Strategies: A logic - automata study
Lecture 1: Exploring structure in strategies

Sujata Ghosh and R. Ramanujam

University of Groningen, The Netherlands

The Institute of Mathematical Sciences, Chennai, India

August 1, 2011

Outline of course

- Lecture 1: Structure of strategies, basic notions, applications in multi-agent systems and other areas.
- Lecture 2: Games of unbounded duration; strategies as finite state transducers; algorithmic questions.
- Lecture 3: Parikh's game logic, extensions to multi-player games, coalitions.
- Lecture 4: ATL and extensions; logics of explicit strategies, automata theoretic decision procedures.
- Lecture 5: Current trends and research directions in reasoning about strategies.

Lectures 1, 3 and 5 will be shared by both of us. Lecture 2 will be by me and Lecture 4 by Sujata.

Course materials

Can be found at the ESSLLI 2011 classroom site:

http://esslli.fmf.uni-lj.si/

You can download recent papers on games and strategies at:

http://www.ai.rug.nl/ sujata/documents.html

Rational players

Underlying assumption of classical game theoretic analysis - players interact in an ideal world.

- Players are perfectly "rational".
- Players have perfect knowledge about all possible strategies.
  - Players have unbounded computational resources.
  - Common knowledge of rationality holds.

Equilibrium theory typically talks about existence of stable strategies.

Large games

When players do not 'know' each other, or when the number of players is large, rationality assumptions need to be re-examined.

- Existence of optimal strategies does not necessarily imply that they will be played.
- Strategy selection is important and interesting to study.
- Players do not come to 'the board' with predecided strategies but start with partial plans that get filled as the game progresses.
- Repeated play calls for different strategisation.

Why bother?

These are philosophical challenges, are they also of computational interest ?

Why bother?

These are philosophical challenges, are they also of computational interest ?

- With the advent of the Internet as the platform of computation, selfish behaviour of agents needs consideration.

Why bother?

These are philosophical challenges, are they also of computational interest ?

- With the advent of the Internet as the platform of computation, selfish behaviour of agents needs consideration.
- Resources need to be shared, leading to situations of cooperation and conflict.

Why bother?

These are philosophical challenges, are they also of computational interest ?

- With the advent of the Internet as the platform of computation, selfish behaviour of agents needs consideration.
- Resources need to be shared, leading to situations of cooperation and conflict.
- Distributed file servers, routers and other such services strategize.

Why bother?

These are philosophical challenges, are they also of computational interest ?

- With the advent of the Internet as the platform of computation, selfish behaviour of agents needs consideration.
- Resources need to be shared, leading to situations of cooperation and conflict.
- Distributed file servers, routers and other such services strategize.
- These are large games.

Rational players

Empirical studies have noted that players deviate from rational behaviour.

Centipede game

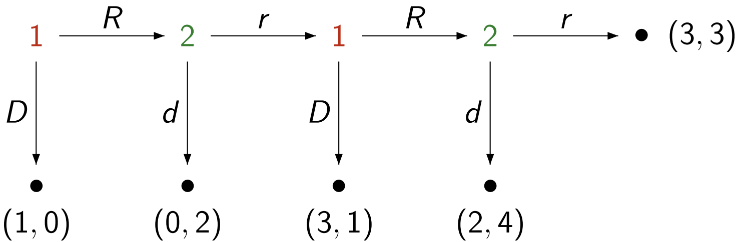

Rational players

Empirical studies have noted that players deviate from rational behaviour.

Centipede game

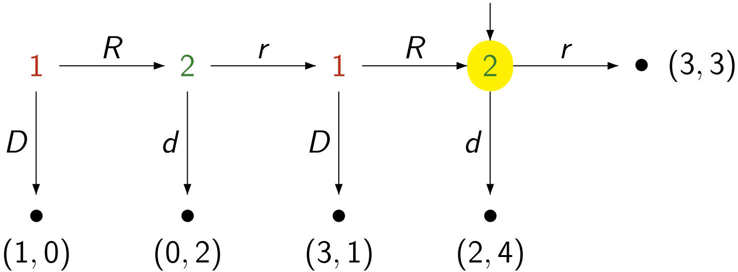

Rational players

Empirical studies have noted that players deviate from rational behaviour.

Centipede game

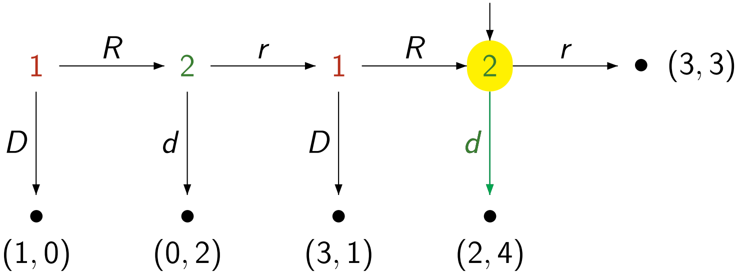

Rational players

Empirical studies have noted that players deviate from rational behaviour.

Centipede game

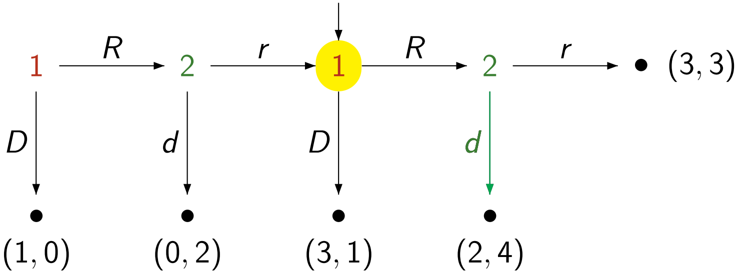

Rational players

Empirical studies have noted that players deviate from rational behaviour.

Centipede game

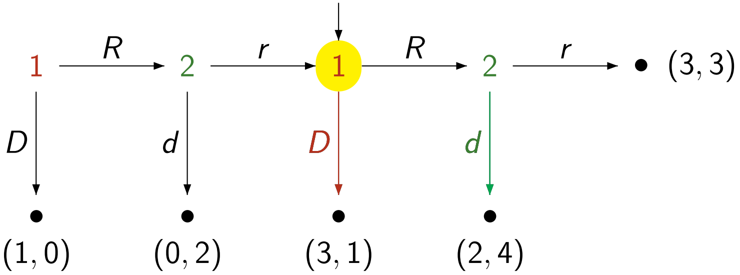

Rational players

Empirical studies have noted that players deviate from rational behaviour.

Centipede game

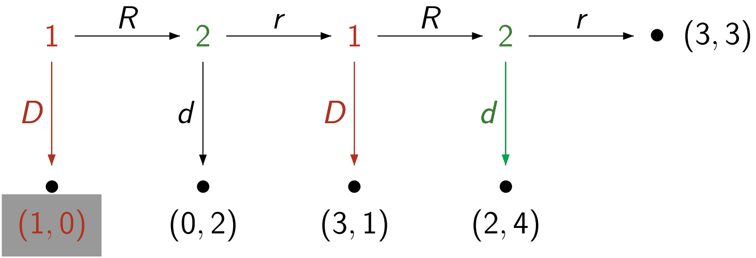

Rational players

Empirical studies have noted that players deviate from rational behaviour.

Centipede game

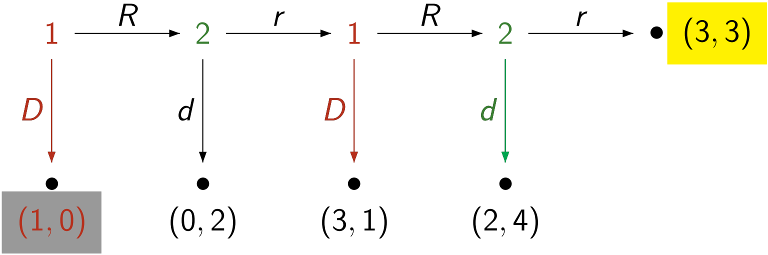

Rational agents

Multi-agent systems: Computational limitations

- Players are resource bounded agents.
  - Agents have limited computational resources.
  - They employ memoryless or bounded memory strategies.

