Linear Filtering, Prediction and Inter-
Math. Statist, 43 (1972), pp. 1935–1941.
ability Ratio Test," Ann. Math. Stat., 36

ass of Coefficients of Divergence of One
Soc., Series B, 28 (1966), pp. 131–142.
Difference of Probability Distributions
itarium Mathematicerium Hungarica, 2

-point Uncertainty Class," IEEE Trans.
374–376.
and Quantized Wiener Filters for p-point
. Sciences and Systems, Princeton Univ.,

sses in Robust Hypothesis Testing and
Conf. Comm., Control and Computing,
4.
ial Capacities," Ann. Statist., 5 (1977),

est Problems for Special Capacities," Z.

of Second-Order Stationary Processes,"

Discrete Time Series," Z. Wahr. verw.

on Stability of Random Variables and

ctrum Estimation: A Review," IEEE
SP-29 (Dec. 1981), pp. 1176–1186.
bust Wiener Filters for Random Signals
nation Theory, IT-29 (July 1983), pp.

Capacities," Ann. Prob., 10 (1982), pp.

and Dynamical Systems, Wiley, New

New York, 1981.
s and the Neyman-Pearson Lemma for
263.
on Classes of Sources Determined by
n Theory, IT-28 (1982), pp. 19–26.
Capacities: The Case of Correlated
Control, AC-31 (1986), pp. 877–878.
lation for Time Series," J. Time Series

nd Robust Time Series Interpolation
982), pp. 185–194.
servers and Regulators for Stochastic
s," IEEE Trans. Automatic Control.

d Prediction for Linear Systems with
oach", IEEE Trans. Automatic Con-

Chapter 10

CONSENSUS IN DISTRIBUTED ESTIMATION[1]

Demosthenis Teneketzis and Pravin Varaiya

ABSTRACT

A team must agree on a common decision to minimize the expected cost. Different team members have different observations relating to the "state of the world," and they may also have different prior beliefs. To reach a consensus they exchange tentative decisions based on their current information. Two questions are discussed: When do the individual estimates converge? If they converge, will a consensus be reached?

1. INTRODUCTION

A team or committee of N people, indexed $i = 1, \ldots, N$, must agree on a common decision u to be selected from a prespecified set U so as to minimize the cost

$$J(\omega, u) \tag{1}$$

Advances in Statistical Signal Processing, vol. 1, pages 361–385

ISBN: 0-89232-570-4

where J is a real valued function of the "state of the world" $\omega \in \Omega$, and the decision u. Initially, different people have different information relating to ω. This is modeled by stipulating that person i observes the value of the random variable $Y_i = Y_i(\omega)$. Everyone knows that i knows Y_i, although j, $j \neq i$, does not know what the value of Y_i actually is. Everyone knows the function J.

Each person has a prior belief concerning ω. We stipulate that i's prior belief is summarized by the probability distribution P^i on $(\Omega, \underline{F})$, where $\underline{F}$ is the σ-field of events. If $P^1 = \cdots = P^N$, we say that the beliefs are **consistent**; otherwise they are **inconsistent**.

Since initially different people have different information, and also because their beliefs may be inconsistent, their estimates of the best decision will also be different. To arrive at a consensus decision it is necessary for them to share information. We suppose that this information is shared by means of the following procedure.

Consider person i. In the first round he makes an estimate $u_i(1)$, which is based on his initial data Y_i, and he communicates this estimate to some or all of the other members. By the time i makes his second estimate, he will have received the estimates of some of the others. More generally, denote by $D_i(t-1)$ the messages received by i from the others before i makes his tth estimate $u_i(t)$. That estimate will be based on Y_i and $D_i(t-1)$. We assume that i communicates all his estimates to a fixed set of the other people, and that there is a message transmission delay of one time unit.

Our aim is to discuss two questions: Will each person's estimate converge as $t \to \infty$? If the individual estimates converge, will they reach a common limit? To formulate these questions mathematically, we need to specify how each person estimates the best decision based on the data available to him. This is done in Section 2. Once this is done, it turns out that the answers depend crucially upon whether the prior beliefs are consistent or inconsistent. The consistent case is considered in Section 3, and the inconsistent case in Section 4. Section 5 outlines some directions for further research.

2. ESTIMATION SCHEMES

Several different estimation schemes have been considered in the literature.

Borkar and Varaiya [2] consider the situation where the committee wants to estimate a random variable X, and they suppose that the tth estimate made by i, $u_i(t)$, is the conditional mean of X given the available data, that is,

$$u_i(t) = E^i\{X \mid Y_i, D_i(t-1)\}. \tag{2}$$

Here E^i denotes expectation with respect to P^i. We see later that the right-hand side of (2) has to be interpreted carefully when the beliefs are inconsistent. For the moment observe that the estimate given by (2) is also the decision that minimizes the (expected value of the) cost function

$$J(\omega, u) := |X(\omega) - u|^2$$

when the information available is $\{Y_i, D_i(t-1)\}$. Aumann [1], and Geanakoplos and Polemarchakis [4] consider the situation in which the group wants to estimate the probability that a particular event $F \in \underline{F}$ has occurred. This is a special case of (2) with $X = 1(F)$. The set Ω of all possible states is finite is [1] and [4]. Tsitsiklis and Athans [7] consider the situation described in the introduction. Sebenius and Geanakoplos [5] discuss the following situation. Suppose there are two people 1, 2, and let $F \in \underline{F}$. Person 1 is allowed to offer a bet, which person 2 may accept or reject. If the bet is accepted and if F is true, that is, if $\omega \in F$, then 2 must pay 1 a fixed sum of money; whereas if F is false, then 1 must pay 2 the same amount. It follows that 1 will make the bet if and only if the conditional mean of $1(\underline{F})$ given the data available to him is *greater* than $\frac{1}{2}$, and 2 will accept the bet if and only if the mean conditioned on the data available to 2 is *less* than $\frac{1}{2}$. This situation can also be reformulated as a special case of Eq. (1) as follows. Take the cost function to be

$$J(\omega, u) := |1(F) - u| \tag{3}$$

and suppose the decision u is restricted to the set $U = \{0, 1\}$. Then person 1 will offer the bet if and only if the decision that minimizes the cost (3) is $u = 1$, and person 2 will accept the bet if and only if the cost minimizing decision is $u = 0$.

Washburn and Teneketzis [8] consider a very general communication framework that includes all the situations mentioned above as special cases. They assume merely that in the tth round, person i selects the decision u according to the rule $u = d(Y_i, D_i(t-1))$. The decision rules d are known to all.

3. CONSISTENT BELIEFS

3.1 The Borkar–Varaiya Estimation Scheme

In Sections 3.1 and 3.2 it is assumed that the beliefs are consistent, that is,

$$P^1 = \cdots = P^N := P, \quad \text{say.}$$

Consider first the case studied by Borkar and Varaiya, in which the tth estimate made by i is the mean of X conditioned on $\{Y_i, D_i(t-1)\}$. (It is assumed that $E|X|<\infty$.) Let $\underline{Y}_i(t-1)$ denote the σ-field generated by the $\{Y_i, D_i(t-1)\}$. Then

$$u_i(t) = E\{X \mid \underline{Y}_i(t-1)\}. \tag{4}$$

Let $\underline{Y}_i(\infty) := V_t \underline{Y}_i(t)$. Since $\underline{Y}_i(t)$ is an increasing sequence, it follows from the martingale convergence theorem that

$$u_i(t) \to u_i(\infty)\, a.s.; \quad u_i(\infty) := E\{X \mid \underline{Y}_i(\infty)\}. \tag{5}$$

Thus the individual estimates do converge.

Next we investigate whether the limiting estimates agree. Suppose i communicates his estimates to j. Then $u_i(t)$ is $\underline{Y}_j(t+1)$-measurable. From (5) it follows that $u_i(\infty)$ is $\underline{Y}_j(\infty)$-measurable, and so,

$$u_i(\infty) = E\{u_j(\infty) \mid \underline{Y}_i(\infty) \cap \underline{Y}_j(\infty)\}. \tag{6}$$

Suppose there is a **communication ring** $i_1, \ldots, i_n = i_1$. This is a not necessarily distinct sequence of persons such that i_k communicates his estimates to i_{k+1}. Then, according to (6), we must have

$$u_{i_k}(\infty) = E\{u_{i_{k+1}}(\infty) \mid \underline{Y}_{i_k}(\infty) \cap \underline{Y}_{i_{k+1}}(\infty)\}, \qquad k = 1, \ldots, n, \tag{7}$$

where $i_{n+1} := i_1$. It is quite easy to show [2, lemma 2] that (7) implies

$$u_{i_1} = \cdots = u_{i_n},$$

so that the asymptotic estimates of the members of a communication ring agree. This suggests the main result of [2].

