Revised version

OPTIMAL USE OF COMMUNICATION RESOURCES

OLIVIER GOSSNER, PENÉLOPE HERNÁNDEZ, AND ABRAHAM NEYMAN

Abstract. We study a repeated game with asymmetric information about a dynamic state of nature. In the course of the game, the better informed player can communicate some or all of his information to the other. Our model covers costly and/or bounded communication. We characterize the set of equilibrium payoffs, and contrast these with the communication equilibrium payoffs, which by definition entail no communication costs.

Keywords: communication, information economics, incomplete information, entropy, information processing.
JEL Classification: C61, C73, D82.

1. Introduction

Communication activities may resolve inefficiencies due to information asymmetries between agents, but are themselves costly, e.g., due to sending and processing costs. The study of optimal trade-offs between the costs and the benefits of communication is to a large extent an open problem, and is the topic of this paper.

Communication equilibria, as proposed by Forges [5] and Myerson [11], extend the rules of a game by adding communication possibilities through arbitrary mechanisms at any stage of a multistage game. This concept captures the largest set of implementable equilibria when no restriction exists on the means of communication between the players.

On the other hand, economic studies like Radner [12] tell us that in an organization like a firm, communication is a costly activity and that a significant amount of resources is devoted to processing and sending

Date: February 11, 2006.
This research was supported in part by Israel Science Foundation grants 382/98 and 263/03 and by the Zvi Hermann Shapira Research Fund. The second author thanks the financial support from the Spanish Ministry of Education under project SEJ 2004-02172/ECON and the Instituto Valenciano de Investigaciones Económicas (Ivie).

information. In these structures, the constant need for information updating entails important costs.

Starting with Forges [6] and Bárány [2], a body of literature, including Urbano and Vila [14], Ben-Porath [3], and Gerardi [7], studies models of decentralized communication. An important conclusion of this literature is that—under various assumptions—all communication equilibria can be implemented through preplay decentralized communication procedures. Hence, decentralized communication schemes can be used without any loss of efficiency if we consider that a finite number of communication stages entails negligible costs compared to the payoffs of the game to be played. Since the costs of communication cannot be explained by considering decentralized communication schemes as opposed to centralized ones, another archetype of costly communication is needed to study the trade-off between the costs and benefits of information transmission.

This paper puts the emphasis on this need for information updating, and studies the communication dynamics in a model where the states of nature evolve through time. One player, the forecaster, has better information than the other player, the agent, about the stream of states of nature. The forecaster may choose to send messages and take actions at any stage, and both components are described as part of the action set of the forecaster.

The payoffs associated to action choices and the limited action set of the forecaster model the costs associated to, and the limitations on, both information processing by the agent and the forecaster and information transmission (e.g., through physical or electronic means).

A repeated game takes place between the forecaster, the agent, and nature. The agent's actions at any stage may depend on all past actions and on all past states of nature. The forecaster's actions may depend on all past actions and on all past states of nature, but also on all future states of nature. Hence, the forecaster's actions include a payoff component (since these actions impact players' payoffs), and an information component (since these actions may inform the agent about future states of nature). At each stage, the agent updates his information using his observation both of the current state of nature and of the forecaster's action.

A specification of players' strategies induces a joint dynamics on the triple (state of nature, forecaster's action, agent's action), called the *action triple*. We study this dynamics through the average distribution Q of this action triple. This distribution contains all expected time average statistics of action triples, and is important for strategic purposes since all expected average payoffs depend on players' strategies through it only.

We characterize the set of distributions Q that are implementable by strategies of the forecaster and the agent. The fact that the information used by the agent cannot exceed the information received leads to an information-theoretic inequality expressed using the Shannon [13] entropy function, and which we call the *information constraint*.

On the one hand, we prove that for all strategies of the forecaster and the agent and for any n, the average distribution during the first n stages fulfills the information constraint. On the other hand, we prove that for any distribution Q that fulfills the information constraint, there exists a pair of pure strategies for the forecaster and the agent such that the long-run average distribution of action triples is Q. Hence, the information constraint fully characterizes the set of implementable distributions.

This result has many implications on the optimal use of communication resources, in both team games and general games.

The cost of communication inefficiencies can be measured in team games, where the (unique) Pareto payoff for the team is the natural solution concept. In the communication equilibrium extension of our model, this Pareto payoff corresponds to the first-best in which both players are perfectly informed of the state of nature at each stage. In our game, the Pareto payoff is in general strictly less than this first-best, and represents a second-best payoff that takes into account the implementation costs of communication processes. Our analysis allows us to compute the optimal payoff and to design optimal strategies for the team that take into account the communication costs and limitations.

In general games, we characterize the set of equilibrium payoffs when players are sufficiently patient. This set is a subset of the set of extensive-form communication equilibrium payoffs, which assume costless and unbounded communication, and it is a superset of the set of "silent" equilibrium payoffs in which no information is transmitted from the forecaster to the agent.

Section 2 presents the model and defines and presents elementary properties of implementable distributions. Section 3 introduces the information constraint and the main results. In Section 4 we show that using mixed or correlated strategies (instead of pure strategies) does not change the analysis. The main results are proved in Sections 5 and 7. Section 6 presents a formulation of the information constraint in terms of relative sizes of sets, which is a fundamental tool in the construction of optimal communication schemes. Section 8 presents applications to team games and to general games, and we conclude with a discussion and extensions in Section 9.

2. The model

We present here the basic version of our model. Most of its assumptions are relaxed in Section 9.

Given a finite set A, $\Delta(A)$ represents the set of probability measures over A, and $|A|$ is the cardinality of A. Random variables are denoted by bold letters.

The finite set of states of nature is denoted by I. There are two players: the forecaster, with finite action set J, and the agent, with finite action set K. The stage payoff functions are $g^f, g^a \colon I \times J \times K \to \mathbb{R}$ for the forecaster and for the agent, respectively, and $g = (g^f, g^a)$. We assume $|J| \geq 2$ so that possibilities of communication from the forecaster to the agent exist.

In the repeated game, the forecaster is informed beforehand of future states of nature. At each stage, the chosen action may depend on past actions, as well as on the sequence of states of nature. A (pure) strategy for the forecaster is thus a sequence $(\sigma_t)_t$ of mappings $\sigma_t \colon I^{\mathbb{N}} \times J^{t-1} \times K^{t-1} \to J$, where σ_t describes the behavior at stage t.

The agent is informed of past realizations of nature and past actions only. A (pure) strategy for the agent is thus a sequence $(\tau_t)_t$ of mappings $\tau_t \colon I^{t-1} \times J^{t-1} \times K^{t-1} \to K$, where τ_t describes the behavior at stage t.

We assume that the sequence $(\mathbf{i}_t)_t$ of states of nature is i.i.d. of stage law μ. A pair of strategies (σ, τ) induces sequences of random variables $(\mathbf{j}_t)_t$ and $(\mathbf{k}_t)_t$ given by $\mathbf{j}_t = \sigma_t((\mathbf{i}_{t'})_{t'}, (\mathbf{j}_1, \ldots, \mathbf{j}_{t-1}), (\mathbf{k}_1, \ldots, \mathbf{k}_{t-1}))$ and $\mathbf{k}_t = \tau_t((\mathbf{i}_1, \ldots, \mathbf{i}_{t-1}), (\mathbf{j}_1, \ldots, \mathbf{j}_{t-1}), (\mathbf{k}_1, \ldots, \mathbf{k}_{t-1}))$. We denote the induced probability distribution over $(I \times J \times K)^{\mathbb{N}}$ by $P_{\mu,\sigma,\tau}$, and the marginal over stage t's action triple by $P^t_{\mu,\sigma,\tau}$. The average distribution up to stage t is $Q^t_{\mu,\sigma,\tau} = \frac{1}{t}\sum_{t'=1}^{t} P^{t'}_{\mu,\sigma,\tau}$.

We say that a distribution $Q \in \Delta(I \times J \times K)$ is *implementable*, respectively t-*implementable*, when there exists a strategy pair (σ, τ) such that $Q^t_{\mu,\sigma,\tau} \to Q$ as $t \to \infty$, respectively $Q^t_{\mu,\sigma,\tau} = Q$, and in this case the strategy pair (σ, τ) *implements*, respectively t-*implements*, the distribution Q.

The set of implementable, respectively t-implementable, distributions is denoted $\mathcal{Q}$, respectively $\mathcal{Q}(t)$. Note the following elementary properties of implementable distributions.

Remark 1. *(1) $\mathcal{Q}(t)$ is closed*[1]*;*
(2) every $Q \in \mathcal{Q}(t)$ is implementable;
(3) $\frac{s}{s+t}\mathcal{Q}(s) + \frac{t}{s+t}\mathcal{Q}(t) \subseteq \mathcal{Q}(s+t)$ and thus also $\mathcal{Q}(t) \subseteq \mathcal{Q}(kt)$;
(4) the Hausdorff distance between $\mathcal{Q}(s)$ and $\mathcal{Q}(s+t)$ is bounded by $\frac{2t}{s+t}$.

These elementary properties imply that

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

Let F_t, respectively F, denote the set of feasible payoff vectors of the t-stage game, respectively of the infinitely repeated game. Note that F_t and F are the linear images of $\mathcal{Q}(t)$ and $\mathcal{Q}$ under the expectation operator: $v \in F_t$ if and only if there is $Q \in \mathcal{Q}(t)$ such that $v = \mathbf{E}_Q g$ and $v \in F$ if and only if there is $Q \in \mathcal{Q}$ such that $v = \mathbf{E}_Q g$.

Interstage-time-dependent discount factors $(\lambda_t)_{t\geq 1}$ lead to a weighted average valuation of a stream of payoffs (g_t) given by $\sum_{t=1}^{\infty}\theta_t g_t$, where $\theta_t = \frac{\prod_{s<t}\lambda_s}{\sum_{u=1}^{\infty}\prod_{s<u}\lambda_s}$. Thus $\theta = (\theta_t)_{t\geq 1}$ is a nonincreasing sequence with

[1]Under the norm distance $\|Q - Q'\| = 2\max_{X\subset I\times J\times K}(Q(X) - Q'(X)) = \sum_{(i,j,k)\in I\times J\times K}|Q(i,j,k) - Q'(i,j,k)|$.

$\sum_{t=1}^{\infty} \theta_t = 1$. The valuation in the commonly used λ-discounted game for some $0 < \lambda < 1$ corresponds to the case where $\theta_t = (1-\lambda)\lambda^{t-1}$, and the valuation in the T-stage finitely repeated game corresponds to the case where $\theta_t = \frac{1}{T}$ for $t \leq T$ and $\theta_t = 0$ for $t > T$.

Long games are characterized by values of θ_1 close to 0. At any stage t, θ_t is the weight of the current payoff whereas $\sum_{t'>t} \theta_t$ is the weight of the future stream of payoffs. Patient players are characterized by values of $\sup_t \frac{\theta_t}{\sum_{t'>t} \theta_t}$ close to 0.

We let E_θ denote the set of Nash equilibria of the game with payoffs evaluated according to the sequence $\theta = (\theta_t)_t$. The key to characterize the limit of E_θ when players are sufficiently patient (Proposition 5) is a characterization *via* the set of implementable distributions of the set of feasible payoff vectors (Corollary 1).

2.1. **Example: Coordination with nature.** We consider a two-player team game in which both players wish to coordinate with nature. $I = J = K = \{0, 1\}$ and μ is uniform. The common payoff function to both players is given by

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

and can be represented by the 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.

Consider the strategy of the forecaster that matches the state of nature at every stage. This strategy conveys no information to the agent about future states of nature. If the agent plays randomly, the average distribution D_0 of action triples up to any stage is in $\mathcal{Q}(1)$ and equals

	0	1
0	$\frac{1}{4}$	$\frac{1}{4}$
1	0	0

$i = 0$

	0	1
0	0	0
1	$\frac{1}{4}$	$\frac{1}{4}$

$i = 1$

The corresponding expected average payoff is $\frac{1}{2}$.

Now consider the strategy of the forecaster that matches nature at even stages, and plays the next state of nature at odd ones. At even stages, the agent is informed of the state of nature by the previous action of the forecaster, and thus can match it. At odd stages, the agent has no information on the state of nature, and we assume he plays randomly. The distribution of action triples at odd stages is

	0	1
0	$\frac{1}{8}$	$\frac{1}{8}$
1	$\frac{1}{8}$	$\frac{1}{8}$

$i = 0$

	0	1
0	$\frac{1}{8}$	$\frac{1}{8}$
1	$\frac{1}{8}$	$\frac{1}{8}$

$i = 1$

and at even stages is

	0	1
0	$\frac{1}{2}$	0
1	0	0

$i = 0$

	0	1
0	0	0
1	0	$\frac{1}{2}$

$i = 1$

The long-run average distribution D_1 is in $\mathcal{Q}(2)$ and equals

	0	1
0	$\frac{5}{16}$	$\frac{1}{16}$
1	$\frac{1}{16}$	$\frac{1}{16}$

$i = 0$

	0	1
0	$\frac{1}{16}$	$\frac{1}{16}$
1	$\frac{1}{16}$	$\frac{5}{16}$

$i = 1$

The corresponding expected payoff is $\frac{5}{8}$.

