

# Lecture 3

## Common Knowledge

*People with the same prior can't "agree to disagree,"  
even on the basis of different information.*

### 1 Context

- So far, we've ignored problems of *private information*
  - In Walrasian equilibrium, it doesn't matter whether other people know your preferences  
Prices come from wherever they come from,  
and all anyone can do is optimize given those prices
  - In Arrow's policy setting, we didn't worry about *learning* each person's preferences  
We assumed they were known,  
and focused on getting from individual preferences to social preferences.
- For the rest of the semester, we'll be looking at settings where people have *private information*
  - Thus, part of the challenge will be understanding their incentives to reveal that information
- First, though, we need a formal framework to think about private information.

- There are lots of types of private information
  - I could have private information about something that’s only relevant to me – like private information about my own preferences.
 

(This still matters strategically to others – it might tell a seller how much money he could try to get out of me, or a competing buyer how much I might bid in an auction – but it only affects others through my actions, not directly.)
  - I could also have private information that’s directly relevant to other peoples’ payoffs
 

I might have inside information about a company – so if I’m trying to sell shares, buyers need to be worried about what I know
- For a lot of applications, we work with a model of private information that’s suited to that particular environment.
- But for today, we’re going to introduce a very general model of information – general enough to (more or less) nest all other models.
- For the next two lectures, we’ll be thinking about interactions between two people – but everything here could be extended to more.

## 2 General Model of Information

### 2.1 Probability Spaces

- We begin with a **finite set  $\Omega$  of states of the world  $\omega$**
- A “state of the world”  $\omega$  is a complete description of the world – basically, a resolution of all uncertainty.
- So, what does it mean to know something? Basically, it means you can distinguish between two states of the world.
- What I know can be described as a **partition of  $\Omega$**  – which states of the world I can tell apart, and which ones I can’t.
- For example, let  $\Omega = \{1, 2, 3\}$

We define my **information partition** as a set of disjoint subsets that make up  $\Omega$ , such as  $\mathcal{P}_1 = \{\{1, 2\}, \{3\}\}$

That means that if the actual state of the world is 3, I know it’s 3; but if the actual state is either 1 or 2, I know that it’s *either* 1 or 2, but I don’t know which.

- At first glance, this sounds like it doesn't capture uncertainty well, but it turns out, it does

We just need to put prior probabilities on each state, and define the states right

Suppose we've caught a murder suspect and want to know whether he actually committed the murder, so we check to see whether his fingerprints match the prints at the scene

There are four states of the world:

$\omega_1$ : he's guilty, and there are fingerprints

$\omega_2$ : he's guilty, but there aren't any fingerprints

$\omega_3$ : he's innocent, and there are fingerprints

$\omega_4$ : he's innocent, and there aren't any fingerprints

The "event" we care about is whether he's guilty – which means what we care about is the probability that the state of the world is *either* 1 or 2

Our information partition, though, is  $\{\{1, 3\}, \{2, 4\}\}$  – we can tell whether there are fingerprints, but not whether he's guilty

There's one other thing we need to make sense of this – a **prior probability distribution** on the different states of the world.

So suppose that...

- With probability 0.4, he's guilty and left fingerprints
- With probability 0.1, he's guilty but there are no fingerprints
- With probability 0.05, he's innocent but there are fingerprints anyway
- With probability 0.45, he's innocent and there are no fingerprints

- Formally, this defines a **probability space** – a set of states  $\Omega$ , a set of events that we might care about  $\mathcal{B}$ , and a prior probability over the states  $p$ . (Each event is just the set of states in which that event has happened – the set of states in which the guy is guilty, for example.)

