


# Review

## Uninformed Search

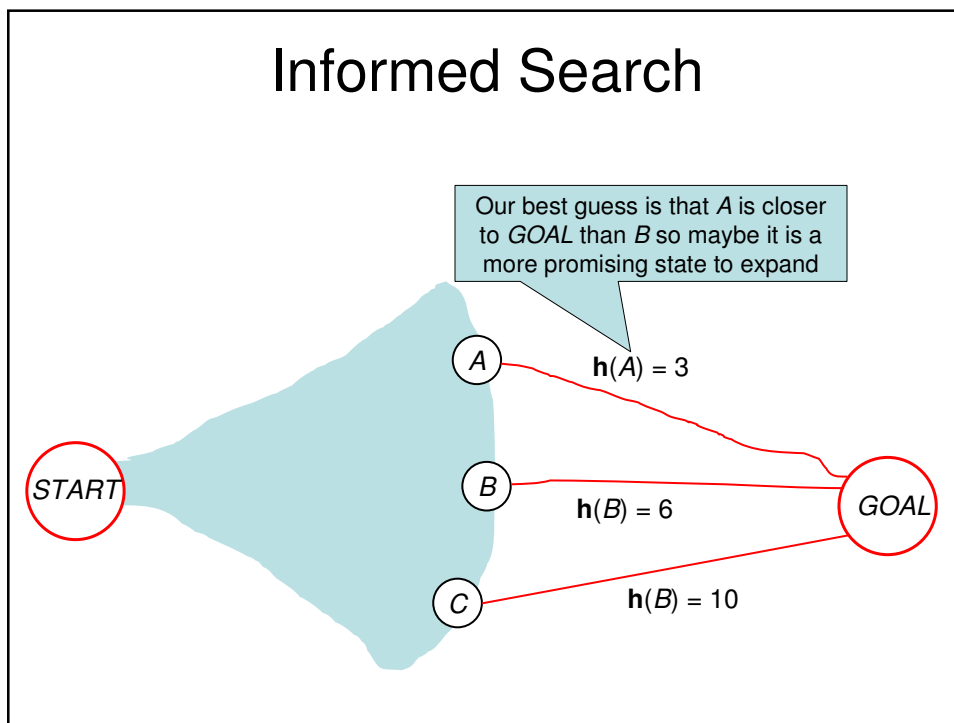


# Complexity

- $N$  = Total number of states
- $B$  = Average number of successors (branching factor)
- $L$  = Length for start to goal with smallest number of steps
- $Q$  = Average size of the priority queue
- $L_{max}$  = Length of longest path from *START* to any state


Algorithm		Complete	Optimal	Time	Space
BFS	Breadth First Search	Y	Y, If all trans. have same cost	$O(\text{Min}(N, B^L))$	$O(\text{Min}(N, B^L))$
BIBFS	Bi- Direction. BFS	Y	Y, If all trans. have same cost	$O(\text{Min}(N, 2B^{L/2}))$	$O(\text{Min}(N, 2B^{L/2}))$
UCS	Uniform Cost Search	Y, If cost > 0	Y, If cost > 0	$O(\log(Q) * \text{Min}(N, B^L))$	$O(\text{Min}(N, B^L))$
PCDFS	Path Check DFS	Y	N	$O(B^{L_{max}})$	$O(B^{L_{max}})$
MEMD FS	Memorizing DFS	Y	N	$O(\text{Min}(N, B^{L_{max}}))$	$O(\text{Min}(N, B^{L_{max}}))$
IDS	Iterative Deepening	Y	Y, If all trans. have same cost	$O(B^L)$	$O(BL)$

# Informed Search



## Informed Search

- Best-First Search: Expand node with minimum  $h(s)$
- A\*: Expand node with minimum  $f(s) = g(s) + h(s)$
- Guaranteed to be optimal if h admissible  
 $h(s) \leq h^*(s)$
- No nodes revisited if monotonic:  
 $h(s) < h(s') + \text{cost}(s, s')$
- IDA\* = Equivalent of Iterative Deepening for A\*  $\rightarrow$  guarantees low memory usage



- (a) Find a goal location for which DFS does much less work than BFS and ID
- (b) Find a goal location for which BFS and ID do much less work than DFS
- (c) Find a goal location for which ID does much less work than BFS and DFS
- (d) Find a goal location for which BFS does much less work than DFS and ID

$$f(s) = (2 - w)g(s) + wh(s)$$

### Most Basic Algorithm: Hill-Climbing (Greedy Local Search)

- $X \leftarrow$  Initial configuration
- Iterate:
  1.  $E \leftarrow Eval(X)$
  2.  $\mathcal{N} \leftarrow Neighbors(X)$
  3. For each  $X_i$  in  $\mathcal{N}$   
 $E_i \leftarrow Eval(X_i)$
  4. If all  $E_i$ 's are lower than  $E$   
Return  $X$   
Else  
 $i^* = \operatorname{argmax}_i (E_i) \quad X \leftarrow X_{i^*} \quad E \leftarrow E_{i^*}$

## Stochastic Search: Randomized Hill-Climbing

- $X \leftarrow$  Initial configuration

- Iterate:

Until when?