Classical game theoretic analysis does not take into account the computational limitations of players.

Structured strategies

Prescriptive theory for resource bounded players.

Strategic choice:

- Depends on observations made during the play.
- Constitutes response to observed behaviour of other players.

Structured strategies

Prescriptive theory for resource bounded players.

Strategic choice:

- Depends on observations made during the play.
- Constitutes response to observed behaviour of other players.

Strategies are better viewed as relations constraining moves rather than complete functions.

Example: Heuristics employed by chess playing programs.

Structured strategies

Question: Can we come up with a framework where strategies are specified as structured objects built in some compositional fashion ?

A possible approach: Logical analysis - Borrow ideas developed in program analysis.

Our attempt

- A syntactic framework where partially specified strategies are composed in a structured fashion.
- Explicate the strategic response of players.
- Independent of the exact depth of the game tree.

Games of unbounded duration

Reasoning in games, not about games.

- The framework is that of extensive form games.
- Players use local, heuristic strategies.
- Locality and bounded memory effectively make this reasoning in a game of unbounded duration.
- Automata theoretic methods help in the analysis.

Rules of the game

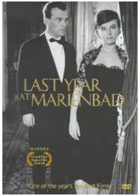

Last year at Marienbad by Alan Resnais

- Character M is always trying to persuade another to play a card game with him.
- Cards arranged in rows of 1, 3, 5, 7.
- They take turns to pick any number of cards from one row. The one who takes the last card loses.
- M is always polite, lets the other start. The other always loses.

Nim

Nim, from the German word nehmen, analysed in 1904 by C. L. Bouton of Harvard University.‘

- Two players I and II, move alternately.
- The game is played with m piles of counters.
- When a player moves, she picks a pile and removes some non-zero many counters from that pile.
- When a player cannot move, he loses (and the other wins).

Ingredients

Every game has three main ingredients:

- The set of players, often $\{I, II\}$. In general, $[n] = \{1, 2, \ldots, n\}$.
- The rules of the game, that specify, at any game position, whose turn it is to move, what moves are applicable, and the resulting new game position after any move.
- Outcomes or winning conditions, that specify at which positions the game is over, and perhaps depending on the course of play, the outcome at those positions.

Game arena

The game arena, that gives the rules of te game, can be envisaged as a finite graph.

- Vertices denote game positions.
- Edges correspond to moves.
- Each vertex is labelled by the player whose turn it is to move.
- Note that winning conditions are not present in the arena.

Game tree

The tree unfolding of a game arena is known as the game in extensive form.

- The root is the initial game position.
- The nodes correspond to paths in the arena from initial position.
- Since the history of play is available, a leaf node can now label the outcome.

Terminology

For the moment, we will consider only two-person zero-sum games of perfect information. $N = \{1, 2\}$.

Terminology

For the moment, we will consider only two-person zero-sum games of perfect information. $N = \{1, 2\}$.

- $\Sigma = \{a_1, a_2, \ldots, a_m\}$ is a finite set of action symbols, representing moves of players, common for both players.

Terminology

For the moment, we will consider only two-person zero-sum games of perfect information. $N = \{1, 2\}$.

- $\Sigma = \{a_1, a_2, \ldots, a_m\}$ is a finite set of action symbols, representing moves of players, common for both players.
- Such games are referred to usually as bipartisan games.

Definitions

- A **game arena** is a graph $\mathcal{G} = (W, \longrightarrow, s_0)$ with $W = \bigcup_{i \in N} W^i \cup \{W^{leaf}\}$.
- For $i \in N$, W^i is the set of *game positions* for player i and W^{leaf} is the set of terminal game positions.
- s_0 is the initial node of the game.
- $\longrightarrow: (W \times \Sigma) \rightarrow W$ is a partial function called the move function.

Plays

Let $\mathcal{G} = (W, \longrightarrow, s_0)$ be a game arena.

- The play of a game can be viewed as placing a token on s_0.
- If player i owns the game position s_0 (i.e $s_0 \in W^i$), then she picks an action 'a' which is enabled for her at s_0 and moves the token to s' where $s_0 \xrightarrow{a} s'$.
- The game then continues from s'.
- Formally, a play in $\mathcal{G}$ is a (possibly infinite) path $\rho : s_0 a_0 s_1 a_1 \cdots$ where $\forall j : s_j \xrightarrow{a_j} s_{j+1}$.

Tree unfolding

Extensive form games are tree unfoldings of game arenas.

- Let $\mathcal{G} = (W, \longrightarrow, s_0)$ be a game arena.
- Let $\mathcal{G}_T = (\pi_W, \rightarrow, \epsilon)$ denote the tree unfolding of the arena $\mathcal{G}$.
- π_W is the set of paths in $\mathcal{G}$ starting from s_0. ϵ is the empty sequence.
- $(u, u') \in \rightarrow$ iff $u' = ua$ and $(u, a, u') \in \longrightarrow$.

Outcomes

The arena $\mathcal{G}$ merely describes the rules by which the game progresses.

- Since we are only interested (for now) in zero-sum games, the outcome is a win for one player and loss for another.

Outcomes

The arena $\mathcal{G}$ merely describes the rules by which the game progresses.

- Since we are only interested (for now) in zero-sum games, the outcome is a win for one player and loss for another.
- Again, for now, we will confine attention to finite games.
- So every leaf node is labelled from $\{1, 2\}$ denoting who wins the play reaching that node.

Outcomes

The arena $\mathcal{G}$ merely describes the rules by which the game progresses.

- Since we are only interested (for now) in zero-sum games, the outcome is a win for one player and loss for another.
- Again, for now, we will confine attention to finite games.
- So every leaf node is labelled from $\{1, 2\}$ denoting who wins the play reaching that node.
- By convention any infinite play is a win for 1 and a loss for 2.

The model

Game model - directed graph where nodes are labelled with players.

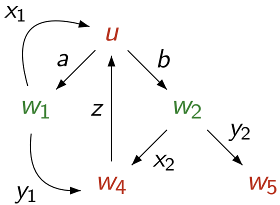

Game arena

The model

Game model - directed graph where nodes are labelled with players.

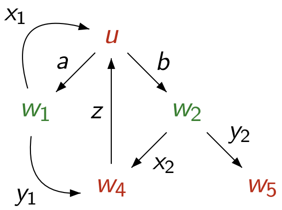

Game arena

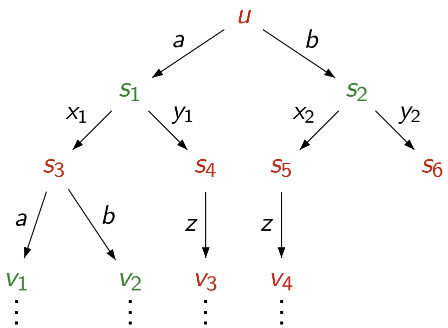

Extensive form game tree

P - countable set of observables

$V : Nodes \rightarrow 2^P$

The model

Strategies of players - subtrees of the game tree.

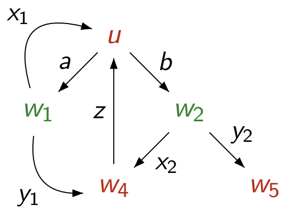

Game arena

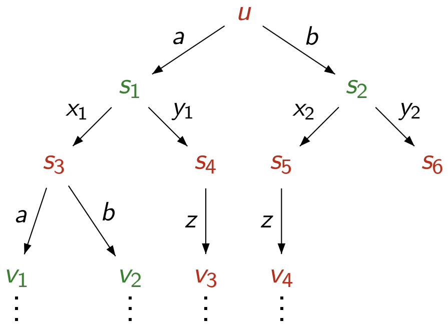

Extensive form game tree

Strategies

Let $\mathcal{G}_T$ denote the tree unfolding of the arena $\mathcal{G}$.

