

Convergence and Asymptotic Agreement in Distributed Decision Problems

JOHN N. TSITSIKLIS AND MICHAEL ATHANS, FELLOW, IEEE

Abstract—We consider a distributed team decision problem in which different agents obtain from the environment different stochastic measurements, possibly at different random times, related to the same uncertain random vector. Each agent has the same objective function and prior probability distribution. We assume that each agent can compute an optimal tentative decision based upon his own observation and that these tentative decisions are communicated and received, possibly at random times, by a subset of other agents. Conditions for asymptotic convergence of each agent's decision sequence and asymptotic agreement of all agents' decisions are derived.

I. INTRODUCTION

CONSIDER the following situation. A set $\{1, \dots, N\}$ of N agents possessing a common model of the world (same prior probabilities) and having the same cost function (common objective) want to make an optimal decision. Each agent bases his decision on a set of observations he has obtained and we allow these observations to be different for each agent. Given this setting, the decisions of the agents will be generally different. Aumann [4] has shown, however, that agreement is guaranteed in the following particular case. If the decision to be made is the evaluation of the posterior probability of some event and if all agents' posteriors are common knowledge, then all agents agree. (In Aumann's terminology, common knowledge of an event means that all agents know it, all agents know that all agents know it, and so on, ad infinitum.)

The situation where each agent's posterior is common knowledge is very unlikely, in general. On the other hand, if agreement is to be guaranteed, posteriors have to be common knowledge. The problem then becomes how to reach a state of agreement where decisions are common knowledge, starting from an initial state of disagreement.

Geanakoplos and Polemarchakis [7] and Borkar and Varaiya [6] gave the following natural solution to the above problem. Namely, agents start communicating to each other their tentative posteriors (or, in the formulation of [6], the conditional expectation of a fixed random variable) and then update their own posterior, taking into account the new information they have received. In the limit, each person's posterior converges (by the martingale convergence theorem) and assuming that "enough" communications have taken place, they all have to converge to a common limit.

The above results hold even when each agent obtains additional raw observations during the adjustment process and when the history of communications is itself random. Similar results were also proved for a detection problem [6].

Manuscript received March 15, 1982; revised August 13, 1982 and February 3, 1983. Paper recommended by T. Basar, Chairman of the Large Scale Systems, Differential Games Committee. This work was supported by the Office of Naval Research under Contract ONR/N00014-77-C-0532 (NR 041-519).

The authors are with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139

A related—and much more general situation—is the subject of this paper; we assume that the agents are not just interested in obtaining an optimal estimate or a likelihood ratio, but their objective is to try to minimize some common cost function, given the available information. (Clearly, if each agent has a different cost function, no agreement is possible even if each agent had identical information.) In this setting, we assume that agents communicate to each other tentative decisions (which initially will be different). That is, at any time, an agent computes an optimal decision given the information he possesses and communicates it to other agents. Whenever an agent receives such a message from another agent, his information essentially increases and he will, in general, update his own tentative decision, and so on. In the sequel we prove that the qualitative results obtained in [6], [7] for the estimation problem (convergence and asymptotic agreement) are also valid for the decision-making problem for several, quite general, choices of the structure of the cost function. However, tentative decisions do not form a martingale sequence and a substantially different mathematical approach is required for the proofs. We point out that estimation problems are a special case of the decision problems studied in this paper, being equivalent to the minimization of the mean square error.

A drawback of the above setting is that each agent is assumed to have an infinite memory. We have implicitly assumed that the knowledge of an agent can only increase with time and, therefore, he has to remember the entire sequence of messages he has received in the past. There is also the implicit assumption that if an agent receives additional raw data from the environment, while the communication process is going on, these data are remembered forever. These assumptions are undesirable, especially if the agents are supposed to model humans, because limited memory is a fundamental component of the bounded rationality behavior of human decision makers [14]. We will therefore relax the infinite memory assumption and allow the agents to forget any portion of their past knowledge. We only constrain them to remember their most recent decision and the most recent message (tentative decision) coming from another agent. (For a particular class of communication protocols, we even allow them to forget their most recent decision.) We then obtain convergence results similar to those obtained for the unbounded memory model, although in a slightly weaker sense.

A particular problem of interest is one in which all random variables are jointly Gaussian and the cost is a quadratic function of an unknown state of the world and the decision. It was demonstrated in [6] that the common limit to which decisions converge (for the estimation problem) is actually the centralized estimate, i.e., the estimate that would be obtained if all agents were to communicate their detailed observations. We prove (Section IV) that the same is true in the presence of memory limitations, provided that each agent never forgets his own raw observations. (That is, he may only forget past tentative decisions sent to him by other agents.) We indicate that for linear quadratic Gaussian (LQG) problems our scheme is essentially a decomposition algorithm for solving static linear estimation problems. As we point out in Section IV, this scheme has certain appealing

features: there is significant parallelism in the computations which matches nicely with the assumed distribution of the data; also, in the course of the algorithm, acceptable estimates are obtained much earlier than the time that would be needed to compute the optimal estimate by centralizing the information. These tentative estimates can be very useful whenever there are strict time limits within which certain decisions have to be made.

We also consider (Section V) a slightly different scheme in which each agent transmits his tentative decision to a coordinator. The latter evaluates a weighted average of the tentative decisions he has received and sends it back to all agents. We show that our results remain valid for this scheme as well, and suggest an economic interpretation in which the coordinator can be viewed as some sort of market mechanism. We also show that making optimal tentative decisions corresponds to Nash strategies for a certain sequential game.

A weak point of the model is that not only does each agent have the same prior information and knows the statistics of the other agents' observations, but also has the same model of the probabilistic mechanism that generates interagent communications. In particular, if this is a deterministic mechanism, an agent must know the precise history of communications between any pair of other agents, a strong requirement. If it is a stochastic mechanism, then there are two possibilities: either the history of communications becomes commonly known on-line (at the expense of additional communications) or each agent will have to make probabilistic inferences about the communications between all other agents. These weaknesses disappear, however, if every tentative decision is broadcast simultaneously to all other agents, at each stage. In that case the history of communications is simple, commonly known, and easy to remember. (This will be the case, for example, if a set of experts with the same objective teleconfer and take turns suggesting what they believe to be the optimal decision.)

Finally, we point out the ways in which our scheme is different from other schemes for distributed decision making or computation: in team decision theory [9] each agent tries to behave optimally, while trying to anticipate the behavior of the other agents; the issue of who implements what component of the decision vector is very important; we are interested instead in consensus and in a common decision, quite independently of implementation issues. In many schemes for distributed computation [2], [5] each agent specializes in updating only those components of the decision vector that have been assigned to him, whereas in our scheme each agent updates the entire decision vector.

Motivation

There are many real world situations in which several agents (or processors) with different on-line information have to cooperate, combine their information, and arrive at a common decision. Examples can be drawn from power, air traffic control, and command and control systems. Consensus on the optimal can be certainly reached by transmitting all information to a predetermined agent (centralization) but there are situations in which this is impractical, due to time and communications limitations. So, it may be preferable to communicate aggregate data. Determining an optimal way to "aggregate" is not a well-posed problem. If constraints are placed on the number of bits to be transmitted, the problem becomes computationally intractable, even in very simple situations where there is only a finite number of possible events and decisions [11]. On the other hand, if a message is allowed to be any real number, all data can be coded in a single message. (This is reminiscent of decentralized control problems in which an agent may observe the decisions of other agents—the so-called control sharing pattern [1], [13].) Any such trick is very sensitive to noise in the channel and is effectively just a more complicated way of centralizing information, which was deemed undesirable in the first place.

The above discussion implies that a particular aggregation of the data should be chosen by means of some ad hoc rule that guarantees that certain desirable characteristics are present. Unless some particular structure on the problem is assumed, the optimal tentative decision given an agent's information seems to be a natural message that an agent could transmit and this is the reason that we have adopted such a framework in this paper. Of course, a demonstration that our scheme leads to approximate consensus with fewer communications and/or computations than direct centralization has to rely on numerical experimentation and the answer will depend on the specific situation.

II. MODEL FORMULATION

In this section we present a mathematical formulation of the model informally described in the Introduction. We start with the general assumptions and later proceed to the development of alternative specialized models to be considered (e.g., memory limitations, particular forms of the cost function, etc.). As far as the description of the sequence of communications and updates goes, we basically adopt the model of Borkar and Varaiya [6] except that time is considered to be discrete. As in [6], events are timed with respect to a common, absolute clock. As far as notation is concerned, we will use subscripts to denote time and superscripts to denote agents.

