האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

OPTIMAL USE OF COMMUNICATION
RESOURCES

by

OLIVIER GOSSNER, PENÉLOPE HERNÁNDEZ
and ABRAHAM NEYMAN

מרכז לחקר הרציונליות

CENTER FOR THE STUDY
OF RATIONALITY

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 with 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.

Classification JEL : 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 extend an open problem, and is the topic of this paper.

Communication equilibria, as proposed by Forges [5] and Myerson [10], 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 [11] tell us that in an organization like a firm, communication is a costly activity and that a significant amount of resources is devoted to it. 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 [13], 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, a new archetype of costly communication is needed.

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 superior information to 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. A repeated game then 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, the forecaster's action, agent's action), called *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 any 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 [12] 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 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.

Our result has implications for the measure of communication inefficiencies, in both in team games and general games.

The set of equilibrium payoffs in our model is in general a proper subset of the set of communication equilibria. This reflects the fact that communication is a costly activity, which consumes part of the player's resources.

The cost of communication inefficiencies can be measured in team games, where the 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 which takes into account the implementation costs of communication processes.

Section 2 presents the model and examples of problems of optimal communication. Section 3 introduces the information constraint and the main results. In Section 4 we prove that using mixed or correlated strategies (instead of pure strategies) will not change the analysis. The

main results are proved in Sections 5 and 6. Section 7 presents applications to team games and to general games, and we conclude with a discussion in Section 8.

2. The model

Given a finite set A, $\Delta(A)$ represents the set of probability measures over A, and $|A|$ is the cardinality of A. Random variables will be 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. 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 all future states of nature. At each stage, the chosen action may depend on past actions, as well as on the whole 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 of states of nature is i.i.d. with stage law μ. Let $(\mathbf{i}_t)_t$ be random variables that represent the sequence of states of nature. 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 $\mathbf{k}_t = \tau_t((\mathbf{i}_1, \ldots, \mathbf{i}_{t-1}), (\mathbf{j}_1, \ldots, \mathbf{j}_{t-1}), (\mathbf{k}_1, \ldots, \mathbf{k}_{t-1}))$. Let $P_{\mu,\sigma,\tau}$ be the induced probability distribution over $(I \times J \times K)^{\mathbb{N}}$, and $P^t_{\mu,\sigma,\tau}$ the marginal over stage t's action triple.

The average distribution up to stage t is $Q^t_{\mu,\sigma,\tau} = \frac{1}{t} \sum_{t'=1}^{t} P^{t'}_{\mu,\sigma,\tau}$. If there is a strategy pair (σ, τ) such that $Q^t_{\mu,\sigma,\tau} \to Q$ as $t \to \infty$ we say

that Q is *implementable* and that the strategy pair σ, τ *implements* the distribution Q.

2.1. **Example: Coordination with nature.** We consider a two-player team game in which both players wish to coordinate with nature. $I = J = K = \{0,1\}$, 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

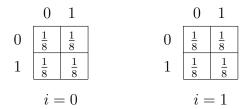|	0	1			0	1
0	1	0		0	0	0
1	0	0		1	0	1

$$i = 0 \qquad\qquad i = 1$$

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 of action triples up to any stage is
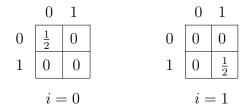
	0	1			0	1
0	$\frac{1}{4}$	$\frac{1}{4}$		0	0	0
1	0	0		1	$\frac{1}{4}$	$\frac{1}{4}$

$$i = 0 \qquad\qquad i = 1$$

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

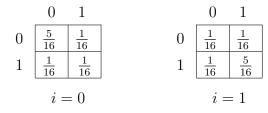Consider now the strategy of the forecaster that matches nature at even stages, and plays the next stage of nature at odd ones. At even stages, the agent is informed of the stage of nature by the previous action of the forecaster, and thus can match it. At odd stages, the agent has no information on the stage 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 is then

	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 $\frac{5}{8}$. A natural question that arises is what is the maximal payoff that corresponds to an implementable distribution.

Our analysis will show that the following distribution

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

$i = 0$

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

$i = 1$

(with corresponding payoff $\frac{4}{5}$) is implementable, while the distribution

	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. Moreover, our analysis enables us to compute the (unique) implementable distribution

that maximizes the corresponding payoff, and describes the implementing strategies. This unique distribution is

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

$$i = 0$$

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

$$i = 1$$

with x satisfying $H(x) + (1 - x)\log_2 3 = 1$, where H is the entropy function.[1] Thus the corresponding payoff equals 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

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

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 with 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

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

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 $H_Q(\mathbf{x}|\mathbf{y})$.

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, considering a triple $(\mathbf{i}, \mathbf{j}, \mathbf{k})$ of random

---

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

variables with joint law Q:

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

From 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 the 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.

Our first result states that, given any pair of strategies (σ, τ), the corresponding average distribution fulfills the information constraint.

Theorem 1. *For every strategy pair (σ, τ) and t, $Q_{\mu, \sigma, \tau}^t$ fulfills the information constraint, and every 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 Q 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.

4. Mixed and correlated strategies

The following result implies that the set of implementable distributions cannot be expanded by considering mixed, or even correlated strategies of the forecaster and the agent.