The previous strategies of the agent and the forecaster are defined over blocks of two stages. Notice that they require the forecaster to know the states of nature only two stages in advance.

Now we show an example of strategies that are defined over blocks of size three. Let $x[r] = (x_1[r], x_2[r], x_3[r])$ denote the sequence of states of nature in the r-th block. Similarly, $y[r]$ and $z[r]$ stand for the sequence of actions of the forecaster and the agent in the r-th block.

The agent's actions in each block are either $(0,0,0)$ or $(1,1,1)$. We design strategies such that in each block after the first, the agent's actions match the state of nature in a majority of stages.

The forecaster's actions in a block signal to the agent the majoritarian state of nature in the next block. This signaling is achieved by playing the majoritarian state of nature of the block $r+1$ in a singled-out stage of block r.

If the actions of the agent match the states of nature at all stages of block r ($x[r] = z[r]$), then the third stage of the block is the one singled out. If the actions of the agent match the sequence of states of nature in exactly two out of three stages, the mismatched stage is the one singled out.

The long-run average distribution D_2 is, by elementary computation,

	0	1
0	$\frac{17}{48}$	$\frac{1}{16}$
1	$\frac{1}{48}$	$\frac{1}{16}$

$i = 0$

	0	1
0	$\frac{1}{16}$	$\frac{1}{48}$
1	$\frac{1}{16}$	$\frac{17}{48}$

$i = 1$

The payoff obtained with D_2 is $\frac{17}{24}$, which is greater than the payoff $\frac{5}{8}$ obtained with D_1. Notice that these strategies require the forecaster to know the states of nature six stages in advance only, and do not rely on the agent observing the states of nature.

Natural questions that arise are what are the implementable distributions that maximize the expected payoff and what are the strategies that implement these distributions.

Using our characterization of the set of implementable distributions, we show in Section 3.1 that the following distribution D_3

	0	1
0	.41	.03
1	.03	.03

$i = 0$

	0	1
0	.03	.03
1	.03	.41

$i = 1$

(with corresponding payoff 0.82) is not implementable, while the distribution D_4

$i = 0$:

	0	1
0	$\frac{2}{5}$	$\frac{1}{30}$
1	$\frac{1}{30}$	$\frac{1}{30}$

$i = 1$:

	0	1
0	$\frac{1}{30}$	$\frac{1}{30}$
1	$\frac{1}{30}$	$\frac{2}{5}$

(with corresponding payoff $\frac{4}{5}$) is implementable.

Moreover, our analysis enables us to compute the (unique) implementable distribution that maximizes the corresponding payoff, and to construct strategies implementing this distribution. The unique payoff-maximizing implementable distribution is

$i = 0$:

	0	1
0	$\frac{x}{2}$	$\frac{1-x}{6}$
1	$\frac{1-x}{6}$	$\frac{1-x}{6}$

$i = 1$:

	0	1
0	$\frac{1-x}{6}$	$\frac{1-x}{6}$
1	$\frac{1-x}{6}$	$\frac{x}{2}$

with x satisfying $H(x) + (1 - x)\log_2 3 = 1$, where H is the entropy function.[2] The corresponding payoff is x, which is approximately 0.81.

3. The information constraint

The entropy of a discrete random variable $\mathbf{x}$ of law p with values in X measures its randomness, and also the quantity of information given by its observation. Its value is

$$H(\mathbf{x}) \quad = \quad -\sum\nolimits_{x\in X} p(\mathbf{x} = x)\log p(\mathbf{x} = x)$$

where the logarithm is taken in basis 2 and $0\log 0 = 0$ by convention.

If $\mathbf{x}, \mathbf{y}$ is a pair of discrete random variables of joint law p and values in $X \times Y$, the entropy of $\mathbf{x}$ given $\mathbf{y}$ measures the randomness of $\mathbf{x}$ given the knowledge of $\mathbf{y}$, or equivalently the quantity of information yielded by the observation of $\mathbf{x}$ to an agent who knows $\mathbf{y}$. Its value is

$$\begin{aligned} H(\mathbf{x}|\mathbf{y}) &= -\sum\nolimits_{x,y\in X\times Y} p(\mathbf{x} = x, \mathbf{y} = y)\log p(\mathbf{x} = x|\mathbf{y} = y) \\ &= -\sum\nolimits_{y\in Y} p(\mathbf{y} = y)\sum\nolimits_{x\in X} p(\mathbf{x} = x|\mathbf{y} = y)\log p(\mathbf{x} = x|\mathbf{y} = y) \end{aligned}$$

When we need to specify explicitly the probability Q of the probability space under which the random variables $\mathbf{x}$ and $\mathbf{y}$ are defined, we shall use the notations $H_Q(\mathbf{x})$ and $H_Q(\mathbf{x}|\mathbf{y})$.

[2]The entropy function H is given by $H(x) = -x\log_2 x - (1 - x)\log_2(1 - x)$ for $0 < x < 1$.

The main property of additivity of entropies states that

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

Let Q be a distribution over $I \times J \times K$. We say that Q fulfills the information constraint when

$$H_Q(\mathbf{i},\mathbf{j}|\mathbf{k}) \geq H_Q(\mathbf{i}) \tag{1}$$

Using the additivity of entropies, the information constraint can be rewritten as

$$H_Q(\mathbf{j}|\mathbf{i},\mathbf{k}) \geq H_Q(\mathbf{i}) - H_Q(\mathbf{i}|\mathbf{k}) \tag{2}$$

The left-hand side of this inequality can be interpreted as the amount of information received by the agent that observes the forecaster's action, $\mathbf{j}$, given the observation of the state of nature $\mathbf{i}$ and his own action $\mathbf{k}$. It is then an amount of information sent by the forecaster to the agent.

The right-hand term of (2) is the difference between the randomness of $\mathbf{i}$ and the randomness of $\mathbf{i}$ given the knowledge of $\mathbf{k}$. It is thus the reduction of uncertainty that $\mathbf{k}$ gives on $\mathbf{i}$, or the amount of information yielded by $\mathbf{i}$ on $\mathbf{k}$. We interpret it as an amount of information used by the agent on the state of nature.

Following this interpretation, the information constraint expresses the fact that the information used by the agent cannot exceed the information received from the forecaster.

Other expressions of the information constraint, such as $H_Q(\mathbf{j},\mathbf{k}|\mathbf{i}) \geq H_Q(\mathbf{k})$ or $H_Q(\mathbf{i},\mathbf{j},\mathbf{k}) \geq H_Q(\mathbf{i}) + H_Q(\mathbf{k})$ are also useful and yield other relevant interpretations in terms of information transmission and information banking.

Theorem 1. *Every implementable distribution fulfills the information constraint. In particular, every t-implementable distribution fulfills the information constraint.*

The next result shows a converse of the previous theorem when the horizon of the game is large.

Theorem 2. *Any distribution that fulfills the information constraint and has marginal μ on I is implementable.*

Together, Theorems 1 and 2 show that the information constraint fully characterizes the set of implementable distributions. It also characterizes the set of feasible payoff vectors as follows.

Corollary 1. *A payoff vector v is feasible if and only if there exists a distribution Q that fulfills the information constraint and has marginal μ on I such that $v = \mathbf{E}_Q g$.*

3.1. **Example continued.** Take up the game of Section 2.1, and for $0 < y < 1$ consider the distribution D_y of action triples

$i = 0$:

	0	1
0	$\frac{y}{2}$	$\frac{1-y}{6}$
1	$\frac{1-y}{6}$	$\frac{1-y}{6}$

$i = 1$:

	0	1
0	$\frac{1-y}{6}$	$\frac{1-y}{6}$
1	$\frac{1-y}{6}$	$\frac{y}{2}$

Conditional on the value 0 or 1 of k, the distributions of (i, j) are

$k = 0$:

	0	1
0	y	$\frac{1-y}{3}$
1	$\frac{1-y}{3}$	$\frac{1-y}{3}$

$k = 1$:

	0	1
0	$\frac{1-y}{3}$	$\frac{1-y}{3}$
1	$\frac{1-y}{3}$	y

Hence, $H_{D_y}(\mathbf{i}, \mathbf{j}|\mathbf{k}) = H(y) + (1 - y) \log 3$.

The distribution D_3 of the example corresponds to $y = 0.82$, and $H_{D_3}(\mathbf{i}, \mathbf{j}|\mathbf{k}) \sim 0.97 < 1$. Hence D_3 does not satisfy the information constraint, and is not implementable by Theorem 1.

The distribution D_4 corresponds to $y = 0.8$, and $H_{D_4}(\mathbf{i}, \mathbf{j}|\mathbf{k}) \sim 1.04 > H_{D_4}(\mathbf{i}) = 1$. Hence D_4 fulfills the information constraint, and is implementable by Theorem 2. Now we construct strategies that implement D_4.

The value of the entropy function $H(x)$, $0 < x < 1$, is a good approximation of $\frac{1}{n} \log \binom{n}{xn}$ for large n. Hence, $H(.8) + .2 \log 3 > H(\mathbf{i}) = 1$ implies $\binom{5n}{n} 3^n > 2^{5n}$ for n sufficiently large and precise computation shows that this holds for all $n \geq 17$.

The left-hand side of this inequality, $\binom{5n}{n} 3^n$, is the number of elements in the set $S(n)$ of all sequences of length $5n$ of the digits $\{1, 2, 3, 4\}$ where exactly $4n$ of the entries are 1. Define a 1-1 map τ from $K^{5n} = \{0, 1\}^{5n}$ into $S(n)$. Choose a positive integer m and consider m consecutive blocks of length $5n$ each. Let $x[r]$ denote the sequence of states of nature in the r-th block, and $y[r]$ and $z[r]$ stand for the sequence of actions of the forecaster and the agent in the r-th block. Define inductively $z[m] = x[m] = y[m]$. Assume that $z[r+1]$ is defined for $1 < r < m$ and define $(y[r], z[r])$ as a function of $x[r]$ and $z[r+1]$ as follows. $y_t[r] = z_t[r] = x_t[r]$ if the t-th coordinate of $\tau(z[r+1])$ is 1, $y_t[r] \neq z_t[r] = 0$ if the t-th coordinate of $\tau(z[r+1])$ is 2, $y_t[r] \neq z_t[r] = 1$ if the t-th coordinate of $\tau(z[r+1])$ is 3, and $y_t[r] = z_t[r] \neq x_t[r]$ if the t-th coordinate of $\tau(z[r+1])$ is 4.

If we set in addition $y[1] = z[2]$ and $z[1]$ arbitrarily, we realize that for $1 \leq r \leq m$ the sequence $z[r]$ is a deterministic function of $(x[r-1], y[r-1], z[r-1])$, and the sequence $y[r]$ is a deterministic function of $x[1], \ldots, x[m]$. Therefore the random (as they depend on the sequence $x[1], \ldots, x[m]$ of states of nature) sequences $z[1], \ldots, z[m]$ and $y[1], \ldots, y[m]$ are implementable by pure strategies of the forecaster and the agent.

Note that for $1 < r < m$ we have $|\{1 \leq t \leq 5n : x_t[r] = y_t[r] = z_t[r]\}| = 4n$, and therefore $4n(m + 1/2) \geq |\{1 \leq t \leq 5nm : x_t = y_t = z_t\}| \geq 4n(m-1)$, and thus for sufficiently large m, the average payoff, $\frac{1}{5nm}|\{1 \leq t \leq 5nm : x_t = y_t = z_t\}|$, is close to 0.8. In addition, $Q(i = 0 = j \neq k)$, $Q(i = 0 = k \neq j)$, etc., are close to $1/30$.

4. Mixed and correlated strategies

The set of t-implementable distributions in pure or mixed strategies is not convex whenever $|K| \geq 2$, and its convex hull is the set of t-implementable distributions in correlated strategies. However, the set of implementable distributions is closed and convex (Remark 2) and therefore also contains the t-implementable distributions in correlated strategies and their limits. Thus, t-implementable distributions using correlated strategies fulfill the information constraint, and distributions that are implementable using correlated strategies are also implementable using pure strategies.

The convexity and the closedness of the set of implementable distributions also follow from Theorems 1 and 2 and from the following lemma, which is useful in the sequel.

