# Games & Adversarial Search B: Alpha-Beta Pruning and MCTS

CS171, Summer 1 Quarter, 2019

Introduction to Artificial Intelligence

Prof. Richard Lathrop

Read Beforehand: R&N 5.3; Optional: 5.5+

UNIVERSITY of CALIFORNIA IRVINE

INFORMATION AND COMPUTER SCIENCES

Alpha-Beta pruning

- Exploit the "fact" of an adversary

- <u>Bad = not better</u> than we already know we can get elsewhere
- If a position is provably bad
  - It's NO USE expending search effort to find out just how bad it is
- If the adversary can force a bad position
  - It's NO USE searching to find the good positions the adversary won't let you achieve anyway

- Contrast normal search:
  - ANY node might be a winner, so ALL nodes must be considered.
  - A* avoids this through heuristics that transmit your knowledge.
  - Alpha-Beta pruning avoids this through exploiting the adversary.
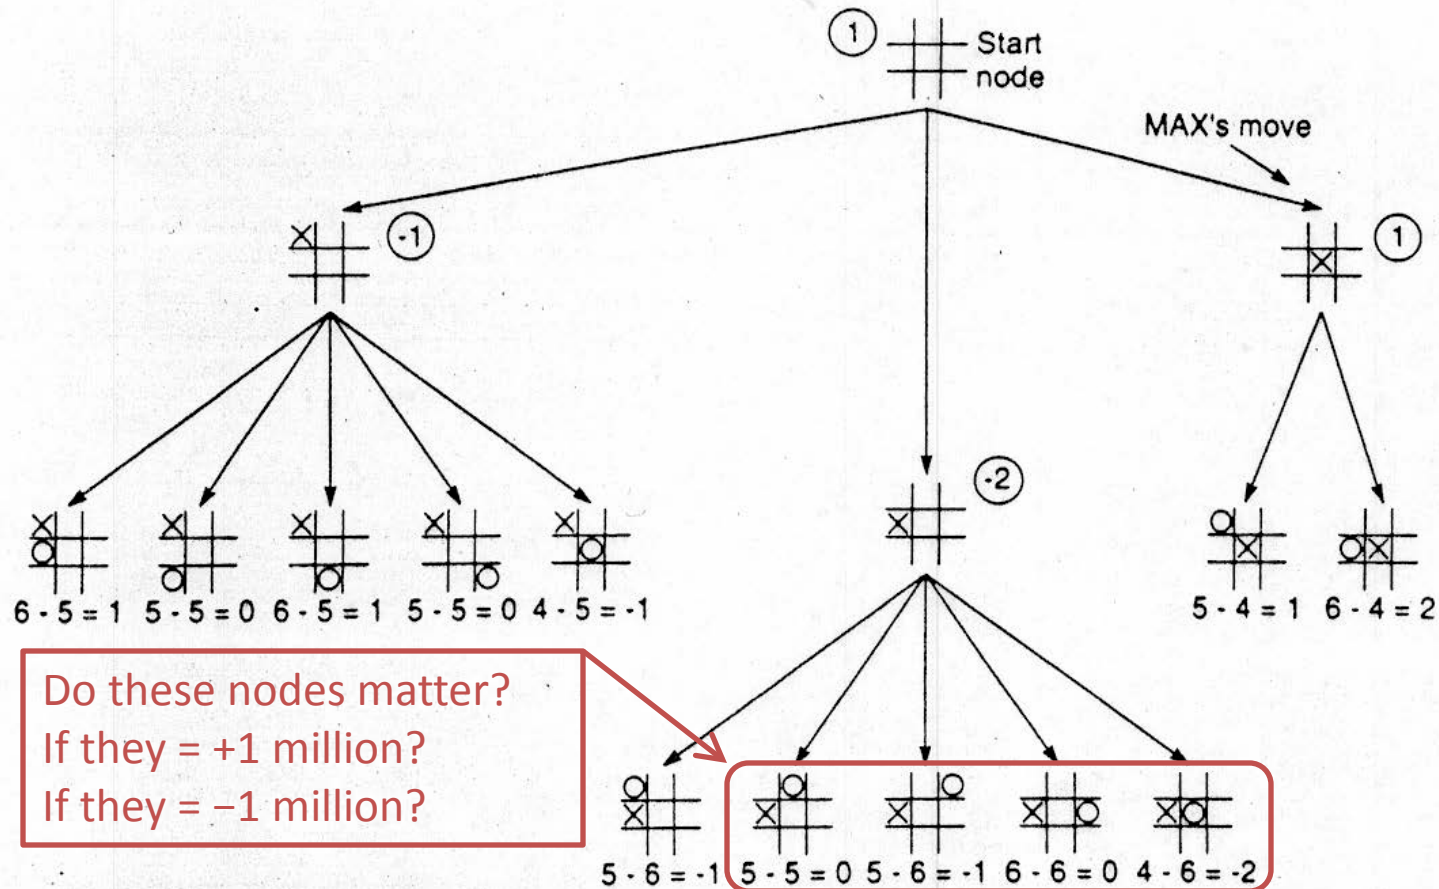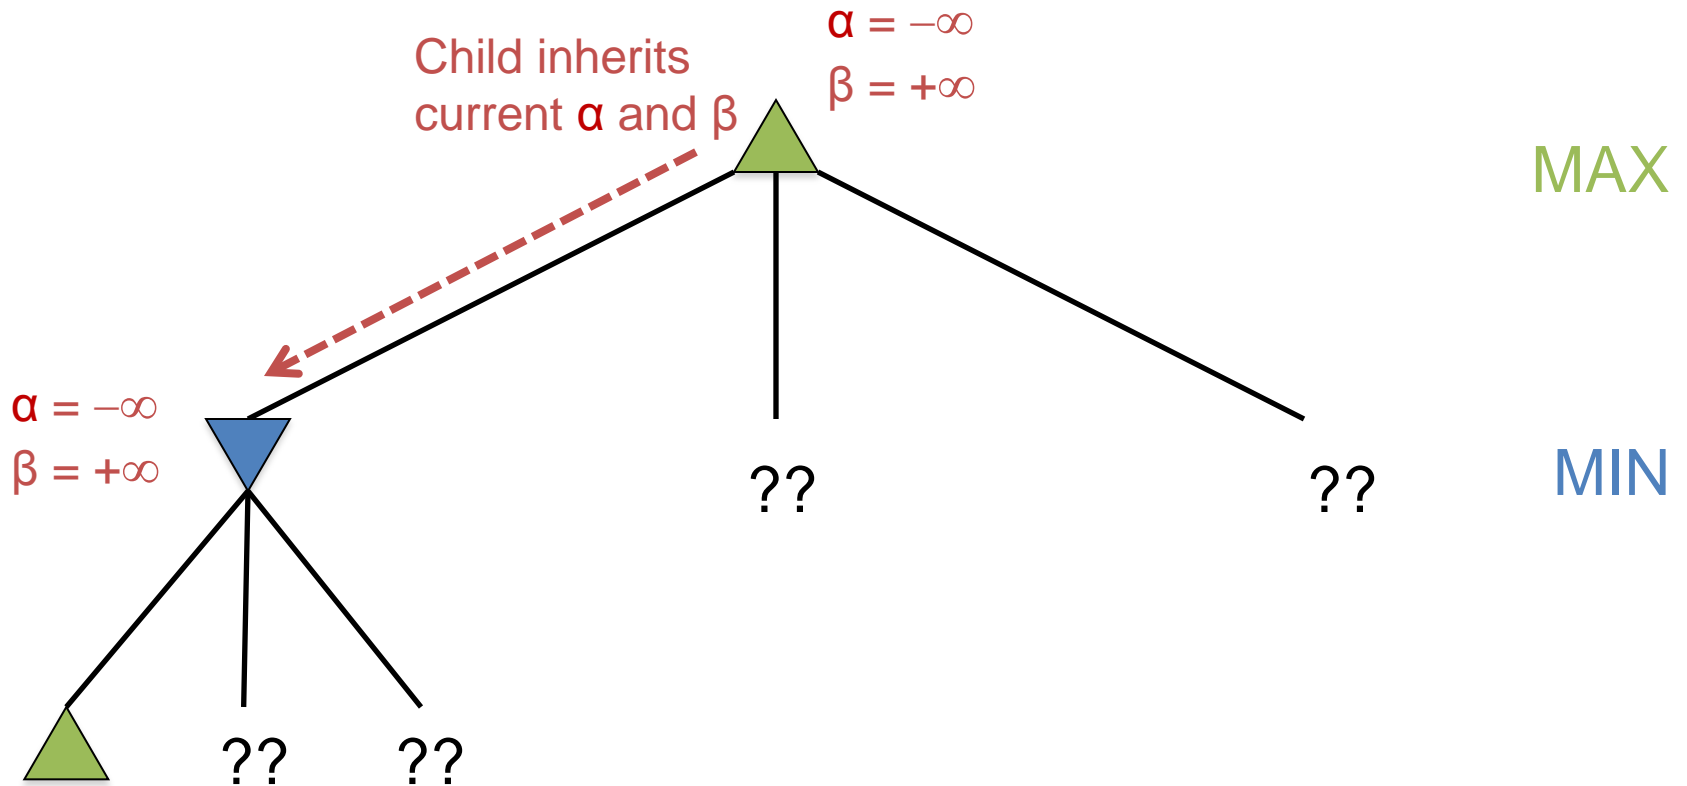
Pruning with Alpha/Beta