Theorem 1. *If the estimates of i are given by* (2), *then each person's estimate converges. Moreover, if everyone in the team is a member of the same communication ring, then the limiting estimates agree.*

Proof. See Appendix A.

A careful study of the preceding argument reveals that the estimation scheme (2) enjoys two properties:

1. The estimate is "continuous" with respect to a monotonically increasing sequence of data, that is (5) holds.
2. The schemes satisfy the "agreement" property (7).

This observation underlies the work of Washburn and Teneketzis [8], which we discuss next.

3.2 The General Case

The estimate (2) is an instance of a decision rule of the form $u_i(t) = d(Y_i, D_i(t-1))$. This suggests the following abstract definitions:

1. A decision rule is a function that associates to any σ-field $\underline{B} \subset \underline{F}$ a $\underline{B}$-measurable random variable $u = d(\underline{B})$ with values in the feasible set U.
2. A decision rule d is said to be **continuous** if for every increasing $\underline{F}_i(1) \subset \underline{F}_i(2) \subset \cdots$ one has

$$\lim_{k\to\infty} d(\underline{F}_i(k)) = d(V_k \underline{F}_i(k)) \qquad a.s.$$

Let d be a decision rule for the team. This rule recursively generates a sequence of estimates (decisions) by:

$$u_i(t) := d(\underline{Y}_i(t-1)), \qquad i = 1, 2, \ldots, n \tag{8}$$

$$\underline{Y}_i(t) = \sigma\{Y_i, D_i(t-1)\}, \qquad i = 1, 2, \ldots, n \tag{9}$$

where $\sigma(Y_i, D_i(t-1)$ denotes the sub-σ-field of $\underline{F}$ generated by Y_i, $D_i(t-1)$. Let $\underline{Y}_i(\infty) = V_t \underline{Y}_i(t)$ and $u_i(\infty) = d(\underline{Y}_i(\infty))$. The following lemma is immediate from the definition of continuity.

Lemma 1. *If d is a continuous decision rule, then for every i*

$$\lim_{t\to\infty} u_i(t) = u_i(\infty). \tag{10}$$

Suppose $i_1, i_2, \ldots, i_{\eta+1} = i$, is a communication ring. Then for each k, $u_{i_k}(t)$ is measurable with respect to $\underline{Y}_{i_k}(t) \cap \underline{Y}_{i_{k+1}}(t+1) \subset Y_{i_{k+1}}(\infty)$. If d is continuous, this yields a chain of conditions:

$$d(\underline{Y}_{i_k}(\infty)) \text{ is } \underline{Y}_{i_{k+1}}\text{-measurable}, \qquad k = 1, 2, \ldots, \eta \quad i_{\eta+1} = i_1.$$

We say that d satisfies the **agreement condition for rings** if for every sequence $(i_1, \underline{F}_{i_1}) \cdots (i_{\eta+1}, F_{i_{\eta+1}}) = (i_1, \underline{F}_{i_1})$, the chain of conditions

$$d(\underline{F}_{i_k}) \text{ is } \underline{F}_{i_{k+1}}\text{-measurable}, \qquad k = 1, 2, \ldots, \eta$$

implies

$$d(\underline{F}_{i_1}) = d(\underline{F}_{i_2}) = \cdots = d(\underline{F}_{i_\eta}).$$

From Lemma 1 and the earlier definition one can prove the main result of Washburn and Teneketzis [8].

Theorem 2. *If the decision rule d is continuous and satisfies the* *agreement condition for rings, and if everyone is a member of the same communication ring, then the individual estimates converge to the same limit.*

Proof. See Appendix B.

If the decision rule d is the conditional expectation as in Eq. (2), then the continuity requirement is just the martingale convergence property; the agreement condition for rings is also satisfied, as shown by Lemma A.3 of Appendix A. Washburn and Teneketzis [8] show also that the agreement condition for rings is satisfied by decision rules that are optimal in the sense defined below.

Proposition 1. *Suppose that d is a decision rule such that $\sigma(d(\underline{F}')) \subset F'$ for all σ-fields $\underline{F}' \subset \underline{F}$. Then d satisfies the agreement condition for rings if and only if there is a partial ordering $\leq$ of the set of functions $\{d(\underline{F}') : \underline{F}' \subset F\}$ such that $d(\underline{F}')$ is the maximum element of $\{d(\underline{G}) : \underline{G} \subset \underline{F}, \sigma(d(\underline{G})) \subset F'\}$ with respect to $\leq$.*

Proof. See Appendix C.

In many cases, as in [2], [7], the partial order relation is defined in terms of a scalar cost function. The following proposition proves that decision rules defined by such cost functions satisfy the agreement condition for rings, provided that the decision includes a tie-breaking rule if the cost function has more than one minima.

Proposition 2. *Suppose that the decision functions take values in a set U. Let L be a real-valued functional of $\underline{F}$-measurable decision functions $\delta : \Omega \to U$. For each $\underline{F}'$ let $D(\underline{F}')$ be the set of $\underline{F}'$-measurable decision functions δ such that $L(\delta) \leq L(\delta')$ for all measurable $\delta' : \Omega \to U$. Assume that U is partially ordered by $\leq'$ and that for each $\underline{F}'$, there is a $\delta \in D(\underline{F}')$ such that $\delta' \in D(\underline{F}')$ implies $\delta'(\omega) \leq' \delta(\omega)$ for all $\omega \in \Omega$. The decision rule $d(\underline{F}')$ that assigns this $\delta \in D(\underline{F}')$ to $\underline{F}'$ satisfies the agreement condition for rings.*

Proof. See Appendix D.

Typically, $L(\delta)$ will have the form

$$L(\delta) = E\{J(\delta(\omega), \omega)\} = E\{J(u, \omega)\}. \tag{11}$$

Under assumptions that guarantee that L has $\underline{F}'$-measurable minima (i.e., $D(\underline{F}') \neq \phi$), Proposition 2 implies that the decision rule minimizing $L(\delta)$, that is,

$$d(\underline{F}_i) = \arg\min_{u \in U} E\{J(\omega, u) | \underline{F}_i\}, \tag{12}$$

satisfies the agreement condition.

In addition, when $\underline{F}_i(t) \uparrow \underline{F}_i$ under fairly weak conditions $E\{J(\cdot, d(\underline{F}_i(t))) | \underline{F}_i(t)\}$ is a convergent supermartingale, converging to $E\{J(\cdot, d(\underline{F}_i)) | \underline{F}_i\}$. Under certain conditions, such as uniform strict convexity of J, U a finite dimensional compact convex set, [7], convergence of the optimal costs implies convergence of the arguments $d(\underline{F}_i(t)) \rightarrow d(\underline{F}_i)$. Thus, the decision rule described by Eq. (12) satisfies the continuity condition. Consequently, Theorem 2 applies to the situation studied by Tsitsiklis and Athans [7].

REMARKS

1. The agreement condition of Theorem 2 is the essence of the agreement problem. The continuity condition is necessary to deal with cases involving infinite measurement σ-fields.
2. In [2, 7, 8] the message exchange model is extended to accommodate the situation where (1) i makes a sequence of observations $Y_i(t)$, $t = 1, 2, \cdots$ and not just the initial observation, and (2) i communicates his estimate $u_i(t)$ to a randomly selected set of the other people, and message transmission takes a random amount of time. These extensions can be easily incorporated in the analysis and Theorems 1 and 2 continue to hold with some minor modifications.

3.3 Common Knowledge

The main feature of the estimation schemes presented in [1], [2], [4], [8] is the following. If all team members use the same decision rule, if everyone in the team is a member of the same communication ring, and if common knowledge decisions agree, then all team members agree on the same decision. The common decision is the decision based on the ultimate common knowledge (common information) of the team members.

Thus, it appears appropriate to define common knowledge at this point, and to show that the definitions of common information given in [1], [2], [8], [9] are essentially equivalent and lead to the same results.

