Last Time:

Defined knowledge, common knowledge, meet (of partitions), and reachability.

Reminders:

- *E is common knowledge at* ω if $\omega \in K_I^{\infty}(E)$.

- *"Reachability Lemma"* $:\omega' \in M(\omega)$ if there is a chain of states $\omega = \omega_0, \omega_1, \ldots \omega_m = \omega'$ such that for each ω_k there is a player *i(k)* s.t. $h_{i(k)}(\omega_k) = h_{i(k)}(\omega_{k+1})$:

- **Theorem:** Event E is common knowledge at ω iff $M(\omega) \subseteq E$.

How does set of NE change with information structure?

Suppose there is a finite number of payoff matrices $u^1,...,u^L$ for finite strategy sets $S_1,...,S_I$

State space Ω, common prior p, partitions H_i, and a map λ so that payoff functions in state ω are $u^{\lambda(\omega)}(.)$; the strategy spaces are maps from H_i into S_i.

When the state space is finite, this is a finite game, and we know that NE is u.h.c. and generically l.h.c. in *p*. In particular, it will be l.h.c. at strict NE.

The "coordinated attack" game

	A	B
A	$8, 8$	$-10, 1$
B	$1, -10$	$0, 0$

u^a

	A	B
A	$0, 0$	$-10, 1$
B	$1, -10$	$8, 8$

u^b

$\Omega = 0,1,2,\ldots.$

In state 0: payoff functions are given by matrix u^b;
In all other states payoff functions are given by u^a.

partitions of Ω
H_1: (0), (1,2), (3,4),… (2n-1,2n)...
H_2 (0,1),(2,3). ..(2n,2n+1)…

Prior p: p(0)=2/3, p(k)= $e(1-e)^{k-1}/3$ for k>0 and $\varepsilon \in (0,1)$.

Interpretation: coordinated attack/email:
Player 1 observes Nature's choice of payoff matrix, sends a message to player 2.

Sending messages isn't a strategic decision, it's hard-coded.

Suppose state is n=2k >0. Then 1 knows the payoffs, knows 2 knows them. Moreover 2 knows that 1 knows that 2 knows, and so on up to strings of length k: $n \in K_I^{n-1}(n > 0)$.

But there is no state at which n>0 is c.k.
(*to see this, use reachability…).*

When it is c.k. that payoff are given by u^a, (*A,A)* is a NE. But..

Claim: the only NE is "play B at every information set.".

Proof: player 1 plays B in state 0 (payoff matrix u^b) since it strictly dominates A.

Let $q = p(0\,|\,(0,1))$, and note that $q > 1/2$.

Now consider player 2 at information set (0,1).

Since player 1 plays B in state 0, and the lowest payoff 2 can get to B in state 1 is 0, player 2's expected payoff to B at (0,1) is at least $8q$.

Playing A gives at most $-10q+8(1-q)$,
and since $q>1/2$, playing B is better.

Now look at player 1 at $h_1=(1,2)$.

Let q'=p(1|1,2), and note that $q'=\dfrac{\varepsilon}{\varepsilon+\varepsilon(1-\varepsilon)}>1/2$.

Since 2 plays B in state 1, player 1's payoff to B is at least 8q';
1's payoff to A is at most -10q'+8(1-q)
so 1 plays B
Now iterate..

Conclude that the unique NE is always B- there is no NE in which at some state the outcome is (A,A).

But (*A*,*A*) is a <u>strict</u> NE of the payoff matrix u^a.
And at large n, there is mutual knowledge of the payoffs to high order- 1 knows that 2 knows that …. n/2 times. So "mutual knowledge to large n" has different NE than c.k.

Also, consider "expanded games" with state space $\Omega = 0,1,....n,...\infty$.

For each small positive ε let the distribution p^{ε} be as above:

$p^{e}(0) = 2/3, \quad p^{e}(n) = e(1 - e)^{n-1}/3$ for $0 < n < \infty$ and $p^{\varepsilon}(\infty) = 0$.

Define distribution p^{*} by $p^{*}(0) = 2/3, p^{*}(\infty) = 1/3$. As $\varepsilon \to 0$, probability mass moves to higher *n*, and there is a sense in which p^{*} is the limit of the p^{ε} as $\varepsilon \to 0$.