- A strategy for player i, $\mu = (W_\mu, \longrightarrow_\mu, s_0)$ is a maximal connected subtree of $\mathcal{G}_T$ where for each player i node, there is a unique outgoing edge and for the other player every move is included.
- That is, for $s \in W_\mu$ the edge relation satisfies the following property:
  - if $s \in W_\mu^i$ then there exists a unique $a \in \Sigma$ such that $s \xrightarrow{a}_\mu s'$, where we have $s \xrightarrow{a}_T s'$.
  - if $s \in W_\mu^j$ $(j \neq i)$, then for each s' such that $s \xrightarrow{a}_T s'$, we have $s \xrightarrow{a}_\mu s'$.

Strategies

Strategies of players - subtrees of the game tree.

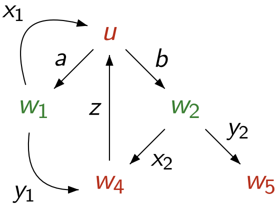

Game arena

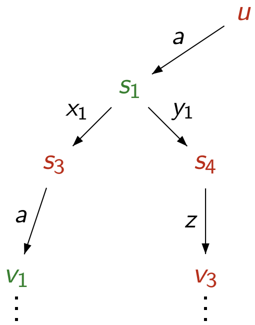

A strategy of player 1

Strategy profiles

A strategy profile is a pair $\langle \mu, \tau \rangle$, that fixes a strategy for each player.

- Note that a strategy profile defines a unique path ρ_μ^τ in the game $\mathcal{G}$.
- This path constitutes a valid *play* if it is of the form $s_0 a_0 \cdots a_{n-1} s_n$ where $s_n = leaf$.
- Every play is either winning or losing for each player. Thus, each strategy profile defines a game outcome.

Winning strategy

Among strategies, some are special.

- A strategy σ for player i is said to be a winning strategy if, every strategy τ of the opponent, the unique play resulting from σ and τ is winning for player i.
- That is, no matter what strategy is adopted by the other player, moving according to the strategy σ guarantees a win for player i.

Solving a game

We can now formalize what it means to analyze a game.

- A win/lose game is said to be determined if, starting from any game position, one of the players has a winning strategy.
- A computer scientist would also like to know which player wins, and to compute her winning strategy.
- We also wish to compute the winning strategy (if it exists).

Backward induction

Zermelo 1913: In every finite extensive form game of perfect information, we can compute whether player i can win (or not).

- Theorem: Backward induction shows who wins, gives a winning strategy in the case of win / lose games, and an NE for general games.
- Note that the game arena for any Nim heap is acyclic and hence the unfolding is a finite tree, so BI applies.

An extensive form game

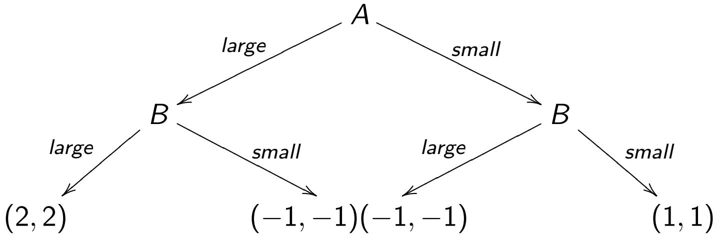

End of course?

Backward induction completely solves finite extensive form games of perfect information, so we might as well go home.

- So the Nim game is solved, isn't it ?

End of course?

Backward induction completely solves finite extensive form games of perfect information, so we might as well go home.

- So the Nim game is solved, isn't it ?
- If we are only interested in existence of winning strategies, this suffices. If we also wish to look at the structure of strategies, this leaves us quite unsatisfied.
- Indeed, in the case of Nim, combinatorial analysis offers more.

Simple cases

Consider the Nim game $(1, 1, 2)$, and you are to move. How would you play ?

Simple cases

Consider the Nim game $(1, 1, 2)$, and you are to move. How would you play ?

- It is easy to see that removing one from the heap with 2 counters is not a good idea.

Simple cases

Consider the Nim game $(1, 1, 2)$, and you are to move. How would you play ?

- It is easy to see that removing one from the heap with 2 counters is not a good idea.
- But removing all counters from that heap works.
- We can call $(1, 1, 2)$ a winning position (for whoever plays) and $(1, 1)$ a losing position (for whoever plays).

Simple cases

Consider the Nim game $(1, 1, 2)$, and you are to move. How would you play ?

- It is easy to see that removing one from the heap with 2 counters is not a good idea.
- But removing all counters from that heap works.
- We can call $(1, 1, 2)$ a winning position (for whoever plays) and $(1, 1)$ a losing position (for whoever plays).
- We can see that $(1, 1, 1)$ is a winning position.

Simple cases

Consider the Nim game $(1, 1, 2)$, and you are to move. How would you play ?

- It is easy to see that removing one from the heap with 2 counters is not a good idea.
- But removing all counters from that heap works.
- We can call $(1, 1, 2)$ a winning position (for whoever plays) and $(1, 1)$ a losing position (for whoever plays).
- We can see that $(1, 1, 1)$ is a winning position.
- Can we see that a k-tuple $(1, 1, \ldots, 1)$ is a winning position iff k is odd ?

The copy strategy

Consider the case of two heaps (m, n).

- Suppose $m = n = 4$, say. Now, whatever move I plays on one heap, II can copy that move on the other heap, thus making the heaps equal again. So this is a losing position for player I.

The copy strategy

Consider the case of two heaps (m, n).

- Suppose $m = n = 4$, say. Now, whatever move I plays on one heap, II can copy that move on the other heap, thus making the heaps equal again. So this is a losing position for player I.
- On the other hand, given heaps of unequal size, player I can equalize them and present II with equal heaps (which is losing for II).

The copy strategy

Consider the case of two heaps (m, n).

- Suppose $m = n = 4$, say. Now, whatever move I plays on one heap, II can copy that move on the other heap, thus making the heaps equal again. So this is a losing position for player I.
- On the other hand, given heaps of unequal size, player I can equalize them and present II with equal heaps (which is losing for II).
- Lemma: For all $m, n \geq 0$. (m, n) is winning iff $m \neq n$.

Subgames

With three or more heaps, analysis becomes difficult: what can we say about $(2, 3, 6)$, for instance ?

Subgames

With three or more heaps, analysis becomes difficult: what can we say about $(2, 3, 6)$, for instance ?

- Observe that every finite extensive form game is of the form 0 or

$$g_1 + g_2 + \ldots + g_m$$

.

- 0 can be thought of as the empty game (in which no player can make any move).
- $g_1, g_2, \ldots, g_m$ are subgames.

Sum of games

Choosing between subgames has an interesting algebraic structure.

- Suppose $g = g_1 + g_2 + \ldots + g_m$.
- Also suppose $h = h_1 + h_2 + \ldots + h_n$.
- Then

$$g + h = (g_1 + h) + \ldots + (g_m + h) + (g + h_1) + \ldots + (g + h_n)$$

.

Sum of games

Choosing between subgames has an interesting algebraic structure.

- Suppose $g = g_1 + g_2 + \ldots + g_m$.
- Also suppose $h = h_1 + h_2 + \ldots + h_n$.
- Then

  $$g + h = (g_1 + h) + \ldots + (g_m + h) + (g + h_1) + \ldots + (g + h_n)$$

  .
- This suggests the notation $1 + 3 + 6$ for the nim game $(1, 3, 6)$.
- When g is a subgame of h, we write $g \leq h$.

Equivalence of games

Once we have game expressions, a natural question arises: when do we consider two games g and g' to be equivalent ?

Equivalence of games

Once we have game expressions, a natural question arises: when do we consider two games g and g' to be equivalent ? (This is only for the analysis of Nim, we will return to these considerations again later.)

- Definition: $g_1 \equiv g_2$ if, for all h, $g_1 + h$ is winning (losing) iff $g_2 + h$ is winning (losing).

Equivalence of games

Once we have game expressions, a natural question arises: when do we consider two games g and g' to be equivalent ? (This is only for the analysis of Nim, we will return to these considerations again later.)

- Definition: $g_1 \equiv g_2$ if, for all h, $g_1 + h$ is winning (losing) iff $g_2 + h$ is winning (losing).
- Sanity check: $\equiv$ is indeed an equivalence relation.

Equivalence of games

Once we have game expressions, a natural question arises: when do we consider two games g and g' to be equivalent ? (This is only for the analysis of Nim, we will return to these considerations again later.)

