# Computational Game Theory in Julia

Eric Pacuit, University of Maryland

Lecture 5

ESSLLI 2023

| a \ b | $L$ | $R$ |
|---|---|---|
| $U$ | 3,1 | 0,0 |
| $D$ | 0,0 | 1,3 |

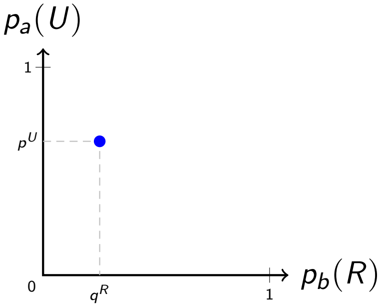

*a*'s current state of indecision

| | | b | |
|---|---|---|---|
| | | $L$ | $R$ |
| a | $U$ | 3,1 | 0,0 |
| | $D$ | 0,0 | 1,3 |

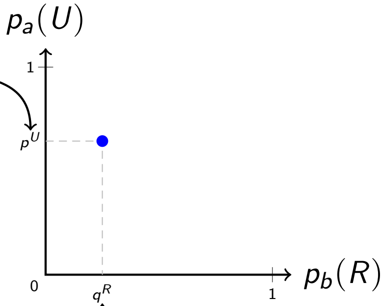


*a*'s current belief about what *b* is going to do

$b$'s current belief about what $a$ is going to do

| | | b | |
|---|---|---|---|
| | | $L$ | $R$ |
| a | $U$ | 3,1 | 0,0 |
| | $D$ | 0,0 | 1,3 |

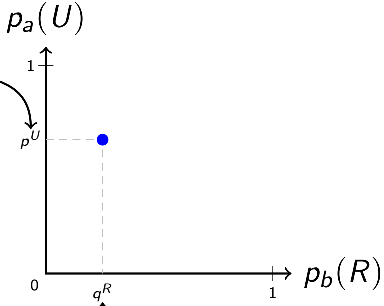


$b$'s current state of indecision

| | | b | |
|---|---|---|---|
| | | $L$ | $R$ |
| a | $U$ | 3,1 | 0,0 |
| | $D$ | 0,0 | 1,3 |

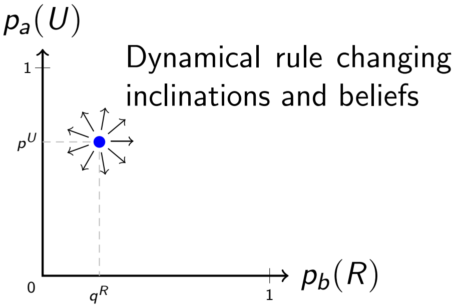

| | | Bob | |
|---|---|---|---|
| | | $L$ | $R$ |
| Ann | $U$ | 2,1 | 0,0 |
| | $D$ | 0,0 | 1,2 |

$\mathbf{P_a} = \langle 0.2, 0.8 \rangle$ and $\mathbf{P_b} = \langle 0.4, 0.6 \rangle$

$$EU(U) = 0.4 \cdot 2 + 0.6 \cdot 0 = 0.8$$
$$EU(D) = 0.4 \cdot 0 + 0.6 \cdot 1 = 0.6$$
$$EU(L) = 0.2 \cdot 1 + 0.8 \cdot 0 = 0.2$$
$$EU(R) = 0.2 \cdot 0 + 0.8 \cdot 2 = 1.6$$
$$SQ_A = 0.2 \cdot EU(U) + 0.8 \cdot EU(D) = 0.2 \cdot 0.8 + 0.8 \cdot 0.6 = 0.64$$
$$SQ_B = 0.4 \cdot EU(L) + 0.6 \cdot EU(R) = 0.4 \cdot 0.2 + 0.6 \cdot 1.6 = 1.04$$

# BoS - Nash Dynamics

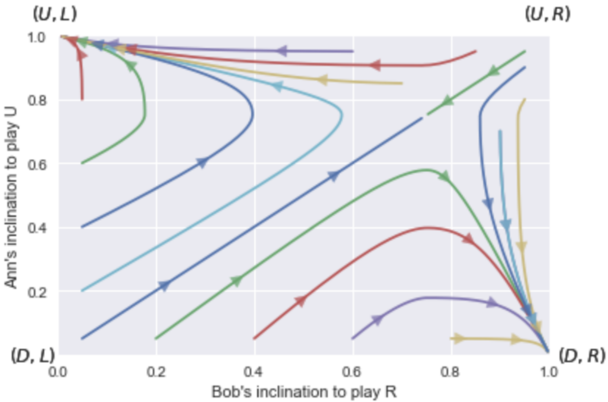

# Bayes dynamics

The **Bayes dynamics**, also called **Darwin dynamics**, transforms $I_i \in \Delta(S_i)$ into a new probability $I_i' \in \Delta(S_i)$ as follows. For each $s \in S_i$:

$$I_i'(s) = I_i(s) + \frac{1}{k} I_i(s) \frac{EU_i(s) - SQ_i}{SQ_i}.$$

where $k > 0$ is the "index of caution".

# BoS - Bayes

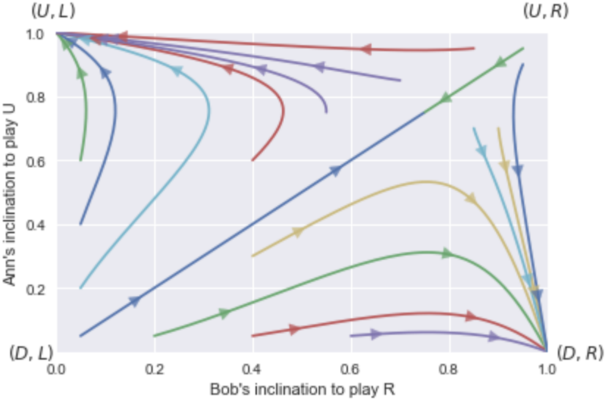

$(U, L)$
$(U, R)$
$(D, L)$
$(D, R)$
1.0
0.8
0.6
0.4
0.2
0.0
0.2
0.4
0.6
0.8
1.0
Ann's inclination to play U
Bob's inclination to play R

# Deliberational Dynamics

"There is nothing in the nature of deliberational dynamics that requires that deliberators be simpleminded, but the illustrations I have chosen...are relatively unsophisticated. These players follow their noses in the direction of the current apparent good, with no real memory of where they have been, no capability of recognizing patterns, and no sense of where they are going." (Skyrms, pg. 152)

# Backward and forward induction reasoning

Backward induction reasoning: player's ignore past behavior and reason only about their opponents' future moves.

