

About

These notes are from Annie Liang's Spring 2018 Topics in Microeconomic Theory seminar (ECON 712). If you find an error, please send me an email at {ianzach+notes[at]seas.upenn.edu} so I can correct it.

Aumann (1976)	1-1
Geanakopolos and Polemarchkis (1982)	2-1
Rubenstein (1989)	3-1
Monderer and Samet (1989)	4-1

Econ 712: Topics in Micro Theory

January 10, 2018

Agreeing to Disagree, Aumann (1976)

Professor Annie Liang

Zach Schutzman

Overview

1. defined “common knowledge”
2. introduced the partitional model (called the Aumann model now)
3. showed that agents cannot “agree to disagree” under a common prior

Idea: If there is a common prior and common knowledge of posterior beliefs (I know that you know that I know... and so on, forever), then the posteriors are identical.

Framework

Let (Ω, \mathcal{B}, p) be a probability space, with (Ω, B) the state space and p a common prior. Denote by \mathcal{P}_i the partition of Player i . At the state $\omega \in \Omega$, Player i learns $P_i(\omega)$, that is, the chunk of her partition which contains ω . Note that we are implicitly assuming that a player cannot believe that the true state of the world is $\omega \in A$ if in fact $\omega \notin A$.

We can illustrate this with an example:

$$\begin{aligned}\Omega &= \{1, 2, 3, 4, 5, 6\} \\ \mathcal{P}_1 &= \{\{1, 2, 3\}, \{4, 5\}, \{6\}\} \\ \mathcal{P}_2 &= \{\{1, 2\}, \{3, 4\}, \{5\}, \{6\}\}\end{aligned}$$

The **join** $\mathcal{P}_1 \vee \mathcal{P}_2$ is the coarsest common refinement of \mathcal{P}_1 and \mathcal{P}_2 .

→ in our example, $\mathcal{P}_1 \vee \mathcal{P}_2 = \{\{1, 2\}, \{3\}, \{4\}, \{5\}, \{6\}\}$

We can think of the join as ‘what players can get by pooling their knowledge’.

The **meet** $\mathcal{P}_1 \wedge \mathcal{P}_2$ is the finest common coarsening of \mathcal{P}_1 and \mathcal{P}_2 .

→ in our example, $\mathcal{P}_1 \wedge \mathcal{P}_2 = \{\{1, 2, 3, 4, 5\}, \{6\}\}$

Given a state $\omega \in \Omega$, an event E is **common knowledge at ω** if E includes the member of the meet which includes ω .

→ The set $\{1, 2, 3, 4, 5\}$ is common knowledge at $\omega = 3$.

Observe that Ω itself is always common knowledge.

Result

Fix an event A . Let q_i be the posterior probability of A given Player i 's information. That is,

$$q_i = \frac{p(A \cap P_i(\omega))}{p(P_i(\omega))}$$

If it is common knowledge at ω that $q_1 = a$ and $q_2 = b$ then $a = b$.

Proof: Let P be the member of the meet containing ω . Write $P = \bigcup_j P_j$ where the P_j are disjoint elements of \mathcal{P}_1 .

Since $q_1 = a$ is common knowledge at ω , it must be that $q_1 = a$ at each partition element P_j . Thus, for all j ,

$$\begin{aligned} a &= p(A \cap P_j)/p(P_j) \\ ap(P_j) &= p(A \cap P_j) \\ a \sum_j p(P_j) &= \sum_j p(A \cap P_j) \\ ap(P) &= p(A \cap P) \end{aligned}$$

By doing the same thing for Player 2, we can get an expression that says $bp(P) = p(A \cap P)$, so $a = b$. ■

★ Note that knowledge of the posterior alone is not sufficient. Consider the following example:

$$\Omega = \{1, 2, 3, 4\}$$

$$\mathcal{P}_1 = \{\{1, 2\}, \{3, 4\}\}$$

$$\mathcal{P}_2 = \{\{1, 2, 3\}, \{4\}\}$$

$$\omega = 2, A = \{1, 4\}, \text{ uniform prior}$$

Player 1's posterior at A is $1/2$, Player 2's is $1/3$. Player 2 knows that Player 1's posterior is $1/2$, but not *why*, as Player 2 doesn't know which of the two chunks she has been told contains ω .