Figure 4.17 Two-ply minimax applied to the opening move of tic-tac-toe.
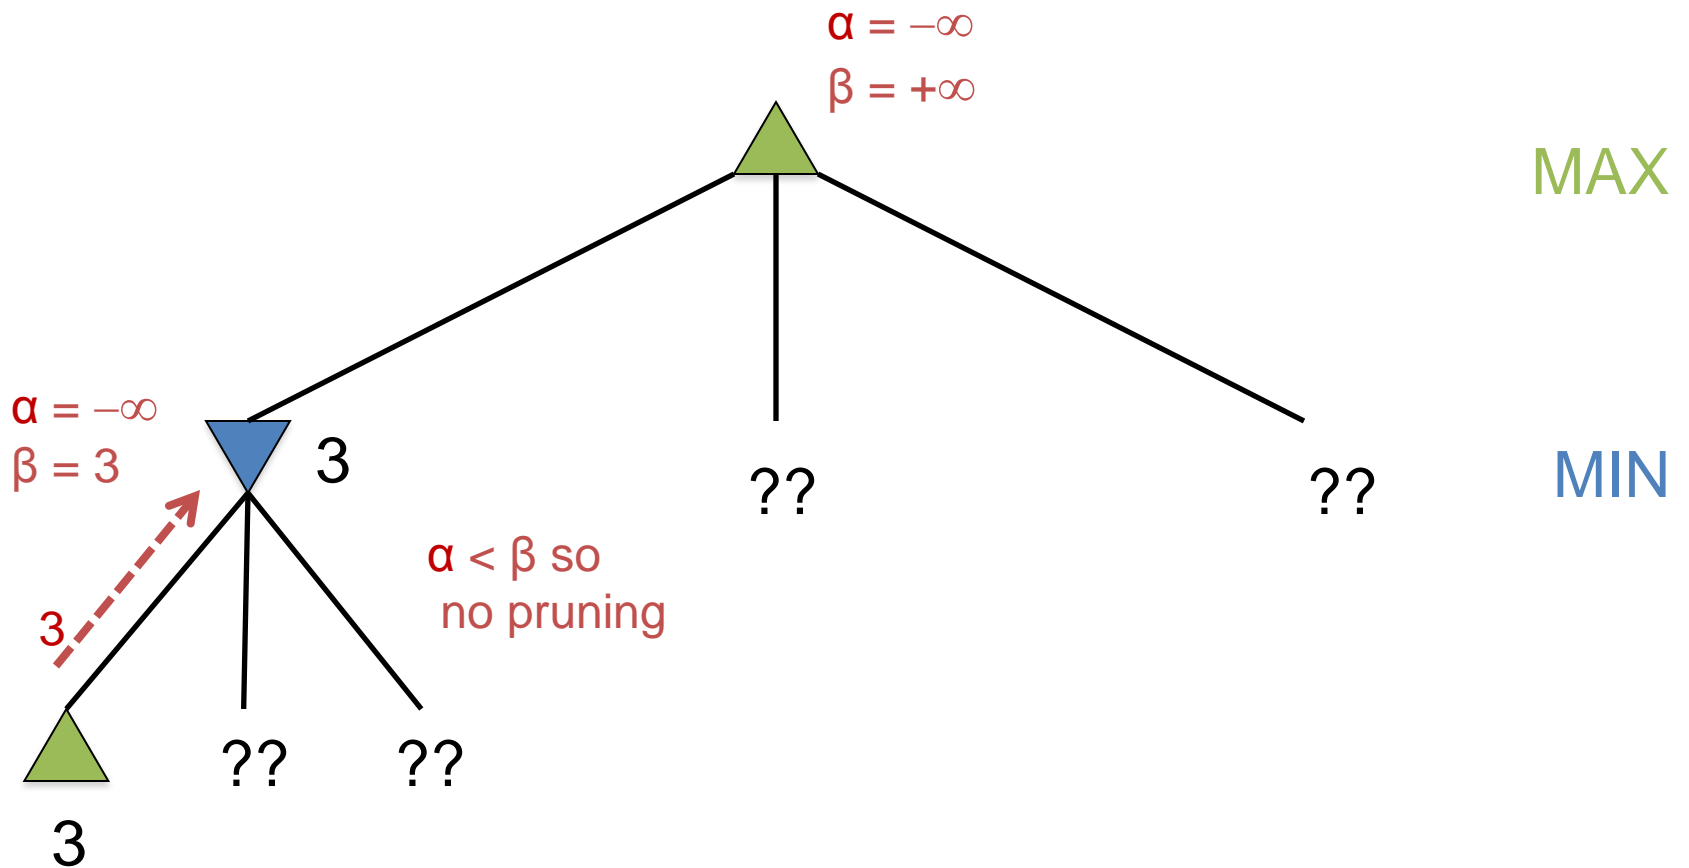
Alpha-Beta Example

Initially, possibilities are unknown: range (α =-∞, β=+∞)
Do a depth-first search to the first leaf.

Alpha-Beta Example

See the first leaf, after MIN's move: MIN updates β

α = −∞
β = +∞

MAX

α = −∞
β = 3   3

??   ??

MIN

α < β so
no pruning

3

3   ??   ??

