CMSC 474, Introduction to Game Theory

18. Bayesian Games & Games of Incomplete Information

Mohammad T. Hajiaghayi
University of Maryland

Introduction

- All the kinds of games we've looked at so far have assumed that everything relevant about the game being played is common knowledge to all the players:
    - the number of players
    - the actions available to each
    - the payoff vector associated with each action vector
- True even for imperfect-information games
    - The actual moves aren't common knowledge, but the game is

- We'll now consider games of **incomplete** (*not **imperfect***) information
    - Players are uncertain about the game being played

Example

- Consider the payoff matrix shown here
  - ε is a small positive constant; Agent 1 knows its value
- Agent 1 doesn't know the values of *a, b, c, d*
  - Thus the matrix represents a ***set*** of games
  - Agent 1 doesn't know which of these games is the one being played

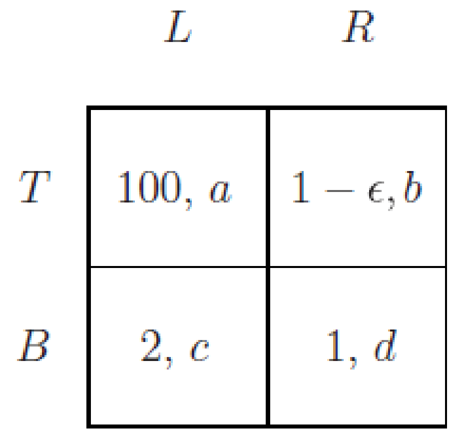

	L	R
T	$100, a$	$1-\epsilon, b$
B	$2, c$	$1, d$

- Agent 1 wants a strategy that makes sense despite this lack of knowledge

- If Agent 1 thinks Agent 2 is malicious, then Agent 1 might want to play a maxmin, or "safety level," strategy
    - minimum payoff of T is 1–ε
    - minimum payoff of B is 1
  - So agent 1's maxmin strategy is B

Bayesian Games

- Suppose we know the set ***G*** of all possible games and we have enough information to put a probability distribution over the games in ***G***
- A **Bayesian Game** is a class of games ***G*** that satisfies two fundamental conditions
- ***Condition 1***:
  - The games in ***G*** have the same number of agents, and the same ***strategy space*** (set of possible strategies) for each agent. The only difference is in the payoffs of the strategies.
- This condition isn't very restrictive
  - Other types of uncertainty can be reduced to the above, by reformulating the problem

Example

- Suppose we don’t know whether player 2 only has strategies L and R, or also an additional strategy C:

Game G_1

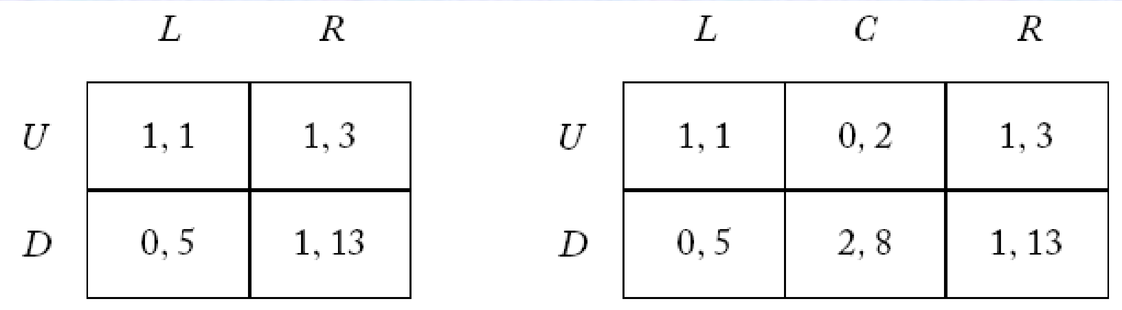

	L	R
U	1, 1	1, 3
D	0, 5	1, 13

	L	C	R
U	1, 1	0, 2	1, 3
D	0, 5	2, 8	1, 13

Game G_2

- If player 2 doesn’t have strategy C, this is equivalent to having a strategy C that’s strictly dominated by other strategies:

Game G_1'

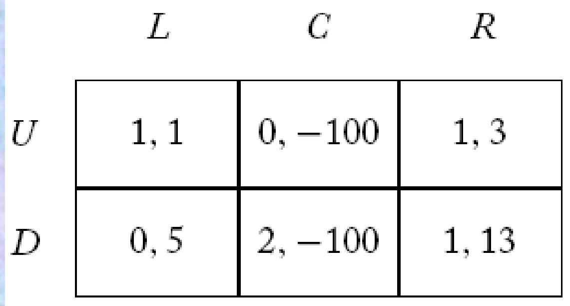

	L	C	R
U	1, 1	0, −100	1, 3
D	0, 5	2, −100	1, 13

  - The Nash equilibria for G_1' are the same as the Nash equilibria for G_1
- We’ve reduced the problem to whether C’s payoffs are those of G_1' or G_2

Bayesian Games

- ***Condition 2 (common prior)*:**
    - The probability distribution over the games in ***G*** is **common knowledge** (i.e., known to all the agents)
- So a Bayesian game defines
    - the uncertainties of agents about the game being played,
    - what each agent believes the other agents believe about the game being played
- The beliefs of the different agents are posterior probabilities
    - Combine the common prior distribution with individual "private signals" (what's "revealed" to the individual players)
- The common-prior assumption rules out whole families of games
    - But it greatly simplifies the theory, so most work in game theory uses it
- There are some examples of games that don't satisfy Condition 2

Definitions of Bayesian Games

- The book discusses three different ways to define Bayesian games
  - All are
    - equivalent (ignoring a few subtleties)
    - useful in some settings
    - intuitive in their own way

- The first definition (Section 7.1.1) is based on information sets
- A Bayesian game consists of
  - a set of games that differ only in their payoffs
  - a common (i.e., known to all players) prior distribution over them
  - for each agent, a partition structure (set of information sets) over the games
- Formal definition on the next page

7.1.1 Definition based on Information Sets

- A **Bayesian game** is a 4-tuple (N,G,P,I) where:
  - N is a set of agents
  - G is a set of N-agent games
  - For every agent i, every game in G has the same strategy space
  - P is a **common prior** over G
    - ***common***: common knowledge (known to all the agents)
    - ***prior***: probability before learning any additional info
  - $I = (I_1, \ldots, I_N)$ is a tuple of partitions of G, one for each agent
    - Information sets

- **Example:**

  G = {Matching Pennies (MP),
  Prisoner's Dilemma (PD),
  Coordination (Crd),
  Battle of the Sexes (BoS)}