Lemma 1. *Given finite sets X and Y, the function $Q \mapsto H_Q(\mathbf{y}|\mathbf{x})$ is concave on the set of probability measures on $X \times Y$.*

Proof. Follows from the concavity of entropy. Let $\bar{Q} = \sum_m \lambda_m Q_m$ be a finite convex combination of distributions over $X \times Y$. Consider a triple of random variables $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}$ such that $P(\boldsymbol{\gamma} = m) = \lambda_m$, and $\boldsymbol{\alpha}, \boldsymbol{\beta}$ has law Q_m conditional on $\boldsymbol{\gamma} = m$. Then

$$\begin{aligned} H_{\bar{Q}}(\mathbf{y}|\mathbf{x}) &= H(\boldsymbol{\beta}|\boldsymbol{\alpha}) \\ &\geq H(\boldsymbol{\beta}|\boldsymbol{\alpha}, \boldsymbol{\gamma}) \\ &= \sum_m \lambda_m H_{Q_m}(\mathbf{y}|\mathbf{x}) \end{aligned}$$

□

5. Proof of Theorem 1

For any pure strategies σ and τ and any stage t,

$$\begin{aligned} \sum_{t'=1}^{t} H_{P^{t'}_{\mu,\sigma\tau}}(\mathbf{i}, \mathbf{j}|\mathbf{k}) &= \sum_{t'=1}^{t} H(\mathbf{i}_t, \mathbf{j}_t|\mathbf{k}_t) \\ &= \sum_{t'=1}^{t} H(\mathbf{i}_t, \mathbf{j}_t, \mathbf{k}_t|\mathbf{k}_t) \\ &\geq \sum_{t'=1}^{t} H(\mathbf{i}_t, \mathbf{j}_t, \mathbf{k}_t|\mathbf{i}_1, \mathbf{j}_1, \mathbf{k}_1, \ldots \mathbf{i}_{t-1}, \mathbf{j}_{t-1}, \mathbf{k}_{t-1}) \\ &= H(\mathbf{i}_1, \mathbf{j}_1, \mathbf{k}_1, \ldots \mathbf{i}_t, \mathbf{j}_t, \mathbf{k}_t) \\ &\geq H(\mathbf{i}_1, \ldots \mathbf{i}_t) = tH(\mu) \end{aligned}$$

where the first inequality follows from the fact that $\mathbf{k}_t$ is a deterministic function of the past, $\mathbf{k}_t = \tau_t \circ (\mathbf{i}_1, \mathbf{j}_1, \mathbf{k}_1, \ldots \mathbf{i}_{t-1}, \mathbf{j}_{t-1}, \mathbf{k}_{t-1})$; the second inequality follows from the fact that $(\mathbf{i}_1, \ldots \mathbf{i}_t)$ is a function of $(\mathbf{i}_1, \mathbf{j}_1, \mathbf{k}_1, \ldots \mathbf{i}_{t-1}, \mathbf{j}_{t-1}, \mathbf{k}_{t-1})$. By Lemma 1,

$$H_{Q^t_{\mu,\sigma\tau}}(\mathbf{i}, \mathbf{j}|\mathbf{k}) \geq \frac{1}{t} \sum_{t'=1}^{t} H_{P^{t'}_{\mu,\sigma\tau}}(\mathbf{i}, \mathbf{j}|\mathbf{k})$$

Hence t-implementable distributions fulfill the information constraint. The result follows from the fact that the maps $Q \mapsto H_Q(\mathbf{i}, \mathbf{j}|\mathbf{k})$ and $Q \mapsto H_Q(\mathbf{i})$ are continuous and therefore the set of distributions Q that obey the information constraint is closed.

6. A Combinatorial View of the Information Constraint

The entropy of the uniform distribution on a finite set A is called the *combinatorial entropy* of A, and is denoted $H(A)$; it is the log to the base two of the number of elements of A. An inequality between the combinatorial entropy of two finite sets is thus an inequality between the number of elements of the sets, and vice versa. In this section we derive a fundamental inequality, (5), that enables us to translate the information constraint, which is an inequality of entropies, to an approximate inequality of combinatorial entropies.

Comparison of sizes of sets, and in particular inequalities derived from (5), play a fundamental role in the proof of Theorem 2. Indeed, if a set A is larger than another set B, then, by pointing to an element in A, the forecaster can signal to the agent which element of B to choose. In addition, the present section introduces notations used in the proof of Theorem 2.

For a finite sequence $\beta = (b_1, \ldots, b_n)$ over a finite alphabet B, $\rho(\beta) \in \Delta(B)$ denotes its empirical distribution ($\rho(\beta)[b] = \frac{1}{n}\sum_{t=1}^{n} \mathbb{I}_{b_t = b}$), and for $\nu \in \Delta(B)$, $T_n(\nu)$ denotes the n-type set of ν, $\{\beta \in B^n : \rho(\beta) = \nu\}$. The set of types is $\mathbb{T}_n(B) = \{\mu \in \Delta(B), T_n(\mu) \neq \emptyset\}$.

The entropy function provides a good approximation of the combinatorial entropy of a nonempty typical set $T_n(\nu)$ (see, for instance, Cover and Thomas [4, Theorem 12.1.3, page 282]). If $\nu \in \Delta(B)$ and $T_n(\nu) \neq \emptyset$, then

$$\frac{2^{nH(\nu)}}{(n+1)^{|B|}} \leq |T_n(\nu)| \leq 2^{nH(\nu)} \tag{3}$$

If $\nu \in \Delta(A \times B)$ has marginals $\nu_A \in \Delta(A)$ on A and ν_B on B, and $\alpha = (a_1, \ldots, a_n) \in T_n(\nu_A)$, the n-type set of ν conditional on α, denoted $T_n(\nu|\alpha)$, is the set of elements in $T_n(\nu)$ whose A-coordinates coincide with α; equivalently, $T_n(\nu|\alpha)$ is identified with $\{\beta \in T_n(\nu_B) : (\alpha, \beta) \in T_n(\nu)\}$, where for $\beta = (b_1, \ldots, b_n)$, $(\alpha, \beta) := (a_1, b_1, \ldots, a_n, b_n)$. The number of elements of $T_n(\nu|\alpha)$ is independent of $\alpha \in T_n(\nu_A)$ and $|T_n(\nu|\alpha)|\ |T_n(\nu_\alpha)| = |T_n(\nu)|$; equivalently, $H(T_n(\nu|\alpha)) + H(T_n(\nu_\alpha)) = H(T_n(\nu))$. The last equality is termed the *additivity of combinatorial entropies.*

The conditional entropy $H_\nu(\mathbf{b}|\mathbf{a}) = H(\nu) - H(\nu_A)$ (where $\nu \in \mathbb{T}_n(A \times B)$) provides a good approximation of the combinatorial entropy of the typical set $T_n(\nu|\alpha)$. If $\nu \in \mathbb{T}_n(A \times B)$ and $\alpha \in T_n(\nu_A)$, using the additivity of combinatorial entropies, inequality (3) implies that $(n+1)^{|A|} 2^{n(H(\nu) - H(\nu_A))} \geq |T_n(\nu|\alpha)| \geq \frac{2^{n(H(\nu)-H(\nu_A))}}{(n+1)^{|A \times B|}}$. Using Lemma 2 (in the Appendix) we have the tighter inequality

$$2^{n(H(\nu)-H(\nu_A))} \geq |T_n(\nu|\alpha)| \geq \frac{2^{n(H(\nu)-H(\nu_A))}}{(\frac{n}{|A|}+1)^{|A \times B|}} \tag{4}$$

Let $Q \in \mathbb{T}_n(I \times J \times K)$, $\varepsilon = \varepsilon_Q = H_Q(\mathbf{i}, \mathbf{j}|\mathbf{k}) - H_Q(\mathbf{i})$, and set $q(n) = n^{|I\times J\times K|}$.

The inequalities (3) and (4) imply[3] that for $x \in T_n(Q_I)$ we have[4]

$$q(n)2^{\varepsilon n} \geq \frac{|T_n(Q|x)|}{|T_n(Q_K)|} \geq \frac{2^{\varepsilon n}}{q(n)} \tag{5}$$

Now we apply the inequality (5) to rewrite the information constraint as an approximate inequality of combinatorial entropies. Note that if $Q \in \Delta(I \times J \times K)$ then Q need not be in $\mathbb{T}_n(I \times J \times K)$. However, for every $Q \in \Delta(I \times J \times K)$ there is $R \in \mathbb{T}_n(I \times J \times K)$ with $\|Q - R\| < 2|I \times J \times K|/n$. Let $Q \mapsto Q^n$ be a map from $\Delta(I \times J \times K)$ into $\mathbb{T}_n(I \times J \times K)$ with $\|Q - Q^n\| \leq C/n$ where $C \geq 2|I \times J \times K|$ is a given constant.

Proposition 1 (The combinatorial information constraint)**.** *The following conditions on a distribution $Q \in \Delta(I\times J\times K)$ are equivalent:*

1) *Q fulfills the information constraint;*
2) *$\exists B \in \mathbb{R}$ such that for $x \in T_n(Q_I^n)$ we have*

$$H(T_n(Q^n|x)) \geq H(T_n(Q_K^n)) + B\log_2 n \tag{6}$$

3) *$\exists B \in \mathbb{R}$ such that for $(x, z) \in T_n(Q_{I\times K}^n)$ we have*

$$H(T_n(Q^n|x, z)) \geq H(T_n(Q_I^n)) - H(T_n(Q_{I\times K}^n|z)) + B\log_2 n$$

Proof. We first prove that 1) and 2) are equivalent. Assume that Q fulfills the information constraint. Thus $H_Q(\mathbf{i}, \mathbf{j}|\mathbf{k}) \geq H_Q(\mathbf{k})$. The inequality $\|Q - Q^n\| \leq C/n \leq 1/2$ implies that $|H_Q(\mathbf{k}) - H_{Q^n}(\mathbf{k})| \leq \frac{2C\log n}{n}$ (for $n \geq |I \times J \times K|$), and $|H_Q(\mathbf{j}, \mathbf{k}|\mathbf{i}) - H_{Q^n}(\mathbf{j}, \mathbf{k}|\mathbf{i})| \leq \frac{4C\log n}{n}$. Therefore, $H_Q(\mathbf{j}, \mathbf{k}|\mathbf{i}) \geq H_Q(\mathbf{k})$ implies that $H_{Q^n}(\mathbf{j}, \mathbf{k}|\mathbf{i}) \geq H_{Q^n}(\mathbf{k}) - \frac{6C\log n}{n}$ and thus, by using inequality (5), for $x \in T_n(Q_I^n)$ we have

$$H(T_n(Q^n|x)) \geq H(T_n(Q_K^n)) - 6C\log_2 n - \log_2 q(n),$$

and thus (6) holds with $B = -6C - |I \times J \times K|$.

If Q violates the information constraint then $H_Q(\mathbf{i}, \mathbf{j}|\mathbf{k}) - H_Q(k) < 0$ and therefore there is $\delta > 0$ such that for sufficiently large n we have

[3]Assuming $n, |I| > 1$.

[4]In fact, the right-hand inequality holds for $q(n) \geq \left(\frac{n+|I|}{|I\times J\times K|}\right)^{|I\times J\times K|}$ and the left-hand inequality holds for $q(n) \geq (n+1)^{|K|}$ and thus in particular the inequalities hold for $q(n) \geq n^{|I\times J\times K|}$.

$H_{Q^n}(\mathbf{i},\mathbf{j}|\mathbf{k}) \leq H_{Q^n}(\mathbf{k}) - \delta$ and thus, by using inequality (5), for $x \in T_n(Q_I^n)$ and n sufficiently large we have $H(T_n(Q^n|x)) \leq H(T_n(Q_K^n)) - \delta n + \log_2 q(n)$ and therefore there is no $B \in \mathbb{R}$ for which inequality (6) holds for all n.

¿From the additivity of entropies, inequality (6) is equivalent to the inequality in condition 3). □

Now we derive consequences of inequality (5) that are used in the proof of Theorem 2. Let $Q \in \mathbb{T}_n(I \times J \times K)$ and $\varepsilon = \varepsilon_Q = H_Q(\mathbf{i},\mathbf{j}|\mathbf{k}) - H_Q(\mathbf{i})$.

Inequality (5) implies that if $\varepsilon > 0$ and n is sufficiently large so that $2^{\varepsilon n} > q(n)$, then

$$|T_n(Q|x)| > |T_n(Q_K)| \quad \text{for } x \in T_n(Q_I); \tag{7}$$

equivalently, $|T_n(Q|x) \cap (J^n \times T_n(Q_K))| > |T_n(Q_K)|$. As shown in Section 7.1, for sufficiently large values[5] of n the set $T_n(Q_K)$ is one of many subsets A of $T_n(Q_K)$ with

$$\forall x \in T_n(Q_I) \quad |T_n(Q|x) \cap (J^n \times A)| > |A| \tag{8}$$

The two terms appearing in inequality (8) have relevant interpretations for the online communication problem between the forecaster and the agent. $T_n(Q|x) \cap (J^n \times A)$ is the set of points $(y,z) \in (J \times K)^n$ with $z \in A$ and $(x,y,z) \in T_n(Q)$. Given $x \in T_n(Q_I)$, $T_n(Q|x) \cap (J^n \times A)$ can be interpreted as a set of messages $(y,z) \in (J \times K)^n$ subject to $z \in A$ and $(x,y,z) \in T_n(Q)$. The set A can be interpreted as a target set of action strings of length n of the agent.