Alpha-Beta Example

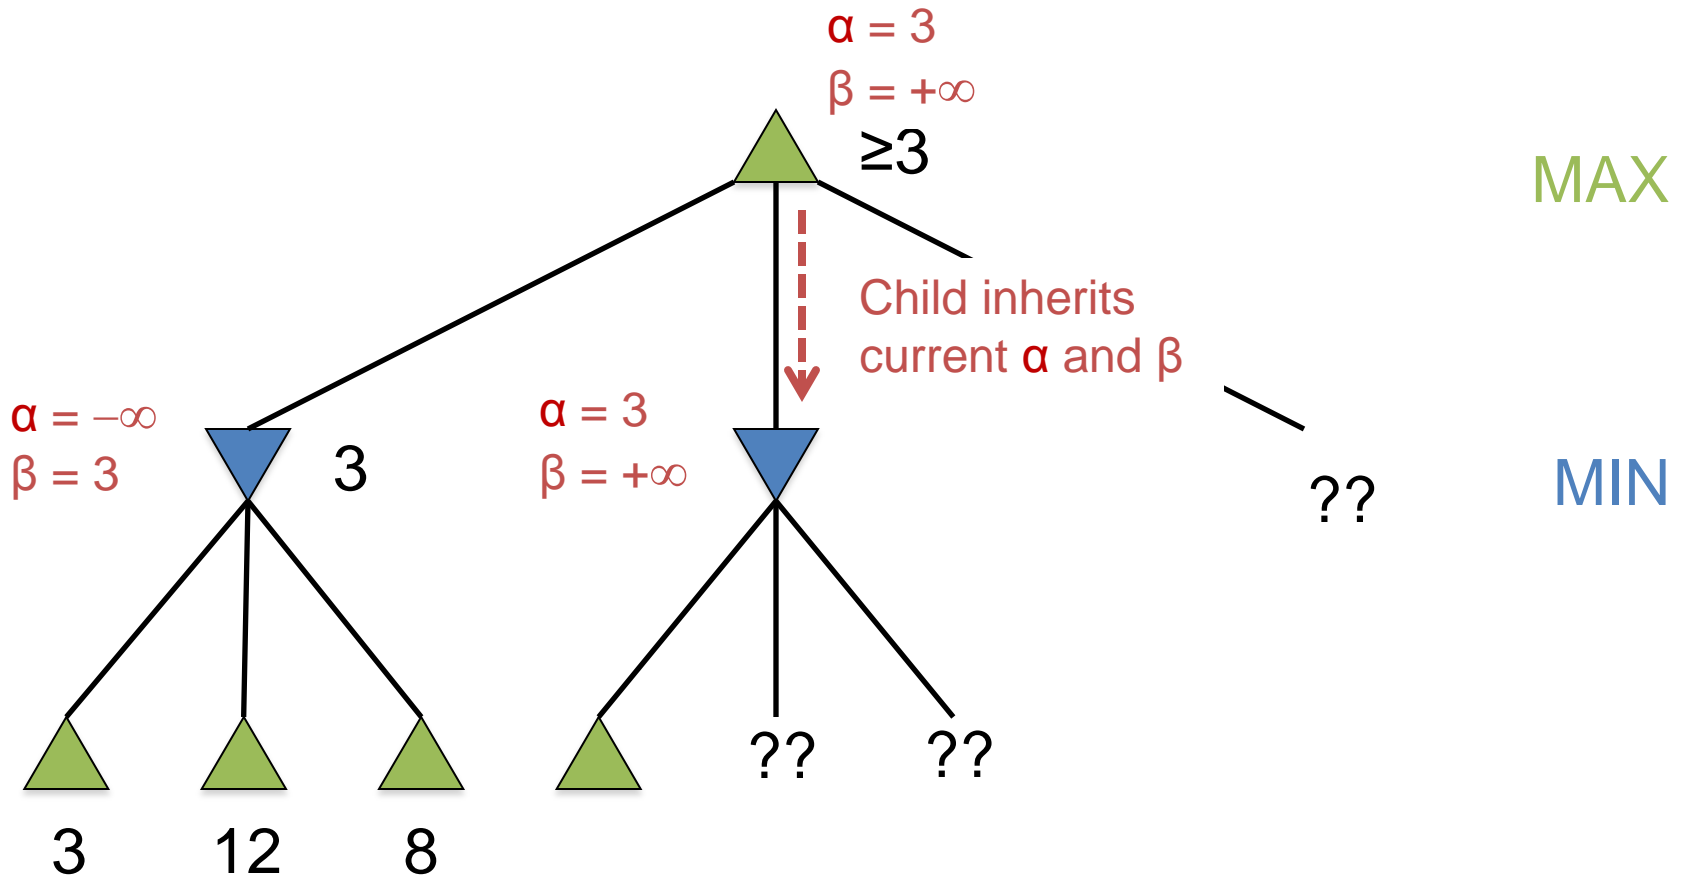See remaining leaves; value is known

Pass outcome to caller; MAX updates α

Alpha-Beta Example

Continue depth-first search to next leaf.
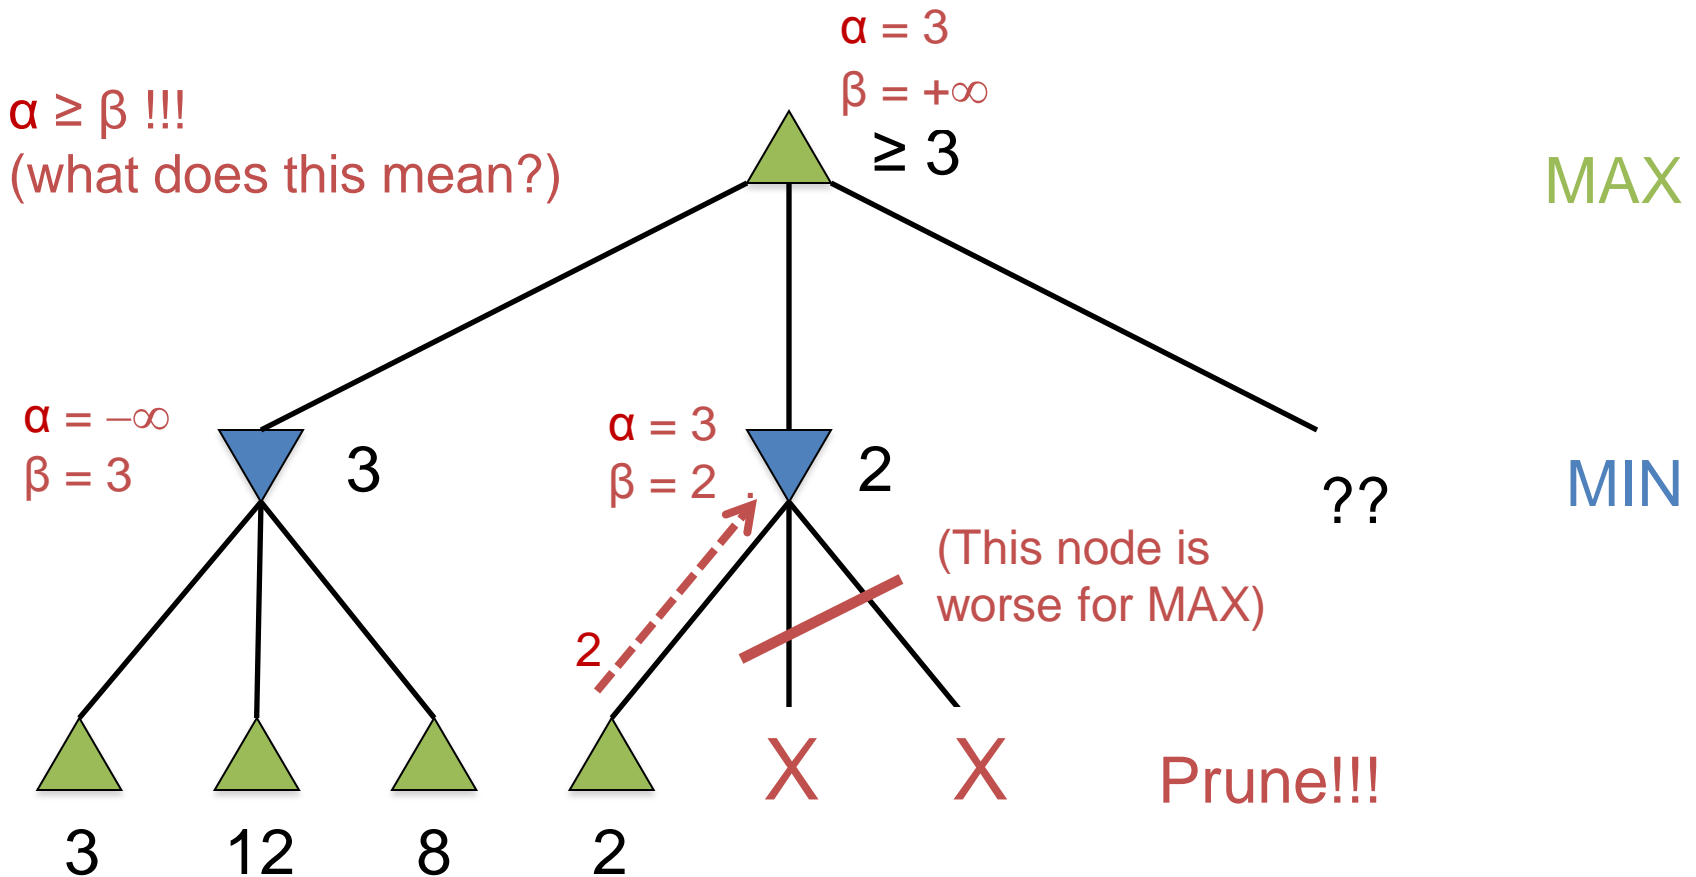
Pass α, β to descendants

Alpha-Beta Example

Observe leaf value; MIN's level; MIN updates β

Prune – play will never reach the other nodes!

α ≥ β !!!
(what does this mean?)

α = 3
β = +∞
≥ 3

MAX

α = −∞
β = 3

3

α = 3
β = 2

2

??

MIN

2

(This node is worse for MAX)

X   X   Prune!!!

3   12   8   2

Alpha-Beta Example

Pass outcome to caller & update caller:

Alpha-Beta Example

Continue depth-first exploration…

No pruning here; value is not resolved until final leaf.

Alpha-Beta Example

Pass outcome to caller & update caller.
Value at the root is resolved.

α = 3
β = +∞
3

MAX level, 3 > 2
α no change

MAX

α = −∞
β = 3

α = 3
β = 2
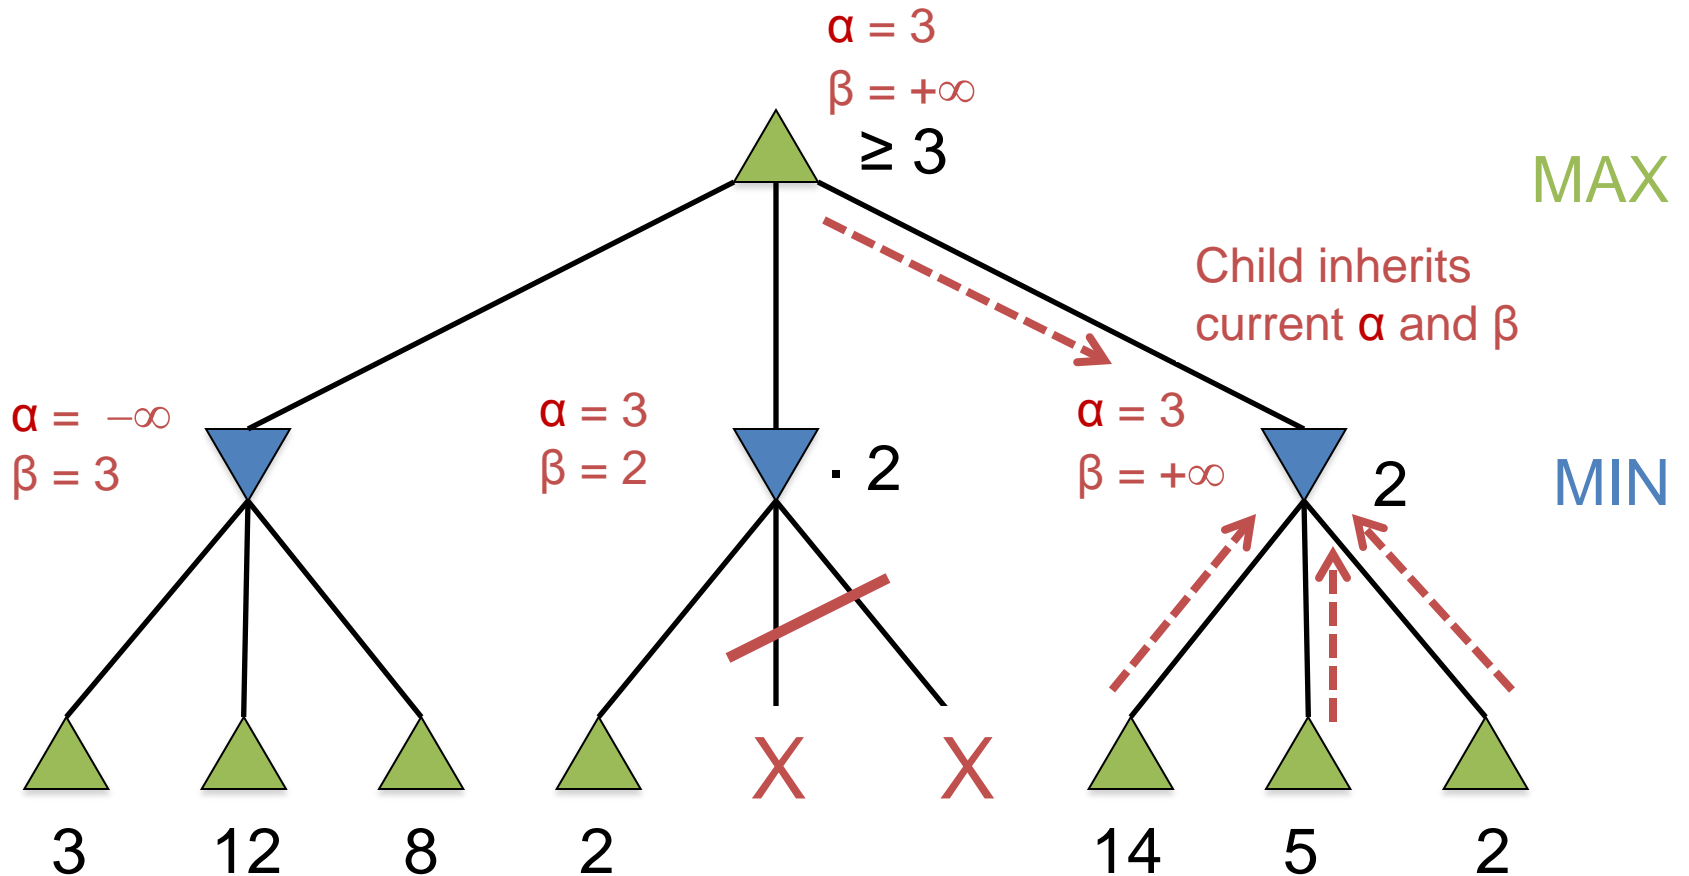· 2

α = 3
β = 2
2

MIN

2

3   12   8   2   X   X   14   5   2

General alpha-beta pruning

- Consider a node n in the tree:

- If player has a better choice at
  - Parent node of n
  - Or, any choice further up!