Aumann [1] represents information by a partition $\underline{P}$ on the sample space Ω. Borkar and Varaiya [2] and Washburn and Teneketzis [8] represent information by σ-fields contained in $\underline{F}$. It can be shown that these two representations are essentially equivalent. The partition $\underline{P}$ is a collection $\{E_1, E_2, \ldots\}$ of mutually disjoint events whose union is the whole sample space. To a partition $\underline{P}$ there corresponds a unique σ-field $\underline{F}$, namely the σ-field generated by the events in $\underline{P}$. Each $E_i \in \underline{P}$ is an atom of $\underline{F}$. If $\underline{P}_1 = \{E_1, E_2, \ldots\}$ and $\underline{P}_2 = \{G_1, G_2, \ldots\}$, then one can define a third partition $\underline{P}_3$ that is the finest partition contained in $\underline{P}_1$ and $\underline{P}_2$ and is denoted by $\underline{P}_1 \wedge \underline{P}_2$. If $\underline{P}_1$ and $\underline{P}_2$ correspond to the σ-fields $\underline{F}_1$ and $\underline{F}_2$, then $\underline{P}_1 \wedge \underline{P}_2$ corresponds to $\underline{F}_1 \wedge \underline{F}_2$. Aumann [1] defines an event E to be common knowledge to team members 1 and 2 (with information $\underline{P}_1$ and $\underline{P}_2$, respectively) at ω if there is an atom $\hat{G} \in \underline{P}_1 \wedge \underline{P}_2$ such that $\omega \in \hat{G} \subset E$. If $\underline{F}_1$ and $\underline{F}_2$ are the σ-fields corresponding to $\underline{P}_1$ and $\underline{P}_2$, respectively, then the definitions of common knowledge at ω given in [2], [8], namely that there is $\hat{G} \in \underline{F}_1 \wedge \underline{F}_2$ and $\omega \in \hat{G} \subset E$, are equivalent to Aumann's definition. Let us say that the event E is common knowledge to the team members 1 and 2 if it is common knowledge at each $\omega \in E$. Then E is common knowledge to 1 and 2 if and only if it belongs to the σ-field generated by $\underline{P}_1 \wedge \underline{P}_2$, namely $\underline{F}_1 \wedge \underline{F}_2$.

Milgrom [9] characterizes common knowledge by

1. associating with each event E another event K_E with the interpretation

$$K_E = \{\omega \in \Omega\colon E \text{ is common knowledge at } \omega\}$$

2. considering the following four conditions:
   - (C1) $K_E \subset E$
   - (C2) $\forall \omega \in K_E, \forall i$, if $\omega \in F_i$, $F_i \subset K_E$
   - (C3) $E_1 \subset E_2 \Rightarrow K_{E_1} \subset K_{E_2}$
   - (C4) $[\forall i, \forall \omega \in E$, if $\omega \in F_i$, $F_i \in E] \Rightarrow E = K_E$.

Condition (C1) asserts that an event E is common knowledge only if it actually occurs. Condition (C2) implies that if E is common knowledge, then every team member knows that E is common knowledge. Conditions (C1) and (C2) imply that E is common knowledge only if E occurs, each team member knows E, each knows that all know E, and so on. Condition (C3) implies that wherever E_1 is common knowledge any logical consequence of E_1 is also common knowledge. Condition (C4) asserts that public events are common knowledge whenever they occur. A public event is defined by the antecedent in (C4); it is an event that if it occurs, it will be known to every team member. Milgrom [9] shows that his characterization of common knowledge is equivalent to Aumann's definition.

The definitions of [1], [2], [8], [9] are equivalent and lead to equivalent agreement conditions. In [1] Aumann notes the simple fact that if the conditional probability functions $P(E|\underline{F}_1)$ and $P(E|\underline{F}_2)$ are common knowledge to 1 and 2 (i.e., they are both $\underline{F}_1 \wedge \underline{F}_2$-measurable), then $P(E|\underline{F}_1) = P(E|\underline{F}_2)$. The following proposition shows that this statement of Aumann is equivalent to the condition for agreement in [8].

Proposition 3. *The conditions*

$$\sigma(d(\underline{F}_1)) \subset \underline{F}_2 \subset \underline{F}_1 \Rightarrow d(\underline{F}_1) = d(\underline{F}_2) \tag{13}$$

of [8] *and*

$$\sigma(d(\underline{F}_1)) \vee \sigma(d(\underline{F}_2)) \subset \underline{F}_1 \cap \underline{F}_2 \Rightarrow d(\underline{F}_1) = d(\underline{F}_2) \tag{14}$$

of [1] *are equivalent.*

Proof. See Appendix E.

In Aumann's terminology, [1], Eq. (14) asserts that if the decisions $d(\underline{F}_1)$ and $d(\underline{F}_2)$ are common knowledge to 1 and 2, then they are equal. Aumann showed that this condition is satisfied by the decision rule

$$d(\underline{F}) = P(E|\underline{F}).$$

It is possible to show [8] that this condition is also satisfied by the Maximum a posteriori (MAP) decision rule. The conditional expectation satisfies condition (13) because of the following fact: If $\underline{G} \subset \underline{F}$ and some version of $E(X|\underline{F})$ is measurable w.r.t. $\underline{G}$, then $E(X|\underline{F}) = E(X|\underline{G})$ with probability one.

The idea of common knowledge proved to be useful in game theory and various areas of mathematical economics. Wilson [11] studied allocation problems under differential information, and defined an efficient allocation in a world of differential information in a way that can be stated succinctly using common knowledge: A contingent allocation f is efficient if there is no other allocation v such that it is common knowledge that all agents prefer v to f. Milgrom [9] and Milgrom and Stokey [10] used the idea of common knowledge to analyze a rational expectations trading model. They showed that when traders exchange a risky security on the basis of private information then they "agree to disagree" (i.e., no trade takes place). Kreps et al. [12] consider finite repetitions of the well-known prisoners' dilemma game. A common observation in experiments involving finite repetitions of the prisoners' dilemma is that players do not always play the single-period dominant strategies, but instead achieve some measure of cooperation. Kreps and his co-authors in [12] show that the lack of common knowledge about one or both players' options, motivation, or behavior can explain the observed cooperation.

4. INCONSISTENT BELIEFS

The analysis is quite different when the beliefs are inconsistent. The discussion in this section is initially based on Teneketzis and Varaiya [6]. Then the results of [6] are extended to the case of a general decision rule d. To keep the notation simple assume there are only two persons, *Alpha* and *Beta*. Initially, Alpha observes the random variable A and Beta observes B. Both wish to estimate the random variable X. We also assume that Ω is finite. The prior probabilities of Alpha and Beta are denoted P^α, P^β, respectively.

For $t = 1, 2, \ldots$ the tth estimate by Alpha(Beta) is denoted $\alpha_t(\beta_t)$. The term α_t is the conditional expectation of X given the observations $A, \beta_1, \ldots, \beta_{t-1}$. After α_t has been calculated it is communicated to Beta whose tth estimate is the conditional expectation of X given $B, \alpha_1, \ldots, \alpha_t$. Once β_t is evaluated it is communicated to Alpha, who incorporates it into the estimate a_{t+1}, and the procedure is repeated.

To complete the specification we assume that the estimation procedures followed by Alpha and Beta are consistent with their own prior models. That is, each assumes the other's model to be the same as his own. Consider Alpha. When he receives Beta's estimate β_{t-1}, Alpha interprets it as if it was based on P^α rather than on P^β. Thus Alpha assumes that Beta's estimate is a realization of the random variable

$$\hat{\beta}_{t-1} := E^\alpha\{X \mid B, \alpha_1, \ldots, \alpha_{t-1}\}.$$

Subsequently, Alpha calculates α_t,

$$a_t := E^\alpha\{X \mid A, \hat{\beta}_1, \ldots, \hat{\beta}_{t-1}\}.$$

Symmetrically, Beta interprets α_t as

$$\hat{a}_t := E^\beta\{X \mid A, \beta_1, \ldots, \beta_{t-1}\},$$

and calculates β_t by

$$\beta_t := E^\beta\{X \mid B, \hat{\alpha}_1, \ldots, \hat{\alpha}_t\}.$$

There is a more revealing description of the functional dependence of these estimates. Suppose a particular realization $\bar{\omega} = (\bar{A}, \bar{B})$ has occurred. Since Alpha observes $\bar{A}$, he concludes that $\bar{\omega} \in \Omega_1^\alpha := \{(A, B) \mid A = \bar{A}\}$, and so his first estimate equals

$$\bar{\alpha}_1 = E^\alpha\{X \mid A = \bar{A}\} = E^\alpha\{X \mid \omega \in \Omega_1^\alpha\}.$$

Alpha transmits the number $\bar{\alpha}_1$ to Beta. Beta interprets it as a realization of the random variable

$$\hat{\alpha}_1 = E^\beta\{X \mid A\},$$

and so he infers that $\bar{\omega} \in \Omega_1^\beta := \{\omega | \hat{\alpha}_1(\omega) = \bar{\alpha}_1, B = \bar{B}\}$, and his first estimate takes the value

$$\bar{\beta}_1 = E^\beta\{X | \omega \in \Omega_1^\beta\}.$$

This value is communicated to Alpha.