## 2.2 Bayes' Rule and Posterior Probabilities

- So now suppose we found fingerprints, and we want to know how likely it is that the guy is guilty
- We use **Bayes' Law**
- Iterated expectations say that if A and B are two different events,

$$\Pr(A \text{ and } B) = \Pr(A)\Pr(B|A)$$

and if we move this around, it gives Bayes' Law, which is

$$\Pr(B|A) = \frac{\Pr(A \text{ and } B)}{\Pr(A)}$$

- So now take a probability space  $(\Omega, \mathcal{B}, p)$ ; an information partition  $\mathcal{P}_1$ ; an element  $P_1 \in \mathcal{P}_1$ ; and an event  $E$ ; Bayes' Law says that

$$\Pr(E|P_1) = \frac{\Pr(E \cap P_1)}{\Pr(P_1)}$$

- Or in our context,

$$\Pr(\textit{guilty}|\textit{fingerprints}) = \frac{\Pr(\textit{guilty and fingerprints})}{\Pr(\textit{fingerprints})}$$

- So if I know the prior probability of each state, I can calculate the posterior probability of an event, given what I've learned; in this case,

$$\Pr(\textit{guilty}|\textit{fingerprints}) = \frac{\Pr(\omega_1)}{\Pr(\omega_1 \cup \omega_3)} = \frac{.4}{.4 + .05} \approx 0.889$$

### 2.3 Who knows who knows who knows what

- Next, we want to think about environments with multiple people, and think not just about who knows what, but what people think about what each other know
- Suppose  $\Omega = \{1, 2, 3, 4, 5, 6, 7, 8\}$ , and there are two of us
- My partition is  $\mathcal{P}_1 = \{\{1, 2, 3\}, \{4, 5\}, \{6, 7, 8\}\}$
- Your partition is  $\mathcal{P}_2 = \{\{1, 2\}, \{3, 4\}, \{5\}, \{6, 7\}, \{8\}\}$
- Suppose the state of the world is 5. Let's think about who knows what.
- I know the state is either 4 or 5, but I don't know which.
- You know the state is 5.

But obviously, I don't know that.

If we think about what I know about what you know: I know that *either* the state is 5, in which case you know it's 5; or the state is 4, and you therefore know that the state is either 3 or 4.

So if we think about *my beliefs about your beliefs* – I know that you know the state is either 3, 4, or 5.

- What about what I know about what you know *about what I know?*

I know that you might think the state is either 3, 4, or 5.

If you know the state is 5, then you know that I know it's either 4 or 5.

But if the state is 4, then you know it's either 3 or 4; which means that you know that *either* I know it's either 4 or 5, *or* you know that I know it's either 1, 2, or 3.

So if we think about what I know that you know that I know: all I know that you know that I know is that the state is either 1, 2, 3, 4, or 5.

- What about you? We already said, you know  $\omega = 5$ .
- So you know that I know it's either 4 or 5.
- So you know I know you know it's either 3, 4, or 5.
- So you know I know you know I know it's either 1, 2, 3, 4, or 5.

## 2.4 Common Knowledge

- Something is **common knowledge** if we both know it's true;  
and I know that you know it's true;  
and you know that I know it's true;  
and I know that you know that I know that you know that I know that you know it's true;  
and so on, for any string of beliefs we put together.
- So something being common knowledge is a pretty high bar – it's a lot more than just both of us knowing something is true.
- In the current example, when the state is 5, you know it's 5, and I know it's either 4 or 5  
So we both know the state is either 4 or 5  
But it's not common knowledge that the state is either 4 or 5 – because as far as I know, you might think the state is 3  
So if a particular event  $E$  occurs only in states 4 and 5, and the state of the world is 5, then the event has occurred, *and we both know the event has occurred*, but it is *not common knowledge* that the event has occurred  
Because I don't know whether you know that I know it's occurred!
- And the big result we'll show today is that common knowledge is very restrictive
- But first, a couple more tools

## 2.5 Refinements and Coarsenings of Information Partitions

- Suppose there are eight states of the world:  $\Omega = \{1, 2, 3, 4, 5, 6, 7, 8\}$
- And suppose my information partition is  $\mathcal{P}_1 = \{\{1, 2, 3\}, \{4, 5\}, \{6, 7, 8\}\}$
- We can make my information better by allowing me to distinguish between more states

This is the same as breaking up some of the elements of my existing partition, leading to

$$\mathcal{P}'_1 = \{\{1, 2\}, \{3\}, \{4, 5\}, \{6, 7, 8\}\}$$

This partition is *finer* than my old one – each element of  $\mathcal{P}'_1$  is a subset of a single element of  $\mathcal{P}_1$ , so I still know everything I used to know, and then more

So  $\mathcal{P}'_1$  is a *refinement* of  $\mathcal{P}_1$ .