- Then n is never reached in play

- So:
  - When that much is known about n, it can be pruned

Player

Opponent   *m*

..
..
..

Player

Opponent   *n*

Recursive α-β pruning (expands on Fig. 5.7)

Alpha-Beta-Search(state)
  alpha = −infty, beta = +infty, act = None
  for each a in Actions(state) do
    val = Min-Value(Result(state, a), alpha, beta)
    if (val > alpha) then alpha = val, act = a
  return act

Simple stub to call recursion functions
Initialize alpha, beta; no move found
Score each action; update alpha & best action

MaxValue(state, al, be)
  if (Cutoff(state)) then return Eval(state)
  val = −infty
  for each a in Actions(state) do
    val = max(val, MinValue(Result(state, a), al, be)
    if (val $\geq$ be) then return val
    al = max(al, val)
  return val

If Cutoff reached, return Eval heuristic
Otherwise, find our best child:
If our options are too good, our min
 ancestor will never let us come this way
Otherwise return the best we can find

MinValue(state, al, be)
  if (Cutoff(state)) then return Eval(state)
  val = +infty
  for each a in Actions(state) do
    val = min(val, MaxValue(Result(state, a), al, be)
    if (val $\leq$ al) then return val
    be = min(be, val)
  return be

If Cutoff reached, return Eval heuristic
Otherwise, find the worst child:
If our options are too bad, our max
 ancestor will never let us come this way
Otherwise return the worst we can find

Effectiveness of α-β Search

- Worst-Case
  - Branches are ordered so that no pruning takes place. In this case alpha-beta gives no improvement over exhaustive search

- Best-Case
  - Each player's best move is the left-most alternative (i.e., evaluated first)
  - In practice, performance is closer to best rather than worst-case

- In practice often get $O(b^{(d/2)})$ rather than $O(b^d)$
  - This is the same as having a branching factor of sqrt(b),
    - since $(sqrt(b))^d = b^{(d/2)}$ (i.e., we have effectively gone from b to square root of b)
  - In chess go from b ~ 35 to b ~ 6
    - permiting much deeper search in the same amount of time

Iterative deepening

- In real games, there is usually a time limit T to make a move

- How do we take this into account?
- Minimax cannot use "partial" results with any confidence, unless the full tree has been searched
  - Conservative: set small depth limit to guarantee finding a move in time < T
  - But, we may finish early – could do more search!

- Added benefit with Alpha-Beta Pruning:
  - Remember node values found at the previous depth limit
  - Sort current nodes so that each player's best move is left-most child
  - Likely to yield good Alpha-Beta Pruning => better, faster search
  - Only a heuristic: node values will change with the deeper search
  - Usually works well in practice
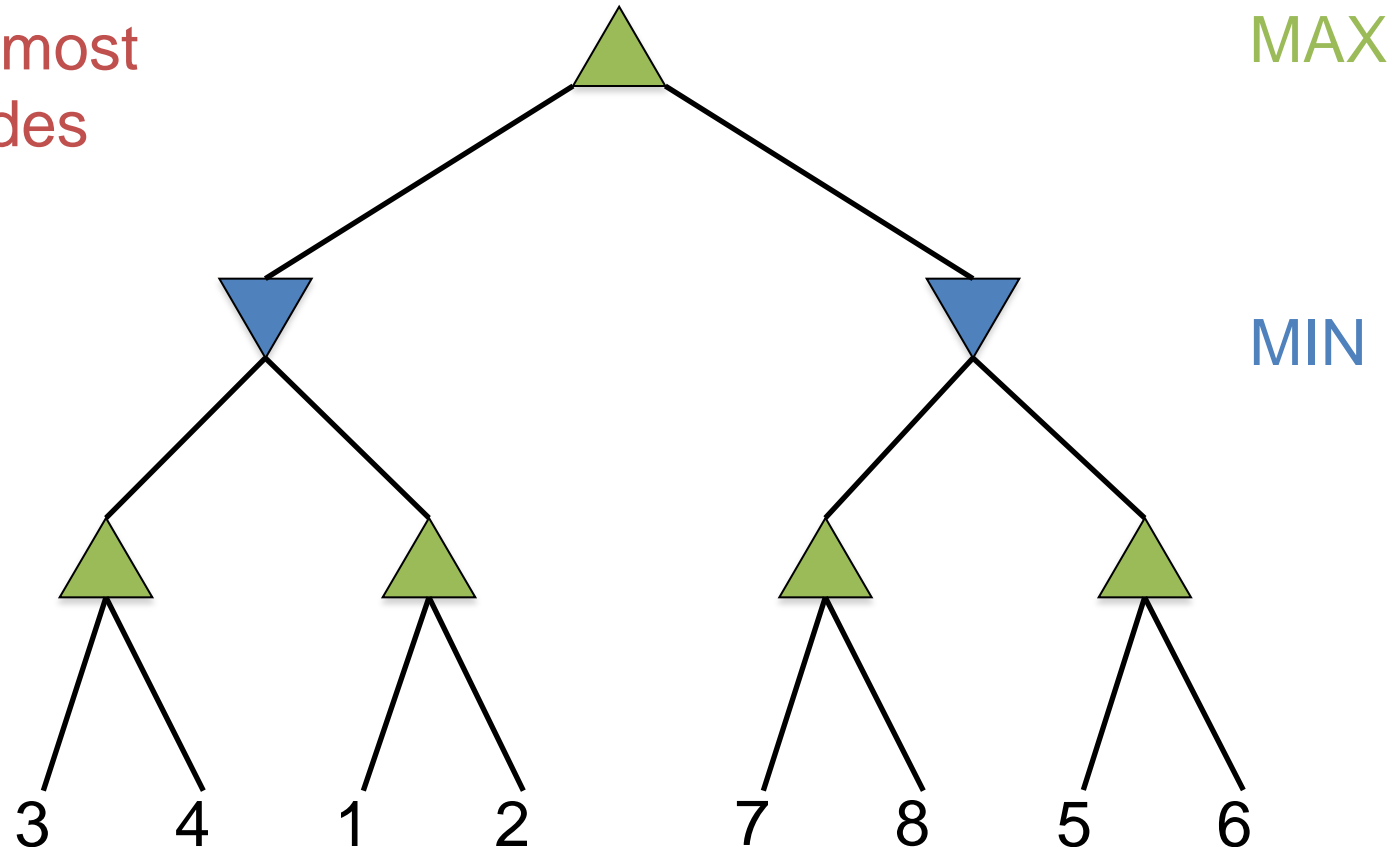
Comments on alpha-beta pruning

- Pruning does not affect final results

- Entire subtrees can be pruned

- Good move ordering improves pruning
  - Order nodes so player's best moves are checked first

- Repeated states are still possible
  - Store them in memory = transposition table

Iterative deepening reordering

Which leaves can be pruned?

None!

because the most favorable nodes are explored last…
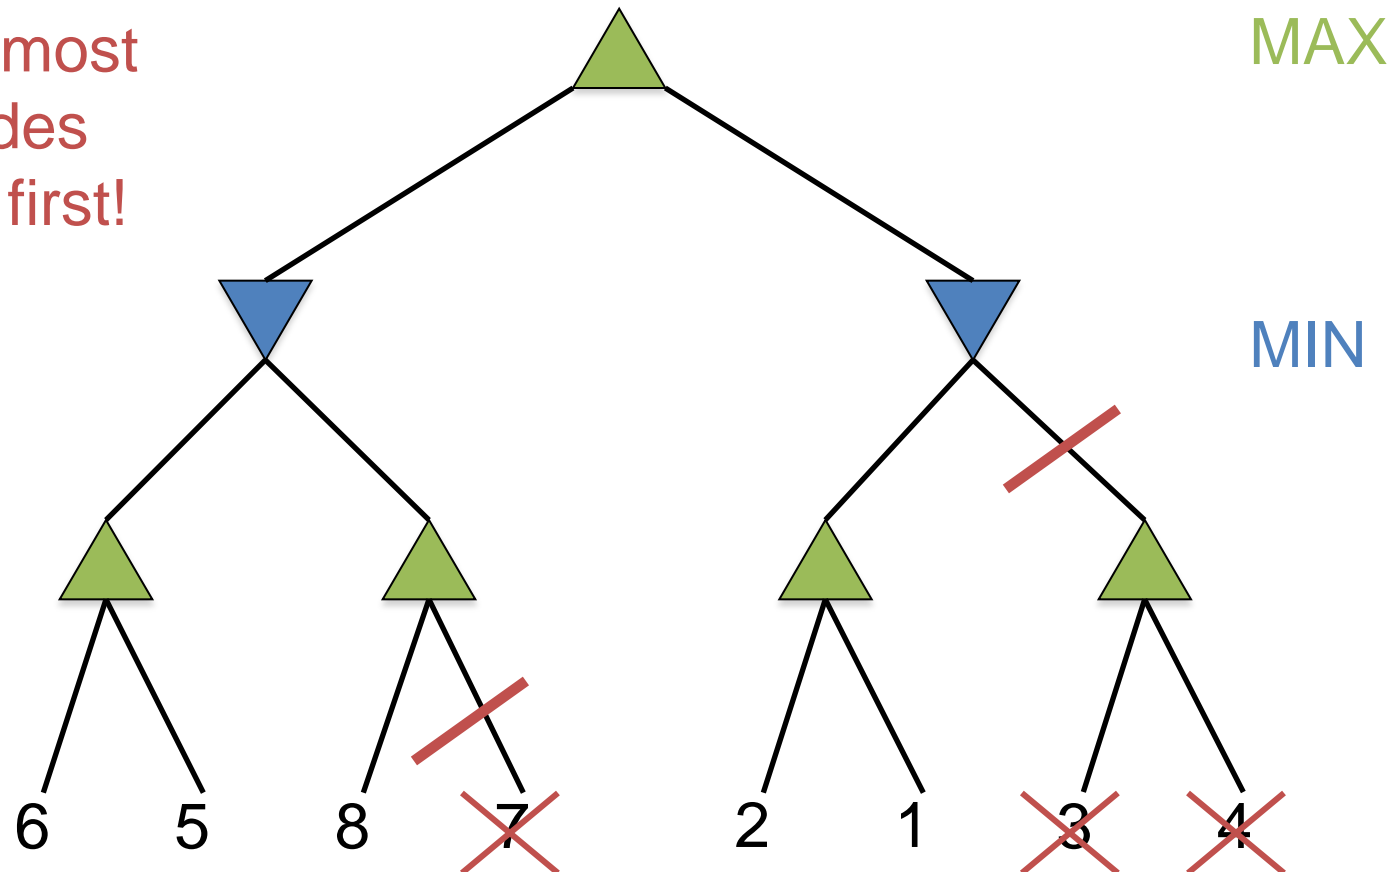
MAX

MIN

3   4   1   2   7   8   5   6

Iterative deepening reordering

Different exploration order: now which leaves can be pruned?