$I_{2,1}$ $I_{2,2}$

$I_{1,1}$

MP $(p = 0.3)$

	L	R
U	2, 0	0, 2
D	0, 2	2, 0

PD $(p = 0.1)$

	L	R
U	2, 2	0, 3
D	3, 0	1, 1

$I_{1,2}$

Crd $(p=0.2)$

	L	R
U	2, 2	0, 0
D	0, 0	1, 1

BoS $(p = 0.4)$

	L	R
U	2, 1	0, 0
D	0, 0	1, 2

Example (Continued)

- G = {Matching Pennies (MP),
  Prisoner's Dilemma (PD),
  Coordination (Crd),
  Battle of the Sexes (BoS)}
- Suppose the randomly chosen game is MP
- Agent 1's information set is $I_{1,1}$
  - 1 knows it's MP or PD
  - 1 can infer **posterior probabilities** for each

$$\Pr[\mathrm{MP}|I_{1,1}] = \frac{\Pr[\mathrm{MP}]}{\Pr[\mathrm{MP}]+\Pr[\mathrm{PD}]} = \frac{0.3}{0.3+0.1} = \frac{3}{4}$$

$$\Pr[\mathrm{PD}|I_{1,1}] = \frac{\Pr[\mathrm{PD}]}{\Pr[\mathrm{MP}]+\Pr[\mathrm{PD}]} = \frac{0.1}{0.3+0.1} = \frac{1}{4}$$

- Agent 2's information set is $I_{2,1}$

$$\Pr[\mathrm{MP}|I_{2,1}] = \frac{\Pr[\mathrm{MP}]}{\Pr[\mathrm{MP}]+\Pr[\mathrm{CrD}]} = \frac{0.3}{0.3+0.2} = \frac{3}{5}$$

$$\Pr[\mathrm{Crd}|I_{2,1}] = \frac{\Pr[\mathrm{Crd}]}{\Pr[\mathrm{MP}]+\Pr[\mathrm{CrD}]} = \frac{0.2}{0.3+0.2} = \frac{2}{5}$$

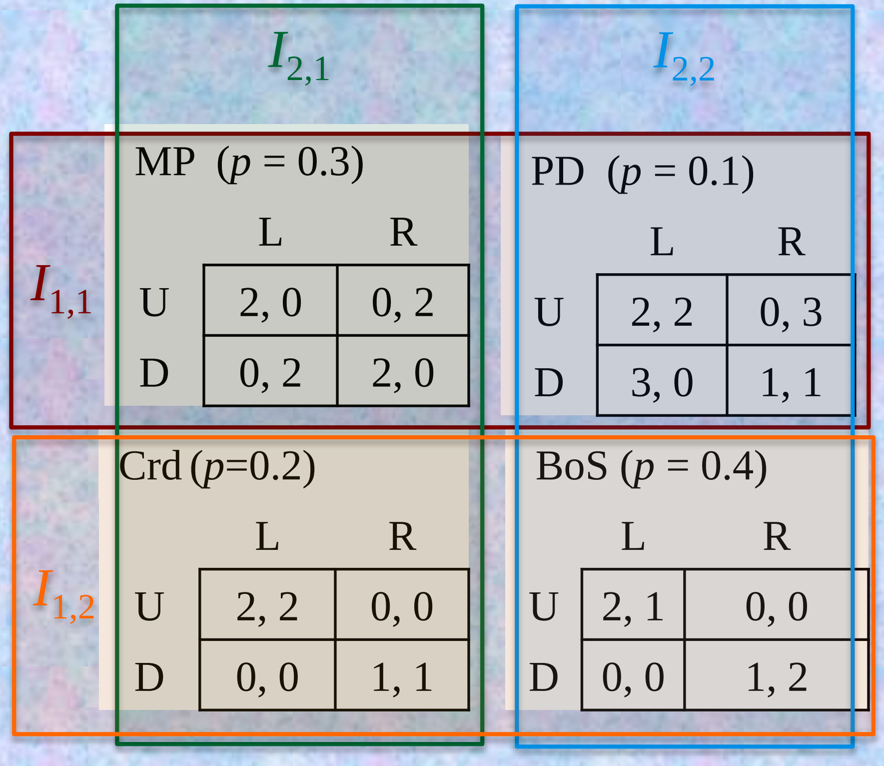

MP (p = 0.3)

	L	R
U	2, 0	0, 2
D	0, 2	2, 0

PD (p = 0.1)

	L	R
U	2, 2	0, 3
D	3, 0	1, 1

Crd (p=0.2)

	L	R
U	2, 2	0, 0
D	0, 0	1, 1

BoS (p = 0.4)

	L	R
U	2, 1	0, 0
D	0, 0	1, 2

7.1.2 Extensive Form with Chance Moves

- **Extensive form with Chance Moves**
  - The book gives a description, but not a formal definition
- Hypothesize a special agent, **Nature**
- Nature has no utility function
  - At the start of the game, Nature makes a probabilistic choice according to the common prior
- The agents receive **individual signals** about Nature’s choice
  - Some of Nature’s choices are “revealed” to some players, others to other players
  - The players receive ***no*** other information
    - In particular, they cannot see each other’s moves

Example

- Same example as before, but translated into extensive form
    - Nature randomly chooses MP, sends signal $I_{1,1}$ to Agent 1, sends signal $I_{2,1}$ to Agent 2

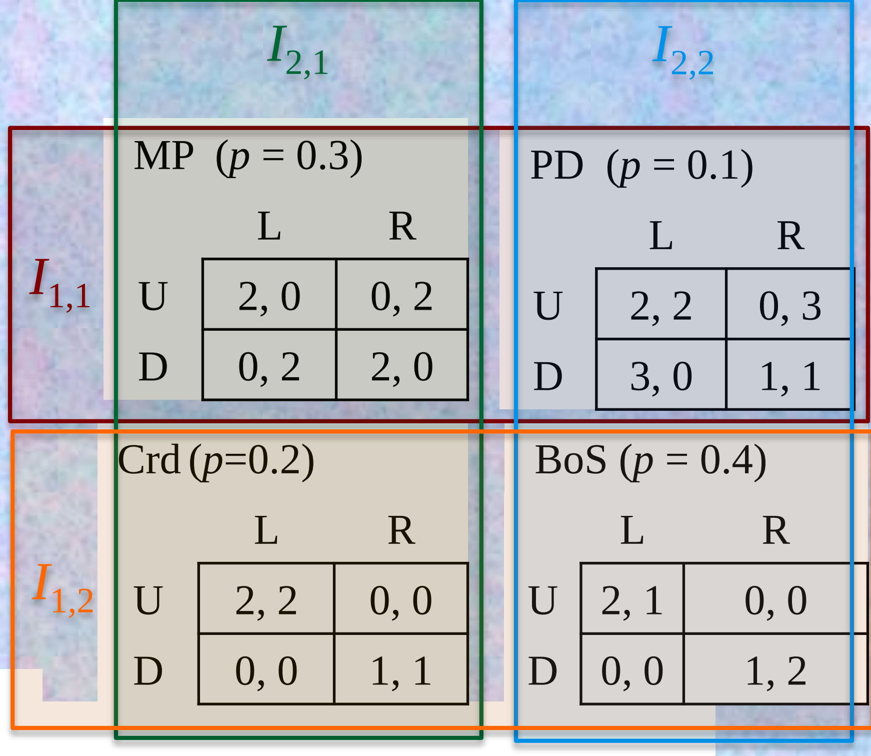

MP (p = 0.3)

