\bar{Q}_I)$ fulfilling (8).

Partition the stages of the game into consecutive blocks of length l. Let $n_r = lr$ for any integer $r \geq 0$. We denote by $\mathbf{x}[r] = (\mathbf{i}[r], \mathbf{j}[r], \mathbf{k}[r])$ the sequence of the states of nature and the actions of the agent and the forecaster in the r block: $(i_t, j_t, k_t)_{t=n_{r-1}+1}^{n_r}$. We also use the notation $\tilde{\mathbf{i}}[r]$ for $\text{Ts}(\mathbf{i}[r])$ and $\tilde{\mathbf{x}}[r] = (\tilde{\mathbf{i}}[r], \mathbf{j}[r], \mathbf{k}[r])$.

To decide upon the strategies in the first block, we choose arbitrarily a sequence of actions for the agent in the first block $\mathbf{k}[1] \in K^l$. In all the following blocks the agent will play according to $\mathbf{k}[r+1] = \text{Com}(\mathbf{j}[r])$.

In any block r the forecaster chooses from the agent's action set A_l a sequence of action $\mathbf{k}'[r+1] \in T_l(\bar{Q}_{I\times K} \mid \tilde{\mathbf{i}}[r+1])$ (this is always possible because A_l is an action set). As Com is a communication function fulfilling (7), the sequence $\mathbf{k}'[r+1] \in A_l$ always has a source under Com in $T_l(\bar{Q} \mid \tilde{\mathbf{i}}[r], \mathbf{k}[r])$ except for $r = 1$ (in the first block we do not have necessarily $\mathbf{k}[1] \in T_l(\bar{Q}_{I\times K} \mid \mathbf{i}[1])$. The forecaster chooses one of those sources as $\mathbf{j}[r]$ and plays it. In the first block the forecaster chooses some $\mathbf{j}[1] \in T_l(\bar{Q}_{I\times J} \mid \tilde{\mathbf{i}}[1])$ not necessarily in $T_l(\bar{Q} \mid \tilde{\mathbf{i}}[1], \mathbf{k}[1])$. The actual way in which the forecaster makes these choices is irrelevant to the proof.

This dictates a pure strategy for the forecaster and the agent. In this strategy we have for each block $r > 1$

$$\mathbf{k}[r] = \mathrm{Com}(\mathbf{j}[r-1]) = \mathbf{k}'[r] \in T_l(\bar{Q}_{I\times K} \mid \tilde{\mathbf{i}}[r])$$
$$\mathbf{j}[r] \in T_l(\bar{Q} \mid \tilde{\mathbf{i}}[r], \mathbf{k}[r])$$

This implies that

$$(\forall r > 1) \quad \tilde{\mathbf{x}}[r] \in T_l(\bar{Q}) \tag{9}$$

Let $\hat{Q}_{n_s}$ be the expectation of the distribution induced by the actions and the states of nature in the first s blocks of the game: $(i_t, j_t, k_t)_{t=1}^{n_s}$ ($\hat{Q}_{n_s} \in \mathcal{Q}(n_s) \subseteq \mathcal{Q}$). Let $\bar{Q}_{n_s}$ be the expectation of the distribution induced by the translation of the states of nature and the actions in the first s blocks: $(\tilde{i}_t, j_t, k_t)_{t=1}^{n_s}$.

Equation (9) implies that $\bar{Q}_{n_s}$ can differ from $\bar{Q}$ only because of the actions in the first block. In fact, it is only the agent's actions which diverge from the typical sequence. The distance between the empirical distributions of $\tilde{\mathbf{x}}[1]$ and $\bar{Q}$ in the $\|\bullet\|_1$ norm is bounded by $\frac{2}{l}\sum_{t=1}^{l} \mathbb{I}_{k_t \neq a_t}$ where $(a_t)_{t=1}^{l}$ is some sequence in $T_l(\bar{Q} \mid \tilde{\mathbf{i}}[1], \mathbf{j}[1])$. This implies that $\|\bar{Q}_{n_s} - \bar{Q}\|_1 \leq 2/s$. So for $s \geq 1/(2\varepsilon)$ we have $\|\bar{Q} - \bar{Q}_{n_s}\|_1 \leq \varepsilon$.