But if we do say that $p^{\varepsilon} \to p^{*}$ we have a failure of lower hemi continuity at a strict NE.

So *maybe* we don't want to say $p^{\varepsilon} \to p^{*}$, and we don't want to use mutual knowledge to large n as a notion of almost common knowledge.

So the questions:

- When should we say that one information structure is close to another?
- What should we mean by "almost common knowledge"?

This last question is related because we would like to say that an information structure where a set of events E is common knowledge is close to another information structure where these events are almost common knowledge.

Monderer-Samet: Player *i r-believes E at* ω if $p(E \mid h_i(\omega)) \geq r$.

$B_i^r(E)$ is the set of all ω where player i r- believes E; this is also denoted ${}^1B_i^r(E)$.

Now do an iterative definition in the style of c.k.:

$^1B_I^r(E) = \cap_i {}^1B_i^r(E)$ *(everyone r-believes E)*

$^nB_i^r(E) = \{w \mid p(^{n-1}B_I^r(E) \mid h_i(w)) \geq r\}$

$^nB_I^r(E) = \cap_i {}^nB_i^r(E)$

E is common r belief at ω if $w \in {}^{\infty}B_I^r(E)$

As with c.k., common r-belief can be characterized in terms of public events:

- An event is a common r-truism if everyone r - believes it when it occurs.

- An event is common r -belief at ω if it is implied by a common r-truism at ω.

Now we have one version of "almost ck" : An event is almost ck if it is common r-belief for r near 1.

MS show that if two player’s posteriors are common r-belief, they differ by at most 2(1-r): so Aumann's result is robust to almost ck, and holds in the limit.

MS also that a strict NE of a game with known payoffs is still a NE when payoffs are "almost ck" - a form of lower hemi continuity.

More formally:
As before consider a family of games with fixed finite action spaces A_i for each player i.
a set of payoff matrices $u^l : A -> R^I$,
a state space W, that is now either finite or countably infinite, a prior p, a map $l : W \to 1,,,L$ such that payoffs at ω are $u(a,w) = u^{l(w)}(a)$.

Payoffs are common r-belief at ω if the event $\{w \mid l(w) = l\}$ is common r belief at ω.

For each λ let σ^λ be a NE for common- knowledge payoffs u^l.
Define s* by $s^*(w) = s^{l(w)}$.
This assigns each w a NE for the corresponding payoffs.
In the email game, one such s^* is
$s^*(0) = (B,B), \quad s^*(n) = (A,A) \forall n > 0$.

If payoffs are c.k. at each ω, then s* is a NE of overall game G. (*discuss)*

<u>Theorem</u>: Monder-Samet 1989

Suppose that for each l , s^{l} is a strict equilibrium for payoffs u^{λ}.
Then for any $e >$ 0 there is $\bar{r} < 1$ and $\bar{q} < 1$ such that for all r Î $[\bar{r}, 1]$ and q Î $[\bar{q}, 1]$,
if there is probability q that payoffs are common r-belief, then there is a NE s of G with
$p(\omega \mid s(\omega) = s^{*}(\omega)) > 1 - \varepsilon$.

Note that the conclusion of the theorem is false in the email game:
there is no NE with an appreciable probability of playing A, even though (A,A) is a strict NE of the payoffs in every state but state 0.

This is an indirect way of showing that the payoffs are never ACK in the email game.

Now many payoff matrices don’t have strict equilibria, and this theorem doesn’t tell us anything about them.

But can extend it to show that if for each state ω, $s^*(\omega)$ is a Nash (but not necessarily strict Nash) equilibrium, then for any $\varepsilon > 0$ there is $\bar{r} < 1$ and $\bar{q} < 1$ such that for all $r \in [\bar{r}, 1]$ and $q \in [\bar{q}, 1]$, if payoffs are common r-belief with probability q, there is an “interim ε equilibria” of G where s^* is played with probability $1 - \varepsilon$.

Interim ε-equilibria:

At each information set, the actions played are within epsilon of maxing expected payoff

$$E(u_i(s_i(w), s_{-i}(w)) \mid h_i(w)) >=$$

$$E(u_i(s_i', s_{-i}(w)) \mid h_i(w)) - \varepsilon$$

Note that this implies the earlier result when s^* specifies strict equilibria.