- We can go the other direction – make my information worse, by lumping together two or more elements of my partition

The partition

$$\mathcal{P}''_1 = \{\{1, 2, 3\}, \{4, 5, 6, 7, 8\}\}$$

is a *coarsening* of  $\mathcal{P}_1$

## 2.6 Meets and Joins

- Suppose my information partition continues to be  $\mathcal{P}_1 = \{\{1, 2, 3\}, \{4, 5\}, \{6, 7, 8\}\}$
- And suppose you also have an information partition,  $\mathcal{P}_2 = \{\{1, 2\}, \{3, 4\}, \{5\}, \{6, 7\}, \{8\}\}$
- Note that  $\mathcal{P}_1$  is neither coarser or finer than  $\mathcal{P}_2$  – they can't be compared in this way

Even though your partition has more pieces, there are some things I know that you don't know – we can't unambiguously say your information is better than mine

For example, consider the event  $E = \{4, 5, 6\}$  – an event that occurs in states 4, 5, and 6, but not otherwise. When  $\omega = 4$ , I know that the event has occurred, but you don't – since for all you know, the state might be 3.

- Now, there are two binary operations we want to define on information partitions, sort of analogous to taking unions and intersections of sets
- Given two information partitions  $\mathcal{P}_1$  and  $\mathcal{P}_2$ , we define their *join*  $\mathcal{P}_1 \vee \mathcal{P}_2$  as the *coarsest common refinement* of the two partitions

- basically, the set of intersections of one of your partition elements with one of my partition elements
- in this example,  $\mathcal{P}_1 \vee \mathcal{P}_2 = \{\{1, 2\}, \{3\}, \{4\}, \{5\}, \{6, 7\}, \{8\}\}$
- this represents what we would know if we shared our information.

- We also define the *meet*  $\mathcal{P}_1 \wedge \mathcal{P}_2$  as the *finest common coarsening*

- this is the finest partition such that each of my elements is contained in a single element of the meet, and each of your elements is contained in a single element of the meet
- in this example, 1, 2 and 3 have to be in the same element of the meet, because they're in the same element of  $\mathcal{P}_1$ ;

3 and 4 need to be together, because they're in the same element of  $\mathcal{P}_2$ ;

and 4 and 5 need to be together, because they're together in  $\mathcal{P}_1$

It turns out, the meet is  $\mathcal{P}_1 \wedge \mathcal{P}_2 = \{\{1, 2, 3, 4, 5\}, \{6, 7, 8\}\}$

- It's not yet clear what this represents intuitively – but it will be soon

### 3 Today's Big Result

- Today's big result is from Aumann, "Agreeing to Disagree," and it says the following:

**Theorem** (Aumann). *If two people have the same priors, and their posteriors for an event  $E$  are common knowledge, then these posteriors are equal.*

- In other words, we can't agree to disagree
  - if it's *common knowledge* that you think the probability of  $E$  given your information is  $x$ ,  
and it's *common knowledge* that I think the probability of  $E$  given my information is  $y$ ,  
then  $x = y$
- To prove this, we first have to figure out what it means for our posteriors to be common knowledge
- If it's common knowledge that you put probability  $x$  on event  $E$ , then this means...
  - Obviously, you have to believe  $\Pr(E) = x$  – which means that  $\Pr(E|P_2) = x$  at the element of  $\mathcal{P}_2$  that contains  $\omega$
  - But also, I have to know that you believe  $\Pr(E) = x$  – which means  $\Pr(E|P_2) = x$  at every element of  $\mathcal{P}_2$  that *I think you might be at*  
Which is every element of  $\mathcal{P}_2$  that intersects the element of  $\mathcal{P}_1$  that contains  $\omega$
  - And also, *you need to know* that I know that you believe  $\Pr(E) = x$   
Which means  $\Pr(E|P_2) = x$  at every element of  $\mathcal{P}_2$  that *you might think* I might think you are at – which is every element of  $\mathcal{P}_2$  that intersects an element of  $\mathcal{P}_1$  that intersects the element of  $\mathcal{P}_2$  that contains  $\omega$
  - And *I need to know* you know I know you believe  $\Pr(E) = x$   
Which means  $\Pr(E|P_2) = x$  at every element of  $\mathcal{P}_2$  that *I might think* you might think I might think you are at  
Which is every element of  $\mathcal{P}_2$  that intersects an element of  $\mathcal{P}_1$  that intersects the element of  $\mathcal{P}_2$  that intersects the element of  $\mathcal{P}_1$  that contains  $\omega$
  - And so on
  - Remember our old example, with  $\mathcal{P}_1 = \{\{1, 2, 3\}, \{4, 5\}, \{6, 7, 8\}\}$  and  
 $\mathcal{P}_2 = \{\{1, 2\}, \{3, 4\}, \{5\}, \{6, 7\}, \{8\}\}$   
If  $\omega = 5$ , I don't know whether the state is 4 or 5, so I don't know whether you're at information partition element  $\{3, 4\}$  or  $\{5\}$