1.  $E \leftarrow Eval(X)$

2.  $X' \leftarrow$  one configuration randomly selected in *Neighbors* ( $X$ )

3.  $E' \leftarrow Eval(X')$

4. If  $E' > E$

$X \leftarrow X'$

$E \leftarrow E'$

Critical change: We no longer select the best move in the entire neighborhood

## Simulated Annealing

1. Do  $K$  times:

- 1.1  $E \leftarrow Eval(X)$

- 1.2  $X' \leftarrow$  one configuration randomly selected in *Neighbors* ( $X$ )

- 1.3  $E' \leftarrow Eval(X')$

- 1.4 If  $E' \geq E$

$X \leftarrow X'; E \leftarrow E';$

Else accept the move with probability

$p = e^{-(E-E')/T} :$

$X \leftarrow X'; E \leftarrow E';$

2.  $T \leftarrow \alpha T$

## Basic GA Outline

- Create initial population  $X = \{X_1, \dots, X_P\}$
- Iterate:
  1. Select  $K$  random pairs of parents  $(X, X')$
  2. For each pair of parents  $(X, X')$ :
    - 1.1 Generate offsprings  $(Y_1, Y_2)$  using crossover operation
    - 1.2 For each offspring  $Y_i$ :
 

Replace randomly selected element of the population by  $Y_i$

With probability  $\mu$ :

Apply a random mutation to  $Y_i$
- Return the best individual in the population

1. Let  $X$  := initial object
2. Let  $E$  := Eval( $X$ )
3. Let  $X'$  := randomly chosen configuration chosen from the moveset of  $X$
4. Let  $E'$  := Eval( $X'$ )
5. Let  $z$  := a number drawn randomly uniformly between 0 and 1
6. If  $E' > E$  or  $\exp(-(E - E')/T) > z$  then
  - $X := X'$
  - $E := E'$
7.  $T := r \times T$
8. If a convergence test is satisfied then halt. Else go to Step 3.

- (a) Normally  $r$ , the temperature decay rate, is chosen in the range  $0 < r < 1$ . How would the behavior of simulated annealing change if  $r > 1$ ?

*The change will always be accepted and we'll do a random walk.*

- (b) Alternatively, how would it change if  $r = 0$ ?

- (c) If we simplified the conditional test in Step 6 to

If  $\exp(-(E - E')/T) > z$  then

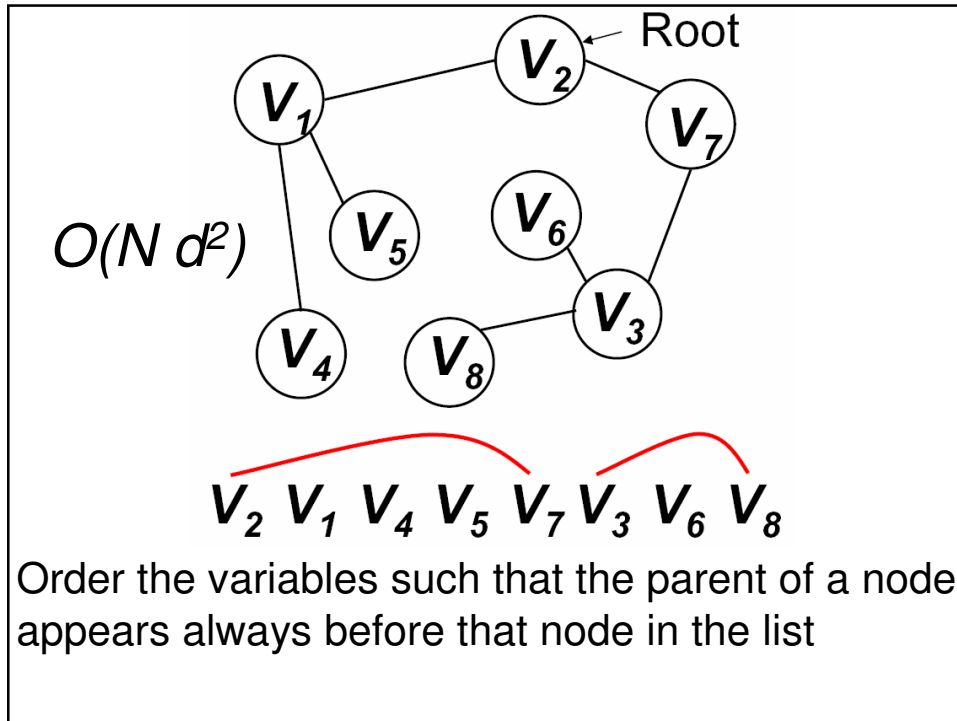
how would the behavior of simulated annealing change?