Forward induction reasoning: player's rationalize past behavior and use it as a basis to form beliefs about their opponents' future moves.

# Backward *versus* Forward Induction

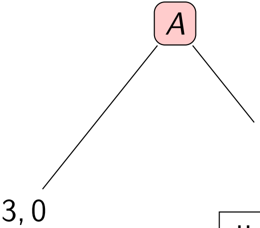


| | | Bob | | |
|---|---|---|---|---|
| | | l | c | r |
| | u | 2, 2 | 2, 1 | 0, 0 |
| Ann | d | 1, 1 | 1, 2 | 4, 0 |

A. Perea. *Backward Induction* versus *Forward Induction Reasoning*. Games, 1, pgs. 168 - 188, 2010.

# Backward *versus* Forward Induction

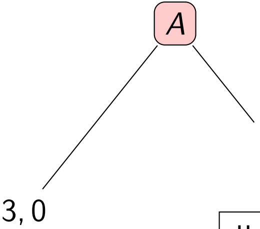


Bob

| | | l | c | r |
|---|---|---|---|---|
| | u | 2, 2 | 2, 1 | 0, 0 |
| Ann | d | 1, 1 | 1, 2 | 4, 0 |

A. Perea. *Backward Induction* versus *Forward Induction Reasoning*. Games, 1, pgs. 168 - 188, 2010.

# Backward *versus* Forward Induction

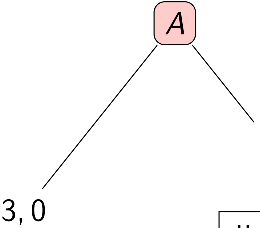


| | | Bob | | |
|---|---|---|---|---|
| | | l | c | r |
| Ann | u | 2, 2 | 2, 1 | 0, 0 |
| | d | 1, 1 | 1, 2 | 4, 0 |

A. Perea. *Backward Induction* versus *Forward Induction Reasoning*. Games, 1, pgs. 168 - 188, 2010.

# Backward *versus* Forward Induction

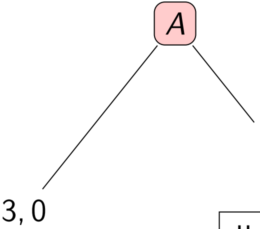


| | | Bob | | |
|---|---|---|---|---|
| | | l | c | r |
| Ann | u | 2, 2 | 2, 1 | 0, 0 |
| | d | 1, 1 | 1, 2 | 4, 0 |

A. Perea. *Backward Induction* versus *Forward Induction Reasoning*. Games, 1, pgs. 168 - 188, 2010.

# Backward *versus* Forward Induction

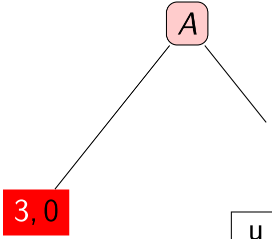


| | | Bob | | |
|---|---|---|---|---|
| | | l | c | r |
| | u | 2, 2 | 2, 1 | 0, 0 |
| Ann | d | 1, 1 | 1, 2 | 4, 0 |


A. Perea. *Backward Induction* versus *Forward Induction Reasoning*. Games, 1, pgs. 168 - 188, 2010.

# Backward *versus* Forward Induction

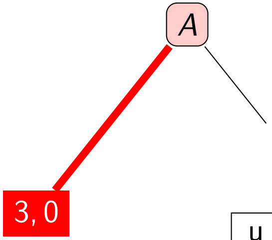


| | | Bob | | |
|---|---|---|---|---|
| | | l | c | r |
| Ann | u | 2, 2 | 2, 1 | 0, 0 |
| | d | 1, 1 | 1, 2 | 4, 0 |

A. Perea. *Backward Induction* versus *Forward Induction Reasoning*. Games, 1, pgs. 168 - 188, 2010.

# Backward *versus* Forward Induction

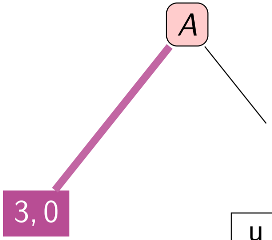


Bob

| | | l | c | r |
|---|---|---|---|---|
| Ann | u | 2, 2 | 2, 1 | 0, 0 |
| | d | 1, 1 | 1, 2 | 4, 0 |

A. Perea. *Backward Induction* versus *Forward Induction Reasoning*. Games, 1, pgs. 168 - 188, 2010.

A. Knoks and EP. *Deliberating between Backward and Forward Induction Reasoning: First Steps*. Proceedings of Theoretical Aspects of Rationality and Knowledge (TARK 2015).