At the beginning of the tth round, Alpha starts with the inference $\bar{\omega} \in \Omega_{t-1}^\alpha$ when he receives the estimate $\bar{\beta}_{t-1}$. He interprets it as a realization of the random variable

$$\hat{\beta}_{t-1} = E^\alpha\{X | B, \alpha_1, \ldots, \alpha_{t-1}\},$$

and so Alpha concludes that $\bar{\omega} \in \Omega_t^\alpha := \{\omega | \omega \in \Omega_{t-1}^\alpha, \hat{\beta}_{t-1}(\omega) = \bar{\beta}_{t-1}\}$. Hence, Alpha's tth estimate takes the value

$$\bar{\alpha}_t = E^\alpha\{X | \omega \in \Omega_t^\alpha\},$$

which is communicated to Beta. Whereupon Beta interprets it as a realization of

$$\hat{\alpha}_t = E^\beta\{X | A, \beta_1, \ldots, \beta_{t-1}\},$$

concludes that $\bar{\omega} \in \Omega_t^\beta := \{\omega | \omega \in \Omega_{t-1}^\beta, \hat{\alpha}_t(\omega) = \bar{\alpha}_t\}$, and evaluates his tth estimate as

$$\bar{\beta}_t = E^\beta\{X | \omega \in \Omega_t^\beta\}.$$

Thus, as expected, the uncertainty diminishes with each exchange, $\Omega_{t+1}^\alpha \subset \Omega_t^\alpha$, $\Omega_{t+1}^\beta \subset \Omega_t^\beta$. From the previous description we also see that if for some k either $\Omega_{k+1}^\alpha = \Omega_k^\alpha$ or $\Omega_{k+1}^\beta = \Omega_k^\beta$, then $\Omega_t^\alpha = \Omega_{k+1}^\alpha$ and $\Omega_t^\beta = \Omega_{k+1}^\beta$ for $t > k+1$. Hence for $t > T$ (which cannot exceed the number of distinct elements in Ω), Ω_t^α and Ω_t^β become constant. These limit sets depend upon the realization ω. Call them $\Omega_*^\alpha(\omega)$ and $\Omega_*^\beta(\omega)$, respectively.

There are two possibilities. The first is that $\Omega_*^\alpha(\omega) = \phi$ and $\Omega_*^\beta(\omega) = \phi$. This happens because at some stage the message $\bar{\beta}_{t-1}$ received by Alpha is "impossible": there is no $\bar{\omega}$ such that $\hat{\beta}_{t-1}(\bar{\omega}) = \bar{\beta}_{t-1}$; or the message $\bar{\alpha}_t$ received by Beta is "impossible": there is no $\bar{\omega}$ such that $\alpha_t(\bar{\omega}) = \bar{\alpha}_t$. Alpha and Beta must realize that their prior models are inconsistent. Let Ω_I be the set of all realizations that lead to this outcome.

The second possibility is that $\Omega_*^\alpha(\omega) \neq \phi$ and $\Omega_*^\beta(\omega) \neq \phi$. In this case for $t > T$ the estimates stop changing: $\hat{\beta}_t(\omega) = \hat{\beta}_*(\omega)$, $\alpha_t(\omega) = \alpha_*(\omega)$, $\hat{\alpha}_t(\omega) = \hat{\alpha}_*(\omega)$, $\beta_t(\omega) = \beta_*(\omega)$. Since for every t, $\hat{\beta}_t(\omega) = \beta_t(\omega)$ and $\hat{\alpha}_t(\omega) = \alpha_t(\omega)$, it follows that

$$\hat{\beta}_*(\omega) = \beta_*(\omega), \qquad \hat{\alpha}_*(\omega) = \alpha_*(\omega).$$

On the other hand, since $\hat{\beta}_t$ and α_t are based on the same model, namely P^α, it follows from Theorem 1 that $\hat{\beta}_*(\omega) = \alpha_*(\omega)$. For the same reason $\hat{\alpha}_*(\omega) = \beta_*(\omega)$. Thus if $\omega \in \Omega_{II} := \Omega - \Omega_I$, there is agreement $\alpha_t(\omega) =$

$\beta_t(\omega)$ for $t > T$. *It is worth emphasizing that this agreement need not be a reflection of the consistency of the two models P^α, P^β. Rather agreement occurs because within each person's model there is sufficient "uncertainty" to permit the reconciliation of the other's messages with his own observation. One might say that agreement could result from two wrong arguments.* We summarize the preceding analysis as follows:

Theorem 3. *The set of events Ω decomposes into two disjoint subsets Ω_I and Ω_{II}. After T exchanges, if $\omega \in \Omega_I$ both agents realize their models are inconsistent, whereas if $\omega \in \Omega_{II}$ the two estimates coincide.*

The result is fragile. In particular, whether a realization ω ends in agreement or in impasse can depend upon the order of communication between Alpha and Beta as demonstrated by the following example.

EXAMPLE. Take $\Omega = [0, 2] \times [0, 3]$, suppose Alpha observes

$$A = \{1(a_1), 1(a_2)\}$$

and Beta observes

$$B = \{1(b_1), 1(b_2), 1(b_3)\}$$

and suppose X is the indicator function of the shaded region as shown in Figure 1. Assume that ω is uniformly distributed under P^α, whereas under P^β,

$$P^\beta(b_1) = \tfrac{2}{12}, \qquad P^\beta(b_2) = \tfrac{3}{12}, \qquad P^\beta(b_3) = \tfrac{7}{12}$$

and within each b_i, ω is uniformly distributed. Suppose $\bar{\omega} \in a_1 \cap b_3$ and that Alpha communicates first. Then

$$\bar{a}_n = E(X \mid \omega \in a_1) = \tfrac{1}{2}.$$

Beta interprets this as a realization of

$$\hat{a}_1 = E(X \mid 1(a_1)1(a_2)).$$

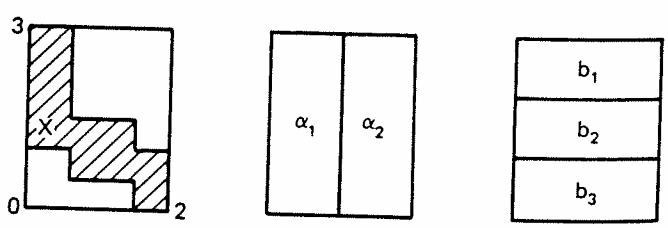

Figure 1

Since $E^{\beta}(x|\ \omega \in a_1) = \frac{5}{12}$, $E^{\beta}(x|\ \omega \in a_2) = \frac{1}{2}$, upon learning that $\bar{a}_1 = \frac{1}{2}$, Beta concludes that $\bar{\omega} \in a_2$, and since he has observed that $\bar{\omega} \in b_3$ his estimate is

$$\bar{\beta}_1 = E^{\beta}(X|\ \omega \in a_2 \cap b_3) = \tfrac{3}{4}.$$

Alpha interprets $\bar{\beta}_1$ as a realization of $E^{\alpha}\ (x|\ \omega \in a_1, B)$. Since

$$E^{\alpha}(x|\ \omega \in a_1 \cap b_1) = \tfrac{1}{2},$$
$$E^{\alpha}(x|\ \omega \in a_1 \cap b_2) = \tfrac{3}{4},$$
$$E^{\alpha}(x|\ \omega \in a_1 \cap b_3) = \tfrac{1}{2},$$

Alpha concludes that $\bar{\omega} \in a_1 \cap b_2$, hence

$$\bar{a}_2 = E^{\alpha}(X|\ \omega \in a_1 \cap b_2) = \tfrac{3}{4}.$$

Evidently, $\bar{\beta}_2 = \bar{\beta}_3 = \cdots = \bar{a}_2 = \bar{a}_3 = \cdots = \frac{3}{4}$ and there is agreement. (Note that Alpha believes that $\omega \in a_1 \cap b_2$, Beta believes that $\bar{\omega} \in a_2 \cap b_3$, and in fact $\bar{\omega} \in a_1 \cap b_3$.)

Now suppose again that $\bar{\omega} \in a_1 \cap b_3$, but this time Beta communicates first. Then

$$\bar{\beta}_1 = E^{\beta}(X|\ \omega \in b_3) = \tfrac{1}{2}.$$

Since

$$E^{\alpha}(X|\ \omega \in b_1) = \tfrac{1}{4},$$
$$E^{\alpha}(X|\ \omega \in b_2) = E^{\alpha}(X|\ \omega \in b_3) = \tfrac{1}{2},$$

upon learning $\bar{\beta}_1 = \frac{1}{2}$, Alpha concludes that $\bar{\omega} \in b_2 \cup b_3$. Then his estimate is

$$\bar{a}_1 = E^{\alpha}(X|\ \omega \in a_1 \cap (b_2 \cup b_3)) = \tfrac{1}{2}.$$

But Beta expects $\bar{a}_1$ to take on the value

$$E^{\beta}(X|\ \omega \in a_1 \cap (b_2 \cup b_3)) = 0.4$$

or

$$E^{\beta}(X|\ \omega \in a_1 \cap (b_2 \cup b_3)) = 0.6.$$

Thus Beta concludes that the models are inconsistent.