We assume that we are given a set $\{1, \dots, N\}$ of N agents, an underlying probability space $(\Omega, \mathcal{F}, \mathcal{P})$, and a real valued cost function $c: \Omega \times U \rightarrow R$, where U is the set of admissible values of the decision variable. It will be useful in the sequel to distinguish between elements of U and U -valued random variables. The letter v will be used to denote elements of U whereas u, w will be used to denote U -valued random variables (measurable functions from Ω to U).

Assumption 1: Either

1.1: U is a finite set, or

1.2: $U = R^n$, for some n .

Assumption 2: The cost function c is nonnegative and jointly measurable in (ω, v) . Moreover, $E[c(v)] < \infty$, $\forall v \in U$. When Assumption 1.2 holds, we assume that there exists a positive and measurable function $A: \Omega \rightarrow R$ such that

$$A(\omega) \|v_1 - v_2\|^2 \leq \frac{1}{2} [c(\omega, v_1) + c(\omega, v_2)] - c\left(\omega, \frac{v_1 + v_2}{2}\right), \quad \forall \omega \in \Omega, \quad \forall v_1, v_2 \in U. \quad (1)$$

(Remark: If we fix $v_1, v_2, v_1 \neq v_2$ and take expectations of both sides of (1), it follows that A is integrable.)

Inequality (1) implies that c is a strictly convex function of v and strict convexity holds in a uniform way, for any fixed $\omega \in \Omega$. It also follows that $c(\omega, v)$ is continuous for any $\omega \in \Omega$. This assumption is satisfied, in particular, if c is twice continuously differentiable in v and its Hessian is positive definite, uniformly in v , for any fixed $\omega \in \Omega$.

We may use the function A , defined in Assumption 2, to define a new measure μ on (Ω, \mathcal{F}) by

$$\mu(B) = \int_B A(\omega) d\mathcal{P}(\omega), \quad B \in \mathcal{F}. \quad (2)$$

This measure will be used in Section III.

We now consider the generic situation facing agent i at some time n . Let $\mathcal{F}_n^i \subset \mathcal{F}$ be a σ -field of events describing the information possessed by agent i at time n . Because of Assumption 2, the conditional expectation $E[c(v) | \mathcal{F}_n^i]$ exists (is finite), is \mathcal{F}_n^i -measurable, and is uniquely determined up to a set of measure zero, for any fixed $v \in U$. Agent i then computes a tentative decision u_n^i that minimizes $E[c(v) | \mathcal{F}_n^i]$. The following lemma (proved in the Appendix) states that u_n^i is well defined and \mathcal{F}_n^i -measurable.

Lemma 1: Under Assumptions 1.2 and 2, there exists an \mathcal{F}_n^i -measurable random variable u_n^i , which is unique up to a set of measure zero, such that

$$E[c(u_n^i) | \mathcal{F}_n^i] \leq E[c(w) | \mathcal{F}_n^i], \quad \text{almost surely,} \quad (3)$$

for any U -valued, \mathcal{F}_n^i -measurable random variable w . The same results are true (except for uniqueness) under Assumptions 1.1 and 2.

We continue with a description of the process of communications between agents. When, at time n , agent i computes his tentative optimal decision u_n^i , he may communicate its realized value (say v_n^i) to any other agent. (If v_n^i is not unique, a particular minimizing v_n^i is selected according to some commonly known rule.) Whether, when, and to which agents v_n^i is to be sent is a random event whose statistics are described by $(\Omega, \mathcal{F}, \mathcal{P})$. In particular, it may depend on the data possessed by agent i at time n . So, we implicitly allow the agents to influence the process of communications, although we do not require this influence to be optimal in any sense. This allows the possibility of signaling additional information, beyond that contained in v_n^i , by appropriately choosing when and to which agents to communicate. We allow the communication delays to be random but finite. We also assume that when an agent receives a message, he knows the identity of the agent who sent it.

We now impose conditions on the number of messages to be communicated in the long run; these conditions are necessary for agreement to be guaranteed. Namely, we require that there is an indirect communication link from any agent to any other agent which is used an infinite number of times. This can be made precise as follows.

Let $A(i)$ be the set of all agents that send an infinite number of messages to agent i , with probability 1. Then, we make the following assumption.

Assumption 3: There is a sequence $m_1, \dots, m_{k+1} = m_1$ of not necessarily distinct agents such that $m_i \in A(m_{i-1})$, $i = 1, 2, \dots, k$. Each agent appears at least once in this sequence.

The main consequence of Assumption 3, which will be repeatedly used, is the following. If $\{h^i: 1, \dots, N\}$ is a set of numbers such that $h^i \leq h^j, \forall j \in A(i), \forall i$, then $h^i = h^j, \forall i, j$.

We continue with a more detailed specification of the operation of the agents. We introduce assumptions on the knowledge \mathcal{F}_n^i which are directly related to the properties of the memory of agent i . An agent may receive (at any time) observations on the state of the world or receive tentative decisions (messages) of other agents. The knowledge of an agent at some time will be a subset (depending on the properties of his memory) of the total information he has received up to that time. We consider four alternative models of memory, formalized with the four assumptions that follow.

Let w_n^i be any message received by agent i at time n . Our most general assumption requires that w_n^i and u_{n-1}^i are remembered at time n .

Assumption 4 (Imperfect Memory): For all n , the σ -field \mathcal{F}_n^i is such that u_{n-1}^i and w_n^i are \mathcal{F}_n^i -measurable.

Assumption 4 can be further weakened if some restrictions are imposed on the communications protocol.

Assumption 5 (Imperfect Memory): For each n there exists a set $I(n)$ of agents such that:

- a) $u_{n-1}^i \in \mathcal{F}_n^j, \quad \forall i \in I(n-1), \quad \forall j \in I(n)$
- b) $\mathcal{F}_n^i = \mathcal{F}_{n-1}^i, \quad \forall i \text{ not in } I(n)$.

Intuitively, $I(n)$ is the set of agents that update their decision at time n . Assumption 5 is satisfied by the following two common communication protocols, provided that agent i may obtain additional observations only at times such that $i \in I(n)$.

Ring Protocol: $I(n) = \{k\}$, where k is the unique integer such that $1 \leq k \leq N$ and $k + mN = n$, for some integer m . Here, exactly one agent updates at any time instance and communicates his tentative decision to the next agent, and so on.

Star Protocol: $I(n) = \{1, \dots, N-1\}$, if n is odd; $I(n) = \{N\}$, if n is even. Here all agents but the last one update simultaneously, communicate to the last agent who updates, and communicates to all other agents, and so on.

Assumption 6 (Own Data Remembered): Let G_n^i be the subfield of \mathcal{F} describing all information that has been observed by agent i up to time n , except for the messages of other agents. We assume that $G_n^i \subset \mathcal{F}_n^i$.

With Assumption 6, we allow the agents to forget the message they received in the past, but they are restricted to remember all their past observations. In this case the total information available to all agents is preserved.

Assumption 7 (Perfect Memory): We let Assumptions 4 and 6 hold and assume that $\mathcal{F}_n^i \subset \mathcal{F}_{n+1}^i, \forall i, n$.

Whenever Assumption 7 holds, we will denote by \mathcal{F}_∞^i the smallest σ -field containing \mathcal{F}_n^i , for all n .

We now define a few special cases of particular interest.

1) *Estimation Problem:* We are given a R^n valued random vector x on $(\Omega, \mathcal{F}, \mathcal{P})$. The objective is to minimize the mean square error. Hence, the cost function is $c(v) = (x - v)^T(x - v)$, where T denotes transpose. It is easy to see that this is a particular case of a strictly convex function covered by Assumption 2, with $A(\omega)$ being a constant.

2) *Static Linear Quadratic Gaussian Decision Problem (LQG):* Let x be an unknown random vector. Let the sequence of transmission and reception times be deterministic. We assume that the random variables observed by the agents are zero mean and, together with x , jointly normally distributed. We allow the total number of observations to be infinite. Let $U = R^n$. The objective is to fix v so as to minimize the expectation of the quadratic cost function $c(v) = v^T R v + x^T Q v$, with $R > 0$. It follows that the optimal tentative decision of agent i at time n is $u_n^i = G E[x | \mathcal{F}_n^i] = E[Gx | \mathcal{F}_n^i]$, where G is a precomputable matrix. If we redefine the unknown vector x to be equal to Gx instead of x , we conclude that we may restrict to estimation problems, without loss of generality.