The distance between the empirical distribution of $\mathbf{x}[r]$ and $\tilde{\mathbf{x}}[r]$ for some r is clearly bounded by $\frac{2}{l}\sum_{t=1}^{l} \mathbb{I}_{i_t \neq \tilde{i}_t}$ (and thus by 2). The translation function fulfills (8) and we deduce for the empirical distribution in each block $r \geq 1$

$$\begin{aligned}
\|\rho(\mathbf{x}[r]) - \rho(\tilde{\mathbf{x}}[r])\|_1 &\leq \Pr_{\mu^{\otimes l}}\Big(\sum_{t=n_{r-1}+1}^{n_r} \mathbb{I}_{i_t \neq \tilde{i}_t} > 8\varepsilon l\Big) \cdot 2 \\
&\quad + \Pr_{\mu^{\otimes l}}\Big(\sum_{t=n_{r-1}+1}^{n_r} \mathbb{I}_{i_t \neq \tilde{i}_t} \leq 8\varepsilon l\Big) \cdot 2 \cdot 8\varepsilon \\
&\leq 2\Big(\frac{|I|^2}{\varepsilon^2 l} + 8\varepsilon\Big) \leq 2\Big(\frac{|I|^2}{\varepsilon^2 m} + 8\varepsilon\Big) \leq 18\varepsilon
\end{aligned}$$

This implies that $\|\hat{Q}_{n_s} - \bar{Q}_{n_s}\|_1 \leq 18\varepsilon$ and by the triangle inequality for every $s \geq 1/(2\varepsilon)$

$$\|Q - \hat{Q}_{n_s}\|_1 \leq \|Q - \bar{Q}\|_1 + \|\bar{Q} - \bar{Q}_{n_s}\|_1 + \|\bar{Q}_{n_s} - \hat{Q}_{n_s}\|_1 \leq (7+1+18)\varepsilon = 26\varepsilon$$

□

7.3 Variations on the Proof

7.3.1 Quasi-optimal Action Sets

Instead of using optimal action sets one can have the same proof unchanged for every series of action sets S_{mn} such that $lim_{n\to\infty}\frac{H(S_{mn})}{n} = mI_{\bar{Q}}(i;k)$. We call such sets quasi-optimal action sets.

The calculation of the actual n-optimal action set is a computationally complex problem, while quasi-optimal action sets, can be generated using a probabilistic approach. The proof in [6] of Proposition 4, which supplies us with quasi-optimal action sets, uses the probabilistic method. A close examination of the proof shows that for a distribution $Q \in \mathbb{T}_m(A \times B)$, the probability that a random set of the same size as in Proposition 4 whose elements are chosen independently and uniformly from $T_{mn}(Q_B)$ is an action set, goes to 1, as n goes to ∞ .

Thus for large blocks we can generate quasi-optimal action sets from random samplings with a high probability of success.

7.3.2 Maximal Action Sets

A different proof can be devised in which maximal action sets $(T_{mn}(\bar{Q}_K))$ are used instead of optimal ones. The proof of the characterization of the implementable distributions for the GHN model in [6] focuses on this method.

In case maximal action sets are used a different but equivalent formulation of the information constraint is used:

$$H_Q(j,k|i) \geq H_Q(k)$$

From this formulation we can see that for large n an approximating strategic distribution $\bar{Q}$ will fulfill $|T_{mn}(\bar{Q}|\mathbf{i})| > |T_{mn}(\bar{Q}_K)|$. So in the GHN model every element of the maximal action set can still be pinpointed by the actions in the previous block but not solely by the forecaster's actions. The agent has to choose his action plan in block r according to both the forecaster's actions and his actions in block $r-1$ (the message set is now $T_{mn}(\bar{Q}|\mathbf{i})$).

To receive the correct message in block r the agent must have played the appropriate sequence of actions in the previous block, and so on. So if we wish for the agent to play some sequence of actions in block r, this defines uniquely in this scheme his actions in all previous blocks and thus the forecaster's as well. The conclusion is that the forecaster must use backtracking from block

r to compute his actions in all previous blocks. For a complete description of this process see [6, Section 7].

The analogy between the GHN model and our USN model in this scheme of actions is the same as with optimal action sets. In order for a strategy of this sort to exist in the USN model we need a communication function from $T_{mn}(\bar{Q})$ to $T_{mn}(\bar{Q}_K)$ whose restriction to $T_{mn}(\bar{Q}|\mathbf{i})$ will be onto $T_{mn}(\bar{Q}_K)$ for every $\mathbf{i} \in T_{mn}(\bar{Q}_I)$. This is equivalent to the existence of a good $|T_{mn}(\bar{Q}_K)|$-coloring of the analogous strategic hypergraph whose vertex set is $T_{mn}(\bar{Q})$ and the set of edges is $\{T_{mn}(\bar{Q}|\mathbf{i})|\mathbf{i} \in T_{mn}(\bar{Q}_I)\}$. The existence of such a good coloring is proved in the same way as in Lemma 2.