Theorem 3. *The set of distributions Q that fulfill the information constraint and have a fixed marginal μ on I is convex.*

The theorem follows from the next lemma. Indeed, observe that $H_Q(\mathbf{i})$ is a constant c over all distributions Q having marginal μ on I, and by setting $X = K$ and $Y = I \times J$ in the lemma it follows that the set of all distributions Q on $I \times J \times K$ with $H_Q(\mathbf{i}, \mathbf{j}|\mathbf{k}) \geq c$ is convex.

Lemma 1. *For any finite set 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

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

$\square$

Any pair of correlated strategies is a distribution over pure strategies. Therefore, the average distribution induced by a pair of correlated strategies is a convex combination the average distributions induced by those pure strategies, and hence fulfills, by Theorem 3, the information constraint. We conclude that for every pair of correlated strategies (σ, τ) and every t, $Q_{\mu,\sigma,\tau}^t$ fulfills the information constraint, and thus every distribution that is implementable by a pair of correlated strategies fulfills the information constraint.

On the other hand, Theorem 2 shows that every distribution Q that fulfills the information constraint and has marginal $Q_I = \mu$ is implementable by pure strategies (and thus in particular by correlated strategies).

5. Proof of Theorem 1

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

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

where the first inequality follows from the fact that k_t is a function of $(\mathbf{i}_1,\mathbf{j}_1,\mathbf{k}_1,\ldots\mathbf{i}_{t-1},\mathbf{j}_{t-1},\mathbf{k}_{t-1})$. By Lemma 1,

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

This completes the proof of the first part of the result. The second part 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. Proof of Theorem 2

Given a distribution Q that fulfills the information constraint and that has marginal μ on X, we construct strategies σ, τ of the forecaster and the agent such that the long-run average distribution induced approximates Q. Strategies are defined over blocks of length n. During each block, the forecaster communicates to the agent which sequence of actions to play during the next block. The sequences of actions for the agent are chosen in a subset A_n of K^n called the set of action plans. The property required on the set of action plans is that, given the sequence $\tilde{x}$ of actions of nature during a block, there exists an element

$\tilde{z}$ in A_n such that the empirical distribution of $(\tilde{x}, \tilde{z})$ during a block is close to $Q_{I \times K}$. The possible set of messages $M_n(\tilde{x}, \tilde{z})$ for the forecaster during a block corresponds to sequences of actions $\tilde{y}$, such that $(\tilde{x}, \tilde{y}, \tilde{z})$ has an empirical distribution close to Q.

Our proof relies on estimates of the sizes of the sets of messages and the set of action plans needed for the forecaster and the agent. The key to the proof is that the information constraint implies that the set of messages has cardinality larger than the set of action plans.

6.1. **Typical sequences.** Given a finite sequence $b = (b_1, \ldots, b_n) \in B^n$ over a finite alphabet B, the *type* $\rho(b)$ of b is its empirical distribution (i.e., $\rho(b)(c) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{I}_{b_i = c}$, for $c \in B$). Given $\mu \in \Delta(B)$, the *type set* of μ is $T^\mu(n) = \{b \in B^n, \rho(b) = \mu\}$. The set of types is $\mathbb{T}_n(B) = \{\mu \in \Delta(B), T_n(\mu) \neq \emptyset\}$. The number of types is $|\mathbb{T}_n(B)| = \binom{n + |B| - 1}{|B| - 1} \leq n^{|B|}$. The following estimates the size of a type set $\mu \in \mathbb{T}_n(B)$ (see, for instance, Cover and Thomas [4, Theorem 12.1.3 page 282]):

$$(3) \qquad \frac{2^{nH(\mu)}}{(n+1)^{|B|}} \leq |T^\mu(n)| \leq 2^{nH(\mu)}$$

Notation 1. *Given two functions $f, g : \mathbb{N} \to \mathbb{R}_+$ we denote by $f \doteq g$ the relation $\lim_{n \to \infty} \frac{\log f(n)}{\log g(n)} = 1$.*

The inequalities (3) imply that for every sequence $\mu_n \in \mathbb{T}_n(B)$ we have

$$|T^{\mu_n}(n)| \doteq 2^{nH(\mu_n)}.$$

Let A and B be two finite sets and $Q \in \Delta(A \times B)$, with marginal distributions Q_A and Q_B on A and B, respectively. Fix $n \in \mathbb{N}$ and $Q \in \Delta(A \times B)$ such that $T^Q(n)$ is nonempty. It follows that the sets $T^{Q_A}(n)$ and $T^{Q_B}(n)$ are also nonempty.