- Definition: $g_1 \equiv g_2$ if, for all h, $g_1 + h$ is winning (losing) iff $g_2 + h$ is winning (losing).
- Sanity check: $\equiv$ is indeed an equivalence relation.
- Sanity check: $g + 0 = g$.

Equivalence of games

Once we have game expressions, a natural question arises: when do we consider two games g and g' to be equivalent ? (This is only for the analysis of Nim, we will return to these considerations again later.)

- Definition: $g_1 \equiv g_2$ if, for all h, $g_1 + h$ is winning (losing) iff $g_2 + h$ is winning (losing).
- Sanity check: $\equiv$ is indeed an equivalence relation.
- Sanity check: $g + 0 = g$.
- Therefore, if $g_1 \equiv g_2$ then g_1 is winning iff g_2 is winning. But the converse is not true.

Equivalence of games

Once we have game expressions, a natural question arises: when do we consider two games g and g' to be equivalent ? (This is only for the analysis of Nim, we will return to these considerations again later.)

- Definition: $g_1 \equiv g_2$ if, for all h, $g_1 + h$ is winning (losing) iff $g_2 + h$ is winning (losing).
- Sanity check: $\equiv$ is indeed an equivalence relation.
- Sanity check: $g + 0 = g$.
- Therefore, if $g_1 \equiv g_2$ then g_1 is winning iff g_2 is winning. But the converse is not true.
- However, all losing games are equivalent, to 0.

The loser's lemma

Lemma: If g is losing then $g \equiv 0$.

The loser's lemma

Lemma: If g is losing then $g \equiv 0$.

- Fix a losing game g.

The loser's lemma

Lemma: If g is losing then $g \equiv 0$.

- Fix a losing game g.
- To prove: for all h, $g + h$ is losing iff h is losing.

The loser's lemma

Lemma: If g is losing then $g \equiv 0$.

- Fix a losing game g.
- To prove: for all h, $g + h$ is losing iff h is losing.
- We prove by induction on h that whenever h is losing, so is $g + h$.

The loser's lemma

Lemma: If g is losing then $g \equiv 0$.

- Fix a losing game g.
- To prove: for all h, $g + h$ is losing iff h is losing.
- We prove by induction on h that whenever h is losing, so is $g + h$.
- Assuming this, the other direction is easy:

The loser's lemma

Lemma: If g is losing then $g \equiv 0$.

- Fix a losing game g.
- To prove: for all h, $g + h$ is losing iff h is losing.
- We prove by induction on h that whenever h is losing, so is $g + h$.
- Assuming this, the other direction is easy: Suppose h is winning. Then there is a move to h' that's losing.

The loser's lemma

Lemma: If g is losing then $g \equiv 0$.

- Fix a losing game g.
- To prove: for all h, $g + h$ is losing iff h is losing.
- We prove by induction on h that whenever h is losing, so is $g + h$.
- Assuming this, the other direction is easy: Suppose h is winning. Then there is a move to h' that's losing. By claim above, $g + h'$ is losing.

The loser's lemma

Lemma: If g is losing then $g \equiv 0$.

- Fix a losing game g.
- To prove: for all h, $g + h$ is losing iff h is losing.
- We prove by induction on h that whenever h is losing, so is $g + h$.
- Assuming this, the other direction is easy: Suppose h is winning. Then there is a move to h' that's losing. By claim above, $g + h'$ is losing. Hence there is a winning move in $g + h$, and hence $g + h$ is winning.

The loser's lemma

We prove by induction on h that whenever h is losing, so is $g + h$.

The loser's lemma

We prove by induction on h that whenever h is losing, so is $g + h$.

- IH1: for all $g' \leq g$, if h is losing, then so is $g' + h$.

The loser's lemma

We prove by induction on h that whenever h is losing, so is $g + h$.

- IH1: for all $g' \leq g$, if h is losing, then so is $g' + h$.
- IH2: for all $h' \leq h$, if h' is losing, then so is $g + h'$.

The loser's lemma

We prove by induction on h that whenever h is losing, so is $g + h$.

- IH1: for all $g' \leq g$, if h is losing, then so is $g' + h$.
- IH2: for all $h' \leq h$, if h' is losing, then so is $g + h'$.
- Every initial move in $g + h$ is either in g or in h. First consider the latter.

The loser's lemma

We prove by induction on h that whenever h is losing, so is $g + h$.

- IH1: for all $g' \leq g$, if h is losing, then so is $g' + h$.
- IH2: for all $h' \leq h$, if h' is losing, then so is $g + h'$.
- Every initial move in $g + h$ is either in g or in h. First consider the latter.
- h is losing, so every move in h to h' is winning, hence there exists a move in h' to h'' which is losing. By IH2, $g + h''$ is losing.

The loser's lemma

We prove by induction on h that whenever h is losing, so is $g + h$.

- IH1: for all $g' \leq g$, if h is losing, then so is $g' + h$.
- IH2: for all $h' \leq h$, if h' is losing, then so is $g + h'$.
- Every initial move in $g + h$ is either in g or in h. First consider the latter.
- h is losing, so every move in h to h' is winning, hence there exists a move in h' to h'' which is losing. By IH2, $g + h''$ is losing.
- g is losing, so every move in g to g' is winning, hence there exists a move in g' to g'' which is losing. By IH2, $g'' + h$ is losing.

The loser's lemma

We prove by induction on h that whenever h is losing, so is $g + h$.

- IH1: for all $g' \leq g$, if h is losing, then so is $g' + h$.
- IH2: for all $h' \leq h$, if h' is losing, then so is $g + h'$.
- Every initial move in $g + h$ is either in g or in h. First consider the latter.
- h is losing, so every move in h to h' is winning, hence there exists a move in h' to h'' which is losing. By IH2, $g + h''$ is losing.
- g is losing, so every move in g to g' is winning, hence there exists a move in g' to g'' which is losing. By IH2, $g'' + h$ is losing.
- Thus every move in $g + h$ is losing, and we are done.

Applications

We now have: if h is losing then for all g, $g + h = g$.

Applications

We now have: if h is losing then for all g, $g + h = g$.

- This means that every losing subgame can be ignored !

Applications

We now have: if h is losing then for all g, $g + h = g$.

- This means that every losing subgame can be ignored !
- Let us use this knowledge to analyse $1 + 2 + 3 + 4 + 5$.

Applications

We now have: if h is losing then for all g, $g + h = g$.

- This means that every losing subgame can be ignored !
- Let us use this knowledge to analyse $1 + 2 + 3 + 4 + 5$.
- Consider $1 + 2 + 3$. Note that removing an entire heap leads to a winning position. Reducing any heap leads to two equal heaps, which is losing and can be ignored. Thus $1 + 2 + 3$ is losing.

Applications

We now have: if h is losing then for all g, $g + h = g$.

- This means that every losing subgame can be ignored !
- Let us use this knowledge to analyse $1 + 2 + 3 + 4 + 5$.
- Consider $1 + 2 + 3$. Note that removing an entire heap leads to a winning position. Reducing any heap leads to two equal heaps, which is losing and can be ignored. Thus $1 + 2 + 3$ is losing.
- We know that $4 + 5$ is winning, so $1 + 2 + 3 + 4 + 5$ is winning.

A principle

Can we get some more general mileage than analysing simple Nim heaps ?

- Question: How do you ensure that you do not lose in a Chess game against a Grandmaster ?