3) *Finite Probability Spaces:* Here we let Ω be a finite set. Then, there exist finitely many σ -fields of subsets of Ω . Strict convexity implies that for each σ -field $\mathcal{F}_0 \subset \mathcal{F}$ and any $\omega \in \Omega$ of positive probability there exists a unique optimal tentative decision. This implies in turn that tentative decisions take values in some finite subset of U , with probability 1. We will therefore assume, without loss of generality, that U is a finite set.

We conclude this section by presenting a simple example that illustrates our scheme, under imperfect memory assumptions, where each agent forgets everything, except for his last decision and the most recent message he has received.

Suppose that we only have two agents, who communicate to each other their tentative decisions at each instant of time. Let $x, z_n^1, z_n^2, n = 1, 2, \dots$ be independent random variables, with known probability distributions. Let $y_n^i = h(x, z_n^i)$ be the observation of agent i at time n . Let $c(x, v)$ be a cost function, satisfying Assumption 2.

The tentative decisions are defined inductively as follows. Suppose that $y_n^i, u_{n-1}^1, u_{n-1}^2$ are known by agent i at time n . He then computes, for each v , the conditional cost $E[c(x, v) | y_n^i, u_{n-1}^1, u_{n-1}^2]$, which is equal to $g_n^i(y_n^i, u_{n-1}^1, u_{n-1}^2, v)$, for some Borel function g_n^i . Finally, he chooses a minimizing v , which is a function of $y_n^i, u_{n-1}^1, u_{n-1}^2$ and this is his tentative decision at time n . Hence, for appropriate functions f_n^1, f_n^2 , we have

$$u_n^i = f_n^i(y_n^i, u_{n-1}^1, u_{n-1}^2), \quad i = 1, 2.$$

If we now define $(\Omega, \mathcal{F}, \mathcal{P})$ to be the product of the probability spaces on which x, z_n^i are defined, and let \mathcal{F}_n^i be the σ -field generated by $y_n^i, u_{n-1}^1, u_{n-1}^2$, Assumptions 3 and 4 are satisfied and the asymptotic properties of the above recursions can be analyzed within our general framework.

III. CONVERGENCE AND AGREEMENT RESULTS

In this section we state and discuss our main results. All proofs can be found in the Appendix. Assumptions 2 and 3 will be assumed throughout the rest of the paper and will not be explicitly mentioned in the statement of each theorem. We start with the least restrictive assumptions on memory.

Theorem 1: We assume that transmissions and receptions are deterministic, that communication delays are bounded, and that the time between consecutive transmissions from agent j to agent i [with $j \in A(i)$] is bounded. Then, under Assumption 1.2 (convex costs) and either Assumption 4 or 5 (imperfect memory):

- a) $\lim_{n \rightarrow \infty} (u_{n+1}^i - u_n^i) = 0$, in probability and in $L_2(\Omega, \mathcal{F}, \mu)$.
- b) $\lim_{n \rightarrow \infty} (u_n^i - u_n^j) = 0$, $\forall i, j$, in probability and in $L_2(\Omega, \mathcal{F}, \mu)$.

Consider the following situation. At time zero, before any observations are obtained, the sequence of transmissions and receptions is selected at random, according to a statistical law which is independent from all observations to be obtained in the future and from $c(v)$, for any $v \in U$. In other words, communications do not carry any information relevant to the decision problem, other than the content of the message being communicated. Suppose that the sequence of communications that has been selected becomes known to all agents. From that point on, the situation is identical with that of deterministic communications. In fact, a moment's thought will show that it is sufficient for the history of communications to become commonly known as it occurs: agent i only needs to know, at time n , what communications have occurred up to that time, so that he can interpret correctly the meaning of the messages he is receiving.

We can formalize these ideas as follows. We are given a product probability space $(\Omega \times \Omega^*, \mathcal{F} \times \mathcal{F}^*, \mathcal{P} \times \mathcal{P}^*)$ where $(\Omega, \mathcal{F}, \mathcal{P})$ describes the decision problem and where $(\Omega^*, \mathcal{F}^*, \mathcal{P}^*)$ describes the communications process. We assume that for each $\omega^* \in \Omega^*$, the resulting process of communications satisfies the assumptions of Theorem 1. Then, note that for each $\omega^* \in \Omega^*$ we obtain a distributed decision problem on $(\Omega, \mathcal{F}, \mathcal{P})$ with deterministic communications. In that case Theorem 2 follows.

Theorem 2: Under Assumption 1.2, either Assumptions 4 or 5, and independent, commonly known communications (as described above), $\lim_{n \rightarrow \infty} (u_{n+1}^i - u_n^i) = \lim_{n \rightarrow \infty} (u_n^i - u_n^j) = 0$, in probability with respect to $\mathcal{P} \times \mathcal{P}^*$.

Strictly speaking, Theorems 1 and 2 do not guarantee convergence of the decisions of each agent. Suppose, however, that the agents operate under the following rule. Fix some small $\gamma > 0$; let the sequence of communications and updates of tentative decisions take place until $|u_n^i - u_n^j| < \gamma$, $\forall i, j$ (small disagreement) and $|u_{n+1}^i - u_n^i| < \gamma$, $\forall i$ (small foreseeable changes in tentative decisions). Then, we obtain the following.

Corollary 1: With the above rule and the assumptions of Theorems 1 or 2, the process terminates in finite time, with probability 1, for any $\gamma > 0$.

When Ω and U are finite, convergence and agreement are obtained after finitely many stages.

Theorem 3: If Ω and U are finite sets, if each agent communicates all the values of v that minimize $E[c(v)|\mathcal{F}_n^i]$ and if Assumption 4 holds, then there exists some positive integer M such that

$$u_M^i = u_M^j, \quad \forall i, j \quad \text{and} \quad u_{M+n}^i = u_M^i, \quad \forall i, \quad \forall n \geq 0, \quad \forall \omega \in \Omega.$$

Strictly speaking, tentative decisions in the above theorem are not elements of U but subsets of U . This is to compensate for the possibility of nonuniqueness of optimal tentative decisions. The equalities appearing in Theorem 3 have to be interpreted, therefore, as equalities of sets.

We now assume that the agents have perfect memory. We

obtain results similar to Theorems 1 and 2 under much more relaxed assumptions on the communications process. Namely, we only need to assume the following.

Assumption 8: Let M_k^{ij} be the k th message sent by agent j to agent i . We assume that when agent i receives M_k^{ij} , he knows that this is indeed the k th message sent to him by agent j .

Remark: This assumption is trivially satisfied if messages arrive at exactly the same order as they are sent, with probability 1.

Theorem 4: Under Assumption 1.2 (convex costs), Assumption 7 (perfect memory), and Assumption 8, there exists a U -valued random variable u^* such that $\lim_{n \rightarrow \infty} u_n^i = u^*$, $\forall i$, in probability and in $L_2(\Omega, \mathcal{F}, \mu)$.

For estimation problems ($u_n^i = E[x|\mathcal{F}_n^i]$), Theorem 4 can be slightly strengthened [5, Theorem 2].

Theorem 5: For estimation problems, under the assumptions of Theorem 4, convergence to u^* takes place with probability 1.

We now consider the case where U is finite but (unlike Theorem 3) Ω is allowed to be infinite. Several complications may arise, all of them due to the fact that optimal decisions, given some information, are not guaranteed to be unique. We discuss these issues briefly, in order to motivate the next theorem.

Suppose that $U = \{v_1, v_2\}$. It is conceivable that $E[c(v_1)|\mathcal{F}_n^i] - E[c(v_2)|\mathcal{F}_n^i]$ is never zero and changes sign an infinite number of times, on a set of positive probability. In that case, the decisions of agent i do not converge. Even worse, it is conceivable that $E[c(v_1)|\mathcal{F}_n^i] > E[c(v_2)|\mathcal{F}_n^i]$ and $E[c(v_1)|\mathcal{F}_n^j] < E[c(v_2)|\mathcal{F}_n^j]$, for all n and for all ω in a set of positive probability, in which case agents i and j disagree forever. It is not hard to show that in both of the above cases $E[c(v_1)|\mathcal{F}_\infty^i] = E[c(v_2)|\mathcal{F}_\infty^i]$, on a set of positive probability and this nonuniqueness is the source of the pathology. The following theorem states that convergence and agreement are still obtained, provided that we explicitly exclude the possibility of nonuniqueness.

Theorem 6: Under Assumption 1.1 (finite U) and Assumption 7 (perfect memory), and if the random variable u^i that minimizes $E[c(w)]$ over all \mathcal{F}_∞^i -measurable random variables is unique up to a set of measure zero, for all i , then $\lim_{n \rightarrow \infty} u_n^i = u^i$, almost surely, and $u^i = u^j$, $\forall i, j$.