7.3.3 A Continuum of Action Sets

In fact, by using the backtracking scheme one can use optimal f_{mn}-fold mn-action sets for a series $(f_{mn})_{n=1}^{\infty}$ that fulfills $\lim_{n\to\infty} = \frac{\log f_{mn}}{mn} = \varphi$ for every $0 \leq \varphi \leq H_{\bar{Q}}(k|i)(= lim_{n\to\infty} \frac{H(T_{mn}(\bar{Q}_{I\times K}|\mathbf{i}))}{mn})$ and large enough block,. where the minimal $\varphi = 0$ corresponds to minimal action sets and the maximal $\varphi = H_{\bar{Q}}(k|i)$ corresponds to the maximal action sets. The appropriate information constraint will be $H_Q(j,k|i) - \varphi \geq H_Q(k) - \varphi$.

Note that by using optimal action sets we can use a strategy without any backtracking, as described in our proof of Theorem 1. So $\varphi = 0$ requires the least amount of computations during the game (in every block each of the players needs only to evaluate a single predefined function either on the forecaster's actions in the previous block or on the states of nature in the next block). The use of maximal action sets, on the other hand, requires backtracking (at the first block the forecaster needs to compute his actions during the whole game), but no computation of an optimal action set before the game.

8 Extensions and Applications

8.1 Imperfect Forecasts

Models with limited forecasting abilities were first discussed in [6, Section 9]. Two cases of limited forecasts are addressed in [6] in regard to the GHN model: finite forecasts and imperfect ones. We shall focus on the second case.

In the imperfect forecast model the forecaster does not receive the infinite sequence of the states of nature $(i_t)_{t=1}^{\infty}$, but rather an infinite sequence of signals $(u_t)_{t=1}^{\infty}$ from a finite signal set U. The signals are drawn independently according to some transition probability distribution ν from I to U.

In the case where the agent observes u_t of previous stages (a GHN model), it was already stated in [6] that a distribution $Q \in \Delta(I \times J \times K)$ with $Q_I = \mu$ is implementable if and only if it is the marginal on $I \times J \times K$ of a distribution $\hat{Q} \in \Delta(I \times U \times J \times K)$ with

$$\hat{Q}(u|i) = \nu[i](u)$$
$$\hat{Q}(i|u, j, k) = \hat{Q}(i|u)$$

and an appropriate information constraint

$$H_{\hat{Q}}(j|u, k) \geq I_{\hat{Q}}(u; k)$$

The first of the three equations above states that the distribution induces the correct transition distribution between the states of nature and the signals. The second equation implies that all the information the players have on the states of nature comes from the signals. The third equation is the standard information constraint in which the state of nature variable was replaced with the signal one: if all the information the players have about the states of nature is the signals then any correlation with nature is achieved through correlation with the signal.

The necessity of these conditions in the GHN model implies their necessity in the USN model (this is proved using the concavity of the entropy function; see [6, Section 5]).

The construction of implementing strategies follows the same lines as in that of a perfect forecast, the only difference being that the states of nature are in most places replaced by the signals.

As with perfect forecasts, every strategy implementable in the GHN model is also implementable in the USN model of this paper (the agent does not observe past signals) by the same technique of full coloring of the appropriate strategic hypergraphs.

8.2 General Signalling Structure

In this section we consider a generalization of the USN model where the agent does not have perfect monitoring of the forecaster's action or the states of nature.

At each stage the agent receives a signal s_t from a finite signal set S. The signal is drawn independently at each stage according to some transition distribution η from $I \times J \times K$. We refer to this model as the USN-signaling model. We'll also discuss the GHN-signaling model in which the agent receives the signal s_t and the past state of nature i_t.

In the USN-signaling model the GHN model corresponds to the case $s_t = (i_t, j_t, k_t)$ and the standard USN model corresponds to $s_t = (j_t, k_t)$.

In [6, Section 9.2.2] the following characterization is given for the GHN-signaling model (the agent observes the signal s_t and the state of nature i_t). A distribution $Q \in \Delta(I \times J \times K)$ with marginal μ over I is implementable if and only if it is the marginal over $I \times J \times K$ of a distribution $\hat{Q} \in \Delta(I \times J \times K \times S)$ with a marginal transition distribution from $I \times J \times K$ to S equal to η ($\hat{Q}(s|i,j,k) = \eta[i,j,k](s)$) for all $(i,j,k,s) \in I \times J \times K \times S$) that fulfills the following information constraint:

$$I_{\hat{Q}}(j; s|i,k) \geq I_{\hat{Q}}(i;k) \tag{10}$$

In the GHN-signaling model (the agent observes s_t and i_t) implementing strategies can be constructed using a similar method to the standard proof in [6] using optimal action sets. In each block the agent can deduce all his actions in the previous block (they depend only on past signals and states of nature that are known to him) and he has knowledge of the past states of nature.