Therefore, given $x \in T_n(Q_I)$, there is a 1-1 map η from A into $T_n(Q|x) \cap (J^n \times A)$ (equivalently, a map η^{-1} from $T_n(Q|x) \cap (J^n \times A)$ onto A), and a point in $\eta(A)$ can be interpreted as a message for the selection of an element of A.

7. Proof of Theorem 2

Fix $Q' \in \Delta(I \times J \times K)$ that satisfies the conditions of Theorem 2. By Remark 2 the set of implementable distributions is closed and by Remark 1 it contains $\mathcal{Q}(t)$. Therefore it suffices to prove that for every

[5] More precisely, $2^{\varepsilon n} > q(n)$ is sufficient.

$\varepsilon > 0$ there exists a strategy profile (σ, τ) and $t = t(\varepsilon, \sigma, \tau)$ such that

$$\|Q^t_{\mu,\sigma,\tau} - Q'\| < \varepsilon$$

Fix $\varepsilon > 0$. By Lemma 4 (in the Appendix) there is $Q \in \Delta(I \times J \times K)$ such that $T_n(Q) \neq \emptyset$ and

$$\|Q - Q'\| < 7\varepsilon \text{ and thus in particular } \|Q_I - \mu\| < 7\varepsilon \tag{9}$$

and

$$H_Q(\mathbf{i}, \mathbf{j}|\mathbf{k}) - H_Q(\mathbf{i}) \geq \varepsilon \tag{10}$$

Let $n > |I|^2/\varepsilon^3$ be a fixed sufficiently large integer; e.g., $2^{\varepsilon n} > n^{|I \times J \times K|}(2 + 2n \ln |I|)$ suffices. The proof associates (given n and Q specified above) to each sequence $x = (x_1, x_2, \ldots)$ of states of nature two sequences $y = (y_1, y_2, \ldots)$ and $z = (z_1, z_2, \ldots)$, of forecaster's actions and agent's actions respectively.

The stages are partitioned into consecutive blocks, the first block is of length $n_1 \geq n \log |K|$ $(\geq nH(Q_K))$ where n_1 is a multiple of n, and the other blocks are of length n. Set $n_0 = 0$ and $n_r = n_1 + (r-1)n$ for $r \geq 1$. We construct the sequences y and z as follows. Denote by $(x[r], y[r], z[r]) = (x_t[r], y_t[r], z_t[r])_t$ (where $1 \leq t \leq n_1$ for $r = 1$ and $1 \leq t \leq n$ for $r > 1$) the action triples in the r-th block (i.e., in stages $n_{r-1} < s \leq n_r$).

We choose for every $x \in I^{\bar{n}}$, where $\bar{n}$ is either n_1 or n, an element $\tilde{x} \in T_{\bar{n}}(Q_I)$ that minimizes the number of coordinates $1 \leq t \leq \bar{n}$ with $x_t \neq \tilde{x}_t$ subject to $\rho(\tilde{x}_1, \ldots, \tilde{x}_n) = Q_I$.

Let $A \subset T_n(Q_K)$ satisfy (8). Then, for every $\tilde{x} \in T_n(Q_I)$ there is a 1-1 map $f_{\tilde{x}}$ from A into the set of all pairs $(y, z) \in T_n(Q|\tilde{x}) \cap (J^n \times A)$. As $|A| \leq |T_n(Q_K)| \leq 2^{n_1}$, there is a 1-1 map f_0 from A into J^{n_1}.

Thus, by backward induction, starting at $r = m$, there are sequences $y[r], z[r]$, $1 < r \leq m$, such that 1) $(\tilde{x}[r], y[r], z[r]) \in T_n(Q)$ for $1 < r \leq m$, 2) $z[r] \in A$ for $1 < r \leq m$, 3) $f_{\tilde{x}[r]}(z[r+1]) = (y[r], z[r])$ for $1 < r < m$, and 4) $y[1] = f_0(z[2])$ and $z[1]$ is an arbitrary element of K^{n_1}.

Conditions 3) and 4) imply the existence of a pure strategy pair (σ, τ) of the forecaster and the agent that generate the actions sequences $y = y[1], y[2], \ldots, y[m]$ and $z = z[1], z[2], \ldots, z[m]$ as a function of the sequence x of states of nature. Explicitly, define a pure strategy pair (σ,τ) of the forecaster and the agent as follows. The agent plays in the first block the sequence $z[1]$; in the second block he plays the sequence $f_0^{-1}(y[1])$ (conditional on $y[1]\in f_0(A)$); in the $(r+1)$-th block, $2\le r<m$, the agent plays the sequence $f^{-1}_{\tilde{x}[r]}(y[r],z[r])$. The forecaster plays in the m-th block a sequence $y[m]\in T_n(Q|\tilde{x}[m],z[m])$; in the r-th block, $1<r<m$, the forecaster plays the J^n-th component of the sequence $f_{\tilde{x}[r]}(z[r+1])$, and in the first block the forecaster plays the sequence $f_0(z[2])$.

Let $Q^{n_m}_\mu$ (respectively, $\tilde{Q}^{n_m}_\mu$) be the expectation of the empirical distribution of the triples (x_t,y_t,z_t) (respectively, $(\tilde{x}_t,y_t,z_t)$) where $1\le t\le n_m$, where $x=(x_1,x_2,\ldots)$ is an i.i.d. sequence with $x_t\sim\mu$. Now we prove that condition 1) above implies that for sufficiently large m we have $\|Q^{n_m}_\mu-Q'\|<26\varepsilon$.

For $1<r\le m$, the norm distance between the empirical distribution of $(\tilde{x}[r],y[r],z[r])$ and $(x[r],y[r],z[r])$ is (obviously ≤ 2 and) bounded by $\frac{2}{n}\sum_{s=1}^n \mathbb{I}_{\tilde{x}_s[r]\neq x_s[r]}$. By Corollary 2 (in the Appendix) the probability that $\sum_{s=1}^n \mathbb{I}_{\tilde{x}_s[r]\neq x_s[r]}\ge 8\varepsilon n$ is $\le |I|^2/(n\varepsilon^2)$, which is $\le\varepsilon$ (as $n>|I|^2/\varepsilon^3$). Therefore $\|Q^{n_m}_\mu-\tilde{Q}^{n_m}_\mu\|<(16+2)\varepsilon=18\varepsilon$ for sufficiently large m. If $\rho(\tilde{x}[r],y[r],z[r])=Q$ for every $1<r\le m$, then for sufficiently large m we have $\|\tilde{Q}^{n_m}_\mu-Q\|<\varepsilon$ and therefore, by the triangle inequality, we conclude that for sufficiently large m we have $\|Q^{n_m}_\mu-Q'\|<(18+1+7)\varepsilon=26\varepsilon$.

This ends the proof of Theorem 2.

We end this section with a discussion on the construction of the strategies of the proof.

7.1. **Choice of the action plan.** The proof of Theorem 2 relies on a set A of action plans for the forecaster (e.g., $T_n(Q_K)$) that satisfies inequality (8).

Let $g(n)=2+2n\ln|I|$, and $1\le f_n\le\frac{|T_n(Q_{I\times K})|}{|T_n(Q_I)|}$. Then $\frac{f_n(g(n)-1)^2}{2g(n)}>nH(Q_I)\ln 2$ and, by Lemma 3 (in the Appendix), there is a subset A of $T_n(Q_K)$ with

$$\frac{|T_n(Q_K)|.|T_n(Q_I)|}{|T_n(Q_{I\times K})|}f_ng(n)\le|A|<1+\frac{|T_n(Q_K)|.|T_n(Q_I)|}{|T_n(Q_{I\times K})|}f_ng(n)$$

such that[6] for every $x \in T_n(Q_I)$ there are at least f_n elements $z \in A$ such that $(x,z) \in T_n(Q_{I\times K})$, i.e., $|T_n(Q_{I\times K}|x) \cap A| \geq f_n$; thus $|T_n(Q|x) \cap (J^n \times A)| \geq f_n |T_n(Q)|/|T_n(Q_{I\times K})|$, which by inequality (5) is no less than $f_n \frac{|T_n(Q_K)|}{|T_n(Q_{I\times K}|x)|} \frac{2^{\varepsilon n}}{q(n)}$. Therefore, if $\varepsilon > 0$ and n is sufficiently large so that $2^{\varepsilon n} > q(n)g(n)$, then A obeys (8).

The combinatorial entropy of A is

$$H(A) = \log_2 f_n + H(T_n(Q_K)) + H(T_n(Q_I)) - H(T_n(Q_{I\times K})) + O(\log n)$$

Hence, the construction of the proof can be implemented using sets of action plans with combinatorial entropies that equal $cn + O(\log n)$ where c is any value between $H_Q(\mathbf{k}) + H_Q(\mathbf{i}) - H_Q(\mathbf{i},\mathbf{k})$ and $H_Q(\mathbf{k})$.

It is worthwhile noting that the implementable strategies described above require a perfect forecast of $n + n_1$ stages when $f_n = 1$ and $n_1 - n + nm$ stages when $f_n = |T_n(Q_{I\times K})|/|T_n(Q_I)|$. Moreover, when $f_n = 1$ the first block can be replaced with a block of size n and thus requires a forecast of $2n$ stages. The strategy corresponding to the minimal value $f_n = 1$ thus requires a smaller forecast, and in addition a smaller dictionary (a function f_x from $T_n(Q|x) \cap (J^n \times A)$ onto A) for specifying the agent's action in the forthcoming block as a function of the sequence of action triples in the last block. As f_n increases the size of the dictionary increases. However, at the other extreme, when $f_n = |T_n(Q_{I\times K})|/|T_n(Q_I)|$ the set A is simply $T_n(Q_K)$ and thus the construction of the set A is made explicit.

An adequate choice of size for the set A may also prove useful in the study of the rate of convergence of the set of t-implementable distributions or the set of distributions implementable with finite forecasts to the full set of implementable distributions.

8. Payoffs and equilibria

In this section we show how for sufficiently long games the information constraint yields characterizations of 1) the set of feasible payoff vectors, 2) the best payoff a team can achieve, and 3) the set of equilibrium payoffs when players are sufficiently patient.

8.1. **Feasible payoffs.** We show that the set F is a good approximation for the set of feasible payoffs of the long games.

[6]Actually, for $f_n > 1$ smaller functions $g(n)$ suffice.

The approximation applies to general interstage-time-dependent discounting, thus in particular to finite repetitions of the game as well as to interstage-time-independent discounting.

For a nonincreasing sequence $\theta = (\theta_t)_t$ of nonnegative numbers summing to 1, let $Q^\theta_{\mu,\sigma,\tau} = \sum_{t=1}^\infty \theta_t P^t_{\mu,\sigma,\tau}$ be the θ-weighted average distribution of action triples. The expectation of the θ-weighted sum of the stage payoffs is

$$\mathbf{E}_{P_{\mu,\sigma,\tau}} \sum\nolimits_{t=1}^\infty \theta_t g(i_t, j_t, k_t) = \mathbf{E}_{Q^\theta_{\mu,\sigma,\tau}} g(i,j,k)$$

Since $\sum_{t=1}^\infty \theta_t P^t_{\mu,\sigma,\tau} = \sum_{t=1}^\infty (\theta_t - \theta_{t+1}) t Q^t_{\mu,\sigma,\tau}$ and $\sum_{t=1}^\infty (\theta_t - \theta_{t+1}) t = 1$, $Q^\theta_{\mu,\sigma,\tau}$ is a convex combination of the family of distributions $\{Q^t_{\mu,\sigma,\tau}\}_{t \geq 1}$. Thus it obeys the information constraint and has marginal μ on I. Therefore, if Σ_f and Σ_a denote the sets of strategies of the forecaster and the agent respectively we have

Proposition 2. *For every nonincreasing sequence $\theta = (\theta_t)$ summing to 1, the set of θ-weighted feasible payoff vectors*

$$F_\theta = \{\mathbf{E}_{P_{\mu,\sigma,\tau}} \sum\nolimits_{t=1}^\infty \theta_t g(i_t, j_t, k_t) : (\sigma, \tau) \in \Sigma_f \times \Sigma_a\}$$

is a subset of F.