	L	R
U	2, 0	0, 2
D	0, 2	2, 0

PD (p = 0.1)

	L	R
U	2, 2	0, 3
D	3, 0	1, 1

Crd (p=0.2)

	L	R
U	2, 2	0, 0
D	0, 0	1, 1

BoS (p = 0.4)

	L	R
U	2, 1	0, 0
D	0, 0	1, 2

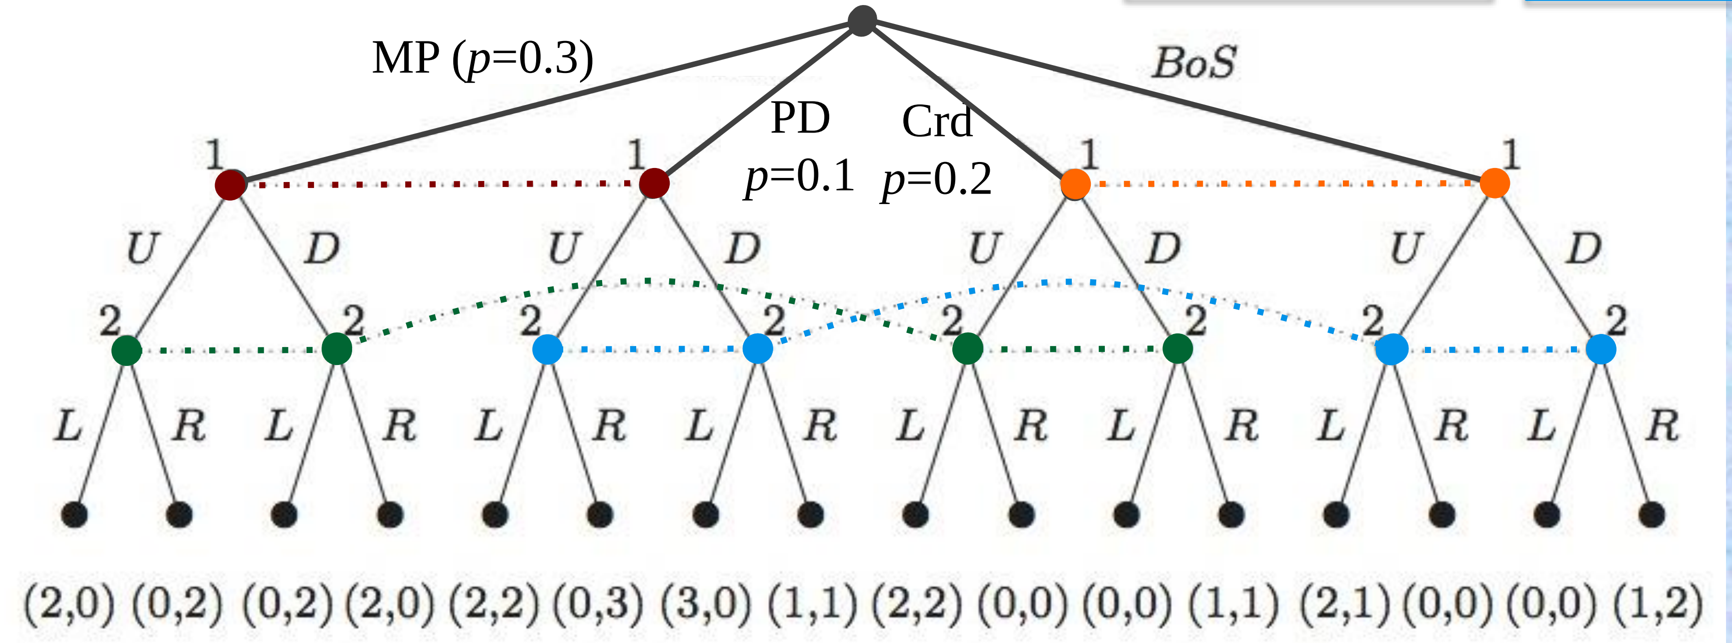

Extensions

- The definition in section 7.1.2 can be extended to include the following:
    - Players sometimes get information about each other’s moves
    - Nature makes choices and sends signals throughout the game
- This allows us to model Backgammon and Bridge

Bridge

- At the start of the game, Nature makes one move
  - The deal of the cards
- Nature signals to each player what that player's cards are
- Each player can always see the other players' moves
  - But imperfect information, since the players can't see each others' hands

