

Consensus in distributed estimation with inconsistent beliefs *

D. TENEKETZIS

ALPHATECH, Inc., 2 Burlington Executive Center, 111
Middlesex Turnpike, Burlington, MA 01803, USA

P. VARAIYA

Department of Electrical Engineering and Computer Sciences and
Electronics Research Laboratory, University of California,
Berkeley, CA 94720, USA

Received 10 February 1984

Two people sequentially revise and exchange their estimates of the same random variable. They may have different models of the underlying probability structure. The two sets of estimates then will converge to the same value, or the people will realize that their beliefs are inconsistent.

Keywords: Distributed estimation, Multiple models, Inconsistent Beliefs, Asymptotic agreement.

1. The problem

Two people, Alpha and Beta, repeatedly calculate and exchange estimates of the same random variable X as follows. Initially, Alpha observes the random variable A and Beta observes B . For $n = 1, 2, \dots$ the n th estimate by Alpha is denoted α_n . It is the conditional expectation of X given the observations $A, \beta_1, \dots, \beta_{n-1}$. After α_n has been calculated it is communicated to Beta whose n th estimate, denoted β_n , is the conditional expectation of X given $B, \alpha_1, \dots, \alpha_n$. Once β_n is evaluated it is communicated to Alpha who incorporates it into the estimate α_{n+1} , and the procedure is repeated.

* Research supported in part by ONR Contract N000014-82-C-0693 and by DOE Contract DEAC01-80-RA50418. Helpful discussion with Dr. Robert Washburn is gratefully acknowledged.

In this setup there are three basic random variables, namely X, A , and B . Alpha's prior model of these variables is given by their joint probability distribution P^α . Beta's prior model is given by P^β . P^α and P^β may be different.

To complete the specification we assume that the estimation procedures followed by Alpha and Beta are consistent with their own prior models. This has two implications. Consider Alpha. When he receives Beta's estimate β_{n-1} , Alpha interprets it as if it were based on the same model as Alpha's, P^α , and not on P^β . That is, Alpha assumes that Beta's estimate is a realization of the random variable

$$\hat{\beta}_{n-1} = E^\alpha\{X | B, \alpha_1, \dots, \alpha_{n-1}\}.$$

Subsequently, Alpha calculates α_n ,

$$\alpha_n = E^\alpha\{X | A, \hat{\beta}_1, \dots, \hat{\beta}_{n-1}\}.$$

Symmetrically, Beta interprets α_n as

$$\hat{\alpha}_n = E^\beta\{X | A, \beta_1, \dots, \beta_{n-1}\},$$

and calculates β_n by

$$\beta_n = E^\beta\{X | B, \hat{\alpha}_1, \dots, \hat{\alpha}_n\}.$$

Our objective is to study how α_n and β_n change. More precisely, we answer the question: Will α_n and β_n agree as n increases? In the case that the two models are the same, $P^\alpha = P^\beta$, an affirmative answer was given by Borkar and Varaiya [2]. Significant variations and extensions of that work have been made by Tsitsiklis and Athans [5], and by Washburn and Teneketzis [6]. In these papers the assumption $P^\alpha = P^\beta$ is maintained; however, the messages exchanged are statistics different from that given by conditional expectations. For earlier work relevant to the question raised above see Aumann [1], Geanakopoulos and Polemarchakis [4], and De Groot [3].

In contrast to the papers cited above our interest here is in the case $P^\alpha \neq P^\beta$. What can happen then is illustrated by two examples.

2. Two examples

Infallible Self. Each believes himself infallible. Alpha assumes $P^\alpha\{X=A\}=1$ and Beta assumes $P^\beta\{X=B\}=1$. Hence

$$\begin{aligned} P^\alpha\{\alpha_n = \hat{\beta}_n = A\} &= 1, \\ P^\beta\{\beta_n = \hat{\alpha}_{n+1} = B\} &= 1, \end{aligned} \quad n = 1, 2, \dots$$

Therefore, when $A \neq B$, agreement is impossible. Indeed, when $A \neq B$, both realize that an 'impossible' event has occurred, or, that their prior models are mutually inconsistent.

Infallible Other. Each believes the other infallible. Alpha assumes $P^\alpha\{X=B\}=1$ and Beta assumes $P^\beta\{X=A\}=1$. Alpha's first estimate is $\alpha_1 = E^\alpha\{X|A\}$ which Beta interprets as

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

and so Beta's first estimate is

$$\beta_1 = E^\beta\{X|B, \hat{\alpha}_1\} = \hat{\alpha}_1 = \alpha_1.$$

Alpha interprets β_1 as

$$\hat{\beta}_1 = E^\alpha\{X|A, \alpha_1\} = X = \beta_1,$$

so that his second estimate

$$\alpha_2 = E^\alpha\{X|A, \hat{\beta}_1\} = \hat{\beta}_1 = \alpha_1$$

is the same as his first estimate, thereby confirming Beta's belief. Thus

$$\alpha_1 = \alpha_2 = \dots = \beta_1 = \beta_2 = \dots = E^\alpha\{X|A\}.$$

In this case there is immediate and lasting consensus. The agreement is not a consequence of consistent beliefs but rather the confirmation of inconsistent models. (Note that if the first estimate was announced by Beta instead of by Alpha, then the consensus estimate would be $E^\beta\{X|B\}$.)