Although the preceding theorems guarantee that (under certain conditions) all agents will agree, nothing has been said concerning the particular decision to which all agents' decisions converge. In particular, it is not necessarily true, as one would be tempted to conjecture, that the limit decision is the optimal centralized solution (that is, the solution to be obtained if all agents were to communicate all their information). On the other hand, the centralized solution is reached for LQG problems, under the perfect memory assumption [6] and is also reached generically for an estimation problem on a finite probability space [7]. This issue will be touched again in the next section.

Robustness with Respect to Communication Noise

Schemes that centralize information by coding, e.g., by using the least significant bits of the allowed messages [1], [13], tend to require high bandwidth and are sensitive to noise in the communication channel. In our scheme, although real numbers are being transmitted (infinite information content), the least significant bits are not as essential. As a result, the qualitative convergence properties of our scheme are retained even if communications of the tentative decisions are assumed to be noisy. We provide a proof of this fact for estimation problems, under the perfect memory assumption.

Suppose, as before, that at random times agent j communicates his optimal tentative decision u_n^j . However, the message received by the other agents is $\hat{u}_n^j = u_n^j + q_n^j$, where q_n^j is a random vector representing the noise in the channel. For simplicity, we assume that the noise vectors are independent, identically distributed.

Theorem 7: Assume noisy communications (as described above). For estimation problems, under Assumption 7 (perfect

memory), there exists a U -valued random variable u^* such that $\lim_{n \rightarrow \infty} u_n^i = u^*$, $\forall i$, with probability 1.

IV. THE LINEAR QUADRATIC GAUSSIAN (LQG) MODEL

In this section we specialize and strengthen some of our results by restricting to the linear quadratic Gaussian model described in Section II. (Recall that any such problem is equivalent to an estimation problem; therefore, $u_n^i = \hat{x}_n^i = E[x | \mathcal{F}_n^i]$, for some random vector x .) Theorems 1, 4, and 5 are applicable. Moreover, the results of [6] guarantee that, under Assumption 7 (perfect memory), u_n^i converges to the optimal centralized estimate, given the information possessed by all agents. The following theorem states that the same is true under the weaker Assumption 6.

Theorem 8: For the LQG problem, under the assumptions of Theorem 1 and Assumption 6 (imperfect memory; own data remembered), $\lim_{n \rightarrow \infty} \hat{x}_n^i = \hat{x}$, in the mean square, where $\hat{x} = E[x | \mathcal{F}_\infty]$ and \mathcal{F}_∞ is the smallest σ -field containing \mathcal{F}_n^i , for all i, n .

Note that Theorem 8 is much stronger than Theorem 1 which was proved for the general case of imperfect memory. We have here convergence to a limit solution which is also guaranteed to be the optimal centralized solution.

Our next result concerns the finite dimensional LQG problem in which the total number of observations is finite. Namely, the smallest σ -field containing \mathcal{F}_n^i for all i, n is generated by a finite number of (jointly Gaussian) random variables. In that case, the centralized solution is going to be reached by all agents in a finite number of stages, provided that all agents have perfect memory.

Theorem 9: For the LQG problem with finitely many observations and under Assumption 7 (perfect memory), the centralized solution is reached by all agents in a finite number of stages.

Theorems 8 and 9 imply that the scheme considered in this paper may be viewed as an algorithm for solving static linear estimation problems, an issue that we discuss below.

The intuitive argument behind Theorem 9 is the following: once an agent has received enough messages, he is able to infer exactly the values of the observations of the other agents (or of some appropriate linear combinations of these observations) and compute the centralized solution himself. So, communicating optimal tentative decisions is in this case just another way of communicating all information to all other agents. This scheme does not seem to have any particular advantages (in terms of communication and computation requirements) over the scheme where each agent communicates all his data directly.

However, the scheme of Theorem 8 (imperfect memory) seems to have some appealing features, as we indicate below. Suppose that we have a single processor who obtains an NM -dimensional vector of observations. He then divides his observations into N M -dimensional vectors that will play the role of the agents of our scheme. Finally, the processor, instead of inverting the $NM \times NM$ covariance matrix to obtain the optimal estimate (which would require $O(N^3M^3)$ operations), he uses the scheme of Theorem 8. At each round there will be one inversion per block of data, that is $O(NM^3)$ operations per round. If, for example, an acceptable estimate is obtained after $O(N)$ rounds, the final objective will have been accomplished with a total of $O(N^2M^3)$ operations, which is one order of magnitude less than the usual algorithm. It is not hard to show that if the noises in observations belonging to different blocks are uncorrelated, agreement is obtained after two rounds only. Accordingly, if the noises in observations in different blocks are weakly correlated, we expect our scheme to be faster than the standard algorithm. We present below some numerical results that support the above statements. So, our scheme leads to a potentially advantageous *decomposition* algorithm for static linear estimation problems. (This algorithm has some conceptual similarities with those suggested in [8].)

We now discuss some issues related to the distributed implementation of the decomposition algorithm, where each block of data actually corresponds to a physically distinct agent

(processor). For any i, n , $\hat{x}_n^i = a_n^i y$, where a_n^i is a row vector and y is the vector of all available observations. When agent j receives \hat{x}_n^i , he must also learn a_n^i , in order to be able to extract information from \hat{x}_n^i . There are two choices: either agent j computes a_n^i , which may be done off-line, or, agent i transmits a_n^i to agent j . Which of the two should be done clearly depends on whether communications or computations are most costly. Whether one of the above two variations can be useful depends on the particularities of the actual situation and its inherent communication and computation limitations. More numerical experience is needed before a definite answer can be given.

Numerical Results

Let x be an unknown scalar, zero mean, random variable to be estimated ($E[x^2] = 5$). Let $y_i = x + w_i$, ($i = 1, \dots, 18$) be the observations. The noises w_i are assumed to be independent of x . (The covariance Σ_w of the noises was randomly generated.) We split the 18-dimensional observation vector into blocks of data (corresponding to distinct agents) and used the decomposition algorithm of Theorem 8. We employed the ring protocol and assumed that at each stage an agent only knows his own observations and the most recent message he received (Assumptions 5 and 6).

Let M_i be the number of observations assigned to agent i . We considered two alternative decompositions: 1) $N = 2$, $M_1 = 10$, $M_2 = 8$; 2) $N = 6$, $M_1 = \dots = M_6 = 3$. We first executed the algorithm using the covariance Σ_w , and then once more using the covariance $\Sigma_w + I$.

The results are presented in Figs. 1 and 2. The horizontal axis denotes stages (each stage corresponds to an update by some agent) and the vertical axis indicates the associated mean square error. The dotted horizontal line indicates the centralized mean square error. The curves $D1$ and $D2$ correspond to the first and second decomposition, respectively. As expected, convergence was much faster when the identity was added to the initial covariance; moreover, the first decomposition converged much faster than the second.

To illustrate the merits of the decomposition algorithm we performed a rough count of operations. We only took matrix inversions into account, assuming that the inversion of an $M \times M$ matrix requires M^3 operations, which is accurate enough for our purposes. With this counting scheme, the centralized algorithm required 5832 operations. The points A, B in the graphs were reached after 4100, 1152 operations, respectively. This leads to the following conclusion: while the first decomposition needs very few stages to converge, it does not have any particular computational advantages. The second decomposition, however, leads to an estimate close to the optimal with much fewer operations than the centralized algorithm.

V. A MODEL INVOLVING A COORDINATOR

In the previous sections we had assumed that for any pair of agents i, j , agent i is allowed to communicate to j . In this section we assume that a particular agent (denoted by the superscript o) has special status and acts like a coordinator. The scheme we envisage is the following: at each instance of time n , agent i evaluates u_n^i which he communicates to the coordinator. The coordinator then combines u_n^1 to u_n^N to produce a tentative decision u_n^o . (We assume that the coordinator has no data of his own.) He then transmits u_n^o to all other agents which accordingly update their decisions. Were the coordinator to combine u_n^1 to u_n^N "optimally," the above scheme would reduce to the one of the previous sections and our past results would apply. We assume, however, that the coordinator simply sets

$$u_n^o = \sum_{i=1}^N a^i u_n^i$$




Fig. 1. Mean-square errors of distributed algorithms with covariance Σ_w ; dashed line represents performance of centralized algorithm.