A. Knoks and EP. *Deliberational dynamics in context*. Proceedings of LOFT (LOFT 2018).

# Forward induction reasoning

| | $L$ | $R$ |
|---|---|---|
| $U$ | $3, 1$ | $0, 0$ |
| $D$ | $0, 0$ | $1, 3$ |

# Forward induction reasoning

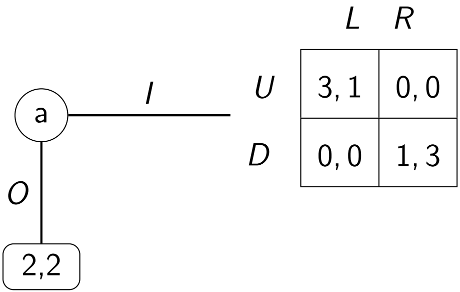

# Forward induction reasoning

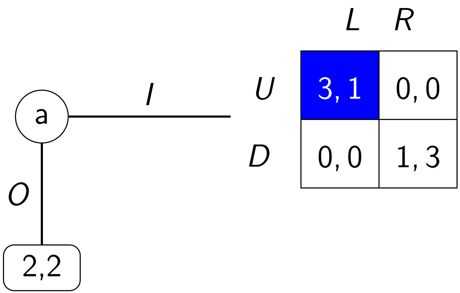

# Forward induction reasoning

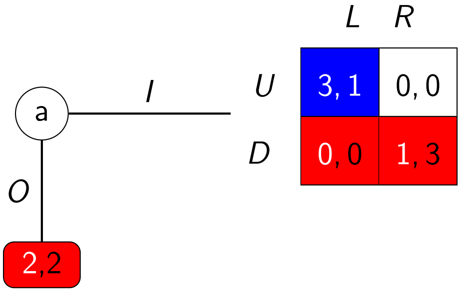

# Forward induction reasoning

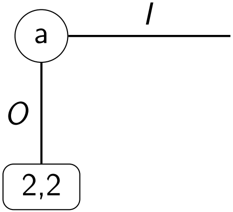


| | $L$ | $R$ |
|---|---|---|
| $U$ | $3, 1$ | $0, 0$ |
| $D$ | $0, 0$ | $1, 3$ |

# Forward induction reasoning

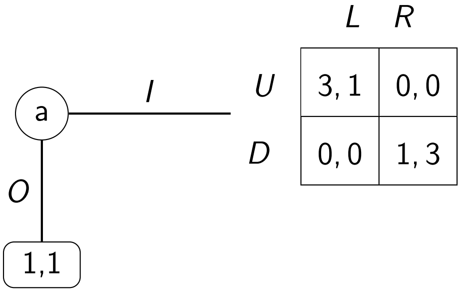

# Forward induction reasoning

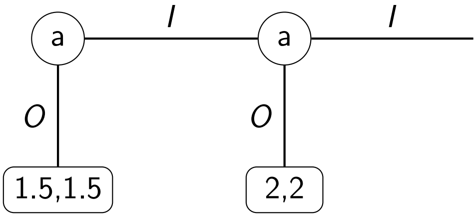


| | $L$ | $R$ |
|---|---|---|
| $U$ | 3, 1 | 0, 0 |
| $D$ | 0, 0 | 1, 3 |

# Forward induction reasoning

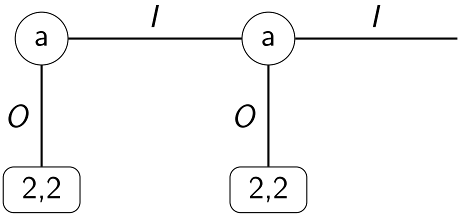


| | $L$ | $R$ |
|---|---|---|
| $U$ | 3, 1 | 0, 0 |
| $D$ | 0, 0 | 1, 3 |

# Forward induction reasoning

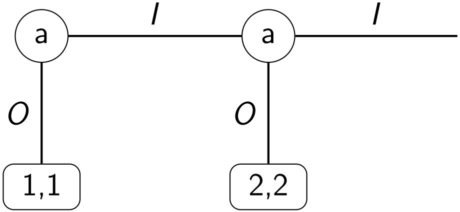


| | $L$ | $R$ |
|---|---|---|
| $U$ | 3, 1 | 0, 0 |
| $D$ | 0, 0 | 1, 3 |

# Forward induction reasoning

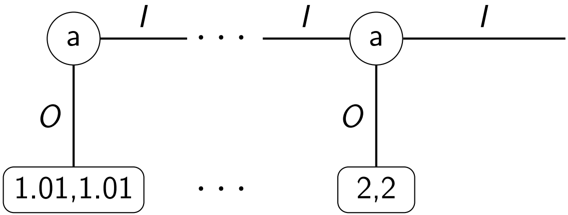


| | $L$ | $R$ |
|---|---|---|
| $U$ | 3, 1 | 0, 0 |
| $D$ | 0, 0 | 1, 3 |

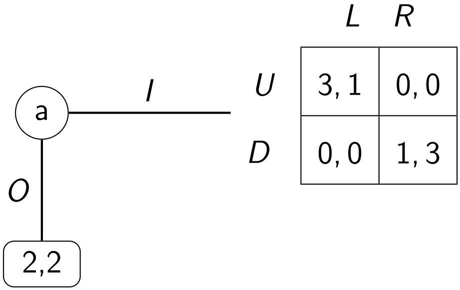


| | $L$ | $R$ |
|---|---|---|
| $O$ | 2, 2 | 2, 2 |
| $IU$ | 3, 1 | 0, 0 |
| $ID$ | 0, 0 | 1, 3 |

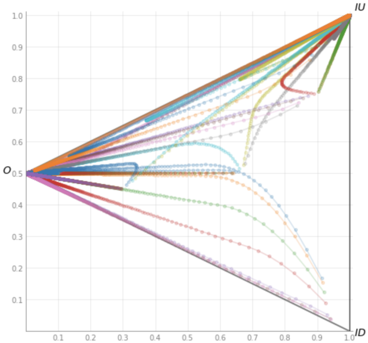


Nash deliberators

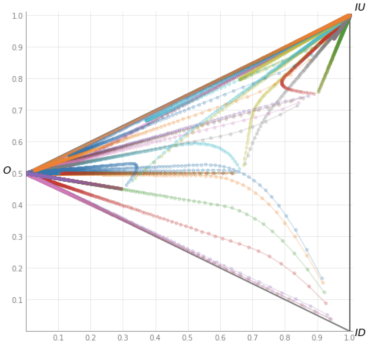


Bayes deliberators

Note that both Bayes and Nash deliberators converge on $(IU, L)$ and $(O, R)$.

Note that both Bayes and Nash deliberators converge on $(IU, L)$ and $(O, R)$.

If Bob is a backward induction reasoner, then he ignores Ann's initial move as he deliberates between $L$ and $R$.

Note that both Bayes and Nash deliberators converge on $(IU, L)$ and $(O, R)$.

If Bob is a backward induction reasoner, then he ignores Ann's initial move as he deliberates between $L$ and $R$.

On the other hand, if Bob is a forward induction reasoner, then, during deliberation, he should assign probability 0 to Ann choosing $I$ then $D$ (since it is strictly dominated by choosing $O$). This belief about Ann's choice does not change during deliberation.

Ann may be uncertain about whether Bob is a forward induction reasoner or a backward induction reasoner (and Bob may assume this about Ann).

Ann may be uncertain about whether Bob is a forward induction reasoner or a backward induction reasoner (and Bob may assume this about Ann).

Ann's belief at time $t+1$ that Bob will play strategy $s$ is:

$$P_a^{t+1}(s) = w_{a,t} I_{b,BI}^t(s) + (1 - w_{a,t}) I_{b,FI}^t(s)$$

where $I_{b,BI}^t$ is Bob's inclinations as a backward induction reasoner at step $t$, $I_{b,FI}^t$ is Bob's inclinations as a backward induction reasoner at step $t$, and $w_{a,t}$ is the weight Ann assigns to Bob being a backward induction reasoner at step $t$.

By varying the weights, we can represent players that are not fully backward induction reasoners or fully forward induction reasoners.

By varying the weights, we can represent players that are not fully backward induction reasoners or fully forward induction reasoners.

1. The weights $w_{a,t}$ are Ann's prior beliefs about whether Bob is a backward induction reasoner or a forward induction reasoner.

2. The weights $w_{a,t}$ are defined by the **context** of the target game.

3. The weights $w_{a,t}$ depend both on the context of the game and the probability that the players are prone to *trembling-hand mistakes*.