Iterative deepening reordering

Order with no pruning; use iterative deepening approach.
Assume node score is the average of leaf values below.

Iterative deepening reordering

Order with no pruning; use iterative deepening approach.
Assume node score is the average of leaf values below.

MAX

For L=2,
switch the order of
these nodes!

6.5

MIN

L=1   2.5   6.5

3   4   1   2   7   8   5   6
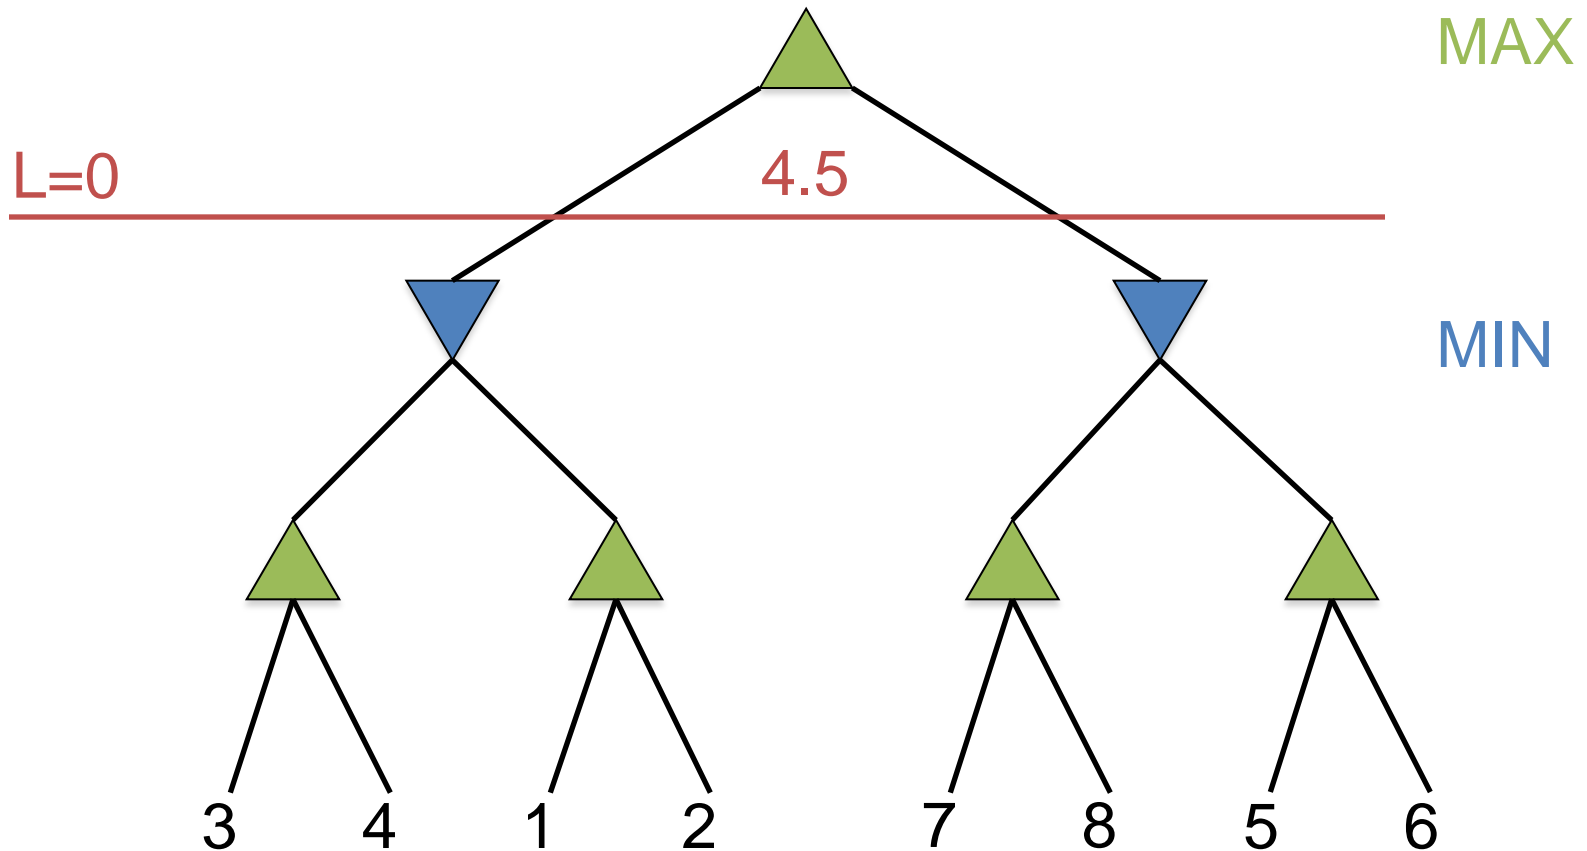
Iterative deepening reordering

Order with no pruning; use iterative deepening approach.
Assume node score is the average of leaf values below.

For L=2,
switch the order of
these nodes!

MAX

MIN

6.5

L=1   6.5   2.5

7   8   5   6   3   4   1   2
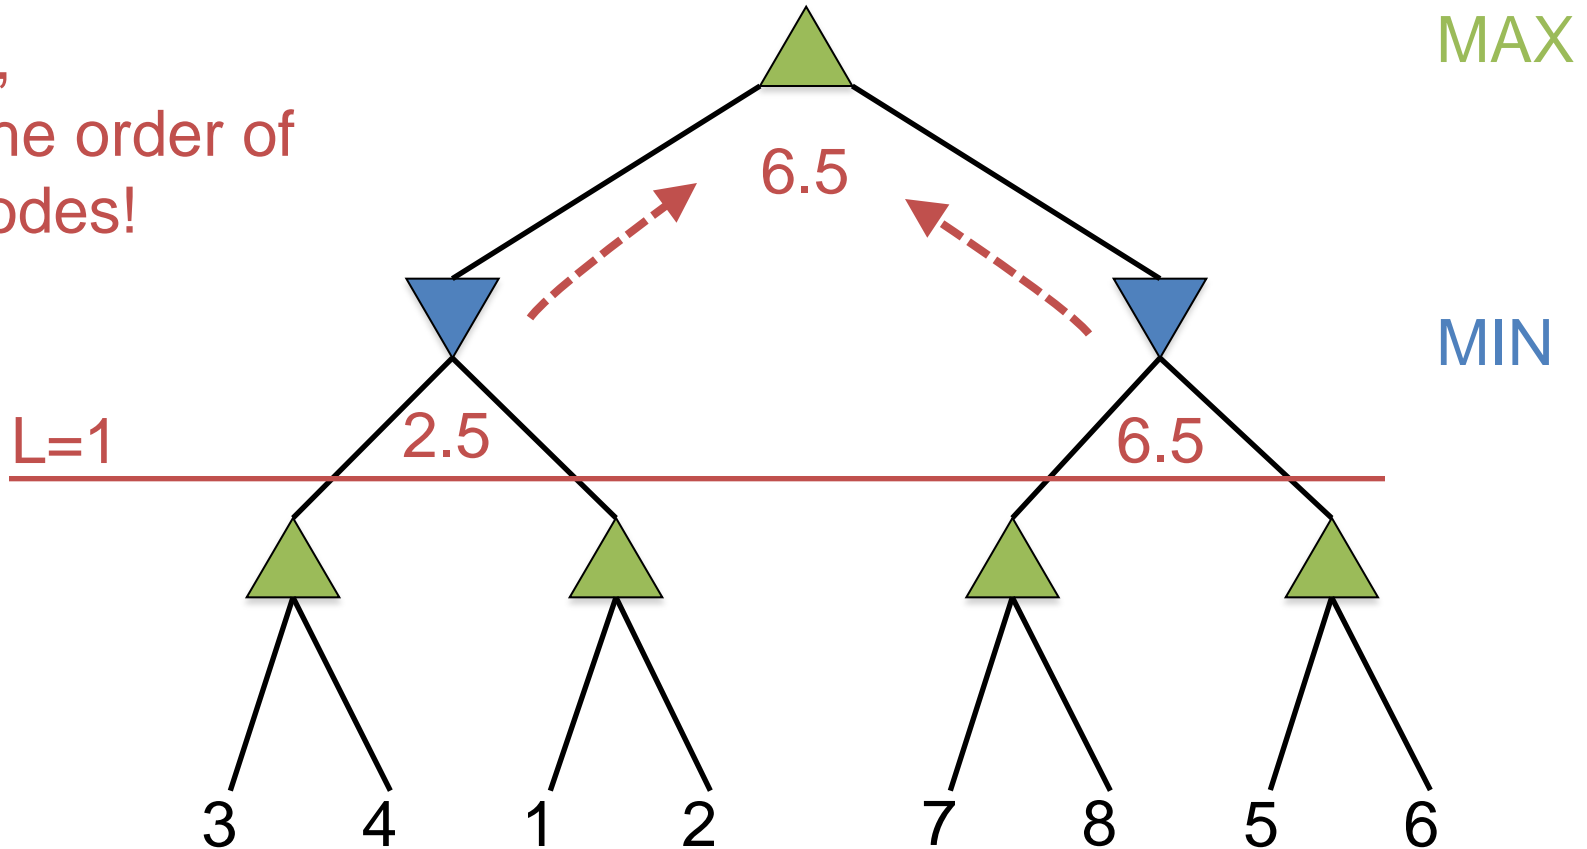
Iterative deepening reordering

Order with no pruning; use iterative deepening approach. Assume node score is the average of leaf values below.

