Aumann's agreement theorem

Tyrrell McAllister

July 7, 2011

Abstract

Aumann's agreement theorem [1] says that, if two Bayesian agents share the same prior and their respective posterior probabilities for some proposition A are "common knowledge", then those posterior probabilities must be the same. We unpack what "common knowledge" means here, and we show why it implies Aumann's result.

1 The Statement

Suppose that you and I both started out using the same prior probability distribution. In other words, to each proposition Q, we both assigned the same prior probability $\text{prob}(Q)$.

But time has passed, and you and I have seen diverse things. Each of us has acquired different bodies of knowledge about the world. As a result, our respective posterior probabilities for some proposition A could be very different.

Now let us make the additional assumption that our respective posterior probabilities for A are *common knowledge*. Aumann's agreement theorem [1] says that this is enough to guarantee that these posterior probabilities are equal. Even though we might know very different things about the subject matter of A, our posterior probabilities for A must nonetheless be the same.

2 What is "Common Knowledge"?

Let p be my posterior probability for A, and let q be your posterior probability. What does it mean to say that our posterior probabilities are "common knowledge"? It's actually a very strong condition, much stronger than just saying that I know your posterior probability and you know mine. What we require is that there be common information C, known to both of us, satisfying the following conditions:

1. The proposition C implies that we both know C. That is, in all possible worlds in which C is true, you and I both condition on C (among other things) to arrive at our posterior probabilities.

2. I would have assigned a posterior probability of p to A, no matter what I had learned in addition to C.

3. You would have assigned a posterior probability of q to A, no matter what you had learned in addition to C.

Let's spell this out a bit more. Typically, neither of us evaluates the probability of A just on the basis of C alone. Each of us has additional information, which, in conjunction with C, constitutes a complete state of knowledge. Let

$$\mathcal{E} = \{C \,\&\, E_1,\, C \,\&\, E_2, \ldots\}$$

be the set of candidates for "everything I know", given C. Similarly, let

$$\mathcal{F} = \{C \,\&\, F_1,\, C \,\&\, F_2, \ldots\}$$

be the set of candidates for "everything you know", given C. We are assuming that the elements of $\mathcal{E}$ are all *a priori* possible, mutually exclusive, and exhaustive, given C. Likewise with the elements of $\mathcal{F}$. Exhaustiveness comes from condition (1) above. Conditions (2) and (3) above amount to saying that

$$\begin{aligned} p &= \text{prob}(A \mid C \,\&\, E_i), \text{ for all } i = 1, 2, \ldots, \text{ and} \\ q &= \text{prob}(A \mid C \,\&\, F_j), \text{ for all } j = 1, 2, \ldots. \end{aligned} \tag{1}$$

3 The Proof

On the one hand, we can decompose C into an exclusive disjunction

$$C \equiv \bigvee_i (C \,\&\, E_i), \tag{2}$$

which gives us

$$\text{prob}(A \mid C) = \sum_i \text{prob}(A \mid C \,\&\, E_i)\, \text{prob}(E_i \mid C).$$

By equation (1) above, conditioning on each disjunct in (2) yields the same posterior probability p. It follows that

$$\text{prob}(A \mid C) = p \sum_i \text{prob}(E_i \mid C) = p.$$

On the other hand, from the exclusive disjunction $C \equiv \bigvee_j (C \,\&\, F_j)$, we similarly get that $\text{prob}(A \mid C) = q$. Thus we find that our posterior probabilities are both equal to $\text{prob}(A \mid C)$, and so to each other.

References

[1] R. J. Aumann, *Agreeing to disagree*, The Annals of Statistics **4** (1976), no. 6, 1236–1239.