## CSP

- Definitions
- Standard search
- Improvements
  - Backtracking
  - Forward checking
  - Constraint propagation
- Heuristics:
  - Variable ordering
  - Value ordering
- Examples
- Tree-structured CSP
- Local search for CSP problems

## Constraint Propagation

- $A$  = queue of active arcs  $(V_i, V_j)$
- Repeat while  $A$  not empty:
  - $(V_i, V_j) \leftarrow$  next element of  $A$
  - For each  $x$  in  $D(V_i)$ :
    - Remove  $x$  from  $D(V_j)$  if there is no  $y$  in  $D(V_i)$  for which  $(x,y)$  satisfies the constraint between  $V_i$  and  $V_j$ .
  - If  $D(V_i)$  has changed:
    - Add all the pairs  $(V_k, V_i)$ , where  $V_k$  is a neighbor of  $V_i$  ( $k$  not equal to  $j$ ) to  $A$



- **Most Constraining Variable**
  - Selecting a variable which contributes to the *largest* number of constraints will have the largest effect on the other variables → Hopefully will prune a larger part of the search
  - This amounts to finding the variable that is connected to the largest number of variables in the constraint graph.
- **Minimum Remaining Values (MRV)**
  - Selecting the variable that has the least number of candidate values is most likely to cause a failure early (“fail-first” heuristic)
- **Least Constraining Value**
  - Choose the value which causes the smallest reduction in the number of available values for the neighboring variables




Consider the perennial problem of scheduling classrooms in Wean Hall. We have 4 instructors ( $I_1, I_2, I_3, I_4$ ) and 3 rooms ( $R_1, R_2, R_3$ ). We need to assign rooms to instructors. We assume that the instructors need the rooms at the following times:

I1: 9am to 11am  
 I2: 10am to 2pm  
 I3: 1pm to 5pm  
 I4: 1pm to 3pm


We assume that a room can be use by only one instructor at a time and that room  $R_3$  is too small for instructor  $I_1$  and that rooms  $R_2$  and  $R_3$  are too small for instructor  $I_3$ .

Variable Instantiated	I1	I2	I3	I4
<i>Initial Domains</i>	R1,R2	R1,R2 R3	R1	R1,R2 R3

## Configuration Space (C-Space)




$q = (x, y, \theta)$   
 $\mathcal{C} = \mathbb{R}^2 \times \text{set of 2-D rotations}$



$q = (q_1, q_2)$   
 $\mathcal{C} = \text{2-D rotations} \times \text{2-D rotations}$

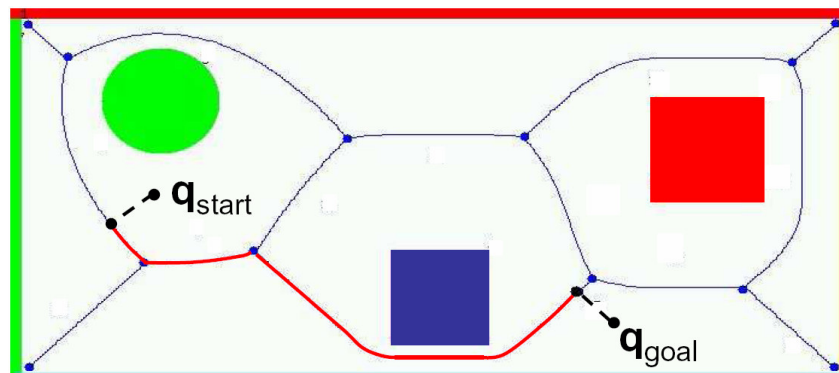
- Configuration space  $\mathcal{C}$  = set of values of  $q$  corresponding to legal configurations of the robot
- Defines the set of possible parameters (the search space) and the set of allowed paths
- Grow the forbidden parts of C-Space  $\rightarrow$  Can assume that robot is a point in C-Space