Iterative deepening reordering

Order with no pruning; use iterative deepening approach.
Assume node score is the average of leaf values below.

Iterative deepening reordering

Order with no pruning; use iterative deepening approach. Assume node score is the average of leaf values below.

MAX

Lots of pruning!
The most favorable nodes are explored earlier.

MIN

6

6

4

6

7

4

5   6   7   8   3   4   1   2

L=3

Longer Alpha-Beta Example

Branch nodes are labelel A..K for easy discussion

$\alpha, \beta, \text{initial values} \longrightarrow \alpha = -\infty$

$\beta = +\infty$

MAX

A   B   C   D   MIN

E   F   G   H   I   J   K   MAX

9   4   9   6

4 5 6 6 1 9 5 4 1 3 2 6 1 3 6 8 6 3

Longer Alpha-Beta Example

Note that cut-off occurs at different depths…

current α, β, passed to kids

$\alpha=-\infty$
$\beta=+\infty$

$\alpha=-\infty$
$\beta=+\infty$
kid=A

$\alpha=-\infty$
$\beta=+\infty$
kid=E

MAX

A   B   C   D   MIN

E   F   G   H   I   J   K   MAX

9   4   9   6

4 5 6 6 1 9 5 4 1 3 2 6 1 3 6 8 6 3

Longer Alpha-Beta Example

*see first leaf,
MAX updates α*

$\alpha=-\infty$
$\beta=+\infty$

MAX

$\alpha=-\infty$
$\beta=+\infty$
kid=A

$\alpha=4$
$\beta=+\infty$
kid=E

A   **B**   **C**   **D**   **MIN**

E   **F**   **G**   **H**   **I**   **J**   **K**   **MAX**

4

4

9   4   9   6

4 5 6 6 1 9 5 4 1 3 2 6 1 3 6 8 6 3
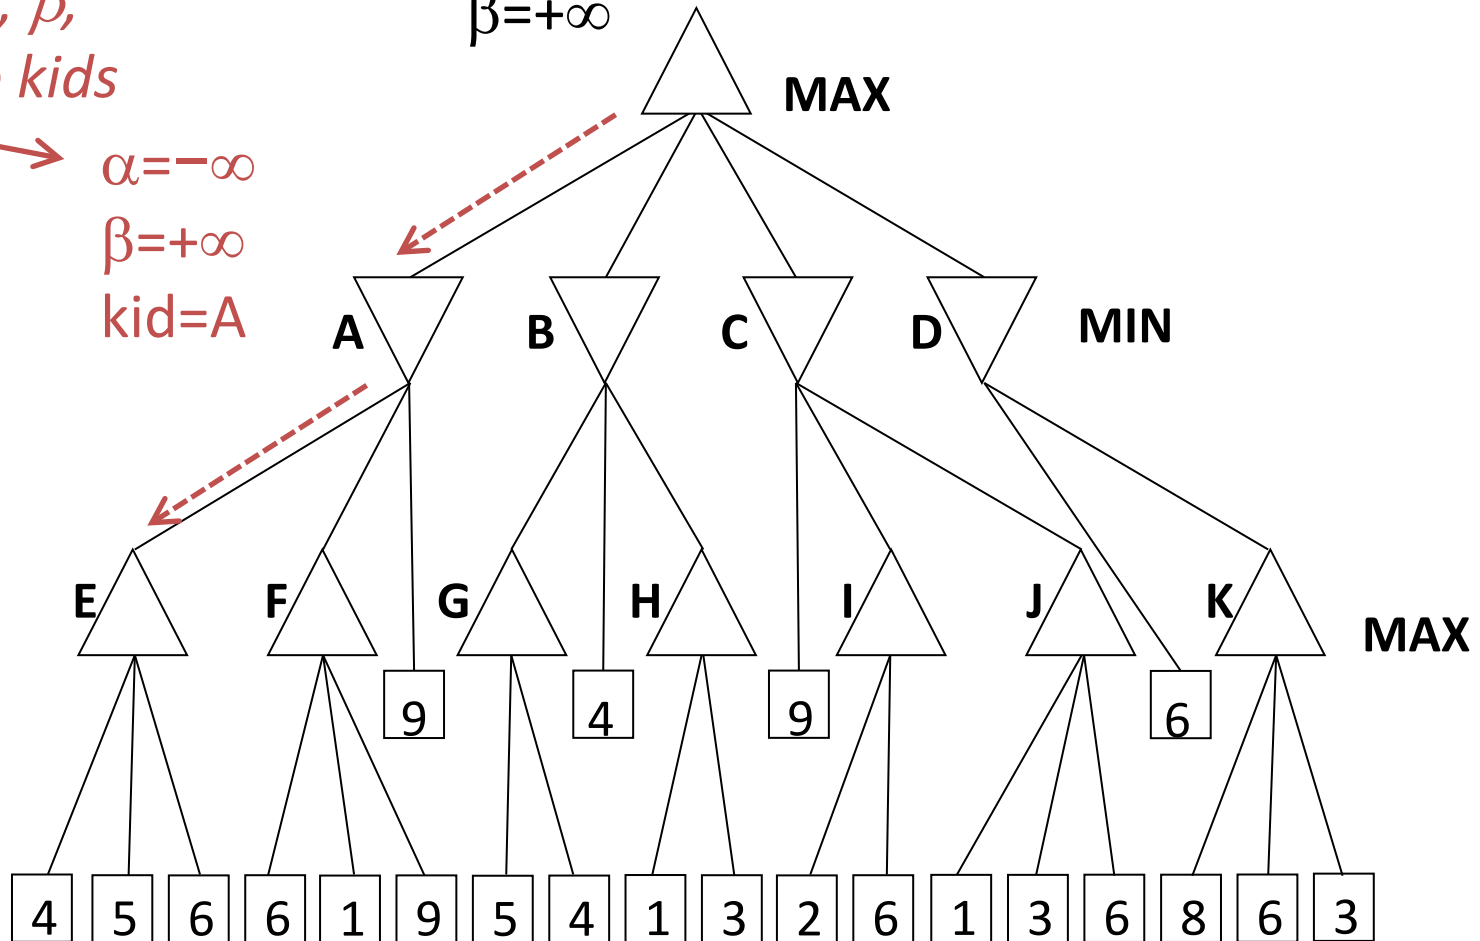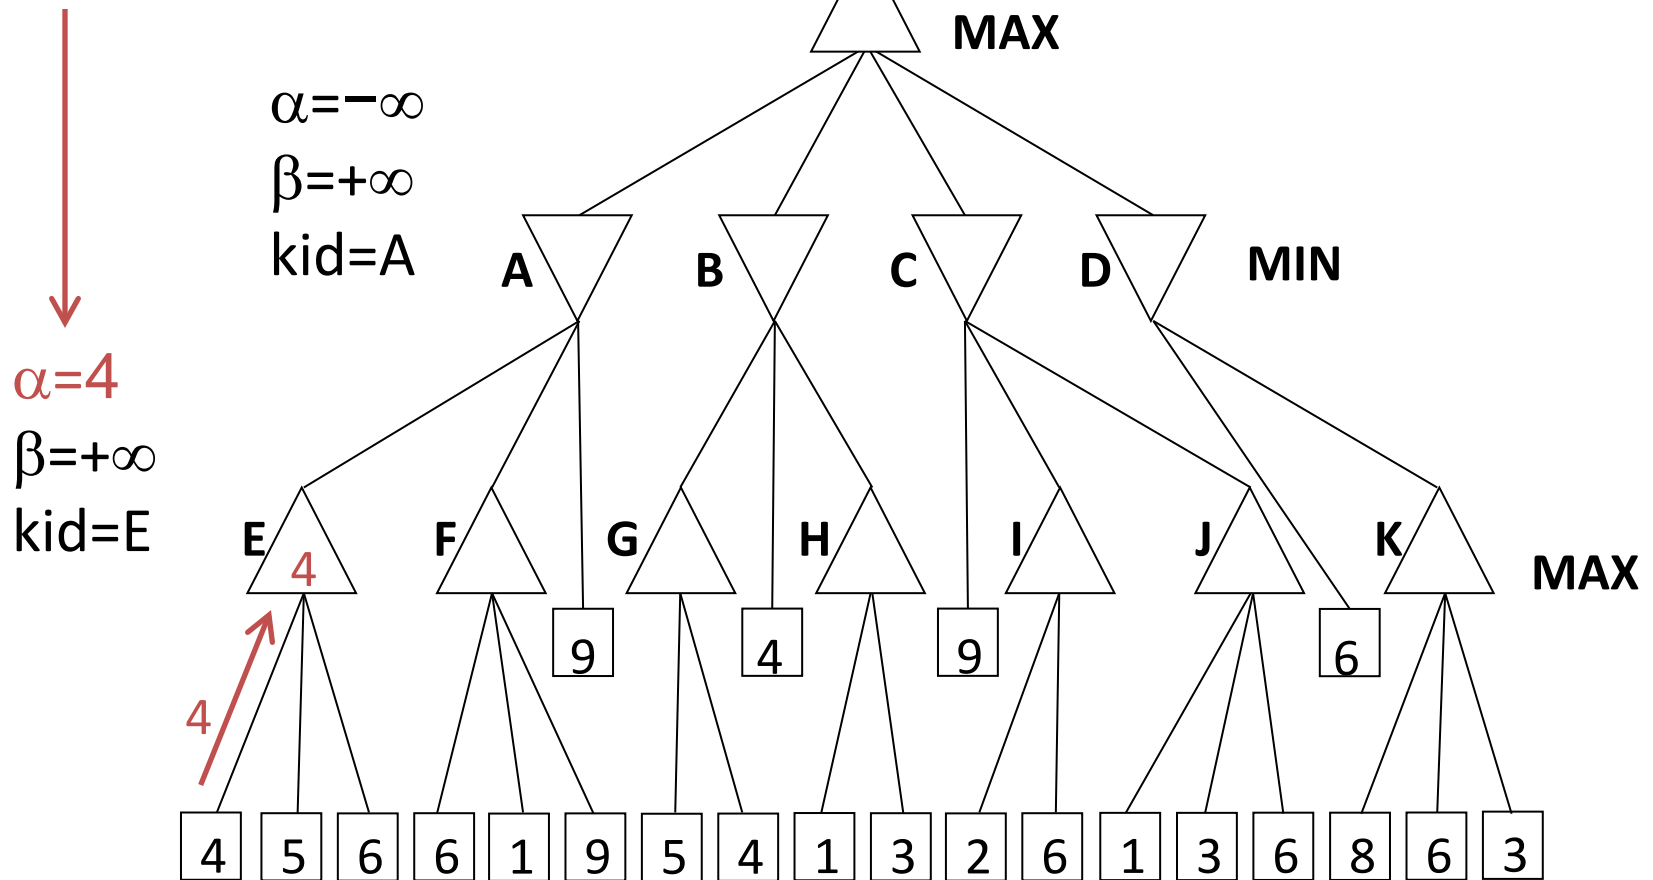
We also are running MiniMax search and recording node values within the triangles, without explicit comment.