The results of Teneketzis and Varaiya [6] can be extended to the case where the decision rule is a general function d, as in Section 3.2. We discuss this case next.

Assume the same model as in Teneketzis and Variaya [6] and suppose the estimates α_t and β_t are generated by the decision rule d given the observations $A, \beta_1, \beta_2, \ldots, \beta_{t-1}$ and $B, \alpha_1, \alpha_2, \ldots, \alpha_{t-1}$, respectively.

Suppose the decision rule d satisfies the **agreement condition**: for all $\underline{G}_1, \underline{G}_2 \subset \underline{F}$,

$$\sigma(d(\underline{G}_2)) \subset \underline{G}_1, \subset \underline{G}_2 \Rightarrow d(\underline{G}_1) = d(\underline{G}_2). \tag{15}$$

Under the previous assumptions one can prove the following result.

Theorem 4. *If Ω is finite and the decision rule d satisfies the agreement condition* (15), *then either the estimates α and β agree after a finite number of communications or Alpha and Beta realize that their models are inconsistent.*

Proof. See Appendix F.

As pointed out in the discussion previously the investigation of convergence and agreement of the estimates can proceed in two steps.

1. Determine what each team member's model predicts about the evolution and the outcome of the estimation process.
2. Examine how these predictions compare with what actually happens during the estimation process.

For finite Ω, the result of Theorem 4 is true for rules that obey the agreement condition for a very simple reason. If a team member's view of the world is consistent with reality, then agreement must result after a finite number of communications, because this is what is predicted by the team member's model; anything else would be inconsistent.

5. CONCLUDING REMARKS

Recall the discussion in Sections 3.1 and 3.2. There a consensus is reached via a sequence of exchanges of tentative decisions. The information available to a person increases with each message exchange and the limiting consensus decision is based on the information common to all in the sense that $d_1(\infty) = \cdots = d_N(\infty)$ is measurable with respect to $\underline{Y}_1(\infty) \cap \cdots \cap \underline{Y}_N(\infty)$. A consensus can also be reached if all people share their initial private data $\underline{Y}_1, \ldots, \underline{Y}_N$. We may call this consensus the **full information** decision. It turns out that the consensus reached by exchanging tentative decisions need *not* coincide with the full information decision. Within a rather simple model, however, Geanakoplos and Polemarchakis [4] have shown that the two decisions are "almost always" the same. It would be worth investigating this in a more general setting.

Secondly, even when the two decisions are the same, it does not follow that all people obtain the full information, that is, it need not be the case that $\underline{Y}_1(\infty) = \sigma\{Y_1, \ldots, Y_N\}$. If $\underline{Y}_1(\infty)$ is a proper subset of $\sigma\{Y_1, \ldots, Y_N\}$, then one could argue that reaching consensus via exchange of tentative decisions requires a transfer of less information than the exchange of all private information. This too is worth further investigation.

Recall now the discussion dealing with the case of inconsistent beliefs. The most interesting finding is that Alpha and Beta can exchange statements about X and eventually agree even when their views are different. Thus paradoxically, the realization that these views are different is only reached when further communication becomes impossible. This raises several basic and knotty issues that need further investigation.

One can readily imagine situations where the most important thing is to determine whether or not the beliefs are inconsistent. In the communication setup of Section 4 the realization that beliefs are inconsistent is fortuitous—it happens only if Alpha and Beta reach an impasse. How should one structure the set of message exchanges so as to expedite the reaching of an impasse?

Suppose now that Alpha and Beta do reach an impasse ($\omega \in \Omega_I$). Our analysis stops at this point, but there are two directions that can be pursued. First, observe that with the realization that their beliefs are different comes the understanding that they have "misread" each other's messages (i.e., they now know that $\hat{\beta}_\eta \not\equiv \beta_\eta$ and $\hat{\alpha}_\eta \not\equiv \alpha_\eta$), and consequently their estimates have been "biased." To eliminate this bias each needs to learn what the other's view is. A straightforward way of permitting such learning is to suppose that from the beginning Alpha admits that Beta's model P^β might be any one of a known set $\underline{P}^\beta$ of models and there is a prior distribution on $\underline{P}^\beta$ reflecting Alpha's initial judgment about Beta's model; a symmetrical structure is formulated for Beta. Within such a framework it seems reasonable to conjecture that each agent will correctly read the other's message and his sequence of estimates will converge. But if their models are different, then the limiting estimates may differ, and a consensus will not emerge.

Suppose, however, that Alpha and Beta want to reach a consensus. To reach a consensus one or both must change their models. One can imagine many different ways in which this can be done. For example, De Groot [3] proposes that each person tell the others what his prior probability is, and he proposes an ad hoc behavioral rule whereby each person adjusts his model to a weighted average of the others' models. This is not very satisfactory in situations where communicating one's prior beliefs is not practicable.

APPENDIX A

Proof of Theorem 1

Convergence of each member's estimates follows from the martingale convergence theorem. The proof of the rest of the theorem proceeds in several steps. Consider two agents i and j and let $G_i(t)$ denote the σ-field generated by the transmission and reception of messages from agent i up to time t. That is,

$$G_i(t) = \sigma\{u_1(1), \ldots, u_1(t-1), \ldots, u_{i-1}(1), \ldots, u_{i-1}(t-1), u_i(1), \ldots, u_i(t), u_{i+1}(1), \ldots, u_{i+1}(t-1), \ldots, u_\eta(1), \ldots, u_\eta(t-1)\}$$

$$G_j(t) = \sigma\{u_1(1), \ldots, u_1(t-1), \ldots, u_{j-1}(1), \ldots, u_{j-1}(t-1), u_j(1), \ldots, u_j(t), u_{j+1}(1), \ldots, u_{j+1}(t-1), \ldots, u_\eta(1), \ldots, u_\eta(2), \ldots, u_\eta(t-1)\}.$$

Define S^{ij} to be the event that agent i sends messages to j infinitely often.

Lemma A1. *Both $u_i(\infty)1(S^{ij})$ and $u_j(\infty)1(S^{ij})$ are common knowledge for $G_i(\infty)$ and $G_j(\infty)$. Moreover,*

$$u_i(\infty)1(S^{ij}) = E(X|\ G_i(\infty) \cap G_j(\infty))1(S^{ij}) \quad a.s. \tag{A.1}$$

and

$$u_i(\infty)1(S) = u_j(\infty)1(S) \tag{A.2}$$

where

$$S = S^{ij} \cap S^{ji}. \tag{A.3}$$

Proof. Since there is a message transmission delay of one unit, it follows that S^{ij} is in $G_i(\infty)$ and $G_j(\infty)$. Since $u_i(t)$ is $G_j(t+1)$-measurable it follows $u_i(\infty)$ is $G_j(\infty)$-measurable. Similarly $u_j(\infty)$ is $G_i(\infty)$-measurable.

Consequently,

$$u_i(\infty)1(S^{ij}) = E(X|\ G_i(\infty) \cap G_j(\infty))1(S^{ij}).$$

Similarly,

$$u_j(\infty)1(S^{ji}) = E(X|\ G_i(\infty) \cap G_j(\infty))1(S^{ji}).$$

Hence,

$$u_i(\infty)1(S) = u_j(\infty)1(S).$$

■

To proceed further we need the following result.

Lemma A2. *Let* $z_1, z_2, \ldots, z_{\eta+1} = z_1$ *be random vectors and*

$\underline{F}_1, \underline{F}_2, \ldots, \underline{F}_\eta$ be σ-fields such that

$$z_i = E(z_{i+1} \mid \underline{F}_i), \qquad i = 1, 2, \ldots, \eta. \tag{A.4}$$

Then

$$z_1 = z_2 = z_3 = \cdots = z_\eta \quad \text{a.s.} \tag{A.5}$$

Proof. We can assume that z_i are scalars, since by applying the same argument to each component we can generalize the result to random vectors. Suppose first that each z_i is square integrable. Since conditional expectation is the best mean square estimate and $z_i = E(z_{i+1}/ F_i)$, it follows that

$$E|z_{i+1}|^2 = E|z_i|^2 + E|z_{i+1} - z_i|^2, \qquad i = 1, 2, \ldots, \eta.$$

Adding the above relations and using $z_{n+1} = z_1$ we get

$$0 = \sum_{i=1}^{\eta} E|z_{i+1} - z_i|^2.$$

Consequently $z_1 = z_2 = z_3 = \cdots = z_\eta$. Thus, Lemma A2 holds for square integrable random variables. To complete the proof of Lemma A2, for any number K let $z_i^K = \min\{z_i, K\}$. Then, by Jensen's inequality, $z_i = E\{z_{i+1} \mid \underline{F}_i\}$ implies

$$z_i^K \geq E\{z_{i+1}^K \mid F_i\}, \quad \forall_i. \tag{A.6}$$

The last inequality implies

$$Ez_1^K \geq Ez_2^K \geq \cdots \geq Ez_\eta^K \geq Ez_{\eta+1}^K = Ez_1^K.$$

Consequently, (A.6) holds with equality.