Outline of proof:

At states Ω^* where some payoff function is common r-belief, specify that players follow s^*. The key is that at these states, each player i r-believes that all other players r-believe the payoffs are common r-belief, so each expects the others to play according to s^*.

Regardless of play in the other states, playing this way is a $4k(1-r)$ best response, where k is a constant that depends on the set of possible payoff functions.

To define play at states in Ω / Ω^* consider an artificial game where players are constrained to play s^* in Ω^*-and pick a NE of this game.

The overall strategy profile is an interim ε-equilibrium that plays like s^* with probability q.

To see the role of the infinite state space, consider the

"truncated email game"

player 2 does not respond after receiving n messages, so there are only 2n states.

When 2n occurs: 2 knows it occurs.

That is, $H_2 = \{(0,1),...(2n-2,2n-1,)(2n)\}$.
$H_1 = \{(0),(1,2),...(2n-1,2n)\}$.

$p\left(2n\,|\,(2n-1,2n)\right) = 1-\varepsilon$, so $2n$ is a "1-ε truism," and thus it is common 1-ε belief when it occurs.

So there is an exact equilibrium where players play A in state 2n.

More generally: on a finite state space, if the probability of an event is close to 1, then there is high probability that it is common r belief for r near 1.

Not true on infinite state spaces…

Lipman, "Finite order implications of the common prior assumption."

His point: there basically aren't any!
All of the "bite" of the CPA is in the tails.

Set up: parameter Q that people "care about"
States $s \in S$,
$f : S \to \Theta$ specifies what the payoffs are at state *s*.
Partitions H_i of S, priors p_i.

Player i's first order beliefs at s: the conditional distribution on Q given s.

For $B \subseteq \Theta$,
$$d_i^1(s)(B) = p_i(s' \mid f(s')\ \hat{\text{I}}\ B \mid h_i(s))$$

Player i's second order beliefs: beliefs about Q and other players' first order beliefs.

$$d_i^2(s)(B) = p_i\left(\{s' \mid (f(s'), d_j^1(s'))\ \hat{\text{I}}\ B\} \mid h_i(s)\right)$$
and so on.

The main point can be seen in his example

Two possible values of an unknown parameter
$q = q_1$ or q_2.

Start with a model w/o common prior, relate it to a model with common prior.

Starting model has only two states $S = \{s_1, s_2\}$.
Each player has the trivial partition- ie no info beyond the prior.

$p_1(s_1) = p_2(s_2) = 2/3.$

example: Player 1 owns an asset whose value is 1 at θ_1 and 2 at θ_2; $f(s_i) = \theta_i$.
At each state, 1's expected value of the asset 4/3, 2's is 5/3, so it’s common knowledge that there are gains from trade.

Lipman shows we can match the players’ beliefs, beliefs about beliefs, etc. to arbitrarily high order in a common prior model.

Fix an integer N. construct the Nth model as follows

State space $S' = S \times \{1, \ldots 2^N\}$

Common prior is that all states equally likely.

The value of θ at (s,k) is determined by the s-component.

Now we specify the partitions of each player in such a way that the beliefs, beliefs about beliefs, look like the simple model w/o common prior.

1's partition: events

$$\{(s_1, 1), (s_1, 2), (s_2, 1)\} \ldots \{(s_1, 2k - 1), (s_1, 2k), (s_2, k)\}$$

for k up to 2^{N-1}; the “left-over” s_2 states go into

$$\{(s_2, 2^{N-1} + 1), \ldots (s_2, 2^N)\}.$$

At every event but the last one, 1 thinks the probability of q_1 is 2/3.

The partition for player 2 is similar but reversed: $\{(s_2, 2k - 1), (s_2, 2k), (s_1, k)\}$ for k up to 2^{N-1}.

And at all info sets but one, player 2 thinks the prob. of q_1 is 1/3.

Now we look at beliefs at the state $(s_1, 1)$.
We matched the first-order beliefs (beliefs about θ) by construction)

Now look at player 1's second-order beliefs.