The copy strategy

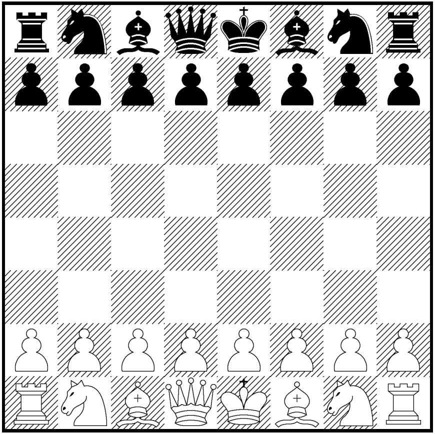

The copy strategy

1. e2–e3

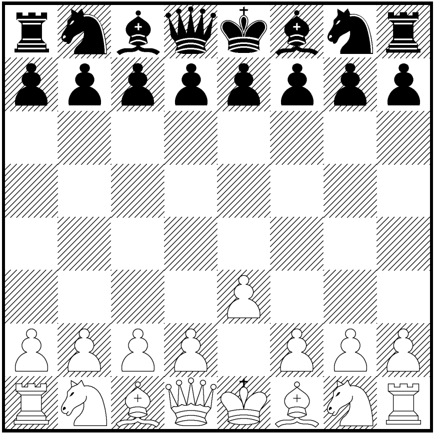

The copy strategy

1. . . . e7–e6

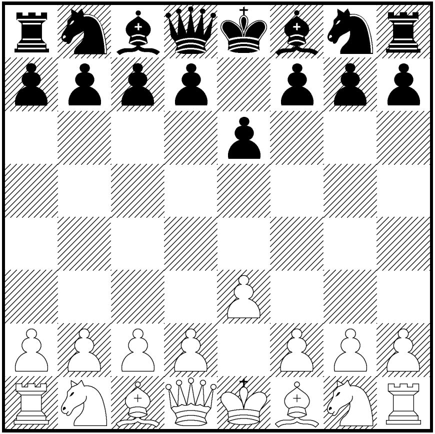

The copy strategy

2. ♕d1–f3

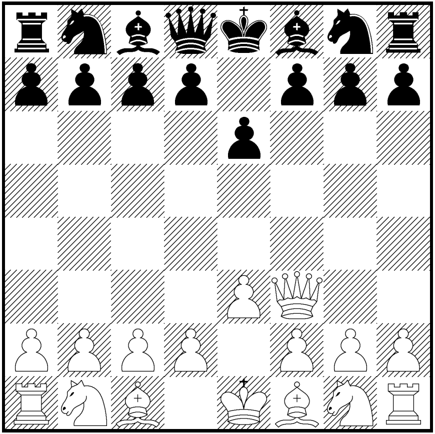

The copy strategy

2. . . . ♛d8–f6

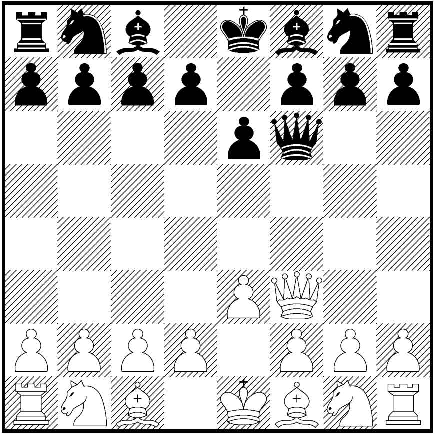

The copy strategy

3. ♕f3–g3

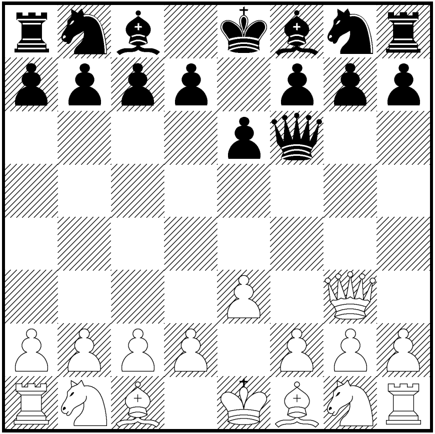

The copy strategy

3. . . . ♕f6–g6

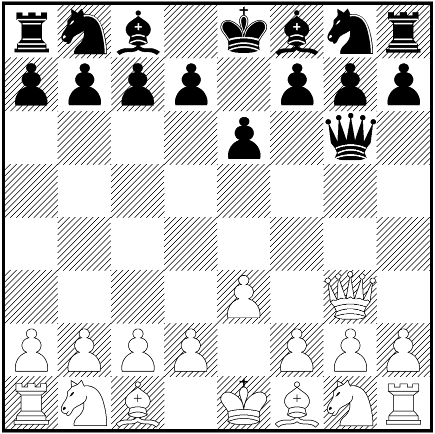

The copy strategy

4. ♘g1–f3

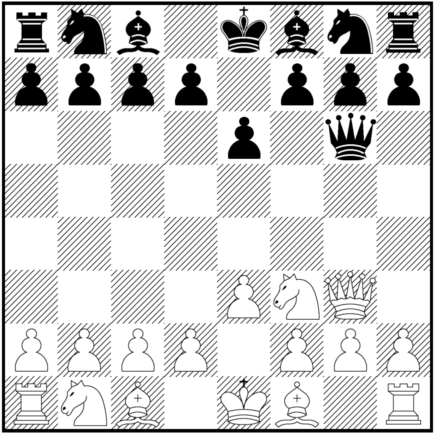

The copy strategy

4. . . . ♘g8–f6

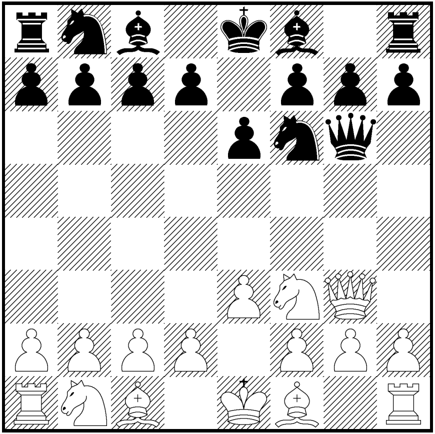

The copy strategy

5. ♔e1–d1

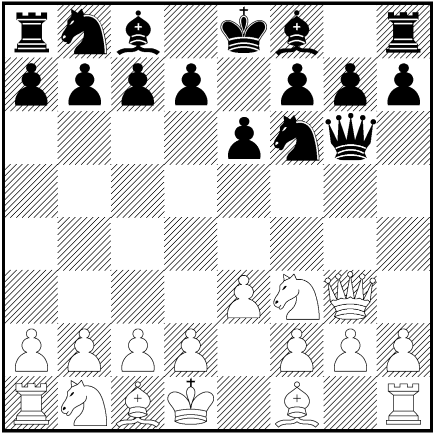

The copy strategy

5. . . . ♔e8–d8

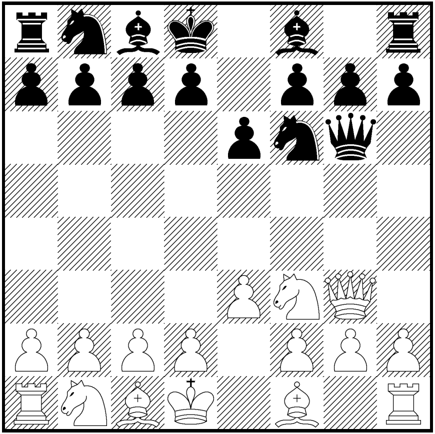

The copy strategy

6. ♗f1–e2

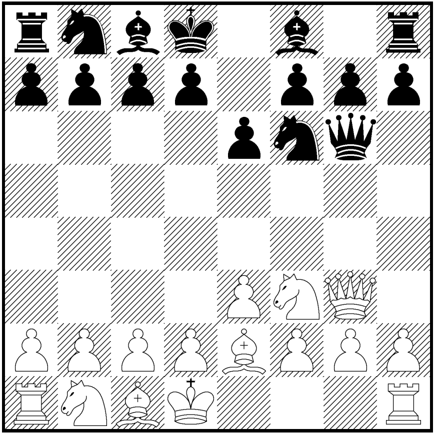

The copy strategy

6. . . . ♝f8–e7

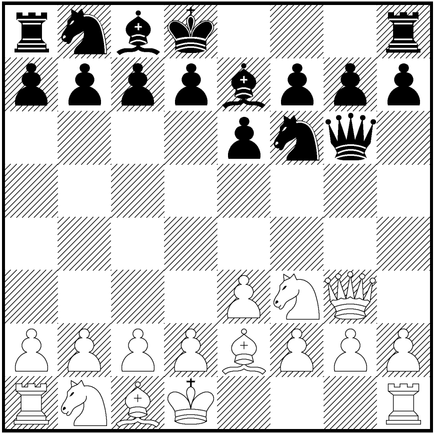

The copy strategy

7. ♖h1–e1

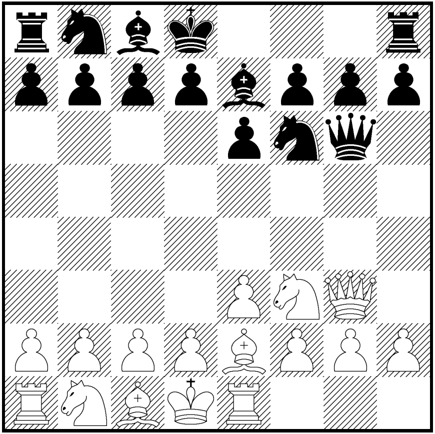

The copy strategy

7. . . . ♖h8–e8

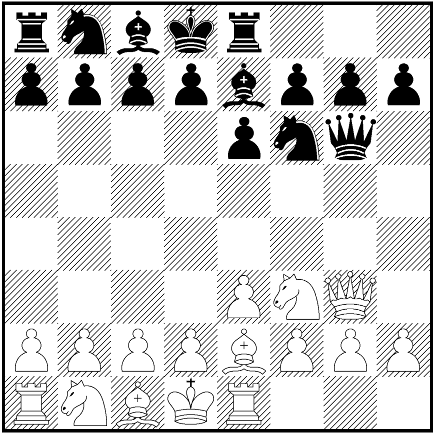

The copy strategy

8. ♘b1–c3

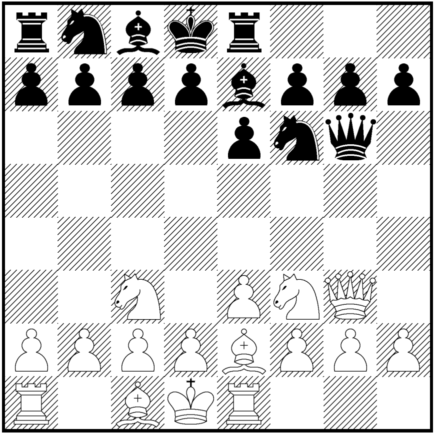

The copy strategy

8. . . . ♘b8–c6

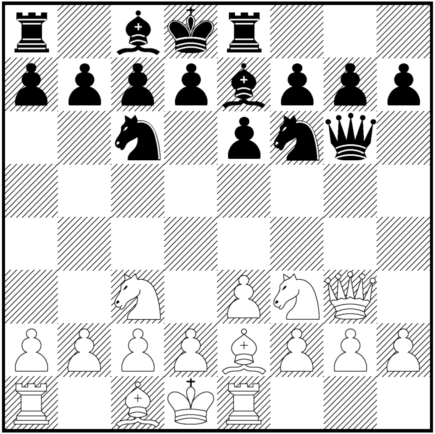

The copy strategy

9. ♘c3–b5

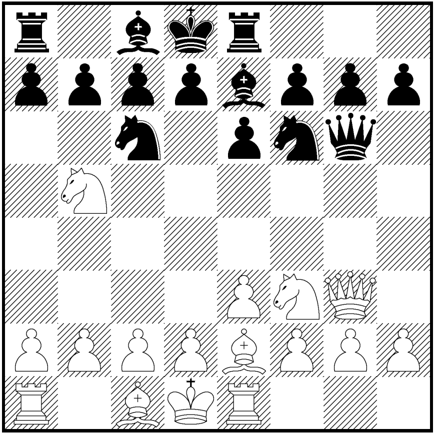

The copy strategy

9. . . . ♘c6–b4

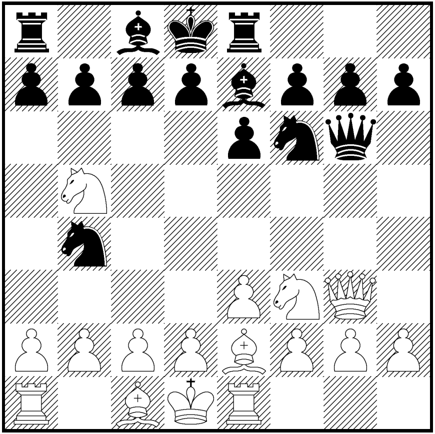

The copy strategy

10. ♕g3×c7#

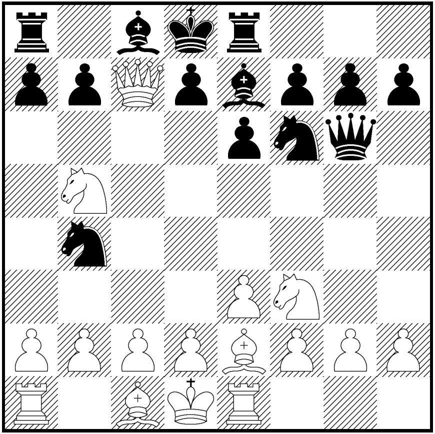

The copycat principle

We can now enunciate an important principle in bipartisan games.

The copycat principle

We can now enunciate an important principle in bipartisan games.

- Lemma $g + g \equiv 0$.

The copycat principle

We can now enunciate an important principle in bipartisan games.

- Lemma $g + g \equiv 0$.
- Given g, assume that the claim holds for $g' \leq g$.

The copycat principle

We can now enunciate an important principle in bipartisan games.

- Lemma $g + g \equiv 0$.
- Given g, assume that the claim holds for $g' \leq g$.
- A move in $g + g$ is of the form $g' + g$.

The copycat principle

We can now enunciate an important principle in bipartisan games.

- Lemma $g + g \equiv 0$.
- Given g, assume that the claim holds for $g' \leq g$.
- A move in $g + g$ is of the form $g' + g$.
  This is winning by the move to the subgame $g' + g'$ which is losing by IH.

The copycat principle

We can now enunciate an important principle in bipartisan games.

- Lemma $g + g \equiv 0$.
- Given g, assume that the claim holds for $g' \leq g$.
- A move in $g + g$ is of the form $g' + g$.
  This is winning by the move to the subgame $g' + g'$ which is losing by IH.