Therefore, for $k_1 > k_2$

$$z_i^{k_1} - z_i^{k_2} = E(z_{i+1}^{k_1} - z_{i+1}^{k_2} \mid \underline{F}_i).$$

Since $z_i^{k_1} - z_i^{k_2}$ is bounded, it is square integrable; therefore

$$z_1^{k_1} - z_1^{k_2} = z_2^{k_1} - z_2^{k_2} = \cdots = z_\eta^{k_1} - z_\eta^{k_2}.$$

Lemma 2 follows by letting $k_1 \to \infty$ and $k_2 \to \infty$. ∎

Lemmas A1 and A2 can now be used to prove the following result.

Lemma A3. *Suppose that*

1. *$i_1, i_2, \ldots, i_{n+1} = i$, form a communication ring for S, and*
2. *$1(S)$ is common knowledge for $G_1(\infty), G_2(\infty), \ldots, G_\eta(\infty)$.*

Then $u_i(\infty)$ agree on S, that is,

$$u_1(\infty)1(S) = u_2(\infty)1(S) = \cdots = u_\eta(\infty)1(S) \quad \text{a.s.} \tag{A.7}$$

Proof. By Lemma A1

$$
\begin{aligned}
u_i(\infty)1(S^{i(i+1)}) &= E(X \mid G_i(\infty) \cap G_{i+1}(\infty)) \\
&= E\{u_{i+1}(\infty) \mid G_i(\infty) \cap G_{i+1}(\infty)\}1(S^{i(i+1)}) \qquad \text{(A.8)} \\
&= E\{u_{i+1}(\infty)1(S^{i(i+1)}) \mid G_i(\infty) \cap G_{i+1}(\infty)\}.
\end{aligned}
$$

By hypothesis (2) $S \subset G_i(\infty)$ and $S \subset S^{i(i+1)}$. Multiplication of both sides of (A.8) by $1(S)$ gives

$$u_i(\infty)1(S) = E\{u_{i+1}(\infty)1(S) \mid G_i(\infty) \cap G_{i+1}(\infty)\}. \qquad \text{(A.9)}$$

Equation (A.9) and Lemma A2 imply

$$u_1(\infty)1(S) = u_2(\infty)1(S) = \cdots = u_\eta(\infty)1(S).$$

Lemma A3 can now be used to prove the following result.

Lemma A4. *Under the hypothesis of Lemma A3*

$$u_i(\infty)1(S) = E\{X \mid G_1(\infty) \cap G_2(\infty) \cap \cdots \cap G_\eta(\infty)\}1(S). \qquad \text{(A.10)}$$

Proof. By Eq. (A.9)

$$u_i(\infty)1(S) = E\{X1(S) \mid G_i(\infty) \cap G_{i+1}(\infty)\}. \qquad \text{(A.11)}$$

By Lemma A3

$$u_i(\infty)1(S) = u_j(\infty)1(S),$$

thus, $u_i(\infty)1(S)$ is common knowledge for $G_1(\infty), G_2(\infty), \ldots, G_\eta(\infty)$. Taking conditional expectation with respect to $G_1(\infty) \cap G_2(\infty) \cap \cdots \cap G_\eta(\infty)$, we obtain

$$
\begin{aligned}
u_i(\infty)1(S) &= E\{X1(S) \mid G_i(\infty) \cap G_2(\infty) \cap \cdots \cap G_\eta(\infty)\} \\
&= E\{X \mid G_1(\infty) \cap G_2(\infty) \cap \cdots \cap G_\eta(\infty)\}1(S)
\end{aligned}
$$

since by hypothesis $1(S)$ is common knowledge for $G_1(\infty)$, $G_2(\infty), \ldots, G_\eta(\infty)$.

■

The assertion of Theorem 1 now follows from Lemma A4, since $1(\Omega)$ is common knowledge for all persons. The estimate of each agent converges to $E(X \mid G_1(\infty) \cap G_2(\infty) \cap \cdots \cap G_\eta(\infty))$.

APPENDIX B

Proof of Theorem 2

The information of team member i is described by the σ-field $\underline{Y}_i$. The σ-fields $\underline{Y}_i$ evolve dynamically as follows:

$$\underline{Y}_i(t+1) = \underline{Y}_i(t) \vee \bigvee_{j\in[i]} \sigma(d(\underline{Y}_j(t))), \qquad (i = 1, 2, \ldots, \eta) \qquad \text{(B.1)}$$ [2]

with initial condition

$$\underline{Y}_i(0) = y_i(0), \qquad (i = 1, 2, \ldots, \eta) \qquad \text{(B.2)}$$

where $[i]$ is the set of team members with whom i communicates either directly or indirectly. By assumption all the team members belong to the same communication ring; thus, (B.1) can be written as

$$\underline{Y}_i(t+1) = \underline{Y}_i(t) \vee \bigvee_{j\neq i} \sigma(d(\underline{Y}_j(t))), \qquad i = 1, 2, \ldots, \eta. \qquad \text{(B.3)}$$

Since $\underline{Y}_i(t) \uparrow \underline{Y}_i(\infty)$, it follows by the continuity of the decision rule d that

$$\lim_{t\to\infty} u_i(t) = u_i(\infty). \qquad \text{(B.4)}$$

Then Eqs. (B.1) and (B.2) imply that for each k, j we have

$$\sigma(d(\underline{Y}_k(t))) \subset \underline{Y}_j(\infty)$$

and

$$\sigma(d(\underline{Y}_k(\infty))) \subset \underline{Y}_j(\infty)$$

Then the agreement condition for rings implies that

$$u_1(\infty) = u_2(\infty) = u_3(\infty) = \cdots = u_\eta(\infty) = d(\bigcap_i \underline{Y}_i(\infty)).$$

APPENDIX C

Proof of Proposition 1

At first we show that if

$$d(\underline{F}') = \max\{d(\underline{G}): \underline{G} \subset \underline{F}, \sigma(d(\underline{G})) \subset \underline{F}'\}, \qquad \text{(C.1)}$$

then d satisfies the agreement condition for rings. Suppose that $\sigma(d(\underline{F}')) \subset \underline{G} \subset \underline{F}'$. Because $\underline{G} \subset \underline{F}'$, $d(\underline{G}) \leq d(\underline{F}')$. Since $\sigma(d(\underline{F}')) \subset \underline{G}$, it is clear that $d(\underline{F}') \in \{d(\underline{H}): \underline{H} \subset \underline{F}, \sigma(d(\underline{H})) \subset \underline{G}\}$. Thus, $d(\underline{F}') \leq d(\underline{G})$. The relation $\leq$ is a partial order; consequently $d(\underline{G}) \leq d(\underline{F}')$ and $d(\underline{F}') \leq d(\underline{G})$ imply $d(\underline{F}') = D(\underline{G})$. Hence, d satisfies the agreement condition for pairs and in particular $d(\sigma(d(\underline{F}'))) = d(\underline{F}')$. Suppose $\underline{F}_1 = \underline{F}_{\eta+1}$ and $\sigma(d(\underline{F}_K)) \subset \underline{F}_{K+1}$ for $1 \leq K \leq \eta$. Then

$$d(\underline{F}_K) = d(\sigma(d(\underline{F}_K))) \leq d(\underline{F}_{K+1})$$

for each K, hence $d(\underline{F}_1) \leq d(\underline{F}_K) \leq d(\underline{F}_{\eta+1}) = d(\underline{F}_1)$, and so $d(\underline{F}_1) = d(\underline{F}_K)$

for all K. This shows that the decision rule d defined by Eq. (C.1) satisfies the agreement condition for rings.