Which means I don't know whether you think I'm at  $\{1, 2, 3\}$  or  $\{4, 5\}$

Which means I don't know whether you think I think you're at  $\{1, 2\}$ ,  $\{3, 4\}$ , or  $\{5\}$

So for your posterior belief to be common knowledge,

it has to be the same at each of those three information sets

- In general, for it to be common knowledge that your posterior belief is  $x$ , we need  $\Pr(E|P_2) = x$  at every  $P_2$  in the element of  $\mathcal{P}_1 \wedge \mathcal{P}_2$  that contains  $\omega$
- And likewise, for it to be common knowledge that my posterior is  $y$ , then  $\Pr(E|P_1) = y$  at every  $P_1$  in the element of  $\mathcal{P}_1 \wedge \mathcal{P}_2$  that contains  $\omega$

- Now we're ready to prove Aumann's result.

### Proof of Aumann's Theorem

- fix a state of the world  $\omega$  and an event  $E$
- let  $P$  be the element of  $\mathcal{P}_1 \wedge \mathcal{P}_2$  that contains  $\omega$  – so in our recent example,  $P$  would be  $\{1, 2, 3, 4, 5\}$
- For it to be *common knowledge* that my posterior is at  $q_1$ , my posterior must be  $q_1$  at every element of  $\mathcal{P}_1$  that's inside  $P$ .
- Now since  $\mathcal{P}_1 \wedge \mathcal{P}_2$  is a coarsening of  $\mathcal{P}_1$ , write  $P$  as the union of some elements of  $\mathcal{P}_1$ ,

$$P = \bigcup_i P^i$$

We just said my posterior probability must be  $q_1$  at each of these elements  $P^i$ , or

$$\Pr(E|P^i) = \frac{\Pr(E \cap P^i)}{\Pr(P^i)} = q_1$$

or

$$\Pr(E \cap P^i) = q_1 \Pr(P^i)$$

But now since the different elements  $P^i$  of my partition are disjoint, and their union is  $P$ , we can sum over  $i$  and get

$$\begin{aligned} \sum_i \Pr(E \cap P^i) &= \sum_i q_1 \Pr(P^i) \\ &\downarrow \\ \Pr(E \cap P) &= q_1 \Pr(P) \\ &\downarrow \\ \frac{\Pr(E \cap P)}{\Pr(P)} &= q_1 \end{aligned}$$

- But if we let  $q_2$  be your prior belief about the probability of  $E$ , and assume that *that* is common knowledge, then we can do the same steps, and show that

$$\frac{\Pr(E \cap P)}{\Pr(P)} = q_2$$

- So if our posterior probabilities  $q_1$  and  $q_2$  on  $E$  are both common knowledge, then  $q_1 = q_2$ , assuming we put the same prior probabilities  $p$  on each state of the world – even if my posterior is based on different information than yours!

Three additional things to note.