Notation 2. *Given sequences $\tilde{a} = (a_1, \ldots, a_n) \in A^n$, $\tilde{b} = (b_1, \ldots, b_n) \in B^n$, and $\tilde{c} = (c_1, \ldots, c_n) \in C^n$, $(\tilde{a}, \tilde{b})$ and $(\tilde{a}, \tilde{b}, \tilde{c})$ represent the points $(a_1, b_1, \ldots, a_n, b_n)$ in $(A \times B)^n$ and $(a_1, b_1, c_1 \ldots, a_n, b_n, c_n)$ in $(A \times B \times C)^n$.*

Notation 3. *Given a probability measure Q over a product set $A \times B$, Q_A and Q_B represent the marginals of Q on A and B, respectively.*

Lemma 2. *For every $\tilde{a} \in T^{Q_A}(n)$, we have*

$$|A|^{|A \times B|} \frac{2^{(H(Q)-H(Q_A))n}}{(n+1)^{|A \times B|}} \leq |\{\tilde{b} \in B^n : (\tilde{a}, \tilde{b}) \in T^Q(n)\}| \leq 2^{(H(Q)-H(Q_A))n}$$

and thus in particular

$$|\{\tilde{b} \in B^n : \rho(\tilde{a}, \tilde{b}) = Q\}| \doteq 2^{(H(Q)-H(Q_A))n}.$$

Proof. The point $\tilde{a} \in A^n$ partitions the set $\{1, \ldots, n\}$ into $|A|$ disjoint subsets N_a, $a \in A$: $N_a = \{1 \leq i \leq n : x_i = a\}$. 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 $\tilde{b} \in B^n$ and a subset N of $\{1, 2, \ldots, n\}$ we denote by $(\tilde{b}|N)$ the N-vector $(\tilde{b}_j)_{j \in N}$. Note that for every $\tilde{b} \in B^n$ we have $\rho(\tilde{a}, \tilde{b}) = Q$ if and only if for every $a \in A$ we have $\rho(\tilde{b}|N_a) = Q^a$. Therefore, it follows from (3) that

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

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

6.2. **Set of action plans.** We now prove for every Q the existence of a sequence of subsets A_n of K^n of size

$$|A_n| \doteq 2^{(H(Q_I)+H(Q_K)-H(Q))n}$$

such that for every $\tilde{x} \in T^{Q_I}(n)$, there exists $\tilde{z} \in A_n$ with $(\tilde{x}, \tilde{z}) \in T^Q(n)$.

Lemma 3. *For $Q \in \mathbb{T}_n(I \times K)$, there exists a subset A_n of $T^{Q_K}(n)$ of size less than or equal to $1 + H(Q_I)n(n+1)^{|I \times K|}2^{(H(Q_I)+H(Q_K)-H(Q))n}$ such that, for every $\tilde{x} \in T^{Q_I}(n)$ there exists $\tilde{z} \in A_n$ with $(\tilde{x}, \tilde{z}) \in T^Q(n)$.*

Proof. Let $Q \in \mathbb{T}_n(I \times K)$. If $H(Q_I) = 0$ we set $A_n = \{\tilde{z}\}$ where $\tilde{z}$ is an arbitrary element of $T^{Q_K}(n)$. We now assume $H(Q_I) > 0$.

Let $(\mathbf{Z}_k)_{k \geq 1}$ be a sequence of $T^{Q_K}(n)$-valued i.i.d. random variables where $\mathbf{Z}_{k,1}$ is uniformly distributed on $T^{Q_K}(n)$. Let $\mathbf{A}_n$ be a random

subset of $T^{Q_K}(n)$ of size

$$\alpha(n) = \lceil H(Q_I)n(n+1)^{|I \times K|}2^{(H(Q_I)+H(Q_K)-H(Q))n}\rceil$$

containing $\{\mathbf{Z}_1, \ldots, \mathbf{Z}_{\alpha(n)}\}$. Denote by P the induced probability over realizations of the $\mathbf{Z}_k$'s. By Lemma 2, the number of elements $\tilde{z} \in T^{Q_K}(n)$ such that $(\tilde{x}, \tilde{z}) \in T^Q(n)$ is no less than $\frac{2^{(H(Q)-H(Q_I))n}}{(n+1)^{|I \times K|}}$. By equation (3), $|T^{Q_K}(n)| \leq 2^{H(Q_K)n}$. Therefore, for any $\tilde{x} \in T^{Q_I}(n)$ and $1 \leq k \leq \alpha(n)$

$$P((x, \mathbf{Z}_k) \in T^Q(n)) \geq \frac{2^{(H(Q_I)+H(Q_K)-H(Q))n}}{(n+1)^{|I \times K|}}$$

¿From this, we deduce that

$$
\begin{aligned}
P(\forall \tilde{z} \in \mathbf{A}_n, (\tilde{x}, \tilde{z}) \notin T^Q(n)) &\leq P(\forall\, 1 \leq k \leq \alpha(n), (\tilde{x}, \mathbf{Z}_k) \notin T^Q(n)) \\
&\leq \left(1 - \frac{2^{(H(Q_I)+H(Q_K)-H(Q))n}}{(n+1)^{|I \times K|}}\right)^{\alpha(n)} \\
&\leq \exp(-nH(Q_I))
\end{aligned}
$$

Hence, the expected number of $\tilde{x} \in T^{Q_I}(n)$ such that $\forall \tilde{z} \in \mathbf{A}_n, (\tilde{x}, \tilde{z}) \notin T^Q(n)$ is at most $\exp(-nH(Q_I))|T^{Q_I}(n)|$, which by equation (3) is bounded by $e^{-nH(Q_I)}2^{nH(Q_I)} < 1$. Therefore, there exists a realization A_n of the random set $\mathbf{A}_n$ that verifies the condition. $\square$