  So $g + g$ is losing and by loser's lemma, equivalent to 0.

The algebraic structure

The class of games with addition forms an abelian group.

The algebraic structure

The class of games with addition forms an abelian group.

- Commutativity and associativity are easy.

The algebraic structure

The class of games with addition forms an abelian group.

- Commutativity and associativity are easy.
- 0 is the neutral element.

The algebraic structure

The class of games with addition forms an abelian group.

- Commutativity and associativity are easy.
- 0 is the neutral element.
- Every element has an inverse, and as we just saw, has itself as inverse.

The algebraic structure

The class of games with addition forms an abelian group.

- Commutativity and associativity are easy.
- 0 is the neutral element.
- Every element has an inverse, and as we just saw, has itself as inverse.

Reminiscent of some very familiar algebraic structures ?

Summing up

This foray into combinatorial games was merely to state the obvious: exploring compositional structure in strategies is worthwhile.

Summing up

This foray into combinatorial games was merely to state the obvious: exploring compositional structure in strategies is worthwhile.

- This is not only interesting, but also required when we consider resource limited players, who can only strategize locally in games.

Summing up

This foray into combinatorial games was merely to state the obvious: exploring compositional structure in strategies is worthwhile.

- This is not only interesting, but also required when we consider resource limited players, who can only strategize locally in games.
- Our model of resource limitation is bounded memory.

What's ahead

- We will be studying reasoning in large games, where a player does not come to the table with pre-decided strategies but starts with partial strategies and composes them.

What's ahead

- We will be studying reasoning in large games, where a player does not come to the table with pre-decided strategies but starts with partial strategies and composes them.
- We will be interested in temporally large games, by which we mean games of unbounded duration, as well as spatially large games where the number of players is large.

What's ahead

- We will be studying reasoning in large games, where a player does not come to the table with pre-decided strategies but starts with partial strategies and composes them.
- We will be interested in temporally large games, by which we mean games of unbounded duration, as well as spatially large games where the number of players is large.
- Such games arise naturally in the study of multi-agent systems.