**Just knowing each others' posteriors is not enough.**

- From Aumann. Consider a set of states  $\Omega = \{1, 2, 3, 4\}$ , with equal prior on each state. My information partition is  $\mathcal{P}_1 = \{\{1, 2\}, \{3, 4\}\}$ , and yours is  $\mathcal{P}_2 = \{\{1, 2, 3\}, \{4\}\}$ . Consider the event  $E = \{1, 4\}$ , and suppose the state is 1.
- Then I know the state is either 1 or 2, with equal probability, and that the event  $\{1, 4\}$  therefore has probability  $\frac{1}{2}$ .
- And you know that that's my posterior – because that would be my posterior in any state!
- You know the state is either 1, 2, or 3, so you put posterior  $\frac{1}{3}$  on the event  $A$ .
- But since I know the state is either 1 or 2, I know that you know it's either 1, 2, or 3 – so I know your posterior too.
- So I know your posterior and you know mine, but they're not equal – and they don't have to be, because they're not common knowledge.

Since you think the state might be 3, you think that I might think it's either 3 or 4 – which means you think that I might think your posterior is  $\frac{1}{3}$ , or I might think it's 1.

## How our posteriors might become common knowledge

- Aumann says that we have to agree if our posteriors are common knowledge, but doesn't explain how that would come to happen
- A different paper – Geanakoplos and Polemarchakis, “We Can't Disagree Forever” – offers one way it could happen
- Suppose I learn whatever I learn and calculate my posterior probability, you learn whatever you learn and calculate yours – and then we do the following.
  - I announce my posterior probability.
  - Based on that new information, you revise your posterior probability, and then announce your new one.
  - Based on that new information, I revise my posterior, and announce my new one.
  - And so on.
- As long as  $\Omega$  is finite, they show that in finite time, our posteriors will become common knowledge – at which point they must be equal.
- They also give a cute example of how, even if our “communication” just consists of saying the same thing over and over for a while, we still eventually converge.
- An example from their paper.  
 $\Omega = \{1, 2, 3, 4, 5, 6, 7, 8, 9\}$ ,  $p(\omega) = \frac{1}{9}$  for each  $\omega \in \Omega$   
 $\mathcal{P}_1 = \{\{1, 2, 3\}, \{4, 5, 6\}, \{7, 8, 9\}\}$   
 $\mathcal{P}_2 = \{\{1, 2, 3, 4\}, \{5, 6, 7, 8\}, \{9\}\}$   
Consider the event  $E = \{3, 4\}$ , and suppose  $\omega = 1$
- Work it out – communication goes  $p = \frac{1}{3}, q = \frac{1}{4}, p = \frac{1}{3}, q = \frac{1}{3}$ .
- If instead  $E = \{1, 5, 9\}$  and  $\omega = 1$ , it would go  $p = \frac{1}{3}, q = \frac{1}{4}, p = \frac{1}{3}, q = \frac{1}{4}, p = \frac{1}{3}, q = \frac{1}{3}$ .

### Common knowledge of posteriors does not imply we know everything.

- Also from Geanakoplos and Polemarchakis
- Since we communicate until our beliefs converge, one might suspect that I learn “all there is to know” from your information, and vice versa
- But this turns out not to be the case.
- Suppose we each get to flip a coin, and the event we’re interested in is that the two coins matched – either both heads or both tails.
- Suppose we each flip our coin, and they both come up heads.
- I believe the probability the two coins match is  $\frac{1}{2}$ .

And you know that. And I know you know it.

It’s common knowledge my posterior is  $\frac{1}{2}$ , because my posterior would be  $\frac{1}{2}$  at either information set.

Same for you – your posterior is  $\frac{1}{2}$ , and that’s common knowledge.

- And of course, since both posteriors are common knowledge, they have to match (Aumann). And since they’re already common knowledge, when I announce that my posterior is  $\frac{1}{2}$ , it doesn’t change your beliefs, and vice versa – our beliefs have already converged.
- But if we shared information, we would know that the two coins matched with probability 1, not probability  $\frac{1}{2}$ .

### References

- Robert Aumann (1976), “Agreeing to Disagree,” *Annals of Statistics*
- Geanakoplos and Polemarchakis (1982), “We Can’t Disagree Forever,” *Journal of Economic Theory*