## Visibility Graphs




- $N$  = total number of vertices of the obstacle polygons
- Naïve:  $O(N^3)$
- Sweep:  $O(N^2 \log N)$
- Optimal:  $O(N^2)$

## Voronoi Diagrams




- Key property: The points on the edges of the Voronoi diagram are the *furthest* from the obstacles
- Idea: Construct a path between  $q_{start}$  and  $q_{goal}$  by following edges on the Voronoi diagram
- (Use the Voronoi diagram as a roadmap graph instead of the visibility graph)

## Cell Decomposition




- Define cells in C-Space
- Mark any cell of the grid that intersects  $\mathcal{C}_{obs}$  as blocked
- Find path through remaining cells by using (for example) A\* (e.g., use Euclidean distance as heuristic)
- Approximate  $\rightarrow$  Easy to compute but not *complete*
- Exact  $\rightarrow$  Hard to compute but complete

## Potential Fields





- Stay away from obstacles: Imagine that the obstacles are made of a material that generate a *repulsive* field
- Move closer to the goal: Imagine that the goal location is a particle that generates an *attractive* field
- Key issue: Local minima
  - Deterministic exploration of local minima
  - Stochastic exploration
- In very high dimension  $\rightarrow$  Randomize traversal of neighbors

# Sampling Techniques

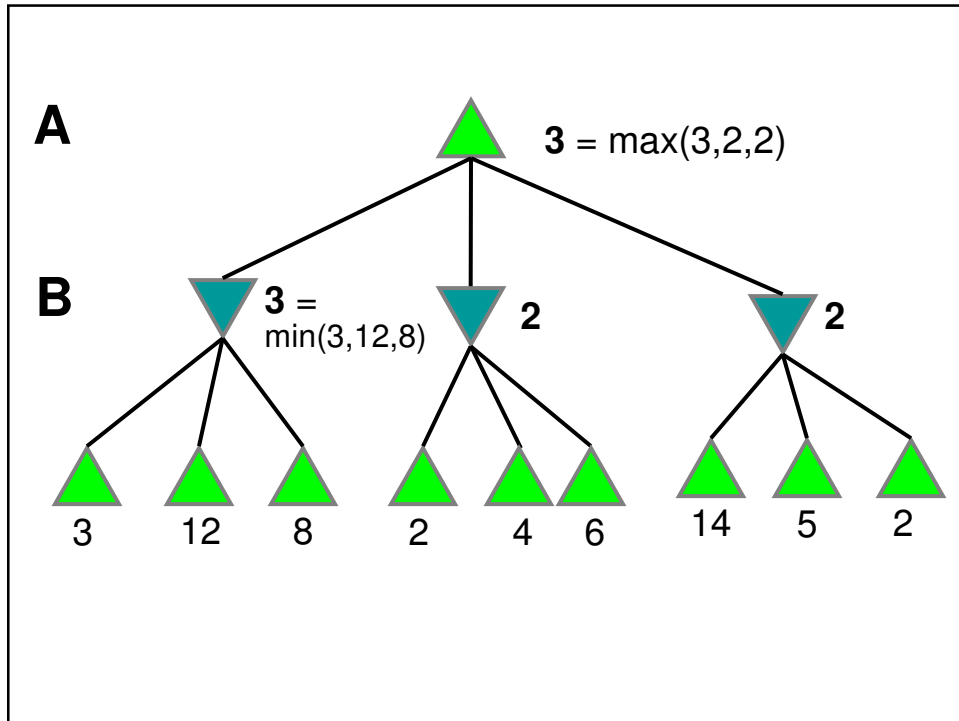


- Tends to explore the space rapidly in all directions
- Does not require extensive pre-processing
- Single query/multiple query problems
- Needs only collision detection test → No need to represent/pre-compute the entire C-space
- For a large class of problems:
  - Prob(finding a path) → 1 exponentially with the number of samples
- *But*, cannot detect that a path does not exist





Games



## Minimax

Minimax ( $s$ )