If in the model with the standard GHN signal $s_t = (i_t, j_t, k_t)$ in block $r+1$ the agent chooses a sequence of actions from an action set $A_{mn}(\bar{Q})$ according to the forecaster's action in the previous block (a sequence from $T_{mn}(\bar{Q}|i[r], k[r])$), then in a general GHN-signaling model the forecaster cannot transmit $j[r]$ to the agent. Instead he can choose a subset of actions $M_{mn} \subseteq T_{mn}(\bar{Q}_{I\times J\times K}|i[r], k[r])$ such that for any $\mathbf{j} \in M_{mn}$ the size of the intersection $(\bigcup_{\mathbf{j}' \in M_{mn}\setminus\{\mathbf{j}\}} T_{mn}(\bar{Q}|i[r], \mathbf{j}', k[r])) \cap T_{mn}(\bar{Q}|i[r], \mathbf{j}, k[r])$ will tend to 0 as $n \to \infty$. If the set M_{mn} is predetermined and known to both the agent and the forecaster, then the agent will be able to guess what sequence $\mathbf{j} \in M_{mn}$ the forecaster played with an arbitrary small error probability for n large enough.

Using the probabilistic method one can easily demonstrate that a maximal sequence of sets $(M_{mn})_{n=1}^{\infty}$ as above (we refer to them as *message sets*) fulfill $\lim_{n\to\infty} \frac{H(M_{mn})}{mn} = I_{\bar{Q}}(j; s|i,k)$, which corresponds to the capacity of the forecaster to transmit information in inequality (10) (see he Channel Capacity Theorem in [3] for a similar problem).

In the case where the signal s_t is independent of the state of nature i_t our analysis using strategic hypergraphs can also be extended to the case of the USN signaling model (the agent does not receive i_t in addition to the signal s_t). The same technique with slight variations proves that the same characterization valid for the GHN-signaling model is also valid for the general USN one if the signal is independent of the state of nature. Denote $M_{mn}(\mathbf{i}) \subseteq T_{mn}(\bar{Q}_{I\times J\times K}|\mathbf{i}, k[r])$ for some $\mathbf{i} \in T_{mn}(\bar{Q}_{I\times K}|k[r])$ to be the maximal message sets as above. The vertex set will be $\bigcup_{\mathbf{i}\in T_{mn}(\bar{Q}_{I\times K}|k[r])} M_{mn}(\mathbf{i})$ and the edge set $\{M_{mn}(\mathbf{i})|\mathbf{i} \in T_{mn}(\bar{Q}_{I\times K}|k[r])\}$. The rest of the analysis is the same as in the proof of Theorem 1 except that now there is some small probability for the agent to err in the interpretation of the forecaster's message, but as this probability can be made arbitrarily small with large enough blocks and $\mathcal{Q}$ is closed we still have a complete proof.

This result is a major step towards solving the general USN-signaling problem. We have extended the information-theoretic techniques to the case of more general signals. The results regarding to the GHN-signaling model in [6] were a first step, but they had a major artificial limitation: the agent has had to receive in addition to the signal the previous states of nature.

References

[1] Alon, N. and J. H. Spencer. *The Probabilistic Method*. Wiley-Interscience, 2nd edition (2000).

[2] C. Berge. *Hypergraphs, Combinatorics of Finite Sets*. North-Holland (1989).

[3] Cover, T.M. and J.M. Thomas. *Elements of Information Theory*. Wiley-Interscience (1991).

[4] Erdős, P. and A. Hajnal (1961). On a property of families of sets. *Acta Mathematica Hungarica*, 12, 87-123.

[5] Gossner, O., P. Hernandez, and A. Neyman (2003). Online matching pennies. Hebrew University, Center for the Study of Rationality, Discussion Paper 316.

[6] Gossner, O., P. Hernandez, and A. Neyman (2005). Optimal use of Communication Resources. *Econometrica* 74, 6 (November, 2006), 1603-1636.

[7] Shannon, C. (1948). A Mathematical Theory of Communication. *Bell Systems Technical Journal*, 27, 379-423; 623-656.