Longer Alpha-Beta Example

*see next leaf,
MAX updates α*
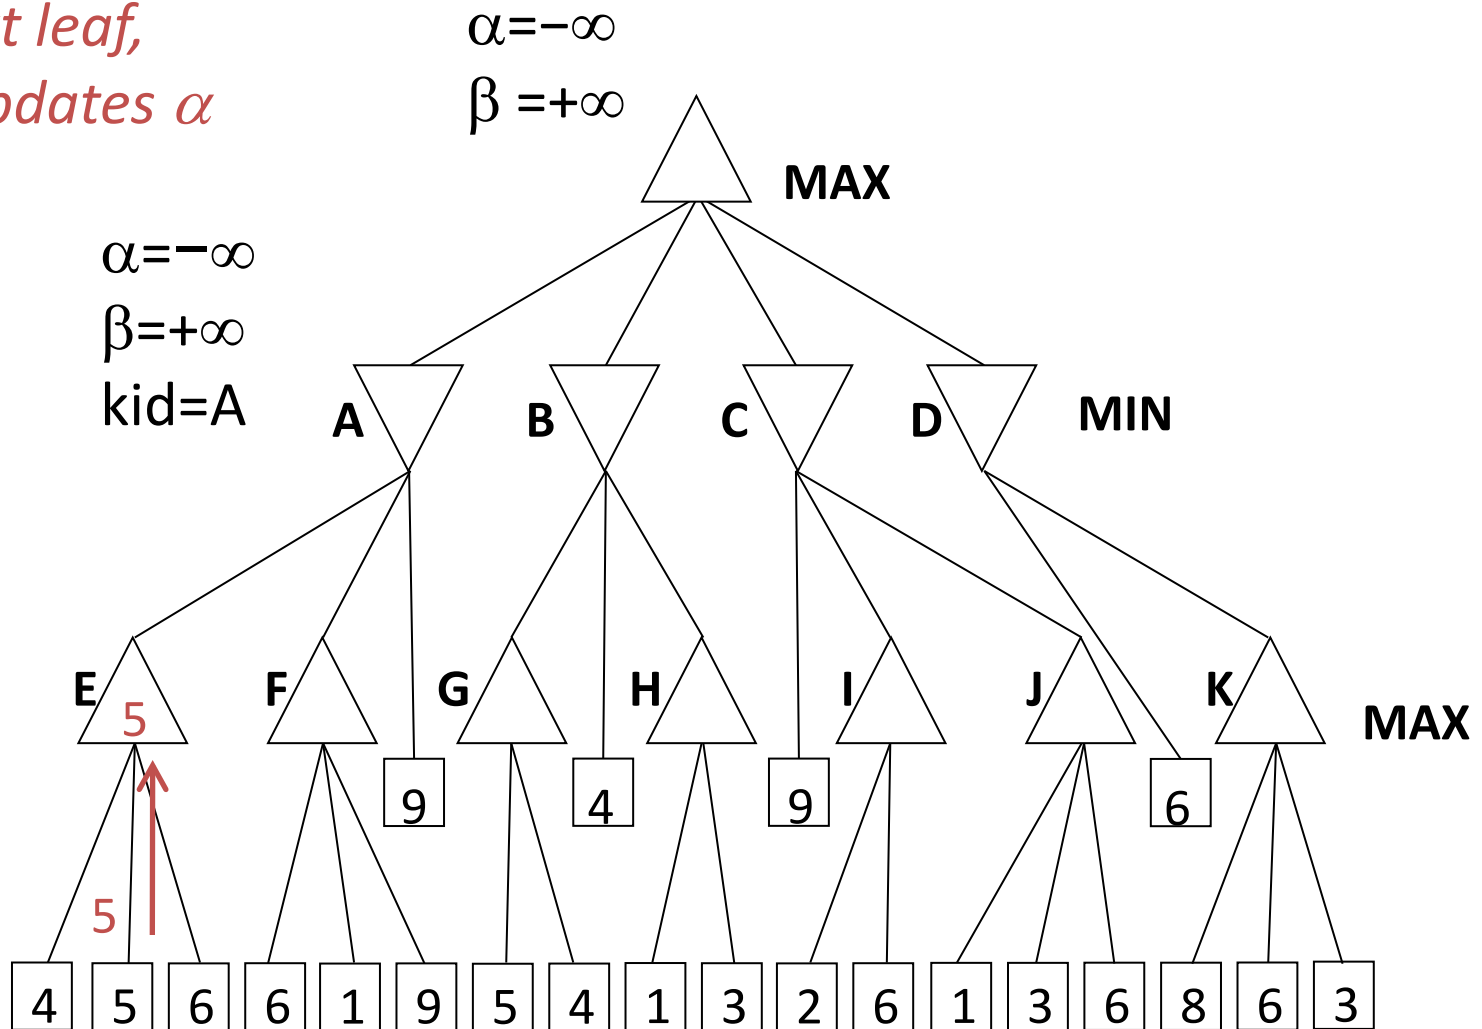
$\alpha = -\infty$
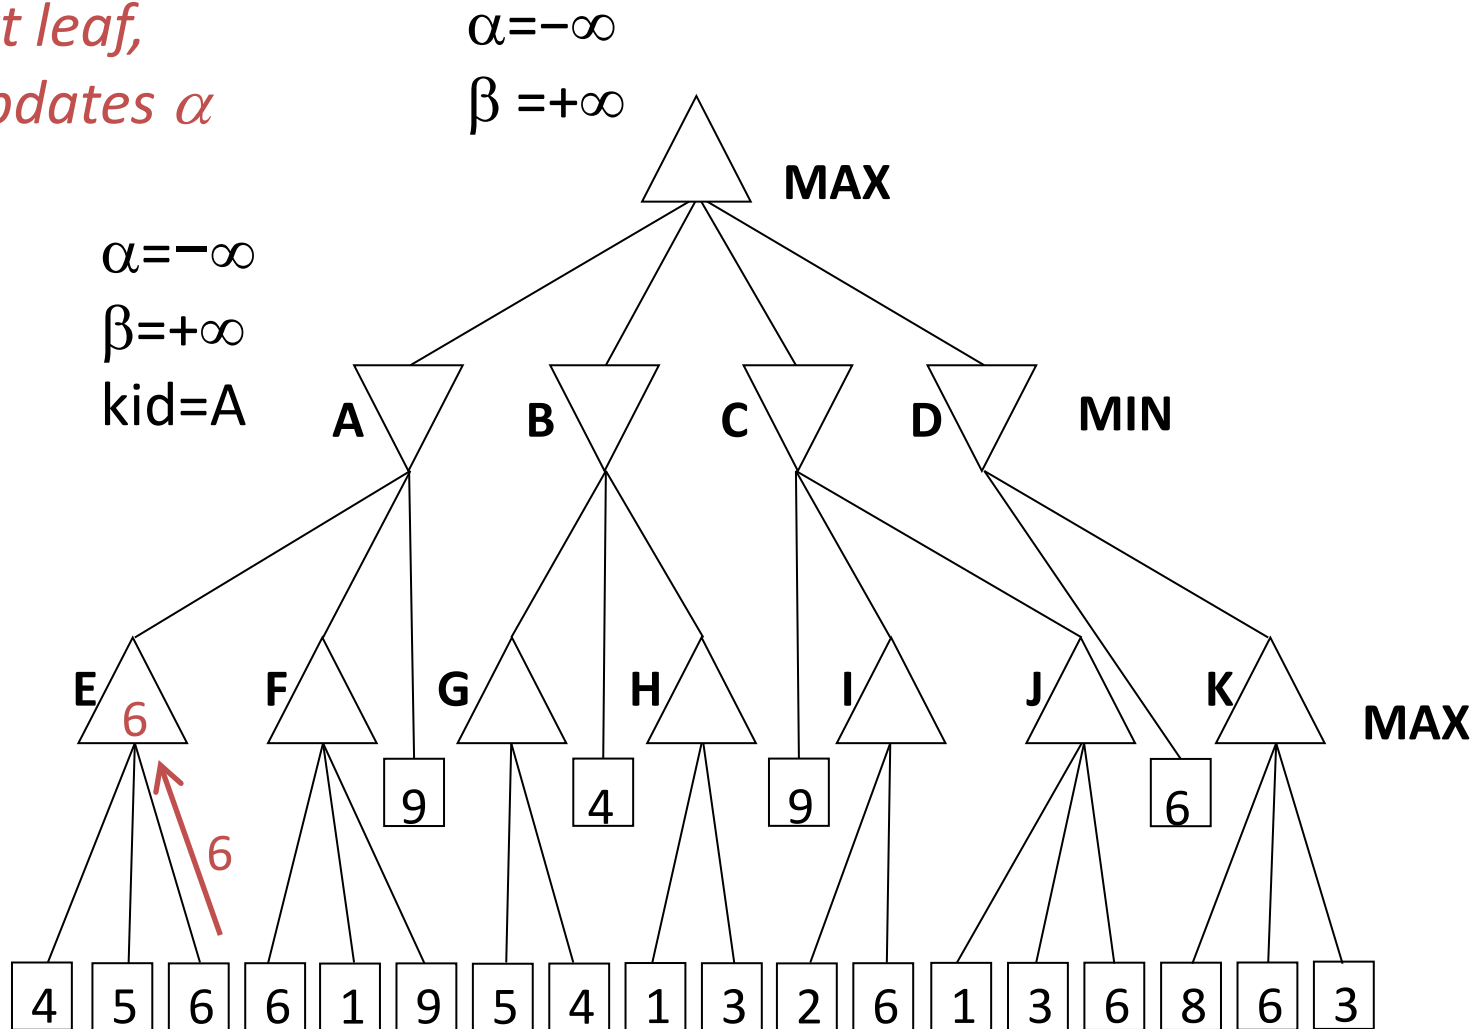$\beta = +\infty$