Fig. 2. Mean-square errors of distributed algorithms with covariance $\Sigma_w + I$; dashed line represents performance of centralized algorithm.

where the coefficients a^i are deterministic, positive, and $\sum_{i=1}^N a^i = 1$. The implicit behavioral assumptions are: 1) the coordinator has no memory; and 2) he need not have a good knowledge of the problem. He only knows how much he can rely on each of the other agents; this is reflected by his choice of the coefficients a^i which may be thought of as a "reliability index" for agent i in the eyes of the coordinator. We then obtain the following.

Theorem 10: The conclusions of Theorems 1, 4, and 8 remain true (under their respective assumptions) with the scheme introduced in this section.

The above scheme can be viewed as a framework for cooperation, where the coordinator simply aids the agents, or, for LQG problems, as another decomposition algorithm. It can be also interpreted, however, from an entirely different point of view. Suppose that the agents are selfish and independent individuals, faced with identical situations, possessing different information and having to make repetitive decisions. They can certainly benefit by observing past decisions of the other agents but assume that this is not possible. They are able, however, to observe a weighted average u_n^o of all decisions made in the last stage, which they take into account for their future actions. The motivation for such a model comes primarily from economics: each agent is a buyer (or seller) in the same market and at each stage he obtains some aggregate information (e.g., the average price) on the transactions that were made in the previous stage. In this sense the "coordinator" simply represents a market mechanism. Our results state that, eventually, an informational equilibrium will be reached. Such an equilibrium has been studied by Radner [12] in a different setting. However, there was no demonstration of an adjustment process that could lead to such an equilibrium. Our scheme provides a model of rational behavior which, if followed by each agent, leads to equilibrium.

Moreover, within such a context (of selfish individuals confronted with identical situations) and for LQG problems with perfect memory, optimal tentative decisions constitute a set of strategies in Nash equilibrium for a certain game. (This is why

optimal tentative decisions can be called "a model of rational behavior.") Let us define the game of interest more precisely.

Let y_n^i be a vector of jointly Gaussian random variables that generate G_n^i , the σ -algebra of events known to agent i at time n if he had received no messages. At each stage, agent i selects a decision u_n^i and incurs (but does not observe) a cost $\alpha^n c(u_n^i)$, where $0 < \alpha < 1$ and c is a quadratic cost function. Then, $u_n^o = \sum a^i u_n^i$ is formed and communicated to all agents. The total cost to agent i is $J^i = \sum_{n=1}^{\infty} \alpha^n E[c(u_n^i)]$. A strategy Γ^i for agent i is a sequence $\{\gamma_n^i, i=1, 2, \dots\}$ of measurable functions such that $u_n^i = \gamma_n^i(y_n^i, u_1^o, \dots, u_{n-1}^o)$. A set $\{\Gamma^1, \dots, \Gamma^N\}$ of strategies is said to be in Nash equilibrium if

$$J^i(\Gamma^1, \dots, \Gamma^{i-1}, \hat{\Gamma}^i, \Gamma^{i+1}, \dots, \Gamma^N) \geq J^i(\Gamma^1, \dots, \Gamma^i, \dots, \Gamma^N)$$

for any strategy $\hat{\Gamma}^i$, for any i . Let $\Gamma = \{\Gamma^i: i=1, \dots, N\}$ denote the particular set of strategies where each agent at each stage plays his optimal tentative decision. (Note that these are linear strategies.) Then Theorem 11 follows.

Theorem 11: Γ is a set of strategies in Nash equilibrium.

VI. CONCLUSIONS

A set of agents with the same objective who start communicating to each other their tentative optimal decisions are guaranteed to agree in the limit. Under certain assumptions, this is true even if the decisions are received in the presence of noise and even if their memory is limited and they are allowed to forget some of their past knowledge. Moreover, they are guaranteed to converge to the optimal centralized decision for linear estimation problems, provided that they do not forget their own observations. This leads to a decomposition algorithm for static linear estimation problems. Similar results are obtained if the agents do not communicate directly but receive messages from a coordinator who evaluates a weighted average of all tentative decisions. In the latter framework, for LQG problems with perfect memory, optimal tentative decisions are Nash strategies for a certain sequential game and admit an economic interpretation.

These results are valid when all agents share the same model of the world. The characterization of the behavior of agents with different models (perceptions) is an open problem.

APPENDIX

Proof of Lemma 1: Let, for notational convenience, $g(\omega, v) = E[c(v)|\mathcal{F}_n^i]$. Under Assumptions 1.2 and 2, $g(\omega, v)$ can be chosen so that it is jointly measurable, strictly convex in v , and converges to infinity as v converges to infinity, for any $\omega \in \Omega$. Hence, $\forall \omega \in \Omega$, the infimum of $g(\omega, v)$ is attained by some $u_n^i(\omega)$, which is unique because of strict convexity.

Let $Q = \{q_k\}_{k=1}^{\infty}$ be a countable dense subset of U . Then, by continuity of g , $\inf_{v \in Q} g(\omega, v) = \inf_{v \in U} g(\omega, v)$, $\forall \omega \in \Omega$. Moreover, $\inf_{v \in Q} g(\omega, v)$ is \mathcal{F}_n^i -measurable. Let $\phi_m(\omega) = q_k$, where k is the smallest index such that

$$g(\omega, q_k) \leq \inf_{v \in U} g(\omega, v) + \frac{1}{m}.$$

Then ϕ_m is \mathcal{F}_n^i -measurable and converges, for each ω to u_n^i . Hence, u_n^i is \mathcal{F}_n^i -measurable. Inequality (3) now follows from the definition of u_n^i and uniqueness is a consequence of strict convexity. The measurability of u_n^i under Assumption 1.1 is trivial. \square

Lemma 2: Let $\{u_n\}, \{w_n\}$ be two sequences of U -valued random variables such that

$$\lim_{n \rightarrow \infty} E \left[c \left(\frac{u_n + w_n}{2} \right) \right] = \lim_{n \rightarrow \infty} E [c(u_n)] = \lim_{n \rightarrow \infty} E [c(w_n)]. \quad (4)$$

If Assumptions 1.2 and 2 hold, then $\lim_{n \rightarrow \infty} (u_n - w_n) = 0$ in $L_2(\Omega, \mathcal{F}, \mu)$ and in probability.

Proof: By Assumptions 1.2 and 2, and (4),

$$\begin{aligned} & \lim_{n \rightarrow \infty} E[A(\omega) \|u_n - w_n\|^2] \\ & \leq \lim_{n \rightarrow \infty} E \left[\frac{c(u_n) + c(w_n)}{2} - c\left(\frac{u_n + w_n}{2}\right) \right] = 0 \quad (5) \end{aligned}$$

which shows that $(u_n - w_n)$ converges to zero in $L_2(\Omega, \mathcal{F}, \mu)$. Therefore, it also converges in measure with respect to μ .

Recall that $A(\omega) > 0$ and $\mu(B) = \int_B A(\omega) d\mathcal{P}(\omega)$, $\forall B \in \mathcal{F}$. Therefore, $\mu(B) = 0$ implies $\mathcal{P}(B) = 0$ and \mathcal{P} is absolutely continuous with respect to μ . Let $B_n^\epsilon = \{|u_n - w_n| \geq \epsilon\}$. Since $(u_n - w_n)$ converges to zero in measure μ , for any $\epsilon > 0$, we have $\lim_{n \rightarrow \infty} \mu(B_n^\epsilon) = 0$ and, by absolute continuity, $\lim_{n \rightarrow \infty} \mathcal{P}(B_n^\epsilon) = 0$, which shows that we have convergence in probability. \square

Proof of Theorem 1: We start with the proof under Assumption 4. Since u_n^i is \mathcal{F}_{n-1}^i -measurable, we have (by the minimizing property of u_{n-1}^i) $E[c(u_{n+1}^i)] \leq E[c(u_n^i)]$. Since c is nonnegative, $E[c(u_n^i)]$ converges to some constant g^i . We also note that $(u_{n+1}^i + u_n^i)/2$ is \mathcal{F}_{n-1}^i -measurable and by taking the limit in the relation

$$\frac{1}{2} E[c(u_{n+1}^i) + c(u_n^i)] \geq E \left[c\left(\frac{u_{n-1}^i + u_n^i}{2}\right) \right] \geq E[c(u_{n-1}^i)]$$

we obtain $\lim_{n \rightarrow \infty} E[c((u_{n+1}^i + u_n^i)/2)] = g^i$. Lemma 2 then yields the first part of the theorem.