# Pure Coordination Game

| | | Bob | |
|---|---|---|---|
| | | $L$ | $R$ |
| Ann | $U$ | 1,1 | 0,0 |
| | $D$ | 0,0 | 1,1 |

# Coordination - Nash deliberators

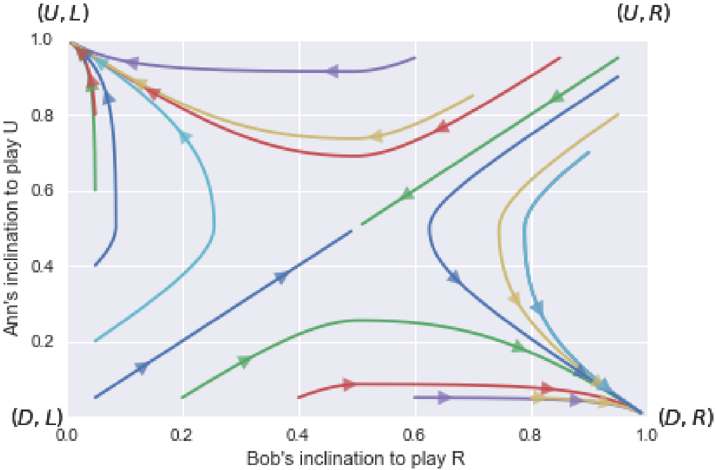

# Imprecise Priors

It is assumed that the players precise states of indecision are common knowledge at the onset of deliberation.

*Imprecise Prior*: Each players prior is a convex set of probability measures over her actions space.

Restrict attention to games with two players where each players has two strategies.

A precise state of indecision for the row player is

$$\mathbf{P}_{row}(t) = \langle p^1_{row}(t), \ldots, p^n_{row}(t) \rangle$$

where $p^j_{row}(t)$ is the probability that row assigns to her strategy $j$ at time $t$.

An imprecise state of indecision has $p^1_{row} = [lp, up]$ and $p^2_{row} = [1 - up, 1 - lp]$. For example, if $p^1_{row} = [0.6, 0.7]$, then $p^2_{row} = [0.3, 0.4]$.

Row (Col) has an expected utility for each probability measure in Col's (Row's) interval. Row (Col) need only compute expected utilities with respect to the endpoints of column's interval.

| | | col | |
|---|---|---|---|
| | | $L$ | $R$ |
| row | $U$ | 0,0 | 1,1 |
| | $D$ | 1,1 | 0,0 |

$$p_{row}^{U}(0) = [0.6, 0.8] \text{ and } p_{col}^{L}(0) = [0.6, 0.9]$$

| | | col | |
|---|---|---|---|
| | | $L$ | $R$ |
| row | $U$ | 0,0 | 1,1 |
| | $D$ | 1,1 | 0,0 |

$$p^U_{row}(0) = [0.6, 0.8] \text{ and } p^L_{col}(0) = [0.6, 0.9]$$

$$EU_{row}(U, 0) = [0.1, 0.4]$$
$$EU_{row}(D, 0) = [0.6, 0.9]$$

| | | col | |
|---|---|---|---|
| | | $L$ | $R$ |
| row | $U$ | 0,0 | 1,1 |
| | $D$ | 1,1 | 0,0 |

$$p^U_{row}(0) = [0.6, 0.8] \text{ and } p^L_{col}(0) = [0.6, 0.9]$$

$$EU_{row}(U, 0) = [0.1, 0.4]$$
$$EU_{row}(D, 0) = [0.6, 0.9]$$

How should you calculate $\mathbf{P}_{row}(1)$ and $\mathbf{P}_{col}(1)$?

1. $p_{row}^{U} = 0.6,\ p_{col}^{L} = 0.6$: $SQ_{row} = 0.30,\ Cov_{row}(U) = 0,\ Cov_{row}(D) = 0.30.$
   $p_{row}^{U}(1) = \frac{0.6+0}{1+0.3} = 0.4615$

2. $p_{row}^{U} = 0.6,\ p_{col}^{L} = 0.9$: $SQ_{row} = 0.40,\ Cov_{row}(U) = 0,\ Cov_{row}(D) = 0.20.$
   $p_{row}^{U}(1) = \frac{0.6+0}{1+0.4} = 0.4286$

3. $p_{row}^{U} = 0.8,\ p_{col}^{L} = 0.6$: $SQ_{row} = 0.32,\ Cov_{row}(U) = 0,\ Cov_{row}(D) = 0.28.$
   $p_{row}^{U}(1) = \frac{0.8+0}{1+0.32} = 0.6061$

4. $p_{row}^{U} = 0.8,\ p_{col}^{L} = 0.9$: $SQ_{row} = 0.20,\ Cov_{row}(U) = 0,\ Cov_{row}(D) = 0.7.$
   $p_{row}^{U}(1) = \frac{0.8+0}{1+0.7} = 0.4706$

$$p_{row}^{U} = [0.4286, 0.6061]$$

- The area of a rectangle of indecision need not be preserved by deliberational dynamics

- The area of a rectangle of indecision need not be preserved by deliberational dynamics
- For example, players may start out with imprecise prior probabilities and deliberation results in point probabilities (E.g., Figure 3.4, 3.5 on pgs. 68, 69)

- The area of a rectangle of indecision need not be preserved by deliberational dynamics
- For example, players may start out with imprecise prior probabilities and deliberation results in point probabilities (E.g., Figure 3.4, 3.5 on pgs. 68, 69)
- The pure mixed strategy in the game of Chicken is not stable for precise probabilities. Starting from $[0.51, 0.49]$, $[0.51, 0.49]$, the orbit explodes to a state of mutual total bewilderment.

- The area of a rectangle of indecision need not be preserved by deliberational dynamics
- For example, players may start out with imprecise prior probabilities and deliberation results in point probabilities (E.g., Figure 3.4, 3.5 on pgs. 68, 69)
- The pure mixed strategy in the game of Chicken is not stable for precise probabilities. Starting from $[0.51, 0.49], [0.51, 0.49]$, the orbit explodes to a state of mutual total bewilderment.
- In matching pennies, the mixed strategy is strongly stable. However, starting from $[0.51, 0.49], [0.51, 0.49]$, the imprecision explodes to cover the whole space (see Figure 3.8, pg. 72)

- The area of a rectangle of indecision need not be preserved by deliberational dynamics
- For example, players may start out with imprecise prior probabilities and deliberation results in point probabilities (E.g., Figure 3.4, 3.5 on pgs. 68, 69)
- The pure mixed strategy in the game of Chicken is not stable for precise probabilities. Starting from $[0.51, 0.49], [0.51, 0.49]$, the orbit explodes to a state of mutual total bewilderment.
- In matching pennies, the mixed strategy is strongly stable. However, starting from $[0.51, 0.49], [0.51, 0.49]$, the imprecision explodes to cover the whole space (see Figure 3.8, pg. 72)
- When analyzed in terms of precise priors, the pure coordination game and Chicken were both seen to be situations in which coordination could arise spontaneously. This is not true when starting with imprecise probabilities.