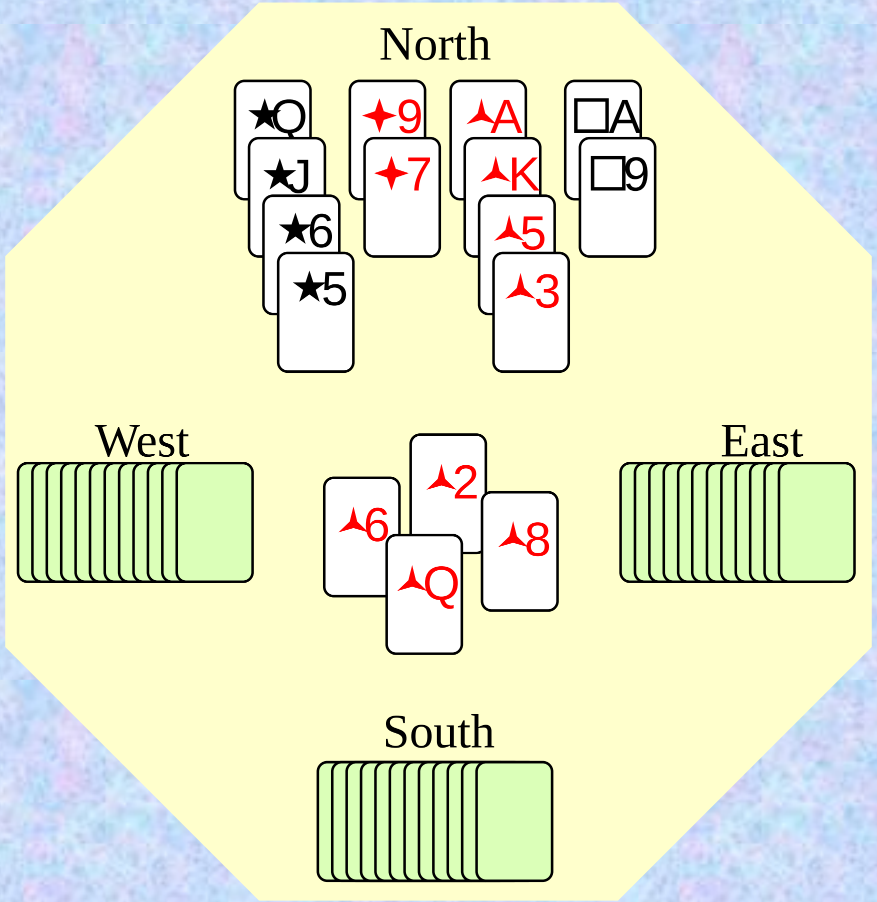

Backgammon

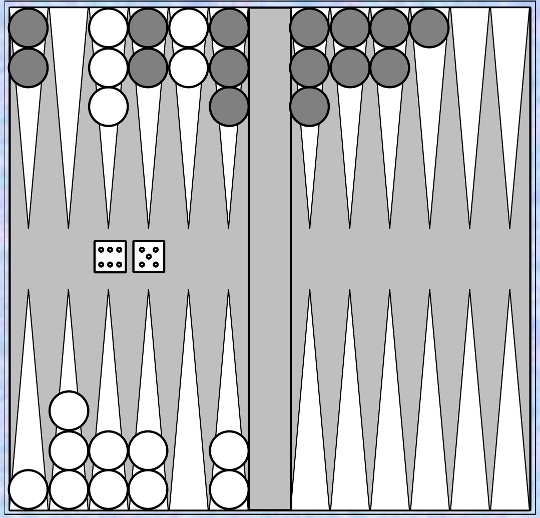

- Nature makes choices throughout the game
  - The random outcomes of the dice rolls
- Nature reveals its choices to both players
  - Both players can see the dice
- Both players always see each other's moves of checkers
- Hence, perfect information

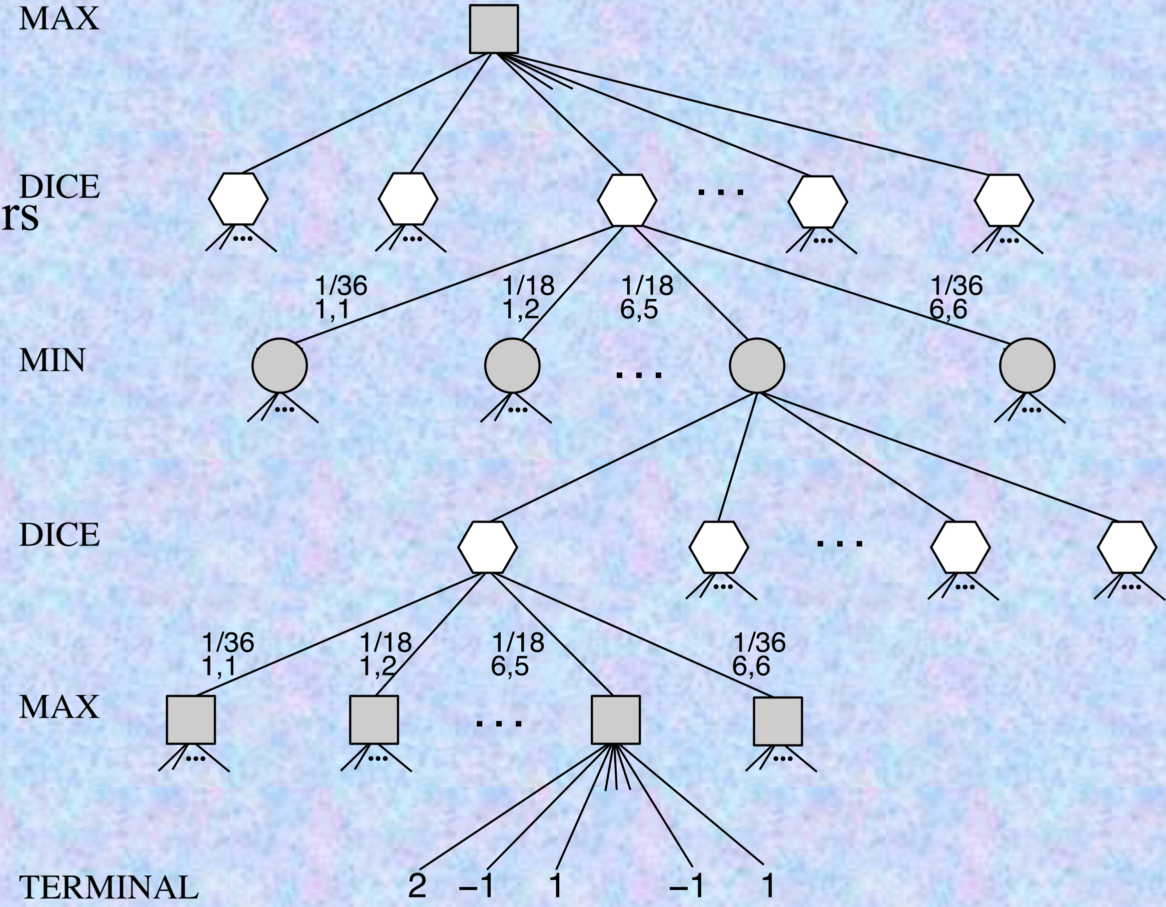

7.1.3 Definition Based on Epistemic Types

- **Epistemic types**
  - Recall that we can assume the only thing players are uncertain about is the game's utility function
  - Thus we can define uncertainty directly over a game's utility function

- **Definition 7.1.2**: a **Bayesian game** is a tuple (N, A, Θ, p, u) where:
  - N is a set of agents;
  - $A = A_1 \times \ldots \times A_n$, where A_i is the set of actions available to player i ;
  - $\Theta = \Theta_1 \times \ldots \times \Theta_n$, where Θ_i is the type space of player i ;
  - $p : \Theta \rightarrow [0, 1]$ is a common prior over types; and
  - $u = (u_1, . . . , u_n)$, where $u_i : A \times \Theta \rightarrow \Re$ is the utility function for player i

- All this is common knowledge among the players
  - And each agent knows its own type

Types

- An agent’s **type** consists of all the information it has that isn’t common knowledge, e.g.,
    - The agent’s actual payoff function
    - The agent’s beliefs about other agents’ payoffs,
    - The agent’s beliefs about *their* beliefs about his own payoff
    - Any other higher-order beliefs

Example

- Agent 1’s possible types: $\theta_{1,1}$ and $\theta_{1,2}$
- 1’s type is $\theta_{1,j}$ $\Leftrightarrow$ 1’s info set is $I_{1,j}$