6.3. **Approximation of probabilities.**

Lemma 4. $\forall \varepsilon > 0 \; \exists N(\varepsilon)$ such that $\forall \tilde{Q} \in \Delta(I \times J \times K) \; \forall n \geq N(\varepsilon)$ $\exists Q \in \mathbb{T}_n(I \times J \times K)$ such that

$$(4) \qquad H_Q(\mathbf{i}, \mathbf{j}|\mathbf{k}) - H_Q(\mathbf{i}) \geq (1 - 3\varepsilon)(H_{\tilde{Q}}(\mathbf{i}, \mathbf{j}|\mathbf{k}) - H_{\tilde{Q}}(\mathbf{i})) + \varepsilon$$

and

$$(5) \qquad \|Q - \tilde{Q}\|_1 < 7\varepsilon$$

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 $R_I^\varepsilon = \tilde{Q}_I$, and the inequality $\log |J| \geq 1$ (which follows from $|J| \geq 2$), we have

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

which implies

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

Let $Q \in \Delta(I \times A)$ with $T^Q(n) \neq \emptyset$ 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

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

which proves (4). In addition

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

$\square$

Lemma 5. *Fix $\nu \in \Delta(I)$ and n such that $T^\nu(n) \neq \emptyset$. For every $\tilde{x} = (x_1, \ldots, x_n) \in I^n$ there is $\tilde{x}' = (x'_1, \ldots, x'_n) \in T^\nu(n)$ such that*

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

Proof. By induction on the integer $d(\tilde{x}) := n\|\rho(\tilde{x}) - \nu\|_1$. If $d(\tilde{x}) = 0$ set $\tilde{x}' = \tilde{x}$. Assume that $d(\tilde{x}) > 0$. There exist elements $i, i' \in I$ such that $\rho(\tilde{x})(i) > \nu(i)$ and $\rho(\tilde{x})(i') < \nu(i')$. 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 $\tilde{x}' \in T_n(\nu)$ such that $|\{t : x''_t \neq x'_t\}| \leq d(x'')$ and therefore $|\{t : x'_t \neq x_t\}| \leq d(x'') + 2 = d(x)$. $\square$

Corollary 1. *Fix $\nu \in \Delta(I)$ and n such that $T^\nu(n) \neq \emptyset$. There exists a map $f \colon I^n \to T^\nu(n)$ such that for $\mu \in \Delta(I)$ we have*

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

Proof. Let $f : I^n \to T^\nu(n)$ be the function that maps $\tilde{x} \in I^n$ to the element $\tilde{x}' \in T^\nu(n)$, as in Lemma 5. For every $i \in I$ we have $|\rho(\tilde{x})(i) - \mu(i)| \geq \frac{\varepsilon}{|I|}$ whenever $|\rho(\tilde{x})(i) - \nu(i)| \geq |\mu(i) - \nu(i)| + \frac{\varepsilon}{|I|}$. Therefore, by using Chebyshev's inequality we have

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

and then

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

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

6.4. **Construction of the strategies.** Fix $Q \in \Delta(I \times J \times I)$ that satisfies the conditions of Theorem 2. Note that it suffices to prove that for every $\varepsilon > 0$ there exists a strategy profile σ, τ and $t(\varepsilon, \sigma, \tau)$ such that

$$\forall\, t \geq t(\varepsilon, \sigma, \tau) \qquad \|Q_{\mu,\sigma,\tau}^t - Q\|_1 < \varepsilon$$

Indeed, from a sequence of strategy profiles that approximate Q closer and closer, one can construct a strategy that implements Q.

Let $Q' \in \Delta(I \times J \times K)$ be as in Lemma 4 such that (4) and (5) hold. By assumption we have $H_Q(\mathbf{i}, \mathbf{j}|\mathbf{k}) \geq H_Q(\mathbf{i})$, and therefore we deduce from (4) that

$$(6) \qquad\qquad H_{Q'}(\mathbf{i}, \mathbf{j}|\mathbf{k}) \geq H_{Q'}(\mathbf{i}) + \varepsilon$$

By Lemma 3 there exists a set of action plans $A_n \subset T^{Q'_K}(n)$ for the agent such that

$$\begin{aligned} |A_n| &\leq\; 1 + H(Q'_I)n(n+1)^{|I \times K|} 2^{(H(Q'_I) + H(Q'_K) - H(Q'))n} \\ &\doteq\; 2^{(H(Q'_I) + H(Q'_K) - H(Q'))n} \end{aligned}$$

and such that for every $\tilde{x} \in T^{Q'_I}(n)$, there exists an action plan $\tilde{z} \in A_n$ of the agent such that $(\tilde{x}, \tilde{z}) \in T^{Q'_{I \times K}}(n)$. For every $(\tilde{x}, \tilde{z}) \in T^{Q'_{I \times K}}(n)$

we set

$$M(\tilde{x}, \tilde{z}) = \{\tilde{y} \in T^{Q'_J}(n) : (\tilde{x}, \tilde{y}, \tilde{z}) \in T^{Q'}(n)\}$$

By Lemma 2 we have

$$|M(\tilde{x}, \tilde{z})| \geq \frac{2^{(H(Q') - H(Q'_{I \times K}))n}}{(n+1)^{|I \times J \times K|}} \doteq 2^{H_{Q'}(\mathbf{j}|\mathbf{i}, \mathbf{k})n}$$

Since by (4) we have $H_{Q'}(\mathbf{j}|\mathbf{i}, \mathbf{k}) > H(Q'_I) + H(Q'_K) - H(Q') + \varepsilon$, for n sufficiently large $|M(\tilde{x}, \tilde{z})| \geq |A_n|$. Therefore there exist maps $m_{\tilde{x}, \tilde{z}}$ from $M(\tilde{x}, \tilde{z})$ onto A_n. In what follows, $m_{\tilde{x}, \tilde{z}}^{-1}$ stands for a function from A_n into $M(\tilde{x}, \tilde{z})$ such that $m_{\tilde{x}, \tilde{z}} \circ m_{\tilde{x}, \tilde{z}}^{-1}$ is the identity on A_n.