Recall that F_t is the set of feasible payoff vectors of the t-stage repeated game, and let F_λ denote the set of feasible payoff vectors of the λ-discounted game. Special cases of the previous proposition are

$$F_t \subseteq F, \quad F_\lambda \subseteq F$$

On the other hand, if Q is implementable there exists a strategy pair (σ, τ) such that for every $\varepsilon > 0$ there exists N so that $\|Q^n_{\mu,\sigma,\tau} - Q\| < \varepsilon$ for every $n \geq N$. Therefore, if $\theta = (\theta_t)_t$ is a nonincreasing sequence summing to 1, then $\|Q^\theta_{\mu,\sigma,\tau} - Q\| = \|\sum_{t=1}^\infty (\theta_t - \theta_{t+1}) t (Q^t_{\mu,\sigma,\tau} - Q)\| \leq 2N\theta_1 + \varepsilon$ (by the triangle inequality and using $0 \leq \sum_{t=1}^N (\theta_t - \theta_{t+1}) t = \sum_{t=1}^N \theta_t - N\theta_{N+1} \leq N\theta_1$), and therefore for sufficiently small θ_1 the distribution $Q^\theta_{\mu,\sigma,\tau}$ is within 2ε of the distribution Q. Therefore,

Proposition 3. *F_θ converges in the Hausdorff metric to F as θ_1 goes to 0. In particular, $F_\lambda \to F$ as $\lambda \to 1$ and $F_t \to F$ as $T \to \infty$.*

8.2. **Team games.** Team games, in which players' preferences are identical, form an adequate setup for the study of inefficiencies due to information asymmetries and communication costs. As shown for instance by Marschak and Radner [10] and by Arrow [1], it is helpful to describe a firm as a team when one focuses on the question of information transmission among its members.

In team games, our model allows us to measure the inefficiencies arising from the need to send and process information. As a benchmark, consider the situation in which both the agent and the forecaster have complete information about the states of nature. In this case, it is possible for both players to choose optimally an action pair at each stage given the current state of nature. The corresponding expected payoff is the best feasible under complete information.

In the game we analyze, both players can use a myopic behavior that seeks to maximize at each stage the payoff of the current stage. In this case, the forecaster's actions are uninformative about the future of the process, and so the agent's belief on the current state of nature is his prior belief. Such behavior rules are not optimal in general. Indeed, in most games the team can secure a better payoff if the forecaster deviates from a myopic maximization rule in order to convey information to the agent. For instance, in the game of Section 2.1, myopic behaviors cannot secure more than $\frac{1}{2}$ whereas the non-myopic strategies that implement the distribution D_1 guarantee $\frac{5}{8}$.

As we see, a good joint behavior for the forecaster must seek to communicate maximal information with the slightest deviation from a stage payoff maximization rule.

The problems of finding optimal strategies for the team and of computing the maximal payoffs that can be achieved in a T-stage game or in a λ-discounted game are difficult ones. Yet these problems are made particularly simple by considering long repeated games and using an approach through the information constraint.

Let v_θ be the maximum payoff for the team when the discount factors are given by $\theta = (\theta_t)_t$, i.e., $v_\theta = \max\{x : (x, x) \in F_\theta\}$, and let $v = \max\{x : (x, x) \in F\}$. Proposition 2 implies that $v_\theta \leq v$ and Proposition 3 implies

Proposition 4. *v_θ goes to v when θ_1 goes to 0. In particular, the maximal payoffs that can be achieved in T-stage games and in λ-discounted games go to v as λ goes to 1 and T goes to ∞.*

We are thus able to characterize the best feasible payoff to the team and its degree of inefficiency compared to the full information case, and to construct strategies that achieve this maximal payoff.

Our model thus applies to the study of the impact of communication costs on team games, which is an important question in the theory of organizations (see van Zandt [15] for a survey). The information constraint does not depend on the specification of payoffs to the team. Since it characterizes the set of implementable distributions, it allows us to write the maximization problem faced by the team in a simple and compact way for any payoff specification.

8.2.1. *Example: bounded communication.* Consider the following team game where the state of nature specifies the matrix, the forecaster is the row player, and the agent is the column player. The sequence of states of nature follows an i.i.d. and uniform process. Payoffs are given by

	1	2	3
1	1	0	0
2	1	0	0

1

	1	2	3
1	0	1	0
2	0	1	0

2

	1	2	3
1	0	0	1
2	0	0	1

3

Thus the payoff to the team depends only on the state of nature and on the agent's action. The forecaster has incentives to send the maximal information to the agent.

Messages do not affect payoffs, but given μ and the size of J, it is not possible to send all the relevant information to the agent. A choice then needs to be made about what information is to be sent, such that only the most important information reaches the agent. This models cheap but bounded communication capacities. For instance, such a situation occurs when either the forecaster or the agent is only capable of processing one binary message per stage.

We illustrate the use of the information constraint in computing the maximal payoff that the team of the forecaster and the agent can obtain. Let Q be an implementable distribution that maximizes the common payoff; i.e., the distribution Q maximizes the probability $Q(i{=}k)$ subject to $H_Q(\mathbf{i},\mathbf{j}|\mathbf{k}) \geq H_Q(\mathbf{i})$. Obviously, by replacing the distribution Q with the product distribution of the uniform distribution U_J on J and the marginal distribution $Q_{I\times K}$ we obtain a distribution $\hat{Q}$ with

$\hat{Q}(i{=}k) = Q(i{=}k)$ and $H_{\hat{Q}}(\mathbf{i},\mathbf{j}|\mathbf{k}) \geq H_{\hat{Q}}(\mathbf{i})$. Therefore we can assume w.l.o.g. that Q is the product distribution $U_J \otimes Q_{I\times K}$, and thus the information constraint is

$$1 + H_Q(\mathbf{i}|\mathbf{k}) \geq H_Q(\mathbf{i}) = \log 3$$

i.e.,

$$H_Q(\mathbf{i}|\mathbf{k}) \geq \log \frac{3}{2}$$

Note that the common payoff depends only on the values of $Q(i{=}k{=}1)$, $Q(i{=}k{=}2)$, and $Q(i{=}k{=}3)$, and equals their sum. By symmetry and by concavity of the map $Q \mapsto H_Q(\mathbf{i}|\mathbf{k})$ (Lemma 1) we can assume w.l.o.g. that $Q(i{=}k{=}1) = Q(i{=}k{=}2) = Q(i{=}k{=}3) = x$. Given this inequality, the conditional entropy $H_Q(\mathbf{i}|\mathbf{k})$ is maximized when $Q(\mathbf{i}{=}i|\mathbf{k}{=}k) = Q(\mathbf{i}{=}i'|\mathbf{k}{=}k)$ for $i \neq i' \neq k \neq i$, hence when $Q(\mathbf{i}{=}i|\mathbf{k}{=}k) = (\frac{1}{3} - x)/2$ for $i \neq k$. In this case $Q_{I\times K}$ is given by

$$Q_{I\times K}(i,k) = \begin{cases} x & \text{if } i = k \\ \frac{\frac{1}{3}-x}{2} & \text{if } i \neq k \end{cases}$$

and thus $H_Q(\mathbf{i}|\mathbf{k}) = H(3x) + 1 - 3x$, which implies that the maximal payoff is the solution v of the equation $H(x) + (1-x) = \log \frac{3}{2}$.

Numerically, $v \sim 0.896$. This has to be compared with the maximal payoff of $\frac{1}{3}$ when the agent is kept uninformed, and with the payoff of 1 when the agent is fully informed.

It follows from the analysis of the next subsection that the set of equilibrium payoffs of the repeated game when players are sufficiently patient goes to the set $\{(y,y) : \frac{1}{3} \leq y \leq v\}$ of all individual rational and feasible payoff vectors.

8.3. **Games with different interests.** Now we consider general payoff functions $g = (g^f, g^a)$.

We compare the set of equilibrium payoffs of our model with the set of *silent* equilibrium payoffs in which no information is transmitted, and with the set of communication equilibrium payoffs.

Call a strategy of the forecaster *silent* if it depends only on past play and on the current state of nature (hence not on future states of nature). When the forecaster uses silent strategies and the agent uses arbitrary strategies, the induced set of feasible payoff vectors is the set of silent feasible payoff vectors given by

$$F^S = \text{ co } \{\mathbf{E}_\mu g(i, \alpha(i), k),\ \alpha\colon I \to J,\ k \in K\}$$

where co stands for the convex hull. At the other extreme, the set of feasible payoff vectors in the extensive-form or normal-form communication extension of the game is

$$F^C = \text{ co } \{\mathbf{E}_\mu g(i, \alpha(i), \beta(i)),\ \alpha\colon I \to J,\ \beta\colon I \to K\}$$

Finally, define the set of feasible payoff vectors with internal communication as F and recall that

$$F = \{\mathbf{E}_Q g(i, j, k) : Q \text{ verifies the information constraint and } Q_I = \mu\}$$

We have the obvious inclusions

$$F^S \subseteq F \subseteq F^C$$

The set F being closed and convex (as the image of the closed and convex set Q by the linear expectation operator), it is defined by its support function

$$x \mapsto \max_{y \in F} \langle x, y \rangle \quad x \in \mathbb{R}^2$$

where $\langle x, y\rangle$ stands for the inner product of x and y. Given $x \in \mathbb{R}^2$, the value $\max_{y \in F}\langle x, y\rangle$ of the support function equals the maximal payoff that the team of the forecaster and the agent can achieve in the team game where the common payoff function equals the inner product $\langle x, (g^f(i, j, k), g^a(i, j, k))\rangle$. Therefore, computing the feasible set F amounts to solving a family of (two-person)[7] team games.

The individually rational level of a player is defined as the best payoff that this player can defend using mixed strategies against every strategy of the other player in long repetitions of the game. For the forecaster, this payoff is

$$v^f = \min_{\alpha \in \Delta(K)} \max_{\beta\colon I \to J} \mathbf{E}_{\mu,\alpha} g^f(i, \beta(i), k)$$

For the agent, this payoff is

$$v^a = \max_{\alpha \in \Delta(K)} \min_{\beta\colon I \to J} \mathbf{E}_{\mu,\alpha} g^a(i, \beta(i), k)$$

[7]In fact, as the implementing strategies in our proof are pure, it follows that solving a family of two-person team games suffices for computing the feasible set of payoff vectors in the model where there are several forecasters and several agents.

The situation is asymmetric between the two players. Indeed, the forecaster possesses a double advantage over the agent. First, he can use his private information concerning the states of nature in order to defend a better payoff against the agent, which results in a higher individually rational level for the forecaster. Second, he can use this information against the agent when punishing him, which induces a lower individually rational payoff for the agent. Let IR be the set of individually rational payoff vectors:

$$IR = \{(x^f, x^a),\ x^f \geq v^f,\ x^a \geq v^a\}$$

The set $F^S \cap IR$ corresponds to the set of equilibrium payoffs of games with sufficiently patient players in which the forecaster uses silent strategies that may depend on the current state of nature, but not on future ones. In these equilibria, the agent is uninformed as to future states of nature.

The set $F^C \cap IR$ is the set of communication equilibrium payoffs of the repeated game with sufficiently patient players. In this case, there are neither restrictions nor costs associated with the communication possibilities.

Finally, $F \cap IR$ is the set of equilibrium payoffs of our original game when players are sufficiently patient, and all communication is internal to the game. Formally:

Proposition 5.

$$E_\theta \rightarrow F \cap IR \ as\ \sup_t \frac{\theta_t}{\sum_{s \geq t} \theta_s} \rightarrow 0$$

In particular, the set of equilibria of the λ-discounted game goes to $F \cap IR$ as λ goes to 1.

Note that all information concerning future states of nature that is sent by the forecaster is eventually verifiable by the agent. Therefore, the proof of the proposition is straightforward (and follows the classical lines of proofs) when the set $F \cap IR$ contains a point that strictly dominates the individual rational payoff vector.

Note that the limit set $F \cap IR$ is convex, but is in general not a polyhedron. It is computed directly from the information constraint, and reflects the costs of communication among the players.

9. Discussion and extensions

In order to preserve maximum transparency, we have tried to keep the basic model of Section 2 as simple as possible. Notably, this has led to greatly simplified assumptions on the forecasting ability, the signalling structure of the game, and the distribution of the process. The aim of this section is to present various extensions and variations of the basic model, and to show how the analysis of implementable distributions through the information constraint can be adapted to these cases. We first discuss relaxations of the perfect and infinite forecast assumption. Second, we examine the impact of the signaling structure of the one-shot game on the set of implementable distributions. Third, we show how autocorrelations of the process of states of nature can reduce the need for information transmission and expand the set of implementable distributions. Next, we illustrate that our main result is robust in the sense that small deviations from the main assumptions lead to a small change in the set of implementable distributions. Finally, we discuss the possibilities of asymmetric information on both sides as an open problem.