- Agent 2’s possible types: $\theta_{2,1}$ and $\theta_{2,2}$
- 2’s type is $\theta_{2,j}$ $\Leftrightarrow$ 2’s info set is $I_{2,j}$

- Joint distribution on the types:

$$\Pr[\theta_{1,1}, \theta_{2,1}] = 0.3; \quad \Pr[\theta_{1,1}, \theta_{2,2}] = 0.1$$

$$\Pr[\theta_{1,2}, \theta_{2,1}] = 0.2; \quad \Pr[\theta_{1,2}, \theta_{2,2}] = 0.4$$

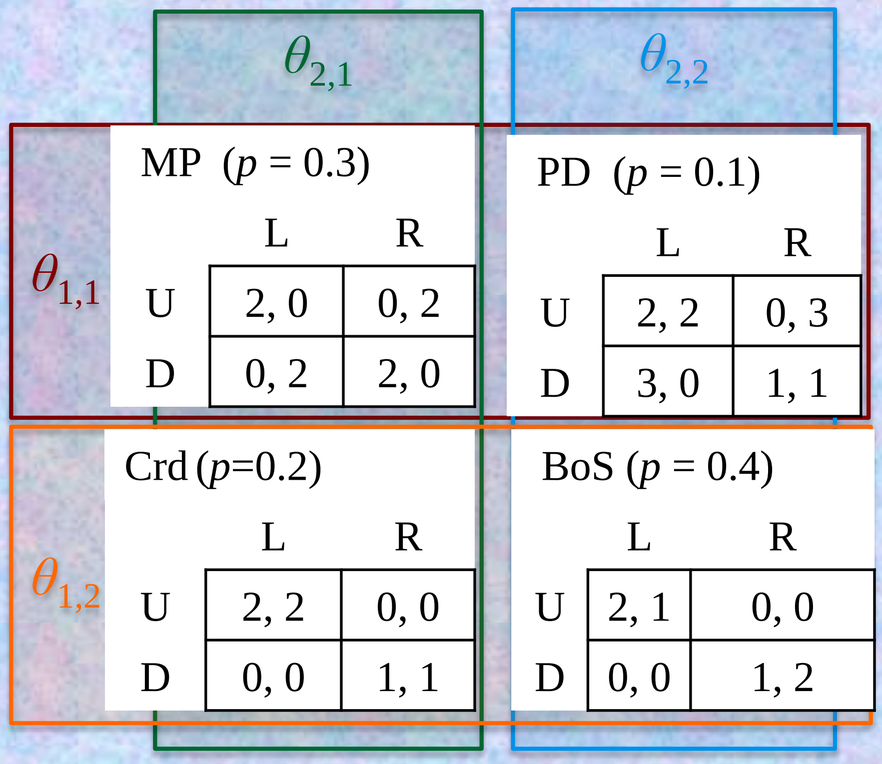

$\theta_{2,1}$ $\theta_{2,2}$

$\theta_{1,1}$

MP (*p* = 0.3)

	L	R
U	2, 0	0, 2
D	0, 2	2, 0

PD (*p* = 0.1)

	L	R
U	2, 2	0, 3
D	3, 0	1, 1

$\theta_{1,2}$

Crd (*p*=0.2)

	L	R
U	2, 2	0, 0
D	0, 0	1, 1

BoS (*p* = 0.4)

	L	R
U	2, 1	0, 0
D	0, 0	1, 2

- Conditional probabilities for agent 1:
  - $\Pr[\theta_{2,1} \mid \theta_{1,1}] = 0.3/(0.3 + 0.1) = 3/4; \quad \Pr[\theta_{2,2} \mid \theta_{1,1}] = 0.1/(0.3 + 0.1) = 1/4$
  - $\Pr[\theta_{2,1} \mid \theta_{1,2}] = 0.2/(0.2 + 0.4) = 1/3; \quad \Pr[\theta_{2,2} \mid \theta_{1,2}] = 0.4/(0.2 + 0.4) = 2/3$

Example (continued)

- The players’ payoffs depend on both their types and their actions
    - The types determine what game it is
    - The actions determine the payoff within that game

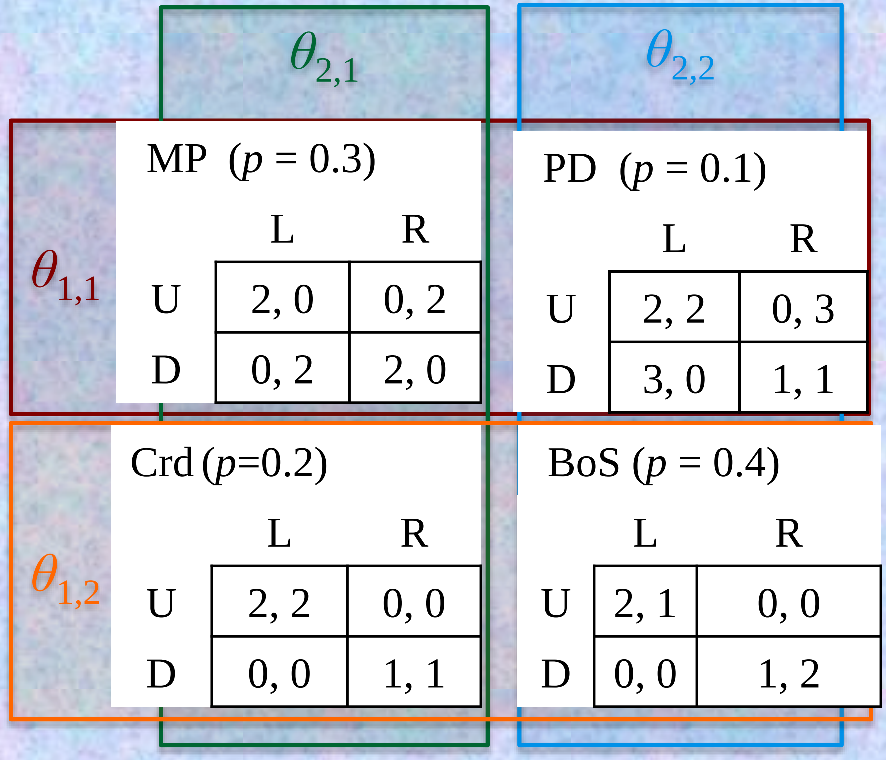

$\theta_{2,1}$ $\theta_{2,2}$

$\theta_{1,1}$

MP (p = 0.3)

	L	R
U	2, 0	0, 2
D	0, 2	2, 0

PD (p = 0.1)

	L	R
U	2, 2	0, 3
D	3, 0	1, 1

$\theta_{1,2}$

Crd (p=0.2)

	L	R
U	2, 2	0, 0
D	0, 0	1, 1

BoS (p = 0.4)