Multi-player games

- More complex structures
- Turns of the players
- Group decisions and coalition formations occur naturally
- Concurrent moves might occur.

Overlapping objectives

- Winning strategies are not appropriate.
- We talk about best-response strategies.
- Non-determinacy comes in.
- Utilities, pay-offs – preference relations
- Winning strategies, best-response ones, equilibrium profiles

A simple case of overlapping objectives

Cake-cutting game (I cut, you choose !!)

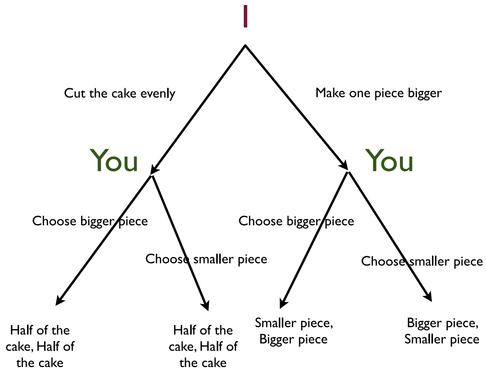

Extensive form game tree

Perfect information multi-player sequential games can be presented in form of trees, with nodes correspond to a finite history of the game, and the edges correspond to moves of the players.

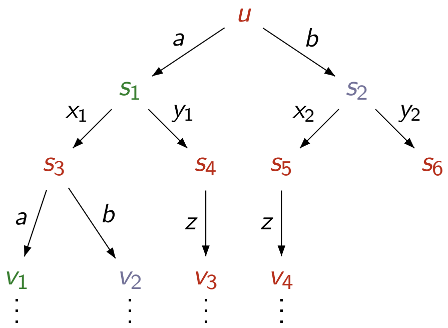

Definitions

Let $N = \{1, 2, \ldots, n\}$ be a finite set of players and $\Sigma = \{a_1, a_2, \ldots, a_m\}$ be a finite set of action symbols, which represent moves of players. We assume that all the actions are available to all the players.

- Formally, we have a tree $\mathcal{G} = (W, \longrightarrow, s_0)$ with $W = \bigcup_{i \in N} W^i \cup \{W^{leaf}\}$.
- For $i \in N$, W^i is the set of *game positions* for player i and W^{leaf} is the set of terminal game positions.
- s_0 is the initial node of the game.
- $\longrightarrow : (W \times \Sigma) \to W$ is a partial function called the move function.

Plays

Let $\mathcal{G} = (W, \longrightarrow, s_0)$ be a game tree.

- The play of a game can be viewed as placing a token on s_0.
- If player i owns the game position s_0 (i.e $s_0 \in W^i$), then she picks an action 'a' which is enabled for her at s_0 and moves the token to s' where $s_0 \xrightarrow{a} s'$.
- The game then continues from s'.
- Formally, a play in $\mathcal{G}$ is a (possibly infinite) path $\rho : s_0 a_0 s_1 a_1 \cdots$ where $\forall j : s_j \xrightarrow{a_j} s_{j+1}$.

Outcomes

The tree $\mathcal{G}$ merely describes the rules by which the game progresses.

- We assume that each player has a preference relation over the set of plays.
- Let $\preceq^i \subseteq (Plays \times Plays)$ be a complete, reflexive, transitive binary relation denoting the preference relation of player i for $i \in N$.
- Then the game G is given as, $G = (\mathcal{G}, \{\preceq^i\}_{i \in N})$.
- Win - lose games: $\mu : Plays \to 2^N$ specifies the outcome. Player i is said to win play u if $i \in \mu(u)$, and is said to lose u otherwise.

Strategies

- A strategy for player i, $\mu = (W_\mu, \longrightarrow_\mu, s_0)$ is a maximal connected subtree of $\mathcal{G}$ where for each player i node, there is a unique outgoing edge and for other players every move is included.
- That is, for $s \in W_\mu$ the edge relation satisfies the following property:
  - if $s \in W_\mu^i$ then there exists a unique $a \in \Sigma$ such that $s \xrightarrow{a}_\mu s'$, where we have $s \xrightarrow{a}_\mathcal{G} s'$.
  - if $s \in W_\mu^j$ $(j \neq i)$, then for each s' such that $s \xrightarrow{a}_\mathcal{G} s'$, we have $s \xrightarrow{a}_\mu s'$.
- Let Ω_i denote the set of all strategies for player i.

Strategy profiles

A strategy profile is a tuple $\langle \mu_1, \ldots \mu_n \rangle$, that fixes a strategy for each player.

- Note that a strategy profile defines a unique path ρ in the game $\mathcal{G}$.
- This path constitutes a valid *play* if it is of the form $s_0 a_0 \cdots a_{n-1} s_n$ where $s_n = leaf$.
- If $\mathcal{G}$ is a win / lose game, the play is either winning or losing for each player. Thus, each strategy profile defines a game outcome.

Overlapping objectives

In the case of games where players have non-binary preferences over outcomes, the notion of winning strategy is not appropriate.

- For a strategy profile $\sigma = \langle \mu_1, \ldots \mu_n \rangle$, let $\sigma_{-i} = \langle \mu_1, \ldots \mu_{i-1}, \mu_{i+1}, \ldots, \mu_n \rangle$ and $(\sigma_{-i}; \mu_i) = \sigma$.
- We say that μ is the best response for σ_{-i} iff $\rho_{(\sigma_{-i};\mu)}$ is a valid play and $\forall \mu' \in \Omega_i$ such that $\rho_{(\sigma_{-i};\mu')}$ is a valid play, we have: $\rho_{(\sigma_{-i};\mu')} \preceq^i \rho_{(\sigma_{-i};\mu)}$.
- A strategy profile $\sigma = \langle \mu_1, \ldots \mu_n \rangle$ is a Nash equilibrium iff for all i, μ_i is the best response for σ_{-i}.

Solving a general game

- We would like to know whether an equilibrium profile of strategies exist.
- We also wish to compute the equilibrium profile.
- Finite games of perfect information: Backward Induction gives an NE.
- Games are sequential: An equilibrium concept acknowledging this notion might be more intuitive.

A simple game

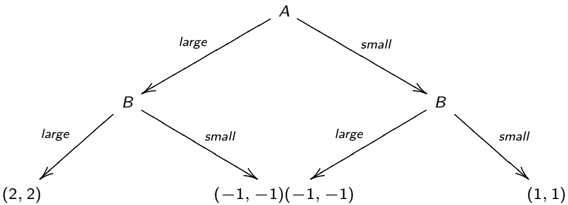

Equilibrium	Strategies	Outcome
NE1	{Large, (Large, Large)}	Both choose large
NE2	{Large, (Large, Small)}	Both choose large
NE3	{Small, (Small, Small)}	Both choose small

Subgame perfection

We need a notion finer than Nash equilibrium.

- A strategy profile that defines a Nash equilibrium for every subgame is called a subgame perfect Nash equilibrium (SPNE).
- In this game, $\{Large, (Large, Small)\}$ is an SPNE, and is indeed the solution given by backward induction.
- Theorem: Backward induction defines an SPNE.

Are we done?

- Infinite games: With ω-regular objectives, existence of NE and SPNE has been shown.
- Large games, structured strategies
- How to play: even in large finite games.