1 thinks there are 3 possible states $(s_1, 1)$, $(s_1, 2)$, $(s_2, 1)$.

At $(s_1, 1)$, player 2 knows $\{(s_1, 1), (s_2, 1), (s_2, 2)\}$.
At $(s_1, 2)$, 2 knows $\{(s_1, 2), (s_2, 3), (s_2, 4)\}$.
At $(s_2, 1)$, 2 knows $\{(s_1, 2), (s_2, 1), (s_2, 2)\}$.

The support of 1's second-order beliefs at $(s_1, 1)$ is the set of 2's beliefs at these info sets.

And at each of them 2's beliefs are (1/3 θ_1, 2/3 θ_2).
Same argument works up to N:

The point is that the N-state models are "like" the original one in that beliefs at some states are the same as beliefs in the original model to high but finite order.
(Beliefs at other states are very different- namely at the states where 1 is sure that state is θ_2 or 2 is sure it's θ_1.)

Conclusion: if we assume that beliefs at a given state are generated by updating from a common prior, this doesn't pin down their finite order behavior. So the main force of the CPA is on the entire infinite hierarchy of beliefs.

Lipman goes on from this to make a point that is correct but potentially misleading: he says that "almost all" priors are close to a common. I think its misleading because here he uses the product topology on the set of hierarchies of beliefs- a.k.a topology of pointwise convergence.

And two types that are close in this product topology can have very different behavior in a NE- so in a sense NE is not continuous in this topology.

The email game is a counterexample. “Product Belief Convergence”:

A sequence of types t_{in} converges to t_i^* if the sequence converges pointwise. That is, if for each k,, $\delta_{k,n}^i \to \delta_{k,*}^i$.

Now consider the expanded version of the email game, where we added the state ∞.

Let t_{in} be the hierarchy of beliefs of player 1 when he has sent n messages, and let $t_{i,*}$ be the hierarchy at the point ∞, where it is common knowledge that the payoff matrix is u^a.

Claim: the sequence t_{in} converges pointwise to $t_{i,*}$.
Proof: At t_{in}, i’s zero-order beliefs assign probability 1 to u^a, his first-order beliefs assign probability 1 to (u^a and j knows it is u^a) and so on up to level n-1. Hence as n goes to infinity, the hierarchy of beliefs converges pointwise to common knowledge of u^a.

In other words, if the number of levels of mutual knowledge go to infinity, then beliefs converge to common knowledge in the product topology. But we know that mutual knowledge to high order is not the same as almost common knowledge, and types that are close in the product topology can play very differently in Nash equilibrium.

Put differently, the product topology on countably infinite sequences is insensitive to the tail of the sequence, but we know that the tail of the belief hierarchy can matter.

Next : B-D JET 93 "Hierarchies of belief and Common Knowledge”.

Here the hierarchies of belief are motivated by Harsanyi's idea of modelling incomplete information as imperfect information.

Harsanyi introduced the idea of a player's "type" which summarizes the player's beliefs, beliefs about beliefs etc- that is, the infinite belief hierarchy we were working with in Lipman's paper.

In Lipman we were taking the state space Ω as given.

Harsanyi argued that given any element of the hierarchy of beliefs could be summarized by a single datum called the "type" of the player, so that there was no loss of generality in working with types instead of working explicitly with the hierarchies.

I think that the first proof is due to Mertens and Zamir. B-D prove essentially the same result, but they do it in a much clearer and shorter paper.

The paper is much more accessible than MZ but it is still a bit technical; also, it involves some hard but important concepts. (*Add hindsight disclaimer…)*

Review of math definitions:

A sequence of probability distributions p^n *converges weakly* to p if

$\int f dp^n \to \int f dp$ for every bounded continuous function f. This defines the *topology of weak convergence*.

In the case of distributions on a finite space, this is the same as the usual idea of convergence in norm.

A metric space X is *complete* if every Cauchy sequence in X converges to a point of X.

A space X is *separable* if it has a countable dense subset.

A *homeomorphism* is a map f between two spaces that is 1-1, and onto (an *isomorphism*) and such that f and f-inverse are continuous.

The *Borel sigma algebra* on a topological space S is the sigma-algebra generated by the open sets. (note that this depends on the topology.)