	L	R
U	2, 1	0, 0
D	0, 0	1, 2

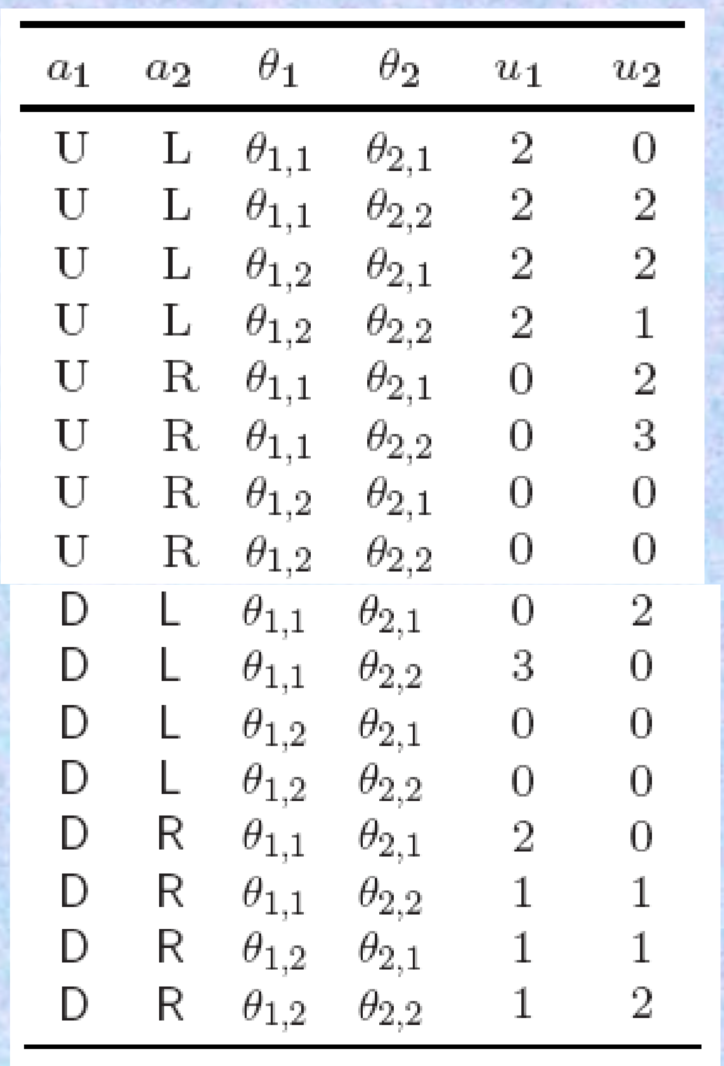

a_1	a_2	θ_1	θ_2	u_1	u_2
U	L	$\theta_{1,1}$	$\theta_{2,1}$	2	0
U	L	$\theta_{1,1}$	$\theta_{2,2}$	2	2
U	L	$\theta_{1,2}$	$\theta_{2,1}$	2	2
U	L	$\theta_{1,2}$	$\theta_{2,2}$	2	1
U	R	$\theta_{1,1}$	$\theta_{2,1}$	0	2
U	R	$\theta_{1,1}$	$\theta_{2,2}$	0	3
U	R	$\theta_{1,2}$	$\theta_{2,1}$	0	0
U	R	$\theta_{1,2}$	$\theta_{2,2}$	0	0
D	L	$\theta_{1,1}$	$\theta_{2,1}$	0	2
D	L	$\theta_{1,1}$	$\theta_{2,2}$	3	0
D	L	$\theta_{1,2}$	$\theta_{2,1}$	0	0
D	L	$\theta_{1,2}$	$\theta_{2,2}$	0	0
D	R	$\theta_{1,1}$	$\theta_{2,1}$	2	0
D	R	$\theta_{1,1}$	$\theta_{2,2}$	1	1
D	R	$\theta_{1,2}$	$\theta_{2,1}$	1	1
D	R	$\theta_{1,2}$	$\theta_{2,2}$	1	2

Strategies

- In principle, we could use any of the three definitions of a Bayesian game
    - The book uses the 3rd one (epistemic types)
- Strategies are similar to what we had in imperfect-information games
    - A **pure strategy** for player *i* maps each of *i*'s types to an action
        - what *i* would play if *i* had that type
    - A mixed strategy s_i is a probability distribution over pure strategies
        - $s_i(a_i \mid \theta_j)$ = Pr[*i* plays action a_j | *i*'s type is θ_j]

- Three kinds of expected utility: *ex post*, *ex interim*, and *ex ante*
    - Depend on what we know about the players' types
- We mainly consider ***ex ante*** in this class (which is simpler than others)
- A ***type profile*** is a vector $\boldsymbol{\theta} = (\theta_1, \theta_2, \ldots, \theta_n)$ of types, one for each agent
    - $\boldsymbol{\theta}_{-i} = (\theta_1, \theta_2, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_n)$
    - $\boldsymbol{\theta} = (\theta_i, \boldsymbol{\theta}_{-i})$

Expected Utility

- Three different kinds of expected utility, depending on what we know about the agents' types
- If we know every agent's type (i.e., the type profile $\boldsymbol{\theta}$)
  - agent *i*'s ***ex post*** expected utility:

$$\boxed{EU_i(\mathbf{s},\boldsymbol{\theta})} = \sum_{\mathbf{a}} \Pr[\mathbf{a} \mid \mathbf{s},\boldsymbol{\theta}]\; u_i(\mathbf{a},\boldsymbol{\theta}) = \sum_{\mathbf{a}} \left(\prod_{j\in N} s_j\left(a_j \mid \theta_j\right)\right) u_i\left(\mathbf{a},\boldsymbol{\theta}\right)$$

- If we only know the common prior
  - agent *i*'s ***ex ante*** expected utility:

$$EU_i(\mathbf{s}) = \sum_{\boldsymbol{\theta}} \Pr[\boldsymbol{\theta}]\; \boxed{EU_i(\mathbf{s},\boldsymbol{\theta})} = \sum_{\theta_i} \Pr[\theta_i]\; EU_i(\mathbf{s},\theta_i)$$

- If we know the type θ_i of one agent *i*, but not the other agents' types
  - *i*'s ***ex interim*** expected utility:

$$EU_i(\mathbf{s},\theta_i) = \sum_{\boldsymbol{\theta}_{-i}} \Pr[\boldsymbol{\theta}_{-i} \mid \theta_i]\; \boxed{EU_i\left(\mathbf{s},(\theta_i,\boldsymbol{\theta}_{-i})\right)}$$

Bayes-Nash Equilibria

- Given a strategy profile $\mathbf{s}_{-i}$, a **best response** for agent *i* is a strategy s_i such that

$$s_i \in \underset{s'_i}{\arg\max}(EU_i\,(s'_i\,, \mathbf{s}_{-i}))$$

- Above, the set notation is because more than one strategy may produce the same expected utility