Conversely, suppose that d satisfies the agreement condition for rings. Define the partial order $\leq$ on $\{d(\underline{F}'): \underline{F}' \subset \underline{F}\}$ as follows: Write $d(\underline{F}_1) \leq d(\underline{F}_2)$ if and only if there is an integer $\eta \geq 1$ and σ-fields $\underline{G}_K \subset \underline{F}$, $1 \leq K \leq \eta$, such that $\sigma(d(\underline{F}_1)) \subset \underline{G}_1$, $\sigma(d(\underline{G}_K)) \subset \underline{G}_{K+1}$ and $d(\underline{G}_\eta) = d(\underline{F}_2)$. It is easy to see that $d(\underline{F}') \leq d(\underline{F}')$ for all $\underline{F}' \subset \underline{F}$ (hence, $\leq$ is reflexive), and that $d(\underline{F}_1) \leq d(\underline{F}_2)$ and $d(\underline{F}_2) \leq d(\underline{F}_3)$ imply $d(\underline{F}_1) \leq d(\underline{F}_3)$ (hence, $\leq$ is transitive). Suppose $d(\underline{F}_1) \leq d(\underline{F}_2)$ and $d(\underline{F}_2) \leq d(\underline{F}_1)$. Then there are σ-fields $\underline{G}_K \subset \underline{F}$, $1 \leq K \leq \eta + m$, such that $\sigma(d(\underline{G}_{\eta+m})) \subset \underline{G}_1$, $\sigma(d(\underline{G}_K)) \subset \underline{G}_{K+1}$, $1 \leq K \leq \eta + m - 1$, $d(\underline{F}_1) = d(\underline{G}_{\eta+m})$, and $d(\underline{F}_2) = d(\underline{G}_\eta)$. The agreement condition implies that $d(\underline{G}_K) = d(\underline{G}_1)$ for all K, therefore $d(\underline{F}_1) = d(\underline{F}_2)$. Consequently, $\leq$ is antisymmetric and so $\leq$ is a partial order. Finally, if $\underline{G} \subset \underline{F}$ and $\sigma(d(\underline{G})) \subset \underline{F}'$, then $d(\underline{G}) \leq d(\underline{F}')$ by definition of $\leq$. Hence, $d(\underline{F}')$ is the maximum element of $\{d(\underline{G}): \underline{G} \subset \underline{F}, \sigma(d(\underline{G})) \subset \underline{F}'\}$ with respect to $\leq$.

APPENDIX D

Proof of Proposition 2

Suppose $\delta_1, \delta_2: \Omega \to U$ are $\underline{F}$-measurable. Define $\delta_1 \leq'' \delta_2$ to mean either that $L(\delta_2) < L(\delta_1)$ or that $L(\delta_2) = L(\delta_1)$ and $\delta_1(\omega) \leq' \delta_2(\omega)$ for all ω. It is easy to see that $\leq''$ so defined partially orders all $\underline{F}'$-measurable decision functions. Suppose that $\underline{F}' \subset \underline{F}$ and δ is an $\underline{F}'$-measurable decision function. Since $d(\underline{F}') \in D(\underline{F}')$, by assumption $L(\delta) \leq L(d(\underline{F}'))$. If $L(\delta) = L(d(\underline{F}'))$, then $\delta \in D(\underline{F}')$ also, and $\delta(\omega) \leq' d(\underline{F}')(\omega)$ for all ω. It follows that $d(\underline{F}')$ maximizes $\{\delta: \sigma(\delta) \in \underline{F}\}$ with respect to $\leq''$. In particular, $d(F')$ maximizes $\{d(\underline{G}): \sigma(d(\underline{G})) \subset \underline{F}', \underline{G} \subset \underline{F}\}$. Thus, Proposition 1 implies that d satisfies the agreement condition for rings.

APPENDIX E

Proof of Proposition 3

Assume that the condition

$$\sigma(d(\underline{F}_1)) \subset \underline{F}_2 \subset \underline{F}_1 \Rightarrow d(\underline{F}_1) = d(\underline{F}_2) \tag{E.1}$$

is true. Then $\sigma(d(\underline{F}_1)) \vee \sigma(d(\underline{F}_2)) \subset \underline{F}_1 \cap \underline{F}_2$ implies $\sigma(d(\underline{F}_1)) \subset \underline{F}_1 \cap \underline{F}_2 \subset \underline{F}_1$, which in turn implies $d(\underline{F}_1) = d(\underline{F}_1 \cap \underline{F}_2)$. Likewise, $d(\underline{F}_2) = d(\underline{F} \cap \underline{F}_2)$.

Hence, the condition

$$\sigma(d(\underline{F}_1)) \vee \sigma(d(\underline{F}_2)) \subset \underline{F}_1 \cap \underline{F}_2 \Rightarrow d(\underline{F}_1) = d(\underline{F}_2)$$

is true.

Conversely, assume that

$$\sigma(d(\underline{F}_1)) \vee \sigma(d(\underline{F}_2)) \subset \underline{F}_1 \cap \underline{F}_2 \Rightarrow d(\underline{F}_1) = d(\underline{F}_2) \qquad \text{(E.2)}$$

is true. Then $\underline{F}_2 \subset \underline{F}_1$ implies $\underline{F}_2 = \underline{F}_1 \cap \underline{F}_2$. Hence, $\sigma(d(\underline{F}_1)) \subset \underline{F}_2$ implies $\sigma(d(\underline{F}_1)) \vee \sigma(d(\underline{F}_2)) \subset \underline{F}_2 = \underline{F}_1 \cap \underline{F}_2$. Because of Eq. (E.2), it follows that $d(\underline{F}_1) = d(\underline{F}_2)$. Thus, the condition

$$\sigma(d(\underline{F}_1)) \subset \underline{F}_2 \subset \underline{F}_1 \Rightarrow d(\underline{F}_1) = d(\underline{F}_2)$$

is true.

APPENDIX F

Proof of Theorem 4

The proof of Theorem 4 proceeds in various steps: First we describe precisely the evolution of the estimation process according to each team member's model, and determine what each member's model predicts. Then, compare these predictions with what happens in reality. Both Alpha and Beta can describe the evolution of the estimation process according to their own view of the world as follows: Let α_t^i and β_t^i be the estimates of Alpha and Beta at time t according to i's perception (i = Alpha, Beta). Then,

$$\alpha_t^i = d^i(A, \beta_1^i, \beta_2^i, \ldots, \beta_{t-1}^i) \qquad \text{(F.1)}$$

$$\beta_t^i = d^i(B, \alpha_1^i, \alpha_2^i, \ldots, \alpha_t^i), \qquad \text{(F.2)}$$

where d^i denotes that the estimates are formed according to the rule d and the probability measure p^i induced by the distribution P^i on Ω. Equations (F.1) and (F.2) considered for all t and for all $\omega \in \Omega$ describe the evolution of the estimation process according to member i's view of the world. To determine what Alpha and Beta predict about the outcome of the estimation process in terms of their own models consider the following σ-fields

$$\underline{F}_t^{iA} = \sigma(A, \beta_1^i, \beta_2^i, \ldots, \beta_{t-1}^i) \qquad (i = \text{Alpha, Beta}) \qquad \text{(F.3)}$$

$$\underline{F}_t^{iB} = \sigma(B, \alpha_1^i, \alpha_2^i, \ldots, \alpha_t^i)$$

The σ-fields $\underline{F}_t^{iA}, \underline{F}_t^{iB}$ describe the view of member i about the information available to Alpha and Beta after the initial observations have been taken and t tentative decisions have been exchanged. The σ-fields $\underline{F}_t^{iA}$ and $\underline{F}_t^{iB}$ evolve dynamically as follows:

$$\underline{F}_{t+1}^{iA} = \underline{F}_t^{iA} \vee \sigma(d^i(\underline{F}_t^{iB}))$$

$$\underline{F}_{t+1}^{iB} = \underline{F}_t^{iB} \vee \sigma(d^i(\underline{F}_{t+1}^{iA})) \qquad \text{(F.4)}$$

with initial condition

$$\underline{F}_0^{iA} = A$$
$$\underline{F}_0^{iB} = B. \tag{F.5}$$

We can view Eqs. (F.4) and (F.5) as a dynamic system defined on the lattice of σ-fields contained in $\underline{F}$. The lattice operations are $\vee$ and $\wedge$, where $\vee$ and $\wedge$ are the join and meet operations on σ-fields. The maximum σ-field is $A \vee B$ and the minimum is the trivial one. $\underline{F}_0 = \{\phi, \Omega\}$. Equation (F.4) generates increasing sequences of σ-fields. Since $A \vee B$ is finite by assumption, and

$$A \subset \underline{F}_t^{iA} \subset \underline{F}_{t+1}^{iA} \subset A \vee B$$
$$B \subset \underline{F}_t^{iB} \subset \underline{F}_{t+1}^{iB} \subset A \vee B, \tag{F.6}$$

it follows that eventually

$$\underline{F}_t^{iA} = \underline{F}^{iA} \tag{F.7}$$

for all $t \geq t_A^i$ (for some $t_A^i \geq 0$) and

$$\underline{F}_t^{iB} = \underline{F}^{iB} \tag{F.8}$$

for all $t \geq t_B^i$ (for some $t_B^i \geq 0$). Let

$$T^i = \max(t_A^i, t_B^i). \tag{F.9}$$

Then, for $t > T^i$

$$\underline{F}^{iA} = \underline{F}^{iA} \vee \sigma(d^i(\underline{F}^{iB}))$$
$$\underline{F}^{iB} = \underline{F}^{iB} \vee \sigma(d^i(\underline{F}^{iA})) \tag{F.10}$$

or equivalently,

$$\sigma(d^i(\underline{F}^{iB})) \subset \underline{F}^{iA}$$
$$\sigma(d^i(\underline{F}^{iA})) \subset \underline{F}^{iB}. \tag{F.11}$$

Consequently,

$$\sigma(d^i(\underline{F}^{iB})) \subset \underline{F}^{iA} \cap \underline{F}^{iB} \subset \underline{F}^{iB}$$

and

$$\sigma(d^i(\underline{F}^{iA})) \subset \underline{F}^{iA} \cap \underline{F}^{iB} \subset \underline{F}^{iA}. \tag{F.12}$$

Since d satisfies the agreement conditions, (F.12) implies that

$$d^i(\underline{F}^{iA}) = d^i(\underline{F}^{iA} \cap \underline{F}^{iB}) = d^i(\underline{F}^{iB}). \tag{F.13}$$

Thus, both Alpha and Beta predict that the estimates will converge and agree after a finite number of steps.