# Coordination

| Ann | | Bob $L$ | Bob $R$ |
|---|---|---|---|
| | $U$ | 1,1 | 0,0 |
| | $D$ | 0,0 | 1,1 |

1. How can convention without communication be sustained? (Lewis)
2. How can convention without communication be generated?

Ann and Bob each have predeliberational probabilities. They can be anything at all. These probabilities are made common knowledge at the start of deliberation.

Ann and Bob each have predeliberational probabilities. They can be anything at all. These probabilities are made common knowledge at the start of deliberation.

You—the philosopher—have some probability distribution over the space of Ann and Bob's initial probabilities. Then you should believe with probability one that the deliberators will converge to one of the pure Nash equilibria.

Ann and Bob each have predeliberational probabilities. They can be anything at all. These probabilities are made common knowledge at the start of deliberation.

You—the philosopher—have some probability distribution over the space of Ann and Bob's initial probabilities. Then you should believe with probability one that the deliberators will converge to one of the pure Nash equilibria.

Precedent and other forms of initial salience may influence the deliberators' initial probabilities, and thus may play a role in determining *which* equilibrium is selected.

# Correlation

Players can improve their expected value by correlating their choices on an "outside signal".

# Correlated Strategies

| | | Player 2 | |
|---|---|---|---|
| | | $l$ | $r$ |
| Player 1 | $u$ | 2, 1 | 0, 0 |
| | $d$ | 0,0 | 1,2 |

- Three Nash equilibria:
  - $(u, l)$: the payoff is $(2, 1)$
  - $(d, r)$: the payoff is $(1, 2)$
  - Mixed Nash Equilibrium: $([\frac{2}{3} : u, \frac{1}{3} : d], [\frac{1}{3} : l, \frac{2}{3} : r])$: the payoff is $(\frac{2}{3}, \frac{2}{3})$

# Correlated Strategies

| | | Player 2 | |
|---|---|---|---|
| | | $l$ | $r$ |
| Player 1 | $u$ | 2, 1 | 0, 0 |
| | $d$ | 0,0 | 1,2 |

- Mixed Nash Equilibrium: $([\frac{2}{3} : u, \frac{1}{3} : d], [\frac{1}{3} : l, \frac{2}{3} : r])$: the payoff is $(\frac{2}{3}, \frac{2}{3})$
- Mixed Strategies: Each player conducts a private, independent lottery to choose their strategy.

# Correlated Strategies

| | | Player 2 | |
|---|---|---|---|
| | | $l$ | $r$ |
| Player 1 | $u$ | 2, 1 | 0, 0 |
| | $d$ | 0,0 | 1,2 |

| | $l$ | $r$ |
|---|---|---|
| $u$ | 0.5 | 0 |
| $d$ | 0 | 0.5 |

- Mixed Nash Equilibrium: $([\frac{2}{3} : u, \frac{1}{3} : d], [\frac{1}{3} : l, \frac{2}{3} : r])$: the payoff is $(\frac{2}{3}, \frac{2}{3})$
- Mixed Strategies: Each player conducts a private, independent lottery to choose their strategy.
- Conduct a *public* lottery: flip a fair coin and follow the strategy $(H \Rightarrow (u, l),\ T \Rightarrow (d, r))$. The expected payoff is $(1.5, 1.5)$.

Two extremes:

1. Completely private, independent lotteries
2. A single, completely public lottery

Two extremes:

1. Completely private, independent lotteries
2. A single, completely public lottery

What about: a public lottery, but reveal only partial information about the outcome to each of the players?

| | | Player 2 | |
|---|---|---|---|
| | | $c$ | $d$ |
| Player 1 | $c$ | 6, 6 | 2, 7 |
| | $d$ | 7, 2 | 0, 0 |

| | $c$ | $d$ |
|---|---|---|
| $c$ | 1/3 | 1/3 |
| $d$ | 1/3 | 0 |

- Three Nash equilibria:
  - $(c, d)$: the payoff is $(2, 7)$; $(d, c)$: the payoff is $(7, 2)$
  - $([\frac{2}{3} : c, \frac{1}{3} : d], [\frac{2}{3} : c, \frac{1}{3} : d])$: the payoff is $(4\frac{2}{3}, 4\frac{2}{3})$

| | | Player 2 | |
|---|---|---|---|
| | | $c$ | $d$ |
| Player 1 | $c$ | 6, 6 | 2, 7 |
| | $d$ | 7, 2 | 0, 0 |

| | $c$ | $d$ |
|---|---|---|
| $c$ | 1/3 | 1/3 |
| $d$ | 1/3 | 0 |

- Three Nash equilibria:
  - $(c, d)$: the payoff is $(2, 7)$; $(d, c)$: the payoff is $(7, 2)$
  - $([\frac{2}{3} : c, \frac{1}{3} : d], [\frac{2}{3} : c, \frac{1}{3} : d])$: the payoff is $(4\frac{2}{3}, 4\frac{2}{3})$
- After conducting the lottery, an outside observer provides Ann with a recommendation to play the first component of the profile that was chosen, and Bob the second component.

| | | Player 2 | |
|---|---|---|---|
| | | $c$ | $d$ |
| Player 1 | $c$ | 6, 6 | 2, 7 |
| | $d$ | 7, 2 | 0, 0 |

| | $c$ | $d$ |
|---|---|---|
| $c$ | 1/3 | 1/3 |
| $d$ | 1/3 | 0 |

- Three Nash equilibria:
  - $(c, d)$: the payoff is $(2, 7)$; $(d, c)$: the payoff is $(7, 2)$
  - $([\frac{2}{3} : c, \frac{1}{3} : d], [\frac{2}{3} : c, \frac{1}{3} : d])$: the payoff is $(4\frac{2}{3}, 4\frac{2}{3})$
- After conducting the lottery, an outside observer provides Ann with a recommendation to play the first component of the profile that was chosen, and Bob the second component.
- The expected payoff is $\frac{1}{3}(6, 6) + \frac{1}{3}(2, 7) + \frac{1}{3}(7, 2) = (5, 5)$

# Correlated Equilibrium

Let $G = \langle N, (S_i)_{i\in N}, (u_i)_{i\in N}\rangle$ be a game.