- A **Bayes-Nash** equilibrium is a strategy profile **s** such that for every s_i in **s**, s_i is a best response to $\mathbf{s}_{-i}$
  - Just like the definition of a Nash equilibrium, except that we're using Bayesian-game strategies

Computing Bayes-Nash Equilibria

- The idea is to construct a payoff matrix for the entire Bayesian game, and find equilibria on that matrix

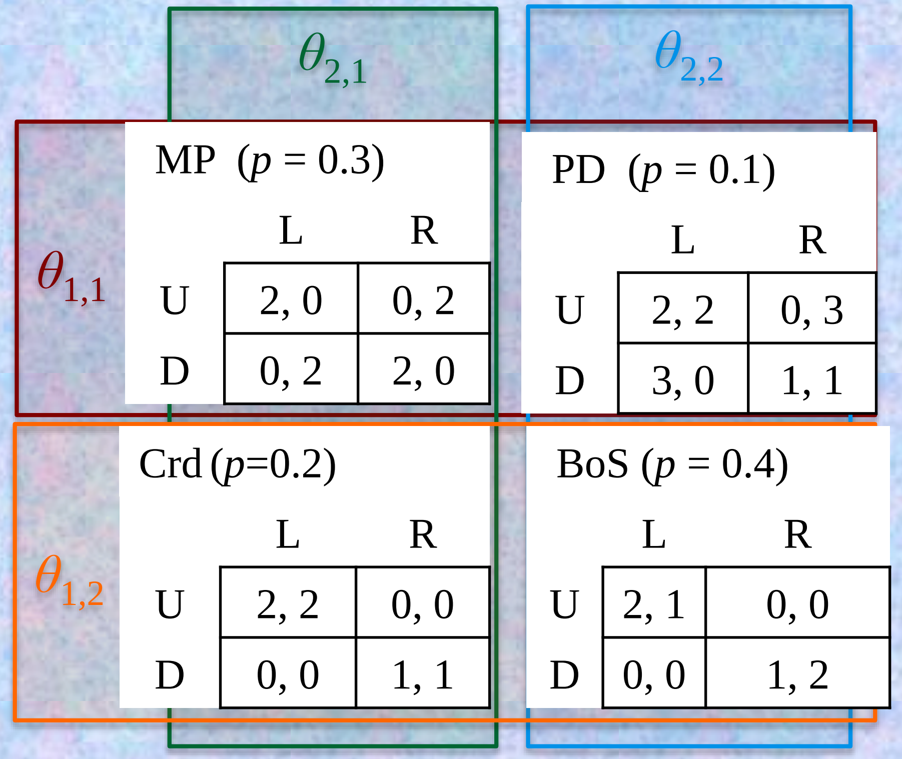

$\theta_{2,1}$ $\theta_{2,2}$

$\theta_{1,1}$

MP (*p* = 0.3)

	L	R
U	2, 0	0, 2
D	0, 2	2, 0

PD (*p* = 0.1)

	L	R
U	2, 2	0, 3
D	3, 0	1, 1

$\theta_{1,2}$

Crd (*p*=0.2)

	L	R
U	2, 2	0, 0
D	0, 0	1, 1

BoS (*p* = 0.4)

	L	R
U	2, 1	0, 0
D	0, 0	1, 2

- First, write each of the pure strategies as a list of actions, one for each type

- Agent 1's pure strategies:
  - UU: U if type $\theta_{1,1}$, U if type $\theta_{1,2}$
  - UD: U if type $\theta_{1,1}$, D if type $\theta_{1,2}$
  - DU: D if type $\theta_{1,1}$, U if type $\theta_{1,2}$
  - DD: D if type $\theta_{1,1}$, D if type $\theta_{1,2}$

- Agent 2's pure strategies:
  - LL: L if type $\theta_{2,1}$, L if type $\theta_{2,2}$
  - LR: L if type $\theta_{2,1}$, R if type $\theta_{2,2}$
  - RL: R if type $\theta_{2,1}$, L if type $\theta_{2,2}$
  - RR: R if type $\theta_{2,1}$, R if type $\theta_{2,2}$

Computing Bayes-Nash Equilibria (continued)

- Next, compute the *ex ante* expected utility for each pure-strategy profile
    - e.g., (note that θ , UU, and LL determine dots)

$$EU_2(UU,LL)=\sum_{\theta}\Pr[\theta]u_2(.,.,\theta)$$

$$=\Pr[\theta_{1,1},\theta_{2,1}]u_2(U,L,\theta_{1,1},\theta_{2,1})$$

$$+\Pr[\theta_{1,1},\theta_{2,2}]u_2(U,L,\theta_{1,1},\theta_{2,2})$$

$$+\Pr[\theta_{1,2},\theta_{2,1}]u_2(U,L,\theta_{1,2},\theta_{2,1})$$

$$+\Pr[\theta_{1,2},\theta_{2,2}]u_2(U,L,\theta_{1,2},\theta_{2,2})$$

$$=0.3(0)+0.1(2)+0.2(2)+0.4(1)$$

$$=1$$

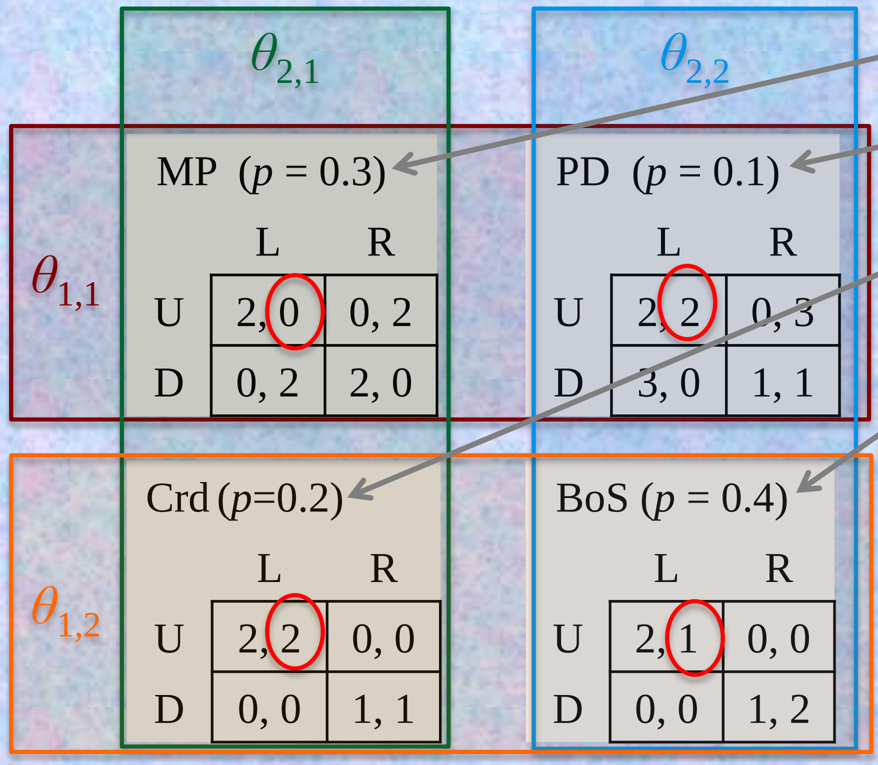

MP (*p* = 0.3)