If  $s$  is terminal

Return  $U(s)$

If next move is  $A$

Return  $\max_{s' \in \text{Succs}(s)} \text{Minimax}(s')$

Else

Return  $\min_{s' \in \text{Succs}(s)} \text{Minimax}(s')$

## Minimax Properties

- *Complete*: If finite game
- *Optimal*: If opponent plays optimally
- *Complexity*: Essentially DFS, so:
  - Time:  $O(B^m)$
  - Space:  $O(Bm)$
  - $B$  = number of possible moves from any state (branching factor)
  - $m$  = depth of search (length of game)
- Pruning ( $\alpha\beta$ ):
  - Guaranteed to find same solution
  - $O(B^{m/2})$  with proper ordering of the nodes  $\rightarrow$  At “A” node, the successor are in order from high to low score

## Non-Deterministic Minimax

Minimax ( $s$ )

If  $s$  is terminal

Return  $U(s)$

If next move is  $A$ : Return

$$\max_{s' \in \text{Succs}(s)} \text{Minimax}(s')$$

If next move is  $B$  Return

$$\min_{s' \in \text{Succs}(s)} \text{Minimax}(s')$$

If chance node Return

$$\sum_{s' \in \text{Succs}(s)} p(s') \text{Minimax}(s')$$

Max value =  
game value = +2

Note that we find the same value and same strategies in both cases. Is that always the case?

	I	II	III	IV	
-1	I	-1	-1	+2	+2
+2	II	+4	+4	+2	+2
+1	III	+5	+1	+5	+1
+1	IV	+5	+1	+5	+1

Min value across each row

$\text{Max}_{\text{Rows } i} \text{Min}_{\text{Columns } j} M(i, j)$

	I	II	III	IV
I	-1	-1	+2	+2
II	+4	+4	+2	+2
III	+5	+1	+5	+1
IV	+5	+1	+5	+1

Max value across each column

+5 +4 +5 +2

Min value =  
game value = +2

$\text{Min}_{\text{Columns } j} \text{Max}_{\text{Rows } i} M(i, j)$

## Minimax vs. Maximin

- Fundamental Theorem I (von Neumann):
  - For a two-player, zero-sum game with perfect information:
    - *There always exists an optimal pure strategy for each player*
    - *Minimax = Maximin*
- *Note: This is a game-theoretic formalization of the minimax search algorithm that we studied earlier.*




# Minimax with Mixed Strategies

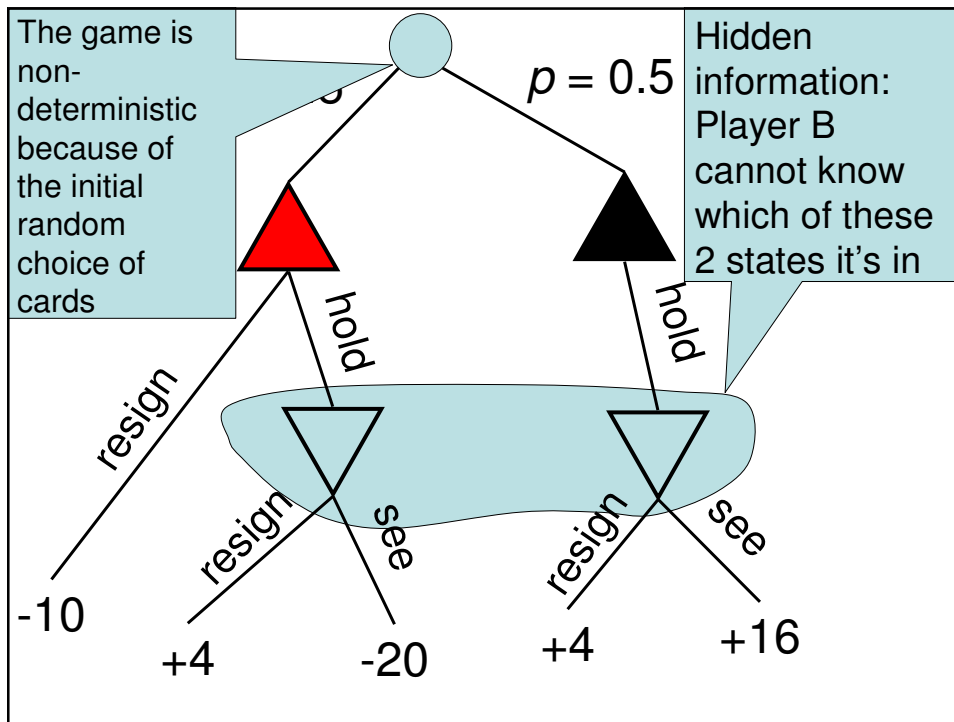
- Theorem II (von Neumann):
  - For a two-player, zero-sum game with hidden information:
    - *There always exists an optimal **mixed strategy***
    - *In addition, just like for games with perfect information, it does not matter in which order we look at the players, minimax is the same as maximin*

$$\max_p \min(p \times m_{11} + (1-p) \times m_{21}, p \times m_{12} + (1-p) \times m_{22}) =$$

$$\min_q \max(q \times m_{11} + (1-q) \times m_{12}, q \times m_{21} + (1-q) \times m_{22})$$