A correlated strategy $n$-tuple in $G$ is a function from a finite probability space $\Gamma$ into $S = S_1 \times \cdots \times S_n$. That is, $f$ is a *random variable* whose values are $n$-tupels of actions.

Chance (according to the probability space $\Gamma$) chooses an element $\gamma \in \Gamma$, then each player is recommended to take action $f_i(\gamma)$.

**Correlated Equilibrium**: A correlated equilibrium in $G$ is a correlated strategy $n$-tuple $f$ such that

$$EU_i(f) \geqslant EU_i(g_i, f_{-i})$$

Imagine an outside observer, who does not know what the players' initial probabilities for the possible actions will be, but rather has his own probability measure over the possible initial states of indecision of the system.

With respect to this probability, the players are at a correlated equilibrium.

Imagine an outside observer, who does not know what the players' initial probabilities for the possible actions will be, but rather has his own probability measure over the possible initial states of indecision of the system.

With respect to this probability, the players are at a correlated equilibrium.

"This correlated equilibrium is a general result of the players' common knowledge and Bayesian dynamic deliberations." (Skyrms, pg. 60)

The same result may be obtained without the outside observer if prior to deliberation the players themselves share the role of the outside observer.

Peter Vanderschraaf and Brian Skyrms. *Learning to take turns*. Erkenntnis 59, pp. 311-348, 2003.

# Impure Coordination

| | | Bob $S_1$ | Bob $S_2$ |
|---|---|---|---|
| Ann | $S_1$ | 1, 1 | 2, 3 |
| Ann | $S_2$ | 3, 2 | 0, 0 |

$(S_1, S_2)$ and $(S_2, S_1)$, and $([S_1 : 0.5, S_2 : 0.5], [S_1 : 0.5, S_2 : 0.5])$
are the Nash equilibria

# Fictitious Play (Brown, 1951)

Fictitious Play: A process by which players who repeatedly play a fixed game, the base game, update their beliefs about each other according to what they actually observe (so their beliefs are the frequencies of observed actions).

# Fictitious Play (Brown, 1951)

Fictitious Play: A process by which players who repeatedly play a fixed game, the base game, update their beliefs about each other according to what they actually observe (so their beliefs are the frequencies of observed actions).

The players are myopically Bayesian rational in that at each play in the sequence, each follows a strategy that maximizes her expected payoff in the base game given her current beliefs.

# Fictitious Play (Brown, 1951)

Fictitious Play: A process by which players who repeatedly play a fixed game, the base game, update their beliefs about each other according to what they actually observe (so their beliefs are the frequencies of observed actions).

The players are myopically Bayesian rational in that at each play in the sequence, each follows a strategy that maximizes her expected payoff in the base game given her current beliefs.

In traditional fictitious play, players enter into the sequence with initial beliefs regarding the strategies the others will follow, and as the sequence progresses they modify their beliefs recursively as a function of the frequency of strategies they follow.

# Fictitious Play (Brown, 1951)

Fictitious Play: A process by which players who repeatedly play a fixed game, the base game, update their beliefs about each other according to what they actually observe (so their beliefs are the frequencies of observed actions).

The players are myopically Bayesian rational in that at each play in the sequence, each follows a strategy that maximizes her expected payoff in the base game given her current beliefs.

In traditional fictitious play, players enter into the sequence with initial beliefs regarding the strategies the others will follow, and as the sequence progresses they modify their beliefs recursively as a function of the frequency of strategies they follow.

There are large interesting classes of games for which fictitious play converges to Nash equilibrium and there are other interesting classes of games for which it does not converge at all.

# Fictitious Play in the Impure Coordination Game

Traditional fictitious play applied to the Impure Coordination game converges to one of three distributions: the distributions that define the strict equilibria $(S_1, S_2)$ and $(S_2, S_1)$, and the and the mixed equilibrium where each player follows $S_1$ and $S_2$ with probability 0.5 each.

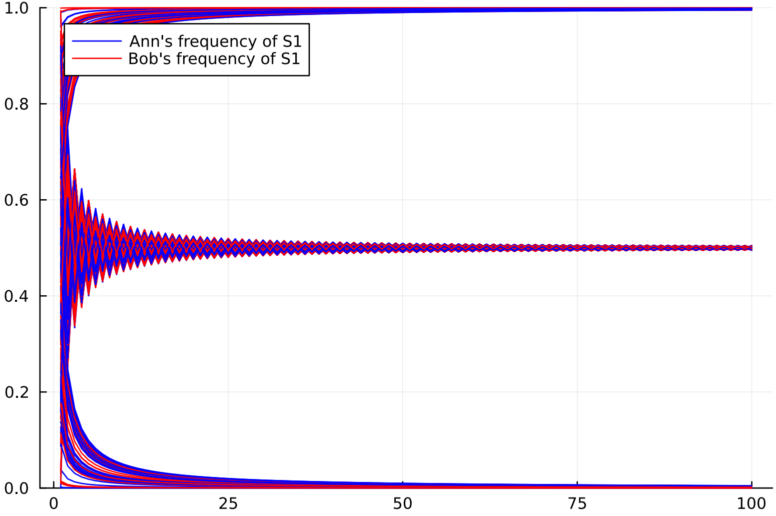
Ann's frequency of S1
Bob's frequency of S1
1.0
0.8
0.6
0.4
0.2
0.0
0
25
50
75
100

But if the players' parameters are such that fictitious play converges to the mixed Nash equilibrium, then their actual sequence of plays oscillates between $(S_1, S_1)$ and $(S_2, S_2)$, so they miscoordinate on every play!

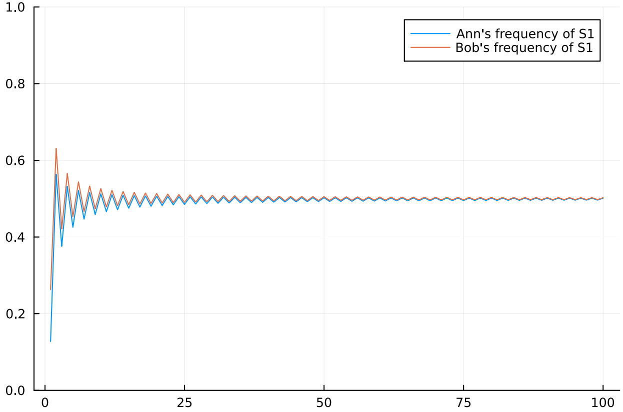

If rational agents start following a pattern of miscoordination, would they not notice this and try to break out of this pattern?