Let $j \in A(i)$. Then there exist sequences $\{m_k\}$ and $\{n_k\}$ of positive integers such that $\lim_{k \rightarrow \infty} m_k = \lim_{k \rightarrow \infty} n_k = \infty$ and m_k, n_k are the times of transmission and reception, respectively, of the k th message from agent j to agent i . Therefore, $u_{m_k}^j$ is $\mathcal{F}_{n_k}^j$ -measurable, for all k , and $E[c(u_{n_k}^i)] \leq E[c(u_{m_k}^j)]$ which shows that $g^i \leq g^j$. Using Assumption 3, we conclude that $g^i = g^j$, $\forall i, j$.

We note that $(u_{n_k}^i + u_{m_k}^j)/2$ is $\mathcal{F}_{n_k}^i$ -measurable and, therefore,

$$\frac{1}{2} E[c(u_{n_k}^i) + c(u_{m_k}^j)] \geq E \left[c\left(\frac{u_{n_k}^i + u_{m_k}^j}{2}\right) \right] \geq E[c(u_{n_k}^i)]. \quad (6)$$

Taking the limit in (6), using Lemma 2 and the boundedness assumptions on the communications, we obtain the second half of the theorem.

We now assume Assumption 5. Let $i(n)$ be a sequence of agents such that $i(n) \in I(n)$, $\forall n$. Then, $u_n^{(n)}$ is $\mathcal{F}_{n-1}^{(n)}$ -measurable and $E[c(u_n^{(n)})]$ is a decreasing sequence. Similarly with the first part of the proof, we conclude that $u_{n+1}^{(n+1)} - u_n^{(n)}$ converges to zero in $L_2(\Omega, \mathcal{F}, \mu)$ and in probability. It follows that $u_{n+1}^i - u_n^i$ and $u_n^i - u_n^j$ converge to zero, for $j \in A(i)$. Using Assumption 3, $u_n^i - u_n^j$ converges to zero, for all i, j . \square

Proof of Theorem 2: Theorem 1 and the discussion preceding the statement of Theorem 2 show that $\lim_{n \rightarrow \infty} (u_{n-1}^i - u_n^i) = 0$, in probability with respect to \mathcal{P} , for all $\omega^* \in \Omega^*$. Let $\chi_n(\omega, \omega^*)$ be the characteristic function of the set $\{(\omega, \omega^*): \|u_{n+1}^i - u_n^i\| < \epsilon\}$. Then,

$$\begin{aligned} \lim_{n \rightarrow \infty} \int \chi_n d(\mathcal{P} \times \mathcal{P}^*) &= \lim_{n \rightarrow \infty} \int \int \chi_n d\mathcal{P} d\mathcal{P}^* \\ &= \int \lim_{n \rightarrow \infty} \int \chi_n d\mathcal{P} d\mathcal{P}^* = 1. \end{aligned}$$

(The first equality holds by the Fubini theorem, the second by the dominated convergence theorem, the third by convergence of $(u_{n-1}^i - u_n^i)$ to zero, with respect to the probability measure \mathcal{P} .) This shows that $u_{n+1}^i - u_n^i$ converges to zero in probability with

respect to $\mathcal{P} \times \mathcal{P}^*$. Similar steps show that $u_n^i - u_n^j$ also converges to zero, in probability. \square

Proof of Corollary 1: By Theorem 1, the $U \times U$ -valued sequence of random variables $(u_n^i - u_n^j, u_{n-1}^i - u_n^j)$ converges to $(0, 0)$ in probability. It therefore contains a subsequence converging to $(0, 0)$, almost surely. Therefore, $\forall \gamma > 0$ and for almost all $\omega \in \Omega$, $\exists n_o$ such that $|u_{n_o}^i - u_{n_o}^j| < \gamma$, $|u_{n_o+1}^i - u_{n_o}^j| < \gamma$ and the termination condition is eventually satisfied with probability 1. \square

Proof of Theorem 3: Because of the finiteness of Ω , there exists a finite (nonrandom) time after which communications (conditioned on past events) are deterministic. We may take that time as the initial time and assume, without loss of generality, that all communications are deterministic.

Let u_n^i be the set of elements of U which are optimal, given \mathcal{F}_n^i . Let w_n^i denote an \mathcal{F}_n^i -measurable random variable such that $w_n^i(\omega) \in u_n^i(\omega)$, $\forall \omega \in \Omega$. (Note that $E[c(w_n^i)]$ is independent of how w_n^i has been selected.) By finiteness of Ω and U , there exist finitely many U -valued random variables and, since $E[c(w_{n+1}^i)] \leq E[c(w_n^i)]$, we conclude that there exists some positive integer T and some g^i such that $E[c(w_n^i)] = g^i$, $\forall n > T$. For any $n > T$, $E[c(w_n^i)] = E[c(w_{n-1}^i)]$ and since w_n^i is \mathcal{F}_{n-1}^i -measurable, w_n^i minimizes $E[c(w)]$ over all \mathcal{F}_{n-1}^i -measurable random variables. Hence, $w_n^i(\omega) \in u_{n-1}^i(\omega)$, $\forall \omega \in \Omega$ which shows that $u_n^i(\omega) \subset u_{n-1}^i(\omega)$, $\forall \omega \in \Omega$. Again, by finiteness of U and Ω , there exists some positive integer M such that $u_{n+1}^i(\omega) = u_n^i(\omega)$, $\forall n > M$, $\forall \omega \in \Omega$, $\forall i$.

If $j \in A(i)$, there exist $m, n > M$ such that w_m^j is \mathcal{F}_n^i -measurable and this shows that $g^j \geq g^i$. By Assumption 3, we obtain $g^i = g^j$, $\forall i, j$. Therefore, w_m^j minimizes $E[c(w)]$ over all \mathcal{F}_n^i -measurable random variables and, therefore, $w_m^j(\omega) \in u_n^i(\omega)$, or, $u_m^j(\omega) \subset u_n^i(\omega)$, $\forall \omega \in \Omega$. Recalling Assumption 3, we obtain $u_m^j(\omega) = u_n^i(\omega)$, $\forall i, j$, $\forall m, n > M$, $\forall \omega$. \square

Lemma 3: Let T be a finite stopping time of an increasing family $\{\mathcal{F}_n\}$ of σ -fields. Let $u_n, n=1, 2, \dots$ be random variables that minimize $E[c(w)|\mathcal{F}_n]$, almost surely, over all \mathcal{F}_n -measurable random variables w . Then, u_T minimizes $E[c(w)]$ over all \mathcal{F}_T -measurable random variables w , where u_T is defined to be equal to u_n if and only if $T = n$.

Proof: Let χ_n be the indicator function of the set $\{\omega: T(\omega) = n\}$. Since T is a stopping time, χ_n is \mathcal{F}_n -measurable. Note that $\chi_n c(u_n) = \chi_n c(u_T)$. Let w be an \mathcal{F}_T -measurable random variable and note that $\chi_n c(w) = \chi_n c(\chi_n w)$ and $\chi_n w$ is \mathcal{F}_n -measurable. Therefore,

$$\begin{aligned} E[\chi_n c(w)|\mathcal{F}_n] &= \chi_n E[c(\chi_n w)|\mathcal{F}_n] \geq \chi_n E[c(u_n)|\mathcal{F}_n] \\ &= E[\chi_n c(u_n)|\mathcal{F}_n] = E[\chi_n c(u_T)|\mathcal{F}_n]. \end{aligned}$$

Taking expectations, we obtain

$$E[\chi_n c(w)] \geq E[\chi_n c(u_T)]$$

and summing over all n 's (and using the monotone convergence theorem to interchange summation and expectation) we obtain $E[c(w)] \geq E[c(u_T)]$. \square

Proof of Theorem 4: Since u_n^i is \mathcal{F}_{n-1}^i -measurable, we have $E[c(u_{n-1}^i)] \leq E[c(u_n^i)]$. Since c is nonnegative, $E[c(u_n^i)]$ converges to some constant g^i . We also note that $(u_n^i + u_{n-m}^i)/2$ is \mathcal{F}_{n-m}^i -measurable. Therefore, $E[c((u_n^i + u_{n-m}^i)/2)] \geq E[c(u_{n-m}^i)] \geq g^i$. Fix some $\epsilon > 0$, and let n be large enough so that $E[c(u_n^i)] \leq g^i + \epsilon$. Then, using Assumption 3, we obtain $E[A(\omega) \|u_{n-m}^i - u_n^i\|^2] \leq \epsilon$, $\forall m \geq 0$. Therefore, $\{u_n^i\}$ is a Cauchy sequence in $L_2(\Omega, \mathcal{F}, \mu)$. By the completeness of L_2 spaces, there exists a U -valued random variable u^i such that $\lim_{n \rightarrow \infty} u_n^i = u^i$, in $L_2(\Omega, \mathcal{F}, \mu)$ and, therefore, in probability, with respect to \mathcal{P} . (The proof of the last implication is contained in the proof of Lemma 2.) Since

$$\begin{aligned} E[E[c(u_{n-1}^i)|\mathcal{F}_{n+1}^i]| \mathcal{F}_n^i] &\leq E[E[c(u_n^i)|\mathcal{F}_{n+1}^i]| \mathcal{F}_n^i] \\ &= E[c(u_n^i)| \mathcal{F}_n^i], \end{aligned}$$

$\{E[c(u_n^i)|\mathcal{F}_n^i], n=1,2,\dots\}$ is a supermartingale, with respect to $\{\mathcal{F}_n^i\}$. Moreover, since for any fixed $v \in U$, $E[c(u_n^i)|\mathcal{F}_n^i] \leq E[c(v)|\mathcal{F}_n^i]$, it is a uniformly integrable supermartingale [10, Theorem T19, p. 90].