9.1. **Limited forecasting abilities.** The assumption of perfect and infinite forecast is relaxed in two ways. First, we can assume that the forecaster is able to make predictions on states of nature a finite number of stages in advance. Second, we introduce possibilities of inaccurate predictions.

9.1.1. *Finite forecasts.* Say that the forecaster has f forecast if, before stage t, the forecaster is informed of $i_t, \ldots, i_{t+f-1}$. Remark that any strategy that is implementable with f forecast is implementable with perfect forecast. Note also that the strategies constructed in the proof of Theorem 2 use f forecast for larger and larger values of f. Hence, the set of implementable distributions with f forecast converges, as f goes to ∞, to the implementable distributions of the basic model.

9.1.2. *Imperfect forecasts.* Now we discuss the case where the forecaster is imperfectly informed of the states of nature. Let S be a set of signals for the forecaster, and let R be a transition probability from I to S. Assume that before the game starts, the forecaster observes a sequence of signals $(s_t)_t$ where each signal s_t is drawn independently according to the probability R_{i_t}. Following the play at stage t the agent observes a stochastic signal that includes the action j_t of the forecaster, and the forecaster observes a stochastic signal that depends on the action triple (i_t, j_t, k_t). The basic model corresponds to the case of perfect monitoring and perfect forecasts (where $S = I$ and $R_i(s) = 1$ if $i = s$ and $= 0$ if $i \neq s$.

In this case, a distribution $Q \in \Delta(J) \times \Delta(I \times K)$ with $Q_I = \mu$ is implementable if and only if there exists a distribution $\hat{Q} \in \Delta(S \times I \times J \times K)$ with $\hat{Q}(s|i) = R_i(s)$ and marginal Q on $I \times J \times K$ such that

(1) $\hat{Q}(i|s, j, k) = \hat{Q}(i|s)$
(2) $H_{\hat{Q}}(\mathbf{j}) \geq H_{\hat{Q}}(\mathbf{s}) - H_{\hat{Q}}(\mathbf{s}|\mathbf{k})$

Condition (2) is the usual information constraint on Q. Condition (1) expresses the fact that all information players have on the current state of nature comes from the signal s of the forecaster.

If the signal to the agent, following the play at stage t, includes s_t in addition to j_t, then a distribution $Q \in \Delta(J \times I \times K)$ with $Q_I = \mu$ is implementable if and only if there exists a distribution $\hat{Q} \in \Delta(S \times I \times J \times K)$ with $\hat{Q}(s|i) = R(i)(s)$ and marginal Q on $I \times J \times K$ such that

(1) $\hat{Q}(i|s, j, k) = \hat{Q}(i|s)$
(2) $H_{\hat{Q}}(\mathbf{j}, \mathbf{s}|\mathbf{k}) \geq H_{\hat{Q}}(\mathbf{s})$

9.2. **Signalling structures.** The basic model assumes that the stage game has perfect monitoring in the sense that both the forecaster and the agent are perfectly informed of the action triple played at each stage.

As a general property, any reduction of the informational content of the signals received by the forecaster or by the agent concerning the action triple results in a reduction of the set of implementable distributions. In other words, all distributions that are implementable with less informative signals are also implementable with more informative ones. This follows from the fact that strategies in the model with less informative signals are also strategies in the model with more informative signals.

Now we discuss the effects of a change either in the forecaster's observation of the agent's action, or in the agent's observation of the forecaster's action, or in the agent's observation of the current state of nature.

In this subsection, it will be useful to use the classical information-theoretic notation $I(\mathbf{a};\mathbf{b})$ to denote the mutual information $H(\mathbf{a}) - H(\mathbf{a}|\mathbf{b})$ $(= H(\mathbf{b}) - H(\mathbf{b}|\mathbf{a}))$ between two random variables a and b. We also use the notation $I_Q(\mathbf{a};\mathbf{b})$ when making explicit the distribution Q of $(\mathbf{a},\mathbf{b})$.

9.2.1. *Imperfect observation of agent's actions.* Assume the forecaster observes at each stage a signal on the agent's actions. The basic model corresponds to the case where this signal is fully informative. Consider the strategies constructed in the proof of Theorem 2. Since the agent uses a pure strategy, which depends on the observed past states of nature and the forecaster's actions only, and since this information is available to the forecaster, the forecaster can reconstitute the past actions of the agent even if the signal received on these actions is completely uninformative. Hence, the designed strategies can still be used. The set of implementable distributions is thus unchanged under the assumption that the forecaster has imperfect monitoring on the agent's actions.

In particular, the characterization of Proposition 4 of the limit Pareto payoff to the team for sufficiently long games (θ_1 arbitrarily close to 0) remains unchanged.

Note however that the set of equilibrium payoffs in the repeated games with different interests is modified. Indeed, some deviations of the agent that are detectable under perfect monitoring may become undetectable under imperfect monitoring.

9.2.2. *Imperfect observation of forecaster's actions.* Consider the situation where the agent observes the states of nature but does not have perfect monitoring of the forecaster's actions. The distribution of the signal $s \in S$ depends on the triple (i, j, k); the conditional distribution of s given (i, j, k) is denoted by $R_{i,j,k}$ $(\in \Delta(S))$.

Given a distribution Q on $I \times J \times K$ we denote by $\hat{Q}$ the distribution on $I \times J \times K \times S$ with marginal Q on $I \times J \times K$ and such that $\hat{Q}(s|i, j, k) = R_{i,j,k}(s)$. Using this notation, Q is implementable if and only if $Q_I = \mu$ and

$$H_{\hat{Q}}(\mathbf{s}|\mathbf{i},\mathbf{k}) - H_{\hat{Q}}(\mathbf{s}|\mathbf{i},\mathbf{j},\mathbf{k}) \geq H_Q(\mathbf{i}) - H_Q(\mathbf{i}|\mathbf{k}) \tag{11}$$

Notice that in the case of the basic model, $H_{\hat{Q}}(\mathbf{s}|\mathbf{i},\mathbf{k}) = H_Q(\mathbf{j}|\mathbf{i},\mathbf{k})$ and $H_{\hat{Q}}(\mathbf{s}|\mathbf{i},\mathbf{j},\mathbf{k}) = 0$ and thus equation (11) particularizes to the information constraint $H_Q(\mathbf{j}|\mathbf{i},\mathbf{k}) \geq I(\mathbf{i};\mathbf{k})$.

The two special cases of the above characterization that are considered below are reformulations of classical results in information theory: Shannon's Noisy Channel Capacity theorem (see e.g. [4, Theorem 8.7.1, p. 198]), and the Rate Distortion theorem for i.i.d. sources (see e.g. [4, Theorem 13.2.1, p. 342]).

First, consider the special case where the distribution $R_{i,j,k}(s)$ depends on $\mathbf{j}$ only and $I = K$. The information constraint (11) for a distribution Q such that $Q(i = k) = 1$ can be expressed as $H_Q(\mathbf{s}) - H_Q(\mathbf{s}|\mathbf{j}) = I_Q(\mathbf{s};\mathbf{j}) \geq H_Q(\mathbf{i})$. Note also that a distribution Q with $Q(i = k) = 1$ is implementable if and only if it is implementable in the variant of the model where the forecaster does not observe the actions of the agent and the agent does not observe the states of nature. Define the capacity of a stochastic signal $\mathbf{s}$ as the maximum over the random variable $\mathbf{j}$ of the mutual information $I(\mathbf{s};\mathbf{j})$. Thus, our result shows that there exists an implementable distribution Q such that $Q(i = k) = 1$ if and only if the capacity of $\mathbf{s}$ exceeds $H(\mathbf{i})$. This result is equivalent to the classical Shannon's Noisy Channel Capacity theorem for i.i.d. sources.

Second, assume that $R_{i,j,k}(s)$ depends on $\mathbf{j}$ only. The information constraint (11) for a distribution $Q \in \Delta(S) \times \Delta(I \times K)$ can be expressed as $I_{\hat{Q}}(\mathbf{s};\mathbf{j}) \geq I_Q(\mathbf{i};\mathbf{k})$. Indeed, for such a product distribution Q we have $H_{\hat{Q}}(\mathbf{s}|\mathbf{i},\mathbf{k}) = H_{\hat{Q}}(\mathbf{s})$ and $H_{\hat{Q}}(\mathbf{s}|\mathbf{i},\mathbf{j},\mathbf{k}) = H_{\hat{Q}}(\mathbf{s}|\mathbf{j})$. Therefore, the left-hand side of inequality (11) equals $I_{\hat{Q}}(\mathbf{s};\mathbf{j})$. Note that $I_{\hat{Q}}(\mathbf{s};\mathbf{j})$ depends only on Q_J and $I_Q(\mathbf{i};\mathbf{k})$ depends only on $Q_{I\times K}$.

Fix $\mu \in \Delta(I)$. Now assume that the payoff function does not depend on j, i.e., $g(i,j,k) = d(i,k)$, and let $R(D)$ be the min of $I_P(\mathbf{i};\mathbf{k})$ when P is a distribution on $I \times K$ such that $P_I = \mu$ and $\mathbf{E}_P d(i,k) \geq D$. Let $\nu \in \Delta(J)$. Our result implies that there exists an implementable distribution $Q \in \Delta(J) \times \Delta(I \times K)$ with $\mathbf{E}_Q d(i,k) \geq D$ and $Q_{I\times J} = \nu \otimes \mu$ if and only if $I_{\hat{Q}}(\mathbf{s};\mathbf{j}) \geq R(D)$. Moreover, the implementability of a distribution $Q \in \Delta(J) \times \Delta(I \times K)$ does not depend on the agent observing the states of nature. This generalizes the Rate Distortion theorem for i.i.d. sources (see e.g. [4, Theorem 13.2.1, p. 342]).

9.2.3. *Unobservable current state of nature.* Now consider the case where the agent observes the forecaster's actions, but is uninformed of the current state of nature.

The characterization of the full set of implementable distributions in this case is beyond the scope of this paper. However, consider the subset R of distributions on $I \times J \times K$ that are the product of a distribution on J and a distribution on $I \times K$.

Following a similar analysis to that of our basic model, one can prove that a distribution $Q \in R$ is implementable if and only if

$$H_Q(\mathbf{j}) \geq I(\mathbf{i}; \mathbf{k})$$

If the agent also does not have perfect monitoring of the forecaster's actions, but receives a signal s as a function of the forecaster's action j, we proceed as in Section 9.2.2. Consider the conditional distribution of s given j by $R_j \in \Delta(S)$. Following the same notation we obtain that a distribution Q that is a product of a distribution on J and a distribution on $I \times K$ is implementable if and only if $Q_I = \mu$ and

$$H_{\hat{Q}}(\mathbf{s}) - H_{\hat{Q}}(\mathbf{s}|\mathbf{j}) \geq I_Q(\mathbf{i}; \mathbf{k}) = H_Q(\mathbf{i}) - H_Q(\mathbf{i}|\mathbf{k})$$

9.3. **State of nature processes.** Now we analyze the variant of the basic model where the states of nature follow a Markov chain. In such cases the distribution of the state of nature in stage $t+1$ is correlated to the distribution of the state in stage t. The adequate element of study is the expected long-run average Q of the distribution of the quadruple (i_{t-1}, i_t, j_t, k_t). Let Q be a distribution on $I \times I \times J \times K$, and $(\mathbf{i}', \mathbf{i}, \mathbf{j}, \mathbf{k})$ have distribution Q.

A Markov chain eventually enters into an ergodic class of states. As players observe past states of nature, they are eventually informed of the ergodic class entered by the chain, and it suffices to study the expected long-run average in the case of an irreducible Markov chain. Let μ be the invariant measure on I and let T denote the transition matrix of the Markov chain. The marginal on $I \times I$ of an implementable distribution Q (on $I \times I \times J \times K$) is deduced from the law of the Markov chain: $Q(\mathbf{i}' = i', \mathbf{i} = i) = \mu(i')T_{i',i}$. It turns out that a distribution Q is implementable if and only if its marginal on $I \times I$ coincides with the marginal imposed by the Markov chain transitions and

$$H_Q(\mathbf{i}, \mathbf{j} \mid \mathbf{k}, \mathbf{i}') \geq H_Q(\mathbf{i} \mid \mathbf{i}') \tag{12}$$

This last condition thus describes the information constraint when the process of states of nature follows a Markov chain.