Econ 712: Topics in Micro Theory

January 10, 2018

We can't disagree forever, Geanakoplos and Polemarchakis (1982)

Professor Annie Liang

Zach Schutzman

Overview

1. points out that the starting point of Aumann (1976) is rather severe. When are posterior beliefs common knowledge? Is that a bad assumption?
2. Asks whether we can get to the same place through a more realistic process - repeated communication of posterior beliefs.
3. Answer: if the partitions are finite, then yes.
4. Repeated communication is generally equivalent to pooling information, but there exist counterexamples.

Example:

$$\Omega = \{1, 2, \dots, 9\}$$

$$\mathcal{P}_1 = \{\{1, 2, 3\}, \{4, 5, 6\}, \{7, 8, 9\}\}$$

$$\mathcal{P}_2 = \{\{1, 2, 3, 4\}, \{5, 6, 7, 8\}, \{9\}\}$$

$$\omega = 1, A = \{3, 4\}, \text{ uniform prior}$$

Player 1 announces her posterior is $1/3$, and Player 2 announces $1/2$. Player 1 knows that the true state is in $\{1, 2, 3, 4\}$, but she already knew this, so her posterior is still $1/3$. Player 2 doesn't learn anything either, as her belief is $1/3$ for $\{1, 2, 3\}$ and for $\{4, 5, 6\}$. Player 1 announces $1/3$ again. But, since she *didn't* announce 1, Player 2 knows that Player 1 thinks it's in $\{1, 2, 3\}$, so Player 2 announces $1/3$ as well. They have arrived at the same posterior and thus pooled their information.

We can describe this communication protocol in an algorithm:

1. Let $P^1 = \{P_1^1, \dots, P_K^1\}$ and $P^2 = \{P_1^2, \dots, P_L^2\}$
2. Player 1 announces initial posterior $q_1^1(\omega) = \frac{p(P^1(\omega) \cap A)}{p(P^1(\omega))}$
3. Player 2 learns that ω is in $\bigcup_{k \in a_1} P_k^1$ where $a_1 = \left\{k : \frac{p(P_k^1 \cap A)}{p(P_k^1(\omega))} = q_1^1(\omega)\right\}$. That is, all of the partitions for which Player 1 would have announced what she did in the previous step.
4. Player 2 announces a revised posterior $q_1^2(\omega) = \frac{p(P^2(\omega) \cap \bigcup_{k \in a_1} P_k^1 \cap A)}{p(P^2(\omega) \cap \bigcup_{k \in a_1} P_k^1)}$
5. Player 1 performs a similar revision, and the process repeats.

Result

The algorithm described converges in no more than $K + L$ announcements, where K and L are the sizes of the players' partitions.

What posteriors does the process converge to? Interestingly, they need not be the same posteriors as in the setting where the players pool information, such as in the following example:

$$\Omega = \{1, 2, 3, 4\}$$

$$\mathcal{P}_1 = \{\{1, 2\}, \{3, 4\}\}$$

$$\mathcal{P}_2 = \{\{1, 3\}, \{2, 4\}\}$$

$$\omega = 1, A = \{1, 4\}, \text{ uniform prior}$$

Each player initially announces a posterior of $1/2$, and nothing is learned. Had they pooled their knowledge, they could know for certain that the state of the world is $\omega = 1$.

Econ 712: Topics in Micro Theory

January 10, 2018

The Electronic Mail Game, Rubenstein (1989)

Professor Annie Liang

Zach Schutzman

Overview

1. How sensitive are strategic predictions to assumptions about knowledge?
2. Demonstrated a case in which predictions are very sensitive. Subsequently, there has been debate over whether this is the “right” notion of common/mutual knowledge.

Framework

Consider a two-player game where each player chooses from actions $\{A, B\}$. There are two states the world could be in, a or b , and each corresponds with a different payoff matrix.

	World a		World b	
	A	B	A	B
A	M, M	$0, -L$	$0, 0$	$0, -L$
B	$-L, 0$	$0, 0$	$-L, 0$	M, M

L and M are arbitrary values satisfying $L > M > 0$ and, letting p denote the probability we are in World a , suppose that $p < 1/2$. Observe that (A, A) is better in World a and (B, B) in World b , but that A is a ‘safe’ action, in that regardless of which world we are in, playing A always has non-negative payoff.

Consider the following communication protocol:

1. Player 1 learns the state of the world
2. If the state is a , nothing happens. If the state is b , her computer sends an email to Player 2, which fails to arrive with probability $\epsilon > 0$.
3. If Player 2 receives an email, he sends one back to Player 1, which fails to send with probability $\epsilon > 0$.
4. This continues until an email fails. That is, the computers automatically send out a new email after receiving one.