We show next that these two examples in a sense bracket the possibilities in general: either the two estimates eventually agree, or both parties realize that their models are incompatible.

3. Analysis

We consider the simple case when the initial observations A, B can take values from finite sets \mathbf{A}, \mathbf{B} respectively. All estimates are functions of

$$\omega = (A, B) \in \mathbf{A} \times \mathbf{B} = \Omega.$$

The estimates are defined sequentially in the following order for $n = 1, 2, \dots$:

$$\begin{aligned} \alpha_n &= E^\alpha\{X|A, \hat{\beta}_1, \dots, \hat{\beta}_{n-1}\}, \\ \hat{\alpha}_n &= E^\beta\{X|A, \beta_1, \dots, \beta_{n-1}\}, \\ \beta_n &= E^\beta\{X|B, \hat{\alpha}_1, \dots, \hat{\alpha}_n\}, \\ \hat{\beta}_n &= E^\alpha\{X|B, \alpha_1, \dots, \alpha_n\}. \end{aligned}$$

There is a more revealing description of the functional dependence of these estimates. Suppose a particular realization $\bar{\omega} = (A, B)$ has occurred. Since Alpha observes \bar{A} , he concludes that

$$\bar{\omega} \in \Omega_1^\alpha = \{(A, B) | A = \bar{A}\}$$

and so his first estimate equals

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

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

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

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

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

and so Alpha concludes that

$$\bar{\omega} \in \Omega_n^\alpha = \{\omega | \omega \in \Omega_{n-1}^\alpha, \hat{\beta}_{n-1}(\omega) = \bar{\beta}_{n-1}\}.$$

Hence Alpha's n th estimate takes the value

$$\bar{\alpha}_n = E^\alpha\{X|\omega \in \Omega_n^\alpha\}$$

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

$$\hat{\alpha}_n = E^\beta\{X|A, \beta_1, \dots, \beta_{n-1}\},$$

concludes that

$$\bar{\omega} \in \Omega_n^\beta = \{\omega | \omega \in \Omega_{n-1}^\beta, \hat{\alpha}_n(\omega) = \bar{\alpha}_n\},$$

and evaluates his n th estimate as

$$\bar{\beta}_n = E^\beta \{ X \mid \omega \in \Omega_n^\beta \}.$$

Thus, as expected, the uncertainty diminishes with each exchange,

$$\Omega_{n+1}^\alpha \subset \Omega_n^\alpha, \quad \Omega_{n+1}^\beta \subset \Omega_n^\beta.$$

From the description above 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

$$\Omega_n^\alpha = \Omega_{k+1}^\alpha \quad \text{and} \quad \Omega_n^\beta = \Omega_{k+1}^\beta \quad \text{for } n > k + 1.$$

Hence for $n > N$ (which cannot exceed the number of distinct elements in Ω), Ω_n^α and Ω_n^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, similar to the ‘infallible self’ example, is that $\Omega_*^\alpha(\omega) = \emptyset$ and $\Omega_*^\beta(\omega) = \emptyset$. This happens because at some stage the message $\bar{\beta}_{n-1}$ received by Alpha is ‘impossible’: there is no $\bar{\omega}$ such that $\hat{\beta}_{n-1}(\bar{\omega}) = \bar{\beta}_{n-1}$; or the message $\bar{\alpha}_n$ received by Beta is ‘impossible’: there is no $\bar{\omega}$ such that $\hat{\alpha}_n(\bar{\omega}) = \bar{\alpha}_n$. 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, similar to the example of the ‘infallible other’, is that $\Omega_*^\alpha(\omega) \neq \emptyset$ and $\Omega_*^\beta(\omega) \neq \emptyset$. In this case for $n > N$ the estimates stop changing:

$$\hat{\beta}_n(\omega) = \hat{\beta}_*(\omega), \quad \alpha_n(\omega) = \alpha_*(\omega), \\ \hat{\alpha}_n(\omega) = \hat{\alpha}_*(\omega), \quad \beta_n(\omega) = \beta_*(\omega).$$

Since for every n ,

$$\hat{\beta}_n(\omega) = \beta_n(\omega) \quad \text{and} \quad \hat{\alpha}_n(\omega) = \alpha_n(\omega),$$

it follows that

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

On the other hand, since $\hat{\beta}_n$ and α_n are based on the same model, namely P^α , it follows from the argument of Borkar and Varaiya [2] 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 $\alpha_n(\omega) = \beta_n(\omega)$ for $n > N$. 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. *The set of events Ω decomposes into two disjoint subsets Ω_I and Ω_{II} . After N 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 the following example demonstrates.