By Corollary 1, for n sufficiently large (e.g., $n \geq |I|^2/\varepsilon^3$), there exists a projection $f \colon I^n \to T^{Q'_I}(n)$ such that

$$P_{\mu^{\otimes \mathbb{N}}}(\frac{1}{n} \sum_{t=1}^{n} \mathbb{1}_{x_t \neq f_t(\tilde{x})}) > 8\varepsilon) < \varepsilon$$

Let $r \colon I^n \to A_n$ such that $(f(\tilde{x}), r(\tilde{x})) \in T^{Q'_{I \times K}}(n)$ for every $x \in I^n$.

Fix a point $(\tilde{x}(0), \tilde{z}(0)) \in T^{Q'_{I \times K}}(n)$ and for $b > 1$ define $(\tilde{x}(b), \tilde{y}(b), \tilde{z}(b))$ to be the coordinates of the play at stages $bn + 1, \ldots, bn + n$, namely,

$$(\tilde{x}(b), \tilde{y}(b), \tilde{z}(b)) = (x_{bn+1}, y_{bn+1}, z_{bn+1}, \ldots, x_{bn+n}, y_{bn+n}, z_{bn+n})$$

The strategy σ of the agent plays at stages $t = 1, \ldots, n$ the sequence of actions $z(0)$. At stages $t = (b+1)n+1, \ldots, (b+2)n$ the agent plays the sequence $z(b + 1) = m_{f(\tilde{x}(b)), \tilde{z}(b)}(\tilde{y}(b))$ if $(f(\tilde{x}(j)), \tilde{y}(j), \tilde{z}(j)) \in T^{Q'}(n)$ and $z(0)$ otherwise. The strategy τ of the forecaster plays at stages $t = 1, \ldots, n$ the sequence of actions $m_{\tilde{x}(0), \tilde{z}(0)}^{-1}(r(\tilde{x}(1)))$, and at stages $t = bn + 1, \ldots, bn + n$ the sequence $m_{f(\tilde{x}(b)), \tilde{z}(b)}^{-1}(r(\tilde{x}(b + 1)))$.

It follows that for every $b \geq 1$ we have $(f(\tilde{x}(b)), \tilde{y}(b), \tilde{z}(b)) \in T^{Q'}(n)$. Hence

$$\|\rho(\tilde{x}(b), \tilde{y}(b), \tilde{z}(b)) - Q'\|_1 \leq \|\rho(\tilde{x}(b)) - Q'_I\|$$

Therefore,

$$P_{\mu^{\otimes n}}(\|\rho(\tilde{x}(b), \tilde{y}(b), \tilde{z}(b)) - Q'\|_1 \geq 8\varepsilon) \leq \varepsilon$$

and as $\|\rho(\tilde{x}(b), \tilde{y}(b), \tilde{z}(b)) - Q'\|_1 \leq 2$ we have

$$\|\mathbf{E}_{\mu^{\otimes n}}(\rho(\tilde{x}(b), \tilde{y}(b), \tilde{z}(b)) - Q'\|_1 \leq 10\varepsilon$$

Hence, for sufficiently large t

$$\|Q^t_{\mu,\sigma,\tau} - Q'\|_1 \leq 11\varepsilon$$

which implies

$$\|Q^t_{\mu,\sigma,\tau} - Q\|_1 \leq 18\varepsilon$$

This ends the proof of Theorem 2.

7. PAYOFFS AND EQUILIBRIA

We show in this section how the information constraint yields characterizations of 1) the set of feasible payoffs, 2) the best payoff a team can achieve, and 3) the set of equilibrium payoffs in repeated games.

7.1. **Feasible payoffs.** Consider a fixed payoff function g defined on action triples ($g : I \times J \times K \to \mathbb{R}^2$) and define the set F by

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

The objective of this section is to demonstrate that the set F is a good approximation for the set of feasible payoffs of the discounted games for large discount factors. The approximation applies not only to interstage-time-independent discounting (i.e., a fixed discount factor), but also for interstage-time-dependent discounting.

The commonly used discounting valuation is obtained by specifying an interstage-time-independent discount factor $0 < \lambda < 1$ and evaluating a stream $(g_t)_t$ of payoffs according to its weighted average $\sum_{t=1}^{\infty} \theta_t g_t$, where $\theta_t = (1-\lambda)\lambda^{t-1}$. The factor $1-\lambda$ is a normalization factor making the sum of θ_t equal 1. Interstage-time-dependent discount factors lead to a weighted average valuation $\sum_{t=1}^{\infty} \theta_t g_t$, where $\theta = (\theta_t)_t$ is a nonincreasing sequence with $\sum_{t=1}^{\infty} \theta_t = 1$.