If rational agents start following a pattern of miscoordination, would they not notice this and try to break out of this pattern?

Indeed, if a pair of agents like Jan and Jill are paired together to play the Impure Coordination game repeatedly, one might expect them to settle into a pattern where they coordinate, not one where they miscoordinate.

If rational agents start following a pattern of miscoordination, would they not notice this and try to break out of this pattern?

Indeed, if a pair of agents like Jan and Jill are paired together to play the Impure Coordination game repeatedly, one might expect them to settle into a pattern where they coordinate, not one where they miscoordinate.

Specifically, one might expect them to take turns between the strict equilibria $(S_1, S_2)$ and $(S_2, S_1)$.

Move from the traditional fictitious play model, where players recursively modify probabilities that their opponents follow given strategies, to a Markov chain model, where players recursively modify the probabilities that they make a *transition* from one state of the game to another. Then we could base our prediction of the next play on the current state and players' transition probabilities.

Consider a Markov chain with $M$ possible states.

$X_{ij}^{T+1}$ is the probability that there is a transition for state $i$ to state $j$ at time $T+1$.

This is determined as follows:

$$X_{ij}^{T+1} = \frac{n_{ij} + \theta_{ij}}{T_i + \sum_{k=1}^{M} \theta_{ik}}$$

- $\theta_{ij} > 0$ is the prior weight of a transition from $i$ to $j$
- $n_{ij}$ is the observed number of transitions from $i$ to $j$
- $T_i$ is the total number of transitions from $i$.

# Impure Coordination

| | | Bob $S_1$ | Bob $S_2$ |
|---|---|---|---|
| Ann | $S_1$ | 1, 1 | 2, 3 |
| Ann | $S_2$ | 3, 2 | 0, 0 |

$s_{ij}$ denotes strategy profile $(S_i, S_j)$ for $i, j \in \{1, 2\}$.

# Impure Coordination

| | | Bob $S_1$ | Bob $S_2$ |
|---|---|---|---|
| Ann | $S_1$ | 1, 1 | 2, 3 |
| Ann | $S_2$ | 3, 2 | 0, 0 |

$s_{ij}$ denotes strategy profile $(S_i, S_j)$ for $i, j \in \{1, 2\}$.

$s_{ij}^T$ means $(S_i, S_j)$ is played in round $T$.

# Impure Coordination

| | | Bob $S_1$ | Bob $S_2$ |
|---|---|---|---|
| Ann | $S_1$ | 1, 1 | 2, 3 |
| Ann | $S_2$ | 3, 2 | 0, 0 |

$s_{ij}$ denotes strategy profile $(S_i, S_j)$ for $i, j \in \{1, 2\}$.

$s_{ij}^T$ means $(S_i, S_j)$ is played in round $T$.

$s_{*j}$ is the strategies that range holding $j$ fixed—e.g., $s_{*2} = \{(S_1, S_2), (S_2, S_2)\}$. Similarly for $s_{j*}$.

# Impure Coordination

| | | Bob $S_1$ | Bob $S_2$ |
|---|---|---|---|
| Ann | $S_1$ | 1, 1 | 2, 3 |
| Ann | $S_2$ | 3, 2 | 0, 0 |

$s_{ij}$ denotes strategy profile $(S_i, S_j)$ for $i, j \in \{1, 2\}$.

$s_{ij}^T$ means $(S_i, S_j)$ is played in round $T$.

$s_{*j}$ is the strategies that range holding $j$ fixed—e.g., $s_{*2} = \{(S_1, S_2), (S_2, S_2)\}$. Similarly for $s_{j*}$.

$\mu_i^T(\cdot)$ is Player $i$'s subjective probability at time $T$

# Transition Probabilities

$$\alpha_{ij,kl} = \mu_1^{T+1}(s_{kl}^{T+1} \mid s_{ij}^{T})$$

$$\beta_{ij,kl} = \mu_2^{T+1}(s_{kl}^{T+1} \mid s_{ij}^{T})$$

# Expected Utilities

$$E_1^{T+1}(u_1(s_{1*}) \mid s_{ij}^T) = u_1(s_{11})(\alpha_{ij,11}^{T+1} + \alpha_{ij,21}^{T+1}) + u_1(s_{12})(\alpha_{ij,12}^{T+1} + \alpha_{ij,22}^{T+1})$$
$$E_1^{T+1}(u_1(s_{2*}) \mid s_{ij}^T) = u_1(s_{21})(\alpha_{ij,11}^{T+1} + \alpha_{ij,21}^{T+1}) + u_1(s_{22})(\alpha_{ij,12}^{T+1} + \alpha_{ij,22}^{T+1})$$

$$E_2^{T+1}(u_2(s_{*1}) \mid s_{ij}^T) = u_2(s_{11})(\beta_{ij,11}^{T+1} + \beta_{ij,12}^{T+1}) + u_2(s_{21})(\beta_{ij,21}^{T+1} + \beta_{ij,22}^{T+1})$$
$$E_2^{T+1}(u_2(s_{*2}) \mid s_{ij}^T) = u_2(s_{12})(\beta_{ij,11}^{T+1} + \beta_{ij,12}^{T+1}) + u_2(s_{22})(\beta_{ij,21}^{T+1} + \beta_{ij,22}^{T+1})$$

$$\boxed{E_1^{T+1}(u_1(s_{1*}) \mid s_{ij}^T)} = u_1(s_{11})(\alpha_{ij,11}^{T+1} + \alpha_{ij,21}^{T+1}) + u_1(s_{12})(\alpha_{ij,12}^{T+1} + \alpha_{ij,22}^{T+1})$$

Player 1's expected utility for playing $S_1$ given that the previous state of play of $(S_1, S_2)$.

$$E_1^{T+1}(u_1(s_{1*}) \mid s_{ij}^{T}) = \boxed{u_1(s_{11})}(\alpha_{ij,11}^{T+1} + \alpha_{ij,21}^{T+1}) + \boxed{u_1(s_{12})}(\alpha_{ij,12}^{T+1} + \alpha_{ij,22}^{T+1})$$

Player 1's utility for the outcome $(S_1, S_2)$ and player 1's utility for the outcome $(S_1, S_2)$.

$$E_1^{T+1}(u_1(s_{1*}) \mid s_{ij}^T) = u_1(s_{11})\boxed{(\alpha_{ij,11}^{T+1} + \alpha_{ij,21}^{T+1})} + u_1(s_{12})(\alpha_{ij,12}^{T+1} + \alpha_{ij,22}^{T+1})$$

The probability of transitioning from a state $(S_i, S_j)$ to a state where player 2 plays $S_1$.

$$E_1^{T+1}(u_1(s_{1*}) \mid s_{ij}^T) = u_1(s_{11})(\alpha_{ij,11}^{T+1} + \alpha_{ij,21}^{T+1}) + u_1(s_{12})\boxed{(\alpha_{ij,12}^{T+1} + \alpha_{ij,22}^{T+1})}$$

The probability of transitioning from a state $(S_i, S_j)$ to a state where player 2 plays $S_2$.