A

B.

(F.5)

a dynamic system defined on the
: lattice operations are $\vee$ and $\wedge$,
eet operations on σ-fields. The
minimum is the trivial one. $\underline{F}_0 =$
sing sequences of σ-fields. Since

$\subset A \vee B$

$\subset A \vee B,$ (F.6)

(F.7)

(F.8)

$t_\beta^i)$. (F.9)

$d^i(\underline{F}^{iB}))$

$d^i(\underline{F}^{iA}))$ (F.10)

$\underline{F}^{iA}$

$\underline{F}^{iB}.$ (F.11)

$\underline{F}^{iB} \subset \underline{F}^{iB}$

$\underline{F}^{iB} \subset \underline{F}^{iA}.$ (F.12)

s, (F.12) implies that

$^{B}) = d^i(\underline{F}^{iB}).$ (F.13)

the estimates will converge and

In reality, the following is happening: At ime $t = 1$ Alpha's estimate is

$$\bar{\alpha}_1 = d^1(\bar{A}) \tag{F.14}$$

(where $\bar{A}$ is Alpha's observation). The message $\bar{\alpha}_1$ is transmitted to Beta. Beta interprets this message according to his own view of the world, that is, he considers that the realization $\tilde{A}$ of A is such that

$$\bar{\alpha}_1 = \alpha_1^B = d^2(\tilde{A}). \tag{F.15}$$

Furthermore, for a consistent interpretation of the data it is required that

$$P^B(\bar{\alpha}_1 = \alpha_1^B) > 0. \tag{F.16}$$

At $t = 2$ *Beta's estimate is*

$$\bar{\beta}_1 = d^2(\bar{B}, \bar{\alpha}_1) \tag{F.17}$$

(where $\bar{B}$ is Beta's observation), and this estimate is transmitted to Alpha who interprets it in terms of his own model, that is, he considers that the realization $\tilde{B}$ of B is such that

$$\bar{\beta}_1 = \beta_1^A = d^1(\tilde{B}, \bar{\alpha}_1). \tag{F.18}$$

For a consistent interpretation of the data it is required that

$$P^A(\beta_1^A = \bar{\beta}_1) > 0. \tag{F.19}$$

In general, when Alpha receives message $\bar{\beta}_K$ he interprets it in terms of his own view of the world, that is, he considers that the realization $\tilde{B}$ of B is such that

$$\bar{\beta}_K = \beta_K^A = d^1(\tilde{\beta}, \bar{\alpha}_1, \bar{\alpha}_2, \ldots, \bar{\alpha}_K) \tag{F.20}$$

Then, Alpha generates

$$\bar{\alpha}_{K+1} = d^1(\bar{A}, \bar{\beta}_1, \bar{\beta}_2, \ldots, \bar{\beta}_K), \tag{F.21}$$

which he sends to Beta. For a consistent interpretation of all the messages received by Alpha and Beta, it is required that at any time t

$$P^A(\beta_l^A = \bar{\beta}_l, \quad 1 \le l \le t) > 0$$

and

$$P^B(\alpha_l^B = \bar{\alpha}_l, \quad 1 \le l \le t) > 0. \tag{F.22}$$

The following result about the evolution of the probabilities of (F.22) is true.

Proposition F.1. *After a finite number of steps s^A, either*

$$P^A(\beta_l^A = \bar{\beta}_l, \quad 1 \le l \le s^A) = 0 \tag{F.23}$$

or

$$P^A(\beta_l^A = \bar{\beta}_l, \quad 1 \le l \le s^A) = 1. \tag{F.24}$$

Moreover, for all $s > s^A$

$$P^A(\beta_l^A = \bar{\beta}_l, \quad 1 \le l \le s) = P^A(\beta_l^A = \bar{\beta}_l, \quad 1 \le l \le s^A). \tag{F.25}$$

Similar results hold for $P^B(\alpha_l^B = \bar{\alpha}_l, \quad 1 \le l \le t)$.

Proof. The result follows directly from the fact that convergence and agreement are predicted to occur in a finite number of steps by both models. The time s^A is given by Eq. (F.1).

■

Based on the previous proposition we can complete the proof of Theorem 4 as follows: If

$$P^A(\beta_l^A = \bar{\beta}_l \quad 1 \le l \le t) > 0$$
$$P^B(\alpha_l^B = \bar{\alpha}_l \quad 1 \le l \le t) > 0$$

are true for all $t < s^A$, s^B, respectively, and Eq. (F.24) is true for both $P^A(\cdot)$ and $P^B(\cdot)$, then because of Eq. (F.13) and the rules by which the messages are interpreted

$$d^A(\underline{F}_l^{AA}) = d^A(\underline{F}_l^{AB}) = d^B(\underline{F}_l^{BB}) = d^B(\underline{F}_l^{BA})$$

for all $l \ge \max(s^A, s^B)$ and the estimates of Alpha and Beta agree asymptotically. If, on the other hand, Eq. (F.23) is true at some time t for either Alpha or Beta, then, at that time Alpha or Beta realize that the sequence of received messages is impossible, or more reasonably, Alpha or Beta must conclude that the two models P^A and P^B are inconsistent.

NOTES

1. Research supported in part by ONR Contract N00014-80-C-0507 and JSEP Contract F49620-79-C-0178.

2. Note that $\vee$ is the join operation on σ-fields: $\underline{F}_1 \vee \underline{F}_2$ is the smallest σ-field containing $\underline{F}_1$ and $\underline{F}_2$.

REFERENCES

[1] Aumann, R. J. "Agreeing to Disagree," *Annals of Statistics*, **4**, (6) (1976), pp. 1236–1239.
[2] Borkar, V., and Varaiya, P. "Asymptotic Agreement in Distributed Estimation," *IEEE Trans. Automatic Control*, **AC-27**, (3) (1982), pp. 650–655.
[3] De Groot, M. H. "Reaching a Consensus," *J. American Statistical Association*, **69**, (345) (Mar. 1974), pp. 118–121.

[4] Geanakoplos, J. D., and Polemarchakis, H. M. "We Can't Disagree Forever." *J. Economic Theory*, **26** (1982), pp. 192–200.

[5] Sebenius, J. K., and Geanakoplos, J. "Don't Bet On It: Contingent Agreements with Asymmetric Information," *J. American Statistical Association*, **78**, (382) (June 1983), pp. 424–236.

[6] Teneketzis, D., and Varaiya, P. "Consensus in Distributed Estimation with Inconsistent Beliefs," *Systems and Control Letters*, **4** (1984), pp. 217–221.

[7] Tsitsiklis, J. N., and Athans, M. "Convergence and Asymptotic Agreement in Distributed Decision Problems," *IEEE Trans. Automatic Control*, **AC-29**, (1) (1984), pp. 42–50.

[8] Washburn, R. B., and Teneketzis, D. "Asymptotic Agreement Among Communicating Decisionmakers," *Stochastics*, **13** (1984), pp. 103–129.

[9] Milgrom, P. "An Axiomatic Characterization of Common Knowledge," *Econometrica*, **49**, 1 (Jan. 1981).

[10] Milgrom, P., and Stokey, N. "Information, Trade and Common Knowledge," Center for Mathematical Studies, Discussion Paper 377R, Northwestern Univ., Evanston, Ill., Nov. 1979.

[11] Wilson, R. "Information, Efficiency and the Core of an Economy," *Econometrica*, **46** (1978), pp. 807–816.

[12] Kreps, D. M., Milgrom, P., Roberts, J., and Wilson, R. "Rational Cooperation in the Finitely-Repeated Prisoners' Dilemma," *J. Economic Theory*, **27** (1982), pp. 245–252.