If $\theta = (\theta_t)_t$ is a nonincreasing sequence of nonnegative numbers summing to 1, then letting $Q = \sum_{t=1}^{\infty} \theta_t P^t_{\sigma,\tau}$, the expectation of the $(\theta_t)_t$-weighted sum of the stage payoffs is

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

As $\sum_{t=1}^{\infty} \theta_t P_{\sigma,\tau}^t = \sum_{t=1}^{\infty} (\theta_t - \theta_{t+1}) t Q_{\sigma,\tau}^t$ and $\sum_{t=1}^{\infty} (\theta_t - \theta_{t+1}) t = 1$ we deduce that Q is a convex combination of $Q_{\sigma,\tau}^t$ and thus obeys the information constraint, and obviously 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 1. *For every nonincreasing sequence $\theta = (\theta_t)$ with $\sum_{t=1}^{\infty} \theta_t = 1$, the set of θ-weighted feasible payoffs,*

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

is a subset of F. In particular, for every $0 < \lambda < 1$, if F_λ is the set of feasible payoffs of the λ-discounted game, i.e., $F_\lambda = F_{\theta(\lambda)}$ where $(\theta(\lambda))_t = (1-\lambda)\lambda^{t-1}$, then

$$F_\lambda \subseteq F$$

On the other hand, if Q is implementable there exists a strategy pair $(\sigma, \tau))$ such that for every $\varepsilon > 0$ there exists N sufficiently large so that $\|Q_{\sigma,\tau}^n - Q\| < \varepsilon$ for every $n \geq N$. Therefore, if $\theta = (\theta_t)_t$ is a nonincreasing sequence summing to 1 and $Q_{\sigma,\tau}^\theta := \sum_t \theta_t P_{\sigma,\tau}^t$, then $\|Q_{\sigma,\tau}^\theta - Q\| = \|\sum_{t=1}^{\infty} (\theta_t - \theta_{t+1}) t (Q_{\sigma,\tau}^t - Q)\| < 2N\theta_1 + \varepsilon$ (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 $Q_{\sigma,\tau}^\theta$ is within 2ε of the distribution Q. Therefore,

Proposition 2. *For every point $(x, y) \in F$ there exists a strategy pair (σ, τ) such that $E_{Q_{\sigma,\tau}^\theta} g(i, j, k)$ converges to (x, y) as θ_1 goes to 0. Thus, F_θ converges (in the Hausdorff metric) to F as θ_1 goes to 0, in particular, $F_\lambda \to F$ as $\lambda \to 1$.*

7.2. **Team games.** Team games, in which players' preferences are identical, form an adequate setup for the study of inefficiencies due the asymmetric information and the communication costs. As shown for instance by Marschak and Radner [9] and by Arrow [1], a firm can be described as a team when one focuses on the question of information transmission between its members.

In team games, our model allows us to measure the inefficiencies arising from the need to transmit information. As a benchmark, consider the situation in which the agent also has complete information about the states of nature. Then, it is possible for the forecaster and the agent to choose optimally an action pair at each stage given the current state of nature. The corresponding expected payoff is the best achievable 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.

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.

Therefore, the problem of finding a rule for the team that maximizes the long-run expected payoff and of computing the exact value that can be achieved is a difficult one. Yet it is made particularly simple by an approach through the information constraint.

The long-run average payoff is the expectation of the stage payoff where the expectation is with respect to the long-run average of action triples. Hence, finding the maximal long-run payoff amounts to finding the implementable distribution Q that maximizes expected payoff.

Let v_λ be the maximum payoff for the team when the discount factor is λ, i.e., $v_\lambda = \max\{x : (x,x) \in F_\lambda\}$, and let $v = \max\{x : (x,x) \in F\}$. Propositions 1 and 2 imply:

Proposition 3. *When λ goes to 1, v_λ goes to v.*

We are thus able to characterize the best achievable 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 [14] 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.

7.2.1. *Example: bounded communication.* The agent makes a decision at each stage in K, and the payoff to the team depends on the state of nature and on the agent's action. The forecaster has incentives to send the maximal information to the agent. Depending on μ and on the size of J, it may or may not be 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.

As an example, 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 states of nature follow an i.i.d. and uniform process. Payoffs of the team 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

Let us illustrate the use of the information constraint in computing the maximal payoff that the team of the forecaster and the agent can approach. 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 $\mathbf{E}_Q g(\mathbf{i}, \mathbf{j}, \mathbf{k}) = \mathbf{E}_{\hat{Q}} g(\mathbf{i}, \mathbf{j}, \mathbf{k})$ 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 $Q(i=k=1) = Q(i=k=2) = Q(i=k=3)$, and let x be this common value. Given these equalities, 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 $i' \neq k \neq i$, hence when $Q(\mathbf{i}=i|\mathbf{k}=k) = (\frac{1}{3} - x)/2$ for $i \neq k$. Therefore, we can assume w.l.o.g. that $Q_{I \times K}$ 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 $H(x) + (1 - x)/2 = \log \frac{3}{2}$.

Numerically, $v \sim 0.8812$. This has to be compared with the maximal payoff of $\frac{1}{3}$ when the agent is uninformed, and with the payoff of 1 when the agent is fully informed.

7.3. **General games.** We now consider general payoff functions g^f, g^a, and let $g = (g^f, g^a)$.