Let $j \in A(i)$. Let m_k, n_k be the times of transmission and reception, respectively, of the k th message from j to i . Because of Assumption 8, m_k and n_k are stopping times of $\{\mathcal{F}_n^j\}, \{\mathcal{F}_n^i\}$, respectively, and since $j \in A(i)$, they are almost surely finite stopping times, for all k . Moreover, $k \leq m_k \leq n_k$ and by the optional sampling theorem [10, Theorem T28, p. 90]

$$E[c(u_k^i)] \geq E[c(u_{n_k}^j)] \geq g^j$$

which shows that $\lim_{k \rightarrow \infty} E[c(u_{n_k}^j)] = g^j$. Similarly, $\lim_{k \rightarrow \infty} E[c(u_{m_k}^j)] = g^j$.

Note that $u_{m_k}^j$ is $\mathcal{F}_{n_k}^j$ -measurable and, by Lemma 3, $E[c(u_{m_k}^j)] \geq E[c(u_{n_k}^j)]$. Taking the limit, we obtain $g^j \geq g^i$, and by Assumption 3, $g^i = g^j = g, \forall i, j$.

We now take the limit of the inequalities

$$\frac{1}{2} E[c(u_{n_k}^i) + c(u_{m_k}^j)] \geq E\left[c\left(\frac{u_{n_k}^i + u_{m_k}^j}{2}\right)\right] \geq E[c(u_{n_k}^i)]$$

to obtain $\lim_{k \rightarrow \infty} E[c((u_{n_k}^i + u_{m_k}^j)/2)] = g$ and, by Lemma 2, $\lim_{k \rightarrow \infty} (u_{n_k}^i - u_{m_k}^j) = 0$, in $L_2(\Omega, \mathcal{F}, \mu)$ and in probability.

We also take the limit of the inequalities

$$\frac{1}{2} E[c(u_{n_k}^i) + c(u_k^i)] \geq E\left[c\left(\frac{u_{n_k}^i + u_k^i}{2}\right)\right] \geq E[c(u_{n_k}^i)]$$

to obtain $\lim_{k \rightarrow \infty} E[c((u_k^i + u_{n_k}^i)/2)] = g^i$ and, by Lemma 2, $\lim_{k \rightarrow \infty} (u_{n_k}^i - u_k^i) = 0$. Similarly, we obtain $\lim_{k \rightarrow \infty} (u_{m_k}^j - u_k^j) = 0$, which shows that $u^i = u^j$, almost surely. \square

Proof of Theorem 6: Fix some $v \in U$ and let $B = \{\omega : u^i(\omega) = v\}$. Then, $E[c(v)|\mathcal{F}_\infty^i] < E[c(v^*)|\mathcal{F}_\infty^i], \forall v^* \in U$, for almost all $\omega \in B$. By the martingale convergence theorem [10, Theorem T17, p. 84], we conclude that, for almost all $\omega \in B$, there exists some $N(\omega)$ such that

$$E[c(v)|\mathcal{F}_n^i] < E[c(v^*)|\mathcal{F}_n^i] \quad \forall n \geq N(\omega).$$

Therefore, $\lim_{n \rightarrow \infty} u_n^i(\omega) = v$, for almost all $\omega \in B$ and, by considering the other elements of U as well, $\lim_{n \rightarrow \infty} u_n^i = u^i$, almost surely.

If $j \in A(i)$, u^j is \mathcal{F}_∞^i -measurable and $E[c(u^j)] \geq E[c(u^i)]$. By Assumption 3, $E[c(u^j)] = E[c(u^j)], \forall i, j$. Therefore, for $j \in A(i)$, u^j minimizes $E[c(w)]$ over all \mathcal{F}_∞^i -measurable random variables and by the assumptions of the theorem, $u^i = u^j$, almost surely. Using Assumption 3 once more, we obtain $u^i = u^j, \forall i, j$. \square

Proof of Theorem 7: Let x be the unknown vector to be estimated. Then $u_n^i = E[x|\mathcal{F}_n^i]$ and converges almost surely to $u^i = E[x|\mathcal{F}_\infty^i]$. Moreover, u^i minimizes $E[c(w)]$ over all \mathcal{F}_∞^i -measurable random variables w . Let $j \in A(i)$ and let m_k be the time of transmission of the k th message from j to i . Note that $(1/M)\sum_{k=1}^M u_{m_k}^j$ is $\mathcal{F}_{n_k}^i$ -measurable (and hence \mathcal{F}_∞^i -measurable) and converges to u^j almost surely. Therefore, u^j is \mathcal{F}_∞^i -measurable and $E[c(u^j)] \leq E[c(u^i)]$ and, by Assumption 3, $E[c(u^i)] = E[c(u^j)], \forall i, j$. The minimizing property of u^i and the strict convexity of the quadratic cost function imply that $u^i = u^j$, almost surely. \square

Proof of Theorem 8: As is usual in linear least squares estimation, we use the setting of Hilbert spaces of square integrable random variables. Let G be a Hilbert space of zero mean, jointly Gaussian random variables on $(\Omega, \mathcal{F}, \mathcal{P})$ such that each component of the unknown vector x and the observations belong to G . The inner product in G is defined by $\langle x, y \rangle = E[xy]$.

For each agent i , let H^i denote the smallest closed subspace of G containing all observations obtained by him. Let H_n^i be the

smallest closed subspace of G containing all observations obtained by agent i up to time n . (Note that H_n^i does not contain all random variables known by agent i at time n because it does not need to contain any of the messages received by agent i .) Note also that, by Assumption 6, $H_n^i \subset H_{n+1}^i \subset H^i$ and that $\sum_{k=1}^N H_n^k$ is the total knowledge available to all agents at time n . The centralized estimate is the projection of x on $\sum_{k=1}^N H^k$. We assume, without loss of generality, that x is a scalar random variable, since each component can be separately estimated.

Let $\hat{x}_n^i = E[x|\mathcal{F}_n^i]$ and $e_n^i = x - \hat{x}_n^i$ and, by the orthogonality of errors and observations, we have $E[xy] = E[\hat{x}_n^i y], \forall y \in H_n^i$. As in the proof of Theorem 1 we have $\|e_{n+1}^i\|^2 \leq \|e_n^i\|^2, \forall n, i$ which implies that

$$\|x\|^2 \geq \|\hat{x}_{n+1}^i\|^2 \geq \|\hat{x}_n^i\|^2. \quad (7)$$

In particular, (7) implies that $\{\hat{x}_n^i\}$ is a norm-bounded sequence. By the weak local sequential compactness of Hilbert spaces [15, p. 126], $\{\hat{x}_n^i\}$ contains a weakly convergent subsequence $\{\hat{x}_{n_l}^i\}$. In other words, there exists an element $\hat{x}_\infty \in G$ such that $\langle y, \hat{x}_{n_l}^i \rangle$ converges to $\langle y, \hat{x}_\infty \rangle, \forall y \in G$. Moreover, $\hat{x}_n^i \in \sum_{k=1}^N H_n^k \subset \sum_{k=1}^N H^k$ and since closed subspaces are also weakly closed [15, Theorem 11, p. 125], $\hat{x}_\infty \in \sum_{k=1}^N H^k$. Now let $y \in H_n^i$. Then, $\langle y, \hat{x}_{n_l}^i \rangle = \langle y, x \rangle$, for all l such that $n_l \geq n$, which implies that $\langle y, \hat{x}_\infty \rangle = \langle y, x \rangle$. Moreover, the sequence of subspaces $\{H_n^i\}$ generates H^i which implies that $\langle y, \hat{x}_\infty \rangle = \langle y, x \rangle, \forall y \in H^i$.