Now we compare the set of implementable distributions under the i.i.d. and the Markov assumptions. Assume that $Q \in \Delta(I \times J \times K)$ has marginal μ on I and verifies the information constraint under the i.i.d. assumption: $H_Q(\mathbf{i}, \mathbf{j}|\mathbf{k}) \geq H_Q(\mathbf{i})$. Let T be the transition of an irreducible Markov chain, and $Q' \in \Delta(I \times I \times J \times K)$ be the law of $(\mathbf{i}', \mathbf{i}, \mathbf{j}, \mathbf{k})$ where

a) $(\mathbf{i}, \mathbf{j}, \mathbf{k})$ have law Q
b) the law of $(\mathbf{i}', \mathbf{i})$ is deduced from the law of the Markov chain: $Q(\mathbf{i}' = i', \mathbf{i} = i) = \mu(i')T_{i',i}$
c) $Q'(\mathbf{j} = j, \mathbf{k} = k|\mathbf{i} = i, \mathbf{i}' = i') = Q(\mathbf{j} = j, k|\mathbf{i} = i)$.

Now we verify that Q' verifies the information constraint under the Markov model. Indeed,

$$\begin{aligned} H_{Q'}(\mathbf{j}|\mathbf{i}, \mathbf{i}', \mathbf{k}) &= H_Q(\mathbf{j}|\mathbf{i}, \mathbf{k}) \\ &\geq H_Q(\mathbf{i}) - H_Q(\mathbf{i}|\mathbf{k}) = H_Q(\mathbf{k}) - H_Q(\mathbf{k}|\mathbf{i}) = H_{Q'}(\mathbf{k}) - H_{Q'}(\mathbf{k}|\mathbf{i}, \mathbf{i}') \\ &\geq H_{Q'}(\mathbf{k}|\mathbf{i}') - H_{Q'}(\mathbf{k}|\mathbf{i}, \mathbf{i}') = H_{Q'}(\mathbf{i}|\mathbf{i}') - H_{Q'}(\mathbf{i}|\mathbf{k}, \mathbf{i}') \end{aligned}$$

where the first and third equalities follow from (a) and (c), the first inequality follows from the information constraint verified by Q, the second inequality follows from the concavity of entropies, and the second and last equality follows from instance from the chain rule of entropies. The obtained inequality is then equivalent to the information constraint for Markov chains (12) applied to Q.

The information constraint is satisfied in the Markov case whenever it is in the i.i.d. case. This shows that the set of implementable distributions is augmented when one takes advantage of the correlations between successive states of nature. This is intuitive since in the Markov case, the need for information transmission is not as important as it is in the i.i.d. case.

9.4. **Information banking.** In the Markov chain case, the distribution of i_t given the sequence of past states $i_1, \dots, i_{t-1}$ is a function of i_{t-1} only. Let ν_i be the distribution of i_t given $i_{t-1} = i$. Define the random partition of $\mathbb{N}$, $\mathbb{N} = \cup \mathbb{N}_i$, where $\mathbb{N}_i$ is the set of all stages t such that $i_{t-1} = i$. For every $i \in I$ and a strategy pair (σ, τ) we define (for every positive integer n) the distribution $Q^{i,n}_{\sigma,\tau}$ as the expected empirical distribution of action triples in stages $t \in \mathbb{N}_i$ with $t \leq n$. The marginal on I of the distribution $Q^{i,n}_{\sigma,\tau}$ is ν_i. Our proof (of the result of the basic model) implies the following: if $Q_i \in \Delta(I \times J \times K)$ verifies the information constraint and has marginal ν_i on I and μ is the invariant distribution of the Markov chain then the distribution $Q = \sum_i \mu(i) Q_i$ is implementable. Indeed, by considering the states in each $\mathbb{N}_i$ separately, the team can collate the strategies that implement Q_i to a strategy pair (σ, τ) in the Markov chain games so that $Q^{i,n}_{\sigma,\tau}$ converges to Q_i and thus as $\frac{|\{t \leq n : t \in \mathbb{N}_i\}|}{n} \to \mu(i)$ as $n \to \infty$ we deduce that $Q = \sum_i \mu(i) Q_i$ is implementable.

Now we verify that the distribution Q' on $I \times I \times J \times K$ defined $Q'(i', i, j, k) = \mu(i') Q_{i'}(i, j, k)$ (and therefore $Q'_{I \times J \times K} = \sum_i \mu(i) Q_i$) verifies the information constraint under the Markov chain model.

$$\begin{aligned} H_{Q'}(\mathbf{i}, \mathbf{j} | \mathbf{k}, \mathbf{i}') &= \sum\nolimits_{i' \in I} \mu(i') H_{Q_{i'}}(\mathbf{i}, \mathbf{j} | \mathbf{k}) \\ &\geq \sum\nolimits_{i' \in I} \mu(i') H_{Q_{i'}}(\mathbf{i}) \\ &= H_{Q'}(\mathbf{i} | \mathbf{i}') \end{aligned}$$

Our characterization of the implementable distribution in the Markov chain process implies however additional implementable distributions. For any real number α, say that a distribution Q on $I \times J \times K$ obeys the α-information-constraint if $H_Q(\mathbf{i}, \mathbf{j} | \mathbf{k}) \geq H_Q(\mathbf{i}) + \alpha$. Note that α can be either positive or negative or zero. The characterization of implementable distributions by the information constraint (12) implies that Q is implementable if and only if there are distributions $Q_i \in \Delta(I \times J \times K)$ with marginals $((Q_i)_I =)$ ν_i on I and constants α_i such that 1) Q_i obeys the α_i-information-constraint, 2) $\sum_i \mu(i) Q_i = Q$, and 3) $\sum_i \mu(i) a_i = 0$.

This comparison of the Markov chain and the i.i.d. cases highlights the need for the forecaster to signal at stages $t \in \mathbb{N}_i$ on states of nature in stages $t \in \mathbb{N}_{i'}$ where $i \neq i'$.

9.5. **Robustness.** The analysis of the previous extensions demonstrates (indirectly) the robustness of our main results to some departures in the assumptions made either on the state of nature process, or on the foresight ability of the forecaster, or on the monitoring and forecasting possibilities of the agent. We wish to comment on robustness when all assumptions are perturbed together and to allow for a wide variety of perturbations. In order to do this, we introduce a generalized version of our model.

We start by describing the dynamics of the states of nature and the signalling structure of the game. A point $\omega = (i_1, i_2, \ldots) \in I^\infty$ is chosen by nature according to some distribution P. Before the game starts, player n ($n = 1, 2$ in the two-player game) observes a random signal s_0^n whose distribution depends on the sequence ω. At stage t player 1 takes action $j_t \in J$ and player 2 takes an action $k_t \in K$. Following the play at stage t player n observes the realization of a random signal[8] s_t^n where the distribution of (s_t^1, s_t^2) depends on the triple (ω, j_t, k_t) of the sequence of states of nature and the action pair (j_t, k_t) of the players and conditional on (ω, j_t, k_t) is independent of all past signals. The payoff vector at stage t depends on the state of nature at stage t and the action pair at stage t.

A strategy of player 1, respectively player 2, specifies the action j_t, respectively k_t, at stage t as a function of all his past information, namely, as a function of $s_0^1, \ldots, s_{t-1}^1$, respectively $s_0^2, \ldots, s_{t-1}^2$.

In order to quantify a small perturbation in this general model we introduce a proper definition to measure such perturbation. First, the stochastic process (I^∞, P) is within δ of an i.i.d. process if there exists $\mu \in \Delta(I)$ and probability distributions $\hat{P}[n]$ on $(I \times I')^m$, where I' is a copy of I and $m = [\delta^{-2}]$, s.t. $(i'_{n+1}, \ldots, i'_{n+m})$ has distribution $\mu^{\otimes m}$, the projection of $\hat{P}[n]$ on the m I-coordinates coincides with P, i.e., $\hat{P}[n]_{I^n}(i_{n+1}, \ldots, i_{n+m}) = P(i_{n+1}, \ldots, i_{n+m} | i_1, \ldots, i_n)$, and

$$\mathbf{E}_{\hat{P}[n]}(\sum_{t=n+1}^{n+m} \mathbb{I}(i_t \neq i'_t)) \leq \delta^{-1} \qquad \forall \text{ sufficiently large } n$$

An example of a process (with values in $I^{\mathbb{N}}$) that is within δ of an i.i.d. process is a nonstationary Markov chain (i.e., with time dependent

[8]In fact, we can assume w.l.o.g. that the signals are moreover deterministic. Indeed, we can "push" all randomness into I; this will however require an infinite set of states.

transitions) where the probabilities $T_t(i', i)$ of transition at stage t from state i' to state i obey $\sum_{i\in I} |T_t(i', i) - \mu(i)| < \delta$ for all sufficiently large t.

Second, we say that the forecaster has δ-*perfect foresight* if for all sufficiently large t the forecaster can guess the future $m := [1/\delta^{-2}]$ states of nature so that the expected number of errors is $\leq 1/\delta$. Formally, there are functions $f_t : (s_0^f, \ldots, s_t^f) \mapsto I^n$, $t \geq 0$, such that

$$\mathbf{E}(\sum_{\ell=1}^{m} \mathbb{I}((f_t)_\ell \neq i_{t+\ell})) \leq 1/\delta \quad \forall \text{ sufficiently large } t$$

The agent has δ-*perfect monitoring* if, for all sufficiently large t, the agent can guess the past $m := [\delta^{-2}]$ (triples of) action profiles so that the expected number of errors is $\leq 1/\delta$. Formally, there are functions $a^t : (s_0^a, \ldots, s_{t-1}^a) \mapsto (I \times J)^{t-1}$, $t \geq 1$, such that

$$\mathbf{E}(\sum_{\ell=1}^{m} \mathbb{I}(a_{t-\ell}^t \neq (i_{t-\ell}, j_{t-\ell}))) \leq \delta^{-1} \quad \forall \text{ sufficiently large } t.$$

The agent has δ-*forecast* if for every $t \geq 1$, every $\omega = (i_1, i_2, \ldots)$ and $\omega' = (i'_1, i'_2, \ldots)$ (in I^∞) with $(i_1, \ldots, i_t) = (i'_1, \ldots, i'_t)$, and every $(j_t, k_t) \in J \times K$, the distribution of s_t^2 given (ω, j_t, k_t) is within[9] δ of its distribution given (ω', j'_t, k'_t).

Finally, we say that the game model Γ is δ-close to the basic model Γ' if the process of states of nature is within δ of the i.i.d. $\mu^{\otimes\mathbb{N}}$, the forecaster has δ-perfect foresight, and the agent has δ-perfect monitoring and δ-forecast.

The robustness theorem states that the set of implementable distributions of a small perturbation of one instance of the basic model is close to the set of implementable distributions of that instance. Formally:

Theorem 3 (The robustness theorem)**.** *Let Γ' be a basic model game. For every $\varepsilon > 0$ there is $\delta > 0$ such that if Γ is δ-close to Γ' then the set of implementable distributions of Γ are within ε of the set of implementable distributions of Γ'.*

Observe that the basic model is the special case where $s_0^1(\omega) = \omega$, $s_0^2(\omega)$ is a constant independent of ω, and $s^n(i_t, j_t, k_t) = (i_t, j_t, k_t)$. The classical model of repeated games with incomplete information is the

[9]If the signal s_t^a takes values in a finite set S^a, then we can use the norm distance between distributions; in the general case we refer to the Kullback-Leibler distance.

special case where $i_t = i_{t+1}$ for all t and $s^n(\omega, j_t, k_t)$ depends only on the triple (i_t, j_t, k_t).

Note, finally, that an important ingredient of the model described above is that the dynamics of states of nature $i \in I$ (where I is the finite set of states of nature) is independent of players' actions. The even more general model, which is not discussed here, enables the transition of states to depend also on players' actions, and generalizes not only the theory of repeated games with incomplete information, but also the theory of stochastic games.

9.6. **Complementary information.** Another important characteristic of the basic model, and of the extensions introduced above, is that all information about future states of nature possessed by the agent is also possessed by the forecaster. One may wish to consider extensions of our models in which both players are partially informed beforehand of the realized sequence of states of nature.

In such cases, sequential communication schemes, in which information is sent back and forth between the players, may be more efficient than simultaneous schemes in which each player sends information independently of the information sent by the other (see e.g. [9]). The characterization of the set of implementable distributions in this model is left as an open problem for future research.

Paris-Jourdan Sciences Economiques, UMR CNRS-EHESS-ENS-ENPC 8545, 48 boulevard Jourdan, 75014 Paris, France,

and

Dpto. Fundamentos del Análisis Económico. Universidad de Alicante, 03071 Alicante, Spain,

and

Institute of Mathematics and Center for Rationality, Hebrew University of Jerusalem, Israel.

Appendix A. Typical sequences

Let A and B be two finite sets.