Games are everywhere

- Interaction and decision-making in multi-agent systems
- Modeling market conditions
- Evolutionary dynamics
- Signaling in communication
- Analyzing theory of mind

Multi-agent systems

- Interactions
- Group decisions
- Coalition formation
- Bargaining - negotiation process
- Auctions - resource allocation

Climate-change negotiations

- Ever-continuing
- Kyoto protocol
- Extensive-form game analysis
- Players: developed countries (Annex B) and the rest of the world (non-Annex B)
- Actions: choice of greenhouse gas mitigation policies

An extensive form game representation

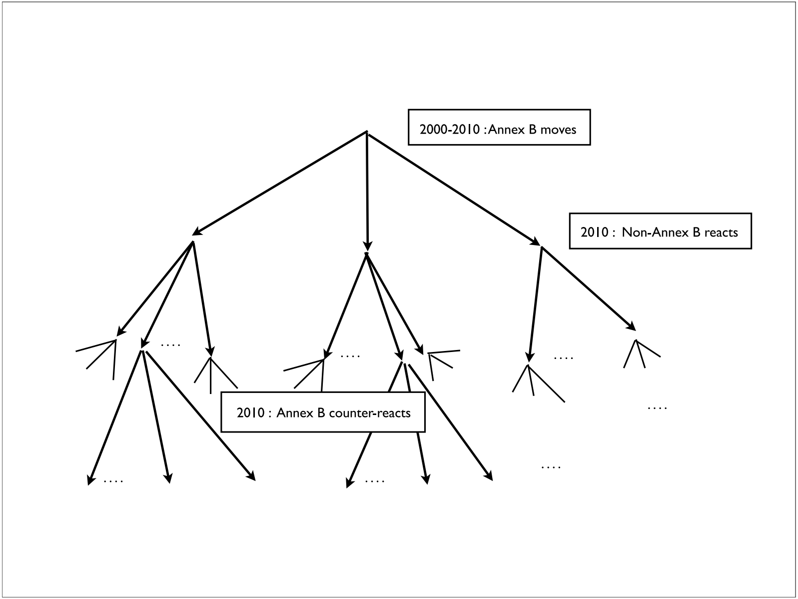

Some issues

- Incentives for free ride: not fully cooperative
- Multi-stage game with conditions to continue
- Can backward induction give a realistic solution?
- Best-response, switching of strategies

A sequential game of coalition formation

- N players: A proposer i proposes a coalition including herself
- Each prospective member responds to the proposal in some order
- If one of the players rejects the proposal, she makes a counteroffer and proposes a coalition including herself
- If all the members accept, the coalition is formed
- All these members then withdraw, and one of the remaining members starts making a proposal.

An extensive form game representation

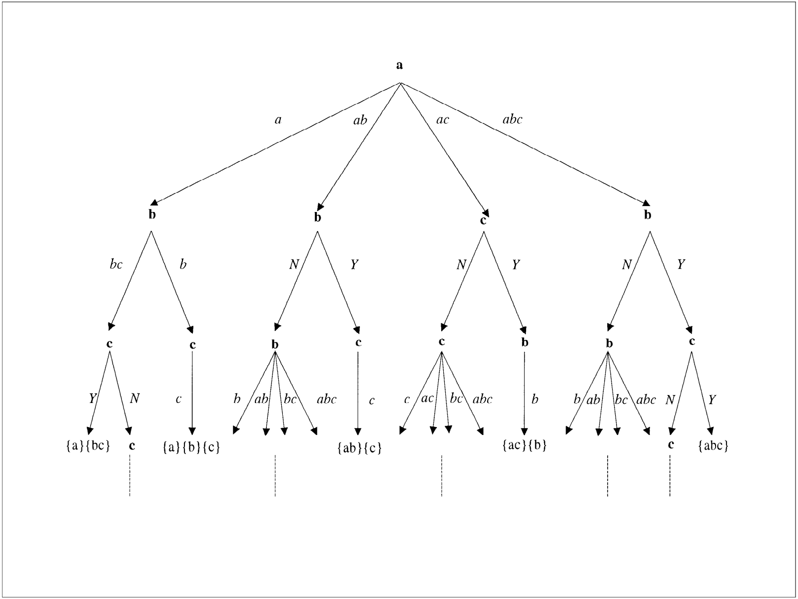

Main issues

- Outcomes are pairs of a set of coalition formed and the period in which it is achieved.
- One who goes on infinitely gets zero payoff.
- The players are committed to the coalitions, once formed.

- Self-interest plays a major role.
- It is not profitable to disagree forever.
- Large number of players get involved. Top-down strategizing makes sense.

Bargaining

- Bargaining with discounts
- Players 1 and 2 bargain about how to split goods whose initial worth is 1 euro, say
- After each round without agreement, the subjective worth of the goods reduces by discount rates δ_1 (for player 1) and δ_2 (for player 2). After t rounds, the goods are worth $(\delta_1)^t$ and $(\delta_2)^t$ respectively.
- At the round t, player 1 (if t is even) or player 2 (if t is odd) makes an offer to split the goods in proportions (x, 1-x), and the other player accepts or rejects it - if rejection occurs, the game continues.

An infinite bargaining game

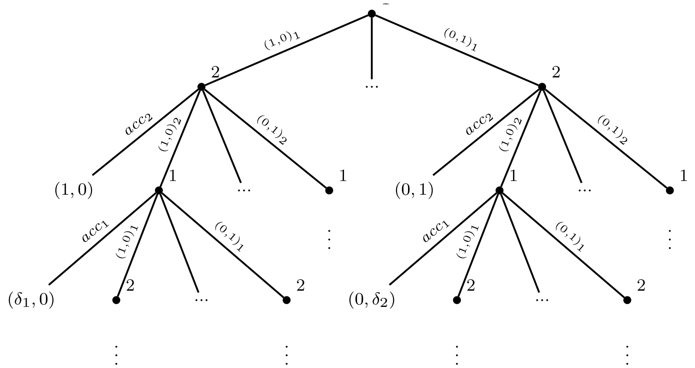

Auctions

- Variety of auction procedures that are in common use for all practical purposes.
- Penny auction, one of the more frequently used auctions in the internet.
- Timed auctions where the user pay-per-bid. Each bid adds a small amount of time to the clock.
- The last bidder to bid before the time runs out typically wins.

An extensive form game representation

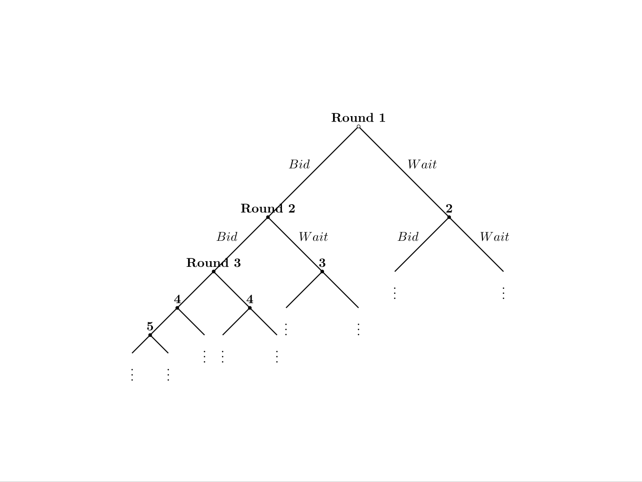

An instance of a penny auction: Each level of the tree corresponds to a round of auction, since all players are selecting the same action

Modeling market conditions

A Classic example: Entry deterrence game

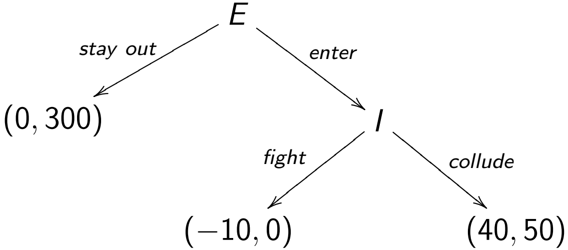

Repeated market games

- Modeling various market conditions
- Incumbent firms facing various potential entrants and strategizing with respect to prior experiences
- Reputation
- Large games, stable strategies

Evolutionary dynamics

- Darwinian natural selection
- Populations of individuals
- Evolutionary stable strategies
- Imitation

Behavioral evolution is sometimes based on imitation, which is an area where the study of extensive form games are indispensable. Study of such structures makes the analyses more tractable, even when the games have a large number of pure strategies.

Signalling in communication

- Rationalizability of communication
- Credibility of messages
- Acquisition of meaning
- Evolution of language

Analysing theory of mind

- Higher order theory of mind
- Experimentations on extensive form games set up
- Cognitive computational models
- Testing and predicting human strategic reasoning

Summing up

- Large games: large temporal structures
- Common methodology of strategizing
- Formation of a library of strategies
- Large number of players
- Different information scenarios