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 Fig. 1. Finally, suppose that ω is uniformly distributed under P^α , whereas

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

and within each b_i , ω is uniformly distributed under P^β .

Suppose that $\bar{\omega} \in a_1 \cap b_3$ and that Alpha communicates first. Then

$$\bar{\alpha}_1 = E^\alpha \{ X \mid \omega \in a_1 \} = \frac{1}{2}.$$

Beta interprets this as a realization of

$$\hat{\alpha}_1 = E^\beta \{ X \mid 1(a_1), 1(a_2) \}.$$

Since

$$E^\beta \{ X \mid \omega \in a_1 \} = \frac{5}{12}, \quad E^\beta \{ X \mid \omega \in a_2 \} = \frac{1}{2},$$

upon learning that $\bar{\alpha}_1 = \frac{1}{2}$, Beta concludes that $\bar{\omega} \in a_2$, and since he has observed that $\bar{\omega} \in b_3$, his estimate is

$$\bar{\beta}_1 = E^\beta \{ X \mid \omega \in a_2 \cap b_3 \} = \frac{3}{4}.$$

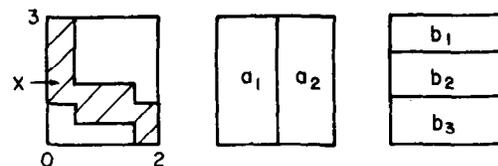


Fig. 1.

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

$$E^\alpha\{X | \omega \in a_1 \cap b_1\} = \frac{1}{2},$$

$$E^\alpha\{X | \omega \in a_1 \cap b_2\} = \frac{3}{4},$$

$$E^\alpha\{X | \omega \in a_1 \cap b_3\} = \frac{1}{4},$$

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

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

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

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

Since

$$E^\alpha\{X | \omega \in b_1\} = \frac{1}{4},$$

$$E^\alpha\{X | \omega \in b_2\} = E^\alpha\{X | \omega \in b_3\} = \frac{1}{2},$$

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

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

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

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

or

$$E^\beta\{X | \omega \in a_2 \cap b_3\} = 0.6.$$

Thus Beta concludes that the models are inconsistent.

4. Concluding remarks

Bertrand Russell once observed that two people could carry on a conversation about London blissfully unaware that their subjective images of London are very different. This is possible, Russell argued, because utterances in English are so ambiguous that each could interpret the other speaker's statements in his own way without realizing that the intended meaning was different.

The point here is similar. Alpha and Beta can exchange statements about X and eventually agree even when their views are different. Paradoxically,

the realization that these views are different is only reached when further communication becomes impossible.

The rudimentary investigation reported here needs to be carried further. First, some 'technical' extensions must be made to include situations when (a) new observations are made in the course of message exchange, (b) these observations are real valued, and (c) messages different from conditional expectations are exchanged. A more challenging problem is to give conditions on the pair of models P^α, P^β which guarantee agreement for all realizations.

There are also more basic and knotty issues. Suppose Alpha and Beta reach an impasse ($\omega \in \Omega_1$). 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 $\hat{\beta}_n \equiv \beta_n$ and $\hat{\alpha}_n \equiv \alpha_n$), 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 \mathcal{P}^β of models and there is a prior distribution on \mathcal{P}^β reflecting Alpha's initial judgement 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. (The necessity for consensus can readily arise in a context where the two parties must agree on a joint decision and such agreement is predicated on a consensus about the expected value of the random outcome of the decision.) To reach a consensus one or both must change their models. One can imagine many different ways in which this can be done.

References

- [1] R.J. Aumann, Agreeing to disagree, *Ann. Statistics* **4** (6) (1976) 1236-1239.
- [2] V. Borkar and P. Varaiya, Asymptotic agreement in distributed estimation, *IEEE Trans. Automat. Control.* **27** (3) (1982) 650-655.

- [3] M.H. De Groot, Reaching a consensus, *J. Amer. Statist. Assoc.* **69** (345) (March 1974).
- [4] J.D. Geanakopoulos and H.M. Polemarchakis, We can't disagree forever, Technical Report No. 277, Institute for Mathematical Studies in the Social Sciences, Stanford University (Dec. 1978).
- [5] J.N. Tsitsiklis and M. Athans, Convergence and asymptotic agreement in distributed decision problems, LIDS-P-1185, Laboratory for Information and Decision Systems, MIT (Mar. 1982).
- [6] R.B. Washburn and D. Teneketzis, Asymptotic agreement among communicating decisionmakers, Alphatech Technical Report TR-145, Burlington MA (Feb. 1983).