	L	R
U	2, 0	0, 2
D	0, 2	2, 0

PD (*p* = 0.1)

	L	R
U	2, 2	0, 3
D	3, 0	1, 1

Crd (*p*=0.2)

	L	R
U	2, 2	0, 0
D	0, 0	1, 1

BoS (*p* = 0.4)

	L	R
U	2, 1	0, 0
D	0, 0	1, 2

Computing Bayes-Nash Equilibria (continued)

- Put all of the *ex ante* expected utilities into a payoff matrix
  - e.g., $EU_2(UU,LL) = 1$
- Now we can compute best responses and Nash equilibria

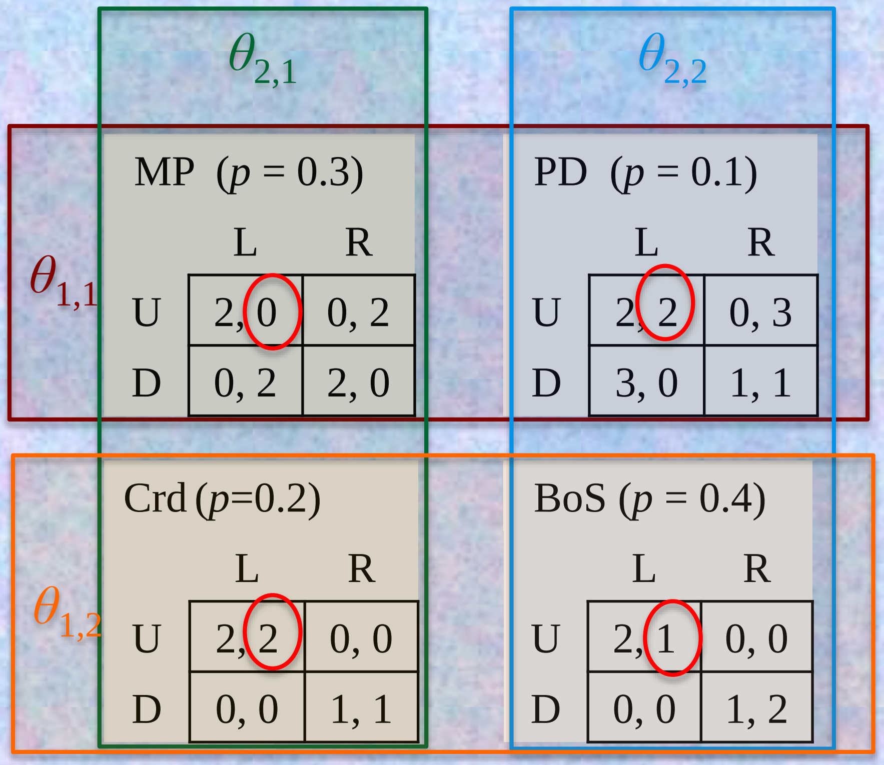

MP (p = 0.3)

	L	R
U	2, 0	0, 2
D	0, 2	2, 0

PD (p = 0.1)

	L	R
U	2, 2	0, 3
D	3, 0	1, 1

Crd (p=0.2)

	L	R
U	2, 2	0, 0
D	0, 0	1, 1

BoS (p = 0.4)

	L	R
U	2, 1	0, 0
D	0, 0	1, 2

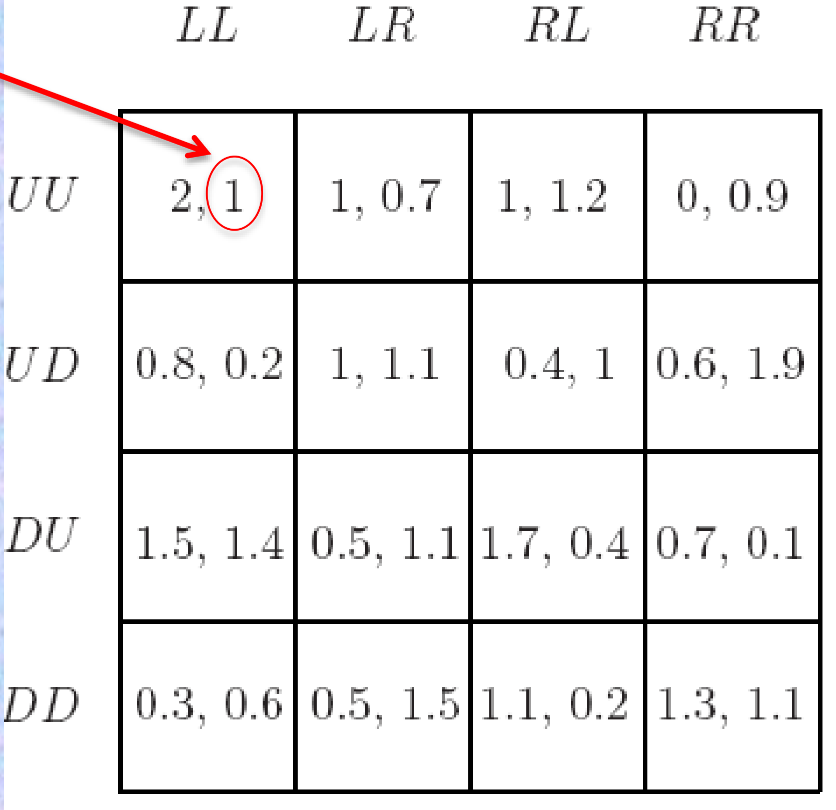

	LL	LR	RL	RR
UU	2, 1	1, 0.7	1, 1.2	0, 0.9
UD	0.8, 0.2	1, 1.1	0.4, 1	0.6, 1.9
DU	1.5, 1.4	0.5, 1.1	1.7, 0.4	0.7, 0.1
DD	0.3, 0.6	0.5, 1.5	1.1, 0.2	1.3, 1.1

Summary

- Incomplete information vs. imperfect information
- Incomplete information vs. uncertainty about payoffs
- Bayesian games (three different definitions)
  - Changing uncertainty about games into uncertainty about payoffs
  - *Ex ante*, *ex interim*, and *ex post* utilities
  - Bayes-Nash equilibria
- Bayesian-game interpretations of Bridge and Backgammon
- Base-Nash instead of Nash