Lemma 2. *Let $Q \in \mathbb{T}_n(A \times B)$. For every $\alpha \in T_n(Q_A)$, we have*

$$|A|^{|A\times B|}\frac{2^{(H(Q)-H(Q_A))n}}{(n+|A|)^{|A\times B|}} \leq |T_n(Q|\alpha)| \leq 2^{(H(Q)-H(Q_A))n}$$

Proof. The point $\alpha \in A^n$ partitions the set $\{1, \dots, n\}$ into $|A|$ disjoint subsets N_a, $a \in A$: $N_a = \{1 \leq i \leq n : a_i = a\}$. For $a \in A$ we denote by Q^a the conditional distribution on B given a, namely, $Q^a(b) = Q(a,b)/\sum_{b\in B} Q(a,b)$. For every point $\beta \in B^n$ and a subset N of $\{1, 2, \dots, n\}$ we denote by $(\beta|N)$ the N-vector $(b_j)_{j\in N}$. Note that for every $\beta \in B^n$ we have $\rho(\alpha, \beta) = Q$ if and only if for every $a \in A$ we have $\rho(\beta|N_a) = Q^a$. Therefore, it follows from (3) that

$$\Pi_{a\in A}\frac{2^{H(Q^a)|N_a|}}{(|N_a|+1)^{|B|}} \leq |\{\beta \in B^n : \rho(\alpha, \beta) = Q\}| \leq \Pi_{a\in A}2^{H(Q^a)|N_a|}.$$

The result follows since $\Pi_{a\in A}(|N_a|+1)^{|B|} \leq (n+|A|)^{|A\times B|}/|A|^{|A\times B|}$ and $\Pi_{a\in A}2^{H(Q^a)|N_a|} = 2^{(H(Q)-H(Q_A))n}$. □

Lemma 3. *Let $Q \in \mathbb{T}_n(A \times B)$, $1 \leq f_n \leq |T_n(Q)|/|T_n(Q_A)|$ and $g(n) > 1$ such that $\frac{f_n(g(n)-1)^2}{2g(n)} > nH(Q_A)\ln 2$. Then, there is a subset S of $T_n(Q_B)$ with*

$$\frac{|T_n(Q_A)|.|T_n(Q_B)|}{|T_n(Q_{A\times B})|}f_n g(n) \leq |S| < 1 + \frac{|T_n(Q_A)|.|T_n(Q_B)|}{|T_n(Q_{A\times B})|}f_n g(n)$$

such that for every $x \in T_n(Q_A)$ there are at least f_n elements $z \in S$ such that $(x, z) \in T_n(Q)$.

Proof. Let m be the least integer $\geq \frac{|T_n(Q_A)|.|T_n(Q_B)|}{|T_n(Q_{A\times B})|}f_n g(n)$. If $m \geq |T_n(Q_B)|$ set $S = T_n(Q_B)$. Otherwise, let S be a random subset of $T_n(Q_B)$ having m elements. Note that for every $x \in T_n(Q_A)$ we have $\mathbf{E}|S \cap T_n(Q|x)| \geq f_n g(n)$. By the large deviation inequality for sampling without replacement,[10] for every $x \in T_n(Q_A)$ we have

$$P(|S \cap T_n(Q \mid x)| \leq f_n) < \exp(-\frac{f_n(g(n)-1)^2}{2g(n)}) < 2^{-nH(Q_A)}$$

and therefore $P(|S \cap T_n(Q \mid x)| \leq f_n)$ times $|T_n(Q_A)|$ is < 1.Therefore the expected number of elements $x \in T_n(Q_A)$ such that $|S \cap T_n(Q \mid x)| \leq f_n$ is < 1. Therefore, there is a subset S of $T_n(Q_B)$ with the desired property. □

[10]E.g., Section 6 and inequality (2.2) of Theorem 1 in Hoeffding [8].

Appendix B. Approximation of probabilities

Lemma 4. *$\forall \varepsilon > 0\, \exists N(\varepsilon)$ such that $\forall\, \tilde{Q} \in \Delta(I\times J\times K)$ with $H_{\tilde{Q}}(\mathbf{i},\mathbf{j}|\mathbf{k})-H_{\tilde{Q}}(\mathbf{i}) \geq 0$ and $\forall\, n \geq N(\varepsilon)\, \exists\, Q \in \mathbb{T}_n(I \times J \times K)$ such that*

$$H_Q(\mathbf{i},\mathbf{j}|\mathbf{k}) - H_Q(\mathbf{i}) \geq \varepsilon \tag{13}$$

and

$$\|Q - \tilde{Q}\|_1 < 7\varepsilon \tag{14}$$

Proof. Let $A := J \times K$. The (real-valued) entropy functions $R \mapsto H_R(\mathbf{i},\mathbf{j}|\mathbf{k})$ and $R \mapsto H_R(\mathbf{i})$ defined on $\Delta(I \times A)$ are continuous, and thus uniformly continuous. Therefore, for every $\varepsilon > 0$ there is $N(\varepsilon) > |I \times A|/\varepsilon$, such that for every $R, R' \in \Delta(I \times A)$ with $\|R - R'\|_1 \leq |I \times A|/N(\varepsilon)$ we have $|H_R(\mathbf{i},\mathbf{j}|\mathbf{k}) - H_{R'}(\mathbf{i},\mathbf{j}|\mathbf{k})| < \varepsilon$ and $|H_R(\mathbf{i}) - H_{R'}(\mathbf{i})| < \varepsilon$. Let R be the product distribution $\tilde{Q}_I \times U_J \times \tilde{Q}_K$ on $I \times A$ where U_J is the uniform distribution over J.

Then $H_R(\mathbf{i},\mathbf{j}|\mathbf{k}) = H_{\tilde{Q}}(\mathbf{i}) + \log|J|$. Let $R^\varepsilon = 3\varepsilon R + (1-3\varepsilon)\tilde{Q}$. Then, using the concavity of the entropy function $R \mapsto H_R(\mathbf{i},\mathbf{j}|\mathbf{k})$ (Lemma 1), the equality $R^\varepsilon_I = \tilde{Q}_I$, and the inequality $\log|J| \geq 1$ (which follows from $|J| \geq 2$), we have

$$H_{R^\varepsilon}(\mathbf{i},\mathbf{j}|\mathbf{k}) \geq (1-3\varepsilon)H_{\tilde{Q}}(\mathbf{i},\mathbf{j}|\mathbf{k}) + 3\varepsilon H(R^\varepsilon_I) + 3\varepsilon$$

which implies

$$H_{R^\varepsilon}(\mathbf{i},\mathbf{j}|\mathbf{k}) - H_{R^\varepsilon}(\mathbf{i}) \geq 3\varepsilon$$

Let $Q \in \Delta(I \times A)$ with $T_n(Q) \neq \emptyset$ and $\|Q - R^\varepsilon\|_1 \leq |I \times A|/n$.

Therefore, for $n \geq N(\varepsilon)$ we have $|H_Q(\mathbf{i},\mathbf{j}|\mathbf{k}) - H_{R^\varepsilon}(\mathbf{i},\mathbf{j}|\mathbf{k})| < \varepsilon$ and $|H_Q(\mathbf{i}) - H_{R^\varepsilon}(\mathbf{i})| < \varepsilon$ and therefore

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

which proves (13). In addition

$$\|Q - \tilde{Q}\|_1 < \|Q - R^\varepsilon\|_1 + \|R^\varepsilon - \tilde{Q}\|_1 \leq 7\varepsilon$$

□

Lemma 5. *Fix $\nu \in \Delta(I)$ and n such that $T_n(\nu) \neq \emptyset$. For every $x = (x_1, \ldots, x_n) \in I^n$ there is $x' = (x'_1, \ldots, x'_n) \in T_n(\nu)$ such that*

$$|\{t : x_t \neq x'_t\}| \leq n\|\rho(x) - \nu\|_1$$

Proof. By induction on the integer $d(x) := n\|\rho(x) - \nu\|_1$. If $d(x) = 0$ set $x' = x$. Assume that $d(x) > 0$. There exist elements $i, i' \in I$ such that $\rho(x)(i) > \nu(i)$ and $\rho(x)(i') < \nu(i')$. Pick $t \in \{1 \leq t' \leq n : x_{t'} = i\}$ and define $\tilde{x}'' \in I^n$ by $x''_k = x_k$ if $k \neq t$, and $x''_t = i'$. It follows that $d(x'') = d(x) - 2$ and therefore by the induction hypothesis there is $x' \in T_n(\nu)$ such that $|\{t : x''_t \neq x'_t\}| \leq d(x'')$ and therefore $|\{t : x'_t \neq x_t\}| \leq d(x'') + 2 = d(x)$. □

Corollary 2. *Fix $\nu \in \Delta(I)$ and n such that $T_n(\nu) \neq \emptyset$. There exists a map $f\colon I^n \to T_n(\nu)$ such that for $\mu \in \Delta(I)$ we have*

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

In particular, if $x \mapsto \tilde{x}$ is a map from I^n to I^n that minimizes the sum $\sum_{1\leq t\leq n} \mathbb{I}_{x_t \neq \tilde{x}_t}$ and $\|\mu - \nu\| \leq 7\varepsilon$ then the probability that $\sum_{1\leq t\leq n} \mathbb{I}_{x_t \neq \tilde{x}_t}$ is $\geq 8\varepsilon n$ is $\leq |I|^2/(\varepsilon n)^2$.

Proof. Let $f : I^n \to T_n(\nu)$ be the function that maps $x \in I^n$ to the element $x' \in T_n(\mu)$, as in Lemma 5. The distribution of x_t, $1 \leq t \leq n$, is μ. Therefore $\rho(x)(i)$ is the sum of the n i.i.d. $\{0, 1/n\}$-valued random variables $\frac{\mathbb{I}_{x_t = i}}{n}$, $1 \leq t \leq n$. Thus, $\mathbf{E}\rho(x)(i) = \mu(i)$, and the variance of $\rho(x)(i)$ is $\leq \mu(i)/n$. For every $i \in I$ we have $|\rho(x)(i) - \mu(i)| \geq \frac{\varepsilon}{|I|}$ whenever $|\rho(x)(i) - \nu(i)| \geq |\mu(i) - \nu(i)| + \frac{\varepsilon}{|I|}$. Therefore, by using Chebyshev's inequality we have

$$P_{\mu^{\otimes n}}(|\rho(x)(i) - \nu(i)| \geq |\mu(i) - \nu(i)| + \frac{\varepsilon}{|I|}) \leq \frac{\mu(i)|I|^2}{\varepsilon^2 n}$$

and then

$$P_{\mu^{\otimes n}}(\|\rho(\tilde{x}) - \nu\|_1 \geq \|\mu - \nu\|_1 + \varepsilon) \leq \frac{|I|^2}{\varepsilon^2 n}$$

Hence the result, since $\sum_{1\leq t\leq n} \mathbb{I}_{x_t \neq f_t(x)} \leq n\|\rho(x) - \nu\|_1$. □

References

[1] Arrow, K. (1985): "Informational Structure of the Firm," *American Economic Review*, 75, 303–307.

[2] Bárány, I. (1992): "Fair Distribution Protocols or How the Players Replace Fortune," *Mathematics of Operations Research*, 53, 327–340.

[3] Ben Porath, E. (2003): "Cheap Talk in Games with Incomplete Information," *Journal of Economic Theory*, 108, 45–71.

[4] Cover, T. M., and J. A. Thomas (1991): *Elements of Information Theory*, Wiley Series in Telecomunications. Wiley, New York.

[5] Forges, F. (1986): "An Approach to Communication Equilibria," *Econometrica*, 54, 1375–1385.

[6] ——— (1990): "Universal Mechanisms," *Econometrica*, 58, 1341–1364.

[7] Gerardi, D. (2004): "Unmediated Communication in Games with Complete and Incomplete Information," *Journal of Economic Theory*, 114, 104–131.

[8] Hoeffding, W. (1958): "Probability Inequalities for Sums of Bounded Random Variables," *Journal of the American Statistical Association*, 301, 13–30.

[9] Kushilevitz, E., and N. Nisan (1996): *Communication Complexity*. Cambridge University Press, New York.

[10] Marschak, J., and R. Radner (1972): *Economic Theory of Teams*. Yale University Press, New Haven.

[11] Myerson, R. B. (1986): "Multistage Games with Communication," *Econometrica*, 54, 323–358.

[12] Radner, R. (1993): "The Organization of Decentralized Information Processing," *Econometrica*, 61, 1109–1146.

[13] Shannon, C. (1948): "A Mathematical Theory of Communication," *Bell System Technical Journal*, 27, 379–423; 623–656.

[14] Urbano, A., and J. Vila (2002): "Computational Complexity and Communication: Coordination in Two-Player Games," *Econometrica*, 70, 1893–1927.

[15] van Zandt, T. (1999): "Decentralized Information Processing in the Theory of Organizations," in *Economic Design and Behavior*, ed. by M. Sertel, vol. 4, pp. 125–160. MacMillan Press Ltd., London.