Now for Brandenburger-Dekel

Two individuals (extension to more is easy)
Common underlying space of uncertainty S (this is called Θ in Lipman)

Assume S is a complete separable metric space. ("*Polish*")

For any metric space, let $\Delta(Z)$ be all probability measures on Borel field of Z, endowed with the topology of weak convergence. (the "weak topology.")

$$X_0 = \Delta(S)$$
$$X_1 = X_0 \times \Delta(X_0)$$
$$X_n = X_{n-1} \times \Delta(X_{n-1})$$

So X_n is the space of n-th order beliefs; a point in X_n specifies (n-1)st order beliefs and beliefs about the opponent's (n-1)st order beliefs.

A *type* for player i is a
$t^i = (d_0^i, d_1^i, d_2^i, \ldots) \in \times_{n=0}^{\infty} \Delta(X_n)$==

T_0 .

Now there is the possibility of further iteration: what about i's belief about j's type? *Do we need to add more levels of i's beliefs about j, or is i's belief about j's type already pinned down by i's type*?

Harsanyi’s insight is that we don't need to iterate further; this is what B-D prove formally.

Coherency: a type is coherent if for every n>=2,
$\mathrm{marg}_{X_{n-2}} d_n = d_{n-1}$.

So the n and (n-1)st order beliefs agree on the lower orders. We impose this because it’s not clear how to interpret incoherent hierarchies..

Let T_1 be the set of all coherent types

<u>Proposition</u> (Brandenburger-Dekel) : There is a homeomorphism between T_1 and D$(S$ ´ $T_0)$.

.

The basis of the proposition is the following

<u>Lemma</u>: Suppose Z_n are a collection of Polish spaces and let
$D = \{(d_1, d_2, \ldots) : d_n \in \Delta(Z_0 \times \ldots Z_{n-1}) \forall n \geq 1,$ and $\mathrm{marg}_{Z_0 \times \ldots \times Z_{n-2}} d_n = d_{n-1}$.

Then there is a homeomorphism

$f : D \to \Delta(\times_{n=0}^{\infty} Z_n)$

This is basically the same as Kolmogorov's extension theorem- the theorem that says that there is a unique product measure on a countable product space that corresponds to specified marginal distributions and the assumption that each component is independent.

To apply the lemma, let $Z_0 = X_0$, and
$Z_n = \Delta(X_{n-1})$.

Then $Z_0 \times \ldots \times Z_n = X_n$ and $\times_0^{\infty} Z_n = S \times T_0$.

If S is complete separable metric than so is $\Delta(S)$.

D is the set of coherent types; we have shown it is homeomorphic to the set of beliefs over state and opponent’s type.

In words: coherency implies that i's type determines i's belief over j's type.

But what about i's belief about j's belief about i's type? This needn’t be determined by i’s type if i thinks that j might not be coherent. So B-D impose “common knowledge of coherency.”

Define $T \times T$ to be the subset of $T_1 \times T_1$ where coherency is common knowledge.

<u>Proposition</u> (Brandenburger-Dekel) : There is a homeomorphism between T and D($S \times T$).

Loosely speaking, this says (a) the “universal type space is big enough” and (b) common knowledge of coherency implies that the information structure is common knowledge in an informal sense: each of i’s types can calculate j’s beliefs about i’s first-order beliefs, j’s beliefs about i’s beliefs about j’s beliefs, etc.

Caveats:

1) In the continuity part of the homeomorphism the argument uses the product topology on types. The drawbacks of the product topology make the homeomorphism part less important, but the *isomorphism* part of the theorem is independent of the topology on T.

2) The space D(S ´ T) that is identified as "universal" depends on the sigma-algebra used on $S \times T$. Does this matter?

Loose ideas and conjectures…

- There can't be an isomorphism between a set X and the power set 2^X, so something about measures as opposed to possibilities is being used.

- The "right topology" on types looks more like the topology of uniform convergence than the product topology. (*this claim isn't meant to be obvious. the "right topology" hasn't yet been found, and there may not be one. But Morris' "Typical Types" suggests that something like this might be true.)*

- The topology of uniform convergence generates the same Borel sigma-algebra as the product topology, so maybe B-D worked with the right set of types after all.