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

Define the set of "silent" feasible payoffs as

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

where co stands for the convex hull. At the other extreme, the set of feasible payoffs under full communication is

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

Finally, define the set of feasible payoffs with internal communication as F; recall that

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

We have the obvious inclusions

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

The set of distributions Q that verifies the information constraint is convex by Theorem 3 and closed, and obviously the set of distributions Q with $Q_I = \mu$ is convex and closed. Therefore the set of feasible payoffs with internal communication F is a closed convex subset of $\mathbb{R}^2$. The closed convex set F 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 game with team games where the common payoff function $g(i, j, k)$ equals the inner product $\langle x, (g^f(i, j, k), g^a(i, j, k)) \rangle$. Therefore, solving for the feasible set F amounts to solving a family of (two-person)^2 team games.

The individually rational level of a player is defined as the best payoff that this player can defend against every strategy of the other player. 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

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

---

[2]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 payoffs 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. The better-informed player possesses a strategic advantage over the less-informed one. Let IR be the set of individually rational payoffs:

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

The set $F^S \cap IR$ corresponds to the set of equilibrium payoffs of games with large discount factors 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 large discount factors. In this case, there are no restrictions nor costs associated with the communication possibilities.

Finally, $F \cap IR$ is the set of equilibrium payoffs of our original game with large discount factors, where all communication is internal to the game. This set is convex, but is in general not a polyhedron. It is computed directly from the information constraint, and reflects the costs of communication between the players.[3]

7.3.1. *Example: Secret Cournot collusion.* Consider a production game in which the forecaster and the agent form a Cournot duopoly, and choose production levels q^p, q^a in finite sets Q^p and Q^a at unit costs c^p, c^a. A state of nature is a pair of positive numbers (A, B). The market inverse demand function is $p = A - B(q^p + q^a)$, and profits are $g^p = p(q^p - c^p)$, $g^a = p(q^a - c^a)$. Since the forecaster has better

---

[3]In the last two cases, all information sent by the forecaster concerning future states of nature is eventually verifiable by the agent. Therefore the truthtelling incentive constraints are not active.

information about future market demand, it may be profitable to share part of this information with the agent. All explicit communication between the firms (through phone lines for instance) is prohibited by law. Nevertheless, nothing prevents the forecaster from transmitting information through his choices of production levels.

8. Discussion and extensions

In order to preserve maximum transparency, we have tried to keep the model of Section 2, henceforth, the basic model, 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 examine the impact of the signaling structure of the one-shot game on the the set of implementable distributions. Second, we discuss relaxations of the perfect and infinite forecast assumption. Third, we discuss information lags. Fourth, 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.

8.1. **Signalling structures.** In the basic model, we have assumed that the stage game has perfect monitoring in the sense that both the forecaster and the agent were 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 more informative signals are also implementable with less informative ones. This follows from the fact that strategies with less informative signals can be mimicked by (mixed) strategies when signals are more informative.

We now 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.

8.1.1. *Unobservable agent's actions.* We discuss here the situation where 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, the forecaster can reconstitute from these the agent's actions. 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 3 of the limit Pareto payoff to the team for a discount factor λ arbitrarily close to 1 remains unchanged.

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

8.1.2. *Noisy signal of forecaster's action.* We contemplate 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 $I \times J \times K \times S$ whose marginal on $I \times J \times K$ equals Q and $\hat{Q}(s|i, j, k)$ equals the probability of the signal s given action triple i, j, k, namely, $\hat{Q}(s|i, j, k) = R_{i,j,k}(s)$. Using this notation, Q is implementable if and only if $Q_I = \mu$ and

$$(7) \qquad H_{\hat{Q}}(\mathbf{s}|\mathbf{i}, \mathbf{k}) - H_{\hat{Q}}(\mathbf{s}|\mathbf{i}, \mathbf{j}, \mathbf{k}) \geq H_Q(\mathbf{i}) - H_Q(\mathbf{i}|\mathbf{k})$$

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 (7) particularizes to the information constraint $H_Q(\mathbf{j}|\mathbf{i}, \mathbf{k}) \geq I(\mathbf{i}; \mathbf{k})$.

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 (7) for a distribution Q such that $Q(i = k) = 1$ becomes $H(\mathbf{s}) - H(\mathbf{s}|\mathbf{j}) = I(\mathbf{s}; \mathbf{j}) \geq H(\mathbf{i})$. 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 that 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 (7) for a distribution $Q \in \Delta(S) \times \Delta(I \times K)$ reduces to $I_{\hat{Q}}(\mathbf{s}; \mathbf{j}) \geq I_Q(\mathbf{i}; \mathbf{k})$. 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 (7) 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 $I_{\hat{Q}}(\mathbf{s}; \mathbf{j}) \geq R(D)$. In addition, 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]).

8.1.3. *Unobservable current state of nature.* Consider now 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 the one 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 forecaster's action j, we proceed as in Section 8.1.2. 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

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

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

8.2.1. *Finite forecasts.* For $n \in \mathbb{N}$, we 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.

8.2.2. *Imperfect forecasts.* We now discuss the case where the forecaster may be imperfectly informed of the states of nature. Let S be a set of signals, and let R be a transition probability from I to S. Assume that before the game starts, the forecaster observes a sequence of signal $(s_t)_t$ 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 is the identity matrix).