Let T_i be the type of Player i , and set it equal to the number of messages Player i ’s computer sent. If $T_1 = T_2 = \infty$, then it is common knowledge that we are in World b . If T_1 and T_2 are both finite and strictly greater than zero, it is mutual, but not common, knowledge that we are in World b .

Result

Proposition 3.1 *There exists a unique Nash equilibrium in which Player 1 plays A in World a. In this equilibrium, both players play A regardless of the number of messages sent.*

Proof: We proceed by induction. Assume that Player 1 will play A if she knows that we are in World a. Denote this by $S_1(0) = A$, the strategy of Player 1 when she is type 0 is A.

In this case, Player 2 thinks that either we are in World a and Player 1 never sent a message at all, or we are in World b but Player 1's first message got lost. The first happens with probability $(1-p)/(1-p+p\epsilon)$ and the second with probability $(p\epsilon)/(1-p+p\epsilon)$.

Then Player 2's expected payoff to A is at least (actually equal to)

$$\frac{M(1-p) + 0(p\epsilon)}{1-p+p\epsilon}$$

and his expected payoff to B is no more than

$$\frac{-L(1-p) + M(p\epsilon)}{1-p+p\epsilon}$$

By our assumptions on p, L, M , the expected payoff to A is strictly greater than the maximum expected payoff to B, so Player 2 plays A in this case, so $S_2(0) = A$.

Now, suppose that $S_i(T_i) = A$ for $T_i < t$. Consider $T_i = t$. There are two possibilities, either Player 1's t^{th} message was lost or Player 2's t^{th} reply was lost. These happen with conditional probabilities $\epsilon/(\epsilon + (1-\epsilon)\epsilon)$ and $((1-\epsilon)\epsilon)/(\epsilon + (1-\epsilon)\epsilon)$, which we will call z and $1-z$, respectively, for ease of notation. Observe that we know that $z > 1/2$.

The expected payoff to A is zero, as we are certainly in World b. The expected payoff to B is at most $z(-L) + (1-z)M$. Because $L > M$ and $z > 1/2$, this is a negative value, so the expected payoff to A is strictly greater than that to B.

Since playing A dominates, $S_1(t) = A$, and a symmetric argument for Player 2 shows that $S_2(t) = A$ as well. ■

Conclusion

1. Is this notion of 'almost common knowledge' reasonable?
2. One objection is that taking the limit as ϵ goes to zero does not yield the same game as in the case where ϵ is actually equal to zero.
3. Formally, interim types are close to types in the product topology, which we will discuss later.
4. Rubenstein's argument is that nevertheless, high T_i is like common knowledge.

Econ 712: Topics in Micro Theory

January 10, 2018

Approximating CK with Common Beliefs, Monderer and Samet (1989)

Professor Annie Liang

Zach Schutzman

Overview

1. What should we mean when we say “almost common knowledge”?
2. Rubenstein (1989) looks at ‘knowledge which is almost-common’
3. This paper looks at ‘common almost-knowledge’

Main Contribution

1. Weakens common knowledge to a notion of *common p-beliefs*
2. Shows that equilibrium predictions under common *p*-belief are close to those under common knowledge when *p* is large

Quick Review of the Standard Model

- I - the set of all agents
- (Ω, Σ, μ) - a probability space
- π_i - the partition of Ω into measurable sets with positive measure of Player i
- \mathcal{F}_i - a σ -field generated by the collection of π_i

The ‘knowledge’ of Player i at the event E means $K_i(E) = \{\omega : \pi_i(\omega) \subset E\}$

We can define *common knowledge* in the following two equivalent ways:

- A) Let C^1 be the set of states at which every agent knows E
Recursively let C^n be the set of states at which every agent knows C^{n-1} .
Let $C(E) = \bigcap_{n \geq 1} C^n$. Then E is *common knowledge at ω* if and only if $\omega \in C(E)$.
- B) Call E *evident knowledge* if $E \in K_i(E)$ for all i . In other words, E is evident knowledge if whenever E occurs, every agent knows that it has. An event F is *common knowledge at ω* if there exists an evident knowledge event E such that $\omega \in E$ and for all agents i , $E \subset K_i(F)$.

New Framework

In the new framework, we'll let I , (Ω, Σ, μ) , and π_i be as before.

Definition 4.2 *Player i **p-believes an event E at ω if $\mu(E|\pi_i(\omega)) \geq p$.***

As a special case, if we take $p = 1$, this is almost, but not exactly the same as, knowledge.