By Theorem 1, $(x_n^i - x_n^j)$ converges in the mean square (and therefore weakly) to zero, which implies that \hat{x}_∞ is also a weak limit point of $\{\hat{x}_{n_l}^j\}$. The same argument as before shows that $\langle y, \hat{x}_\infty \rangle = \langle y, x \rangle, \forall y \in H^j, \forall j$. Therefore, $\langle y, \hat{x}_\infty \rangle = \langle y, x \rangle, \forall y \in \sum_{k=1}^N H^k$. But this is exactly the condition that \hat{x}_∞ is the centralized estimate, given the observations of all agents. So, $\{\hat{x}_n^i\}$ has a unique weak limit point, which is the same for all i and coincides with the centralized estimate.

It only remains to show that \hat{x}_n^i converges to \hat{x}_∞ strongly (in the mean square). We know from [15, p. 120] that $\|\hat{x}_\infty\| \leq \liminf_{n \rightarrow \infty} \|\hat{x}_n^i\|$. On the other hand,

$$\|x\|^2 - \|\hat{x}_\infty\|^2 = \|x - \hat{x}_\infty\|^2 \leq \|x - \hat{x}_n^i\|^2 = \|x\|^2 - \|\hat{x}_n^i\|^2$$

which shows that $\|\hat{x}_\infty\| \leq \limsup_{n \rightarrow \infty} \|\hat{x}_n^i\|$. Therefore, $\|\hat{x}_\infty\| = \lim_{n \rightarrow \infty} \|\hat{x}_n^i\|$ and by [15, Theorem 8, p. 124], we conclude that $\lim_{n \rightarrow \infty} \|\hat{x}_n^i - \hat{x}_\infty\|^2 = 0$. \square

Proof of Theorem 9: We use again the Hilbert space formalism of the previous proof. Let G_n^i be the subspace of G describing the knowledge of agent i at time n (both his observations and the messages he has received). By Assumption 7, we have $G_n^i \subset G_{n+1}^i \subset G$. Since G can be chosen to be finite dimensional, there exists some M (depending on the sequence of communications but deterministic) such that $G_{M+n}^i = G_M^i, \forall n \geq 0, \forall i$. Equivalently, $\hat{x}_{M+n}^i = \hat{x}_M^i, \forall n \geq 0, \forall i$, and by Theorem 8, $\hat{x}_M^i = \hat{x}_M^j = \hat{x}_\infty, \forall i, j$. \square

Proof of Theorem 10: Since u_n^i is \mathcal{F}_{n+1}^i -measurable, it follows that

$$\begin{aligned} E[c(u_{n+1}^i)] &\leq E[c(u_n^i)] = E\left[c\left(\sum_{j=1}^N a^j u_n^j\right)\right] \\ &\leq \sum_{j=1}^N a^j E[c(u_n^j)]. \end{aligned} \quad (8)$$

Also, since u_n^i is \mathcal{F}_{n+1}^i -measurable, $E[c(u_n^i)]$ decreases and converges to some g^i . Taking the limit in (8) we conclude that $g^i = g^j, \forall i, j$. From this point on, the proofs of Theorems 1, 4, and 8 (with minor modifications) are valid and establish the desired conclusions. \square

Proof of Theorem 11: Let u_n^0, u_n^i be the coordinator messages and decisions, respectively, when all agents use Γ^i . Let \mathcal{F}_n^i be the smallest σ -algebra generated by y_n^i and u_1^0, \dots, u_{n-1}^0 . Now sup-

pose that a particular agent, say agent 1, uses a strategy $\hat{\Gamma}^1$ different from Γ^1 , while all other agents use Γ^i . Let \hat{u}_n^o, \hat{u}_n^1 be the coordinator messages and decisions resulting from the application of this new set of strategies.

We proceed by induction. Clearly, \hat{u}_1^1 is \mathcal{F}_1^1 -measurable. Assume that \hat{u}_n^o and $\hat{u}_1^o, \dots, \hat{u}_{n-1}^o$ are \mathcal{F}_n^1 -measurable. Then, because of the linearity of the Γ^i 's, $u_n^i - \hat{u}_n^i$ ($i \neq 1$), is linear in $u_1^o - \hat{u}_1^o, \dots, u_{n-1}^o - \hat{u}_{n-1}^o$, and hence \mathcal{F}_n^1 -measurable. Since u_n^o is \mathcal{F}_{n+1}^1 -measurable and $\mathcal{F}_n^1 \subset \mathcal{F}_{n+1}^1$ (perfect memory), \hat{u}_n^o is \mathcal{F}_{n+1}^1 -measurable, and so is \hat{u}_{n+1}^1 . The induction shows that this is true for all n .

Therefore, by the minimizing property of u_n^1 , $E[c(u_n^1)] \leq E[c(\hat{u}_n^1)]$, $\forall n$ and this completes the proof. \square

ACKNOWLEDGMENT

The authors would like to thank Prof. R. Dudley of the Mathematics Department at M.I.T. for suggesting the proof of Lemma 1.

REFERENCES

- [1] M. Aoki, "On decentralized linear stochastic control problems with quadratic cost," *IEEE Trans. Automat. Contr.*, vol. AC-18, pp. 243-250, June 1973.
- [2] K. J. Arrow and L. Hurwicz, "Decentralization and computation in resource allocation," in *Essays in Economics and Econometrics*, R. W. Pfouts, Ed. Chapel Hill, NC: Univ. North Carolina Press, 1969, pp. 34-104.
- [3] M. Athans, "The expert team of experts approach to C^3 organizations," *IEEE Contr. Syst. Mag.*, vol. 2, Sept. 1982.
- [4] R. J. Aumann, "Agreeing to disagree," *Ann. Statist.*, vol. 4, no. 6, pp. 1236-1239, 1976.
- [5] D. P. Bertsekas, "Distributed dynamic programming," *IEEE Trans. Automat. Contr.*, vol. AC-27, pp. 610-616, June 1982.
- [6] V. Borkar and P. Varaiya, "Asymptotic agreement in distributed estimation," *IEEE Trans. Automat. Contr.*, vol. AC-27, pp. 650-655, June 1982.
- [7] J. D. Geanakoplos and H. M. Polemarchakis, "We can't disagree forever," *Inst. Math. Studies in Social Sci.*, Stanford University, Stanford, CA, Tech. Rep. 277, 1978.

- [8] A. J. Laub and F. N. Bailey, "An iterative coordination approach to decentralized decision problems," *IEEE Trans. Automat. Contr.*, vol. AC-23, pp. 1031-1036, Dec. 1978.
- [9] J. Marschak and R. Radner, *Economic Theory of Teams*. New Haven, CT: Yale Univ. Press, 1972.
- [10] P. A. Meyer, *Probability and Potentials*. Waltham, MA: Blaisdell, 1966.
- [11] C. H. Papadimitriou and J. N. Tsitsiklis, "On the complexity of designing distributed protocols," submitted for publication.
- [12] R. Radner, "Rational expectations equilibrium: Generic existence and the information revealed by prices," *Econometrica*, vol. 47, pp. 655-678, May 1979.
- [13] N. R. Sandell, Jr. and M. Athans, "Solution of some nonclassical LQG stochastic decision problems," *IEEE Trans. Automat. Contr.*, vol. AC-19, pp. 108-115, Apr. 1974.
- [14] H. A. Simon, *The Sciences of the Artificial*, 2nd ed. Cambridge, MA: M.I.T. Press, 1981.
- [15] K. Yosida, *Functional Analysis*. New York: Springer-Verlag, 1980.



John N. Tsitsiklis was born in Thessaloniki, Greece, on February 18, 1958. He received the B.S. degrees in electrical engineering and mathematics in 1980, the M.S. degree in electrical engineering in 1981, and the Ph.D. degree in electrical engineering in 1983, all from the Massachusetts Institute of Technology, Cambridge.

Currently he is Assistant Professor of Electrical Engineering at Stanford University, Stanford, CA. His research interests are in decentralized decision making, estimation and computation, as

well as stochastic control.

Dr. Tsitsiklis is a member of Phi Beta Kappa, Sigma Xi, and the Mathematical Association of America.

Michael Athans (S'58-M'61-SM'69-F'73), for a photograph and biography, see this issue, p. 8.