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

8.3. **Information lags.** We now introduce possibilities of delays in the observation by the agent of the state of nature and of the action of the forecaster. For $f \in \mathbb{N}$, we say that the agent has f-*delayed perfect monitoring* if the agent observes (i_t, j_t) just before the play at stage $t + f + 1$. (Thus 0-delayed perfect monitoring corresponds to perfect monitoring). Note that if the agent has f-delayed perfect monitoring then a strategy of the agent specifies the agent's action k_{t+f+1} at stage t as a function of $(i_1, j_1, k_1, \ldots, i_t, j_t, k_t, k_{t+1}, \ldots, k_{t+f})$.

Our proof shows that the set of implementable distributions of this variant with f-perfect monitoring coincides with the set of implementable distributions of the basic model.

More generally, given a function $f : \mathbb{N} \to \mathbb{N}$ we say that the agent has f-*delayed perfect monitoring* if the agent observes the (i_t, j_t) just before the play at stage $t + f(t) + 1$. Our proof shows, in addition, that if $f(t) = o(t)$ then the set of implementable distributions of the variant of the basic where the agent has f-perfect monitoring coincides with the set of implementable distributions of the basic model.

8.4. **State of nature processes.** In a forthcoming paper we analyze the variant of the basic model where the states of nature follow a Markov chain. In that case 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

$$(8) \qquad\qquad H_Q(\mathbf{i}, \mathbf{j} \mid \mathbf{k}, \mathbf{i}') \geq H_Q(\mathbf{i} \mid \mathbf{i}')$$

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

We now 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) $Q'(\mathbf{j} = j, \mathbf{k} = k|\mathbf{i} = i, \mathbf{i}' = i') = Q(\mathbf{j} = j, k|\mathbf{i} = i)$.

We now 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, for instance, from the chain rule of entropies. The obtained inequality is then equivalent to the information constraint for Markov chains (8) applied to Q.

The information constraint is satisfied in the Markov case whenever it is in the i.i.d. case.

On the other hand, consider the Markov chain with state $I = \{0, 1\}$ moving from state i to state $1 - i$, i.e., alternates between the states 0 and 1. The invariant distribution is $1/2, 1/2$. If $J = K = I$ then the distribution Q on $I \times J \times K$ with $Q_I(0) = 1/2$ and $Q(i = j = k) = 1$ is implementable. However, it does not satisfy the information constraint of the basic model.

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.

In the Markov chain case, the distribution of i_t given the sequence of past states $i_1, \ldots, i_{t-1}$ 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 $Q_{\sigma,\tau}^{i,n}$ 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 $Q_{\sigma,\tau}^{i,n}$ 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 $Q_{\sigma,\tau}^{i,n}$ 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.

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

Our characterization of the implementable distribution in the Markov chain process implies however additional implementable distributions. For every 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 (8) implies that Q is implementable if and only if there are distribution $Q_i \in \Delta(I \times J \times K)$ with marginals $((Q_i)_I =) \nu_i$ 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 case and the i.i.d. highlight the need for the forecaster signaling at stages $t \in \mathbb{N}_i$ on states of nature in stages $t \in \mathbb{N}_{i'}$ where $i \neq i'$.

8.5. **Robustness.** We now prove that our characterization of implementable distributions is robust to small departures from the assumptions made in the basic model on the state of nature process, the foresight ability of the forecaster, and the monitoring and forecasting possibilities of the agent. Notice that the analysis of the previous extensions demonstrate (indirectly) robustness to some specific departures in each of these assumptions separately. We wish here to show robustness when all assumptions are perturbed together and to allow for a wide variety of perturbations. In order to do this, we introduce a general model that extends the basic model in several directions: the process of states of nature is arbitrary, both players (not only the forecaster) receive signals about future states of nature, the signalling structure of the stage game is arbitrary, and stage signals depend also on future states of nature.

We first describe the dynamics of 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 p ($p = 1, 2$ in the two-player game) observes a random signal s_0^p whose distribution depends on the infinite 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 p observes the realization of a random signal[4] s_t^p where the distribution of (s_t^1, s_t^2) depends on the triple (ω, j_t, k_t) of the infinite 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 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 $s_0^2, \ldots, s_{t-1}^2$.

---

[4]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 of nature.

We now introduce the notions that allow us to describe instances of the general model whose implementable distributions are close to those of the basic model. 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

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

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

In the basic model, the forecaster has perfect and infinite forecast. As an approximation of this assumption, we say that the forecaster has δ-*perfect forecast* 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

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

We approximate the assumption of perfect monitoring of the action triple by the agent by saying that 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

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

To define closeness to the hypothesis of absence of forecasting abilities for the agent, say that the agent has δ-*no-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[5] 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 forecast, the agent has δ-perfect monitoring and δ-no-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 4 (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 P is i.i.d. $s_0^1(\omega) = \omega$, $s_0^2(\omega)$ 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 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).

Remark 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 of nature 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.

8.6. **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

---

[5]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.

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 a case, 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.,[8]). The characterization of the set of implementable distributions in this model is left as an open problem for future research.

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.

[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*, forthcoming.

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

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

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

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

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

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

[14] 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.

CERAS, UMR CNRS 2036

Universidad de Alicante

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