# Expected Utilities

$$E_1^{T+1}(u_1(s_{1*}) \mid s_{ij}^T) = u_1(s_{11})(\alpha_{ij,11}^{T+1} + \alpha_{ij,21}^{T+1}) + u_1(s_{12})(\alpha_{ij,12}^{T+1} + \alpha_{ij,22}^{T+1})$$
$$E_1^{T+1}(u_1(s_{2*}) \mid s_{ij}^T) = u_1(s_{21})(\alpha_{ij,11}^{T+1} + \alpha_{ij,21}^{T+1}) + u_1(s_{22})(\alpha_{ij,12}^{T+1} + \alpha_{ij,22}^{T+1})$$

$$E_2^{T+1}(u_2(s_{*1}) \mid s_{ij}^T) = u_2(s_{11})(\beta_{ij,11}^{T+1} + \beta_{ij,12}^{T+1}) + u_2(s_{21})(\beta_{ij,21}^{T+1} + \beta_{ij,22}^{T+1})$$
$$E_2^{T+1}(u_2(s_{*2}) \mid s_{ij}^T) = u_2(s_{12})(\beta_{ij,11}^{T+1} + \beta_{ij,12}^{T+1}) + u_2(s_{22})(\beta_{ij,21}^{T+1} + \beta_{ij,22}^{T+1})$$

|  |  | Bob |  |
| --- | --- | --- | --- |
|  |  | $S_1$ | $S_2$ |
| Ann | $S_1$ | 1, 1 | 2, 3 |
|  | $S_2$ | 3, 2 | 0, 0 |

$$\alpha_{12,21}^{T+1} = \alpha_{21,12}^{T+1} = 1 \qquad \beta_{12,21}^{T+1} = \beta_{21,12}^{T+1} = 1$$

| | | Bob $S_1$ | Bob $S_2$ |
|---|---|---|---|
| Ann | $S_1$ | 1, 1 | 2, 3 |
| Ann | $S_2$ | 3, 2 | 0, 0 |

$$\alpha_{12,21}^{T+1} = \alpha_{21,12}^{T+1} = 1 \qquad \beta_{12,21}^{T+1} = \beta_{21,12}^{T+1} = 1$$

$$E_1^{T+1}(u_1(s_{1*}) \mid s_{12}^T) = 1(0+1) + 2(0+0) = 1$$
$$E_1^{T+1}(u_1(s_{2*}) \mid s_{12}^T) = 3(0+1) + 0(0+0) = 3$$

| | | Bob | |
|---|---|---|---|
| | | $S_1$ | $S_2$ |
| Ann | $S_1$ | 1, 1 | 2, 3 |
| | $S_2$ | 3, 2 | 0, 0 |

$$\alpha_{12,21}^{T+1} = \alpha_{21,12}^{T+1} = 1 \qquad \beta_{12,21}^{T+1} = \beta_{21,12}^{T+1} = 1$$

$$E_1^{T+1}(u_1(s_{1*}) \mid s_{12}^{T}) = 1(0+1) + 2(0+0) = 1$$
$$E_1^{T+1}(u_1(s_{2*}) \mid s_{12}^{T}) = 3(0+1) + 0(0+0) = 3$$

$$E_2^{T+1}(u_1(s_{*1}) \mid s_{12}^{T}) = 1(0+0) + 2(1+0) = 2$$
$$E_2^{T+1}(u_1(s_{*2}) \mid s_{12}^{T}) = 3(0+0) + 0(1+0) = 0$$

# Alternating equilibrium

| Ann \ Bob | $S_1$ | $S_2$ |
|---|---|---|
| $S_1$ | 1, 1 | 2, 3 |
| $S_2$ | 3, 2 | 0, 0 |

# Alternating equilibrium

| | | Bob $S_1$ | Bob $S_2$ |
|---|---|---|---|
| Ann | $S_1$ | 1, 1 | 2, 3 |
| Ann | $S_2$ | 3, 2 | 0, 0 |

# Markov Fictitious Play

$$\mu_1^{T+1}(s_{kl}^{T+1} \mid S_{ij}^{T}) = \alpha_{ij,kl}^{T+1} = \frac{n_{ij,kl}^{T} + \lambda_1 \alpha_{ij,kl}^{0}}{n_{ij}^{T} + \lambda_1}$$

where $\lambda_1 \alpha_{ij,kl}^{0}$ is the prior weight for the initial beliefs.

# Simulation

At the start of a run, the initial transitions transition probabilities are sampled randomly from the uniform probability distribution over $[0, 1]$

$\lambda_1$ and $\lambda_2$ are $10 * |x|$ where $x$ is sampled from a normal distribution with mean 0 and variance 1.

Each simulation for a game consisted of $10,000$ independent runs of 500 rounds each.

We record the percentage of convergence to alternative equilibrium, non-alternating equilibrium or other.

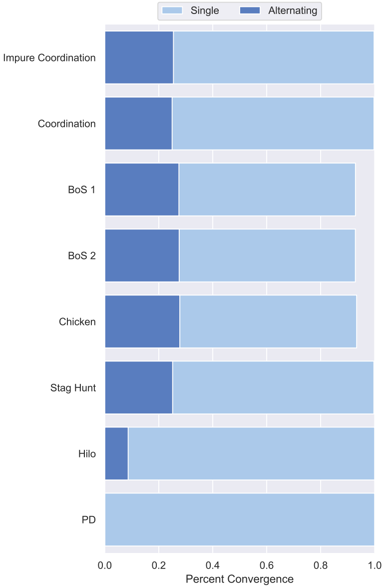

Single
Alternating
Impure Coordination
Coordination
BoS 1
BoS 2
Chicken
Stag Hunt
Hilo
PD
0.0
0.2
0.4
0.6
0.8
1.0
Percent Convergence

Thank you!!

https://syl1.gitbook.io/julia-language-a-concise-tutorial/

https://juliadynamics.github.io/Agents.jl/stable/

https://pacuit.org/esslli2023/game-theory-julia/