- *Note: This is a direct generalization of the minimax result to mixed strategies.*





**Player B**

↔

	Resign	See
<b>Player A</b>	Resign	
	Hold	

↕

- Generate the matrix form of the game (be careful: It's not a deterministic game)
- Find the optimal mixed strategy
- Find the expected payoff for Player A

Strategies

<b>I</b>	<b>3,0</b>	<b>4,1</b>	<b>5,9</b>	<b>5,6</b>
<b>II</b>	<b>5,3</b>	<b>5,8</b>	<b>9,7</b>	<b>9,0</b>
<b>III</b>	<b>2,3</b>	<b>8,4</b>	<b>6,2</b>	<b>6,3</b>
<b>IV</b>	<b>3,8</b>	<b>3,1</b>	<b>2,3</b>	<b>4,5</b>

## Pure Strategy Nash Equilibrium

$$u_i(s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_n^*) \leq u_i(s_1^*, \dots, s_{i-1}^*, s_i^*, s_{i+1}^*, \dots, s_n^*)$$

$$s_i^* = \arg \max_{s_i} u_i(s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_n^*)$$

- Does not always exist
- Is not always unique
- For continuous games:

$$\frac{\partial u_i}{\partial s_i}(s_1^*, \dots, s_n^*) = 0$$

## Mixed Strategy Nash Equilibrium

$$\bar{u}_A(p, q^*) \leq \bar{u}_A(p^*, q^*)$$

$$p^* = \arg \max_p \bar{u}_A(p, q^*)$$

- Player A plays strategy  $s_i$  with probability  $p_i$
- For non-zero games  $\rightarrow$  Mixed Nash equilibrium *always* exists
- Is not always unique

$t_B = \text{meet}$		<b>Hockey</b>	<b>Movie</b>
	<b>Hockey</b>	<b>+2,+1</b>	<b>0,0</b>
	<b>Movie</b>	<b>0,0</b>	<b>+1,+2</b>
$t_B = \text{avoid}$		<b>Hockey</b>	<b>Movie</b>
	<b>Hockey</b>	<b>+2,0</b>	<b>0,+2</b>
	<b>Movie</b>	<b>0,+1</b>	<b>+1,0</b>

$P(t_A = \text{meet} \mid t_B = \text{meet}) = 1$     $P(t_B = \text{meet} \mid t_A = \text{meet}) = 1/2$   
 $P(t_A = \text{meet} \mid t_B = \text{avoid}) = 1$     $P(t_B = \text{meet} \mid t_A = \text{avoid}) = 1/2$   
 $P(t_A = \text{avoid} \mid t_B = \text{meet}) = 0$     $P(t_B = \text{avoid} \mid t_A = \text{meet}) = 1/2$   
 $P(t_A = \text{avoid} \mid t_B = \text{avoid}) = 0$     $P(t_B = \text{avoid} \mid t_A = \text{avoid}) = 1/2$

Payoff if Player A knows that Player B is of type  $t_B$

Probability that Player B is indeed of type  $t_B$

$$\bar{u}_A = \sum_{\substack{\text{all possible types} \\ t_B \text{ of Player B}}} u_A(s_A(t_A), s_B(t_B)) P(t_B | t_A)$$

Since Player A does not know Player B's type, it has to sum over all possible types to get the expected value

		D	E	F
Player 1	A	0, 1	3, 5	2, 1
	B	6, 3	1, 3	5, 2
	C	4, 2	3, 4	7, 7

		Player 2	
		C	D
Player 1	A	2	0
	B	0	1

Now consider a very different game. Two companies, A and B, both make elbow warmers. The more they spend on advertising, the more sales they get, but there are diminishing returns. A's advertising somewhat helps B, and B's advertising somewhat helps A. the exact fomulas are

$$\begin{aligned}
 P_a &= \overbrace{\log(2a + b)}^{\text{revenue}} - \overbrace{a}^{\text{expense}} \\
 P_b &= \log(a + 2b) - b
 \end{aligned}$$

where

- a = # dollars that A spends on advertising
- b = # dollars that B spends on advertising
- $P_a$  = profit to A
- $P_b$  = profit to B