MAX

$\alpha = -\infty$
$\beta = +\infty$
kid=A

A   B   C   D   MIN

$\alpha = 5$
$\beta = +\infty$
kid=E

E 5   F   G   H   I   J   K   MAX

9   4   9   6

5

5

4 5 6 6 1 9 5 4 1 3 2 6 1 3 6 8 6 3

Longer Alpha-Beta Example

see next leaf, MAX updates α

$\alpha = -\infty$
$\beta = +\infty$

MAX

$\alpha = -\infty$
$\beta = +\infty$
kid=A

A B C D **MIN**

$\alpha = 6$
$\beta = +\infty$
kid=E

E F G H I J K **MAX**

6

6

9 4 9 6

4 5 6 6 1 9 5 4 1 3 2 6 1 3 6 8 6 3

Longer Alpha-Beta Example

return node value,
MIN updates β

α=−∞
β =+∞

MAX

α=−∞

β=6

kid=A

A 6   **B**   **C**   **D**   **MIN**

6

E 6   **F**   **G**   **H**   **I**   **J**   **K**   **MAX**

9   4   9   6

4 5 6 6 1 9 5 4 1 3 2 6 1 3 6 8 6 3

Longer Alpha-Beta Example

current α, β, passed to kid F

α=−∞
β =+∞

MAX

α=−∞
β=6
kid=A

α=−∞
β=6
kid=F

A 6 **B** **C** **D** **MIN**

E 6 **F** **G** **H** **I** **J** **K** **MAX**

9   4   9   6

4 5 6 6 1 9 5 4 1 3 2 6 1 3 6 8 6 3

Longer Alpha-Beta Example

*see first leaf,
MAX updates α*

$\alpha = -\infty$
$\beta = +\infty$

MAX

$\alpha = -\infty$
$\beta = 6$
kid=A

A 6   **B**   **C**   **D**   **MIN**

$\alpha = 6$
$\beta = 6$
kid=F

E 6   **F** 6   **G**   **H**   **I**   **J**   **K**   **MAX**

6

9   4   9   6

4 5 6 6 1 9 5 4 1 3 2 6 1 3 6 8 6 3

Longer Alpha-Beta Example

$\alpha \geq \beta$!!
Prune!!

$\alpha = -\infty$
$\beta = +\infty$

MAX

$\alpha = -\infty$
$\beta = 6$
kid=A

$\alpha = 6$
$\beta = 6$
kid=F

A 6   **B**   **C**   **D**   **MIN**

E 6   **F** 6   **G**   **H**   **I**   **J**   **K**   **MAX**

9   4   9   6

4 5 6 6 X X 5 4 1 3 2 6 1 3 6 8 6 3

Longer Alpha-Beta Example

*return node value,
MIN updates β,
no change to β*

$\alpha = -\infty$
$\beta = +\infty$

MAX

$\alpha = -\infty$
$\beta = 6$
kid = A

A 6   **B**   **C**   **D**   **MIN**

6

E 6   **F** 6   **G**   **H**   **I**   **J**   **K**   **MAX**

9   4   9   6

4 5 6 6 X X 5 4 1 3 2 6 1 3 6 8 6 3

If we had continued searching at node F, we would see the 9 from its third leaf. Our returned value would be 9 instead of 6. But at A, MIN would choose E(=6) instead of F(=9). Internal values may change; root values do not.

Longer Alpha-Beta Example

Longer Alpha-Beta Example

return node value, → α=6

MAX updates α   β =+∞

6

MAX

A 6   B   C   D   MIN

6

E 6   F 6   G   H   I   J   K   MAX

9   4   9   6

4   5   6   6   X X   5   4   1   3   2   6   1   3   6   8   6   3

Longer Alpha-Beta Example

current α, β,
passed to kids

$\alpha=6$
$\beta=+\infty$
kid=B

$\alpha=6$
$\beta=+\infty$
kid=G

$\alpha=6$
$\beta=+\infty$

6 MAX

A 6   B   C   D   MIN

E 6   F 6   G   H   I   J   K   MAX

9   4   9   6

4 5 6 6 **X X** 5 4 1 3 2 6 1 3 6 8 6 3

Longer Alpha-Beta Example

see first leaf,
MAX updates α,
no change to α

α=6
β =+∞

6   **MAX**

α=6
β=+∞
kid=B

α=6
β=+∞
kid=G

A 6   **B**   **C**   **D**   **MIN**

E 6   **F** 6   **G** 5   **H**   **I**   **J**   **K**   **MAX**

9   5   4   9   6

4   5   6   6   X   X   5   4   1   3   2   6   1   3   6   8   6   3

5

Longer Alpha-Beta Example

α=6
β =+∞

see next leaf,
MAX updates α,
no change to α

α=6
β=+∞
kid=B

MAX

α=6
β=+∞
kid=G

6

A 6   **B**   **C**   **D**   **MIN**

E 6   **F** 6   **G** 5   **H**   **I**   **J**   **K**   **MAX**

9   4   9   6

4

4   5   6   6   **X**   **X**   5   4   1   3   2   6   1   3   6   8   6   3

Longer Alpha-Beta Example

return node value, MIN updates β

$\alpha=6$
$\beta=+\infty$

$\alpha=6$
$\beta=5$
kid=B

6 MAX

A 6   B 5   C   D   MIN

5

E 6   F 6   G 5   H   I   J   K   MAX

9   4   9   6

4 5 6 6 X X 5 4 1 3 2 6 1 3 6 8 6 3

Longer Alpha-Beta Example

$\alpha \geq \beta$!!
Prune!!

$\alpha=6$
$\beta =+\infty$

6 MAX

$\alpha=6$
$\beta=5$
kid=B

A 6   B 5   C   D   MIN

E 6   F 6   G 5   H ?   I   J   K   MAX

9   X   9   6

4 5 6 6 X X 5 4 X X 2 6 1 3 6 8 6 3

Note that we never find out, what is the node value of H? But we have proven it doesn't matter, so we don't care.

Longer Alpha-Beta Example

return node value,
MAX updates α,
no change to α

$\alpha=6$

$\beta=+\infty$

Longer Alpha-Beta Example

current α, β,
passed to kid=C

α=6
β=+∞
kid=C

α=6
β =+∞

6   MAX

A 6   B 5   C   D   MIN

E 6   F 6   G 5   H ?   I   J   K   MAX

9   X   9   6

4 5 6 6 X X 5 4 X X 2 6 1 3 6 8 6 3

Longer Alpha-Beta Example

*see first leaf,
MIN updates β*

$\alpha=6$
$\beta=9$
kid=C

$\alpha=6$
$\beta=+\infty$

6 **MAX**

A 6   B 5   C 9   D   **MIN**

9

9

E 6   F 6   G 5   H ?   I   J   K   **MAX**

9   **X**   9   6

4 5 6 6 **X** **X** 5 4 **X** **X** 2 6 1 3 6 8 6 3

Longer Alpha-Beta Example

Longer Alpha-Beta Example

Longer Alpha-Beta Example

Longer Alpha-Beta Example

return node value, MIN updates β

$\alpha=6$
$\beta =+\infty$

$\alpha=6$
$\beta=6$
kid=C

6   **MAX**

A 6   B 5   C 6   D   **MIN**

6

E 6   F 6   G 5   H ?   I 6   J   K   **MAX**

9   X   9   6

4 5 6   6 X X 5 4 X X 2 6   1 3 6 8   6 3

Longer Alpha-Beta Example

$\alpha \geq \beta$!!
Prune!!

α=6
β =+∞

α=6
β=6
kid=C

6 MAX

A 6   B 5   C 6   D   MIN

E 6   F 6   G 5   H ?   I 6   J ?   K   MAX

9   X   9   6

4 5 6 6 X X 5 4 X X 2 6 X X X 8 6 3

Longer Alpha-Beta Example

return node value,
MAX updates α,
no change to α

$\alpha = 6$

$\beta = +\infty$

Longer Alpha-Beta Example

Longer Alpha-Beta Example

see first leaf,
MIN updates β

$\alpha=6$
$\beta =+\infty$

6 **MAX**

$\alpha=6$
$\beta=6$
kid=D

A 6   **B** 5   **C** 6   **D** 6   **MIN**

E 6   **F** 6   **G** 5   **H** ?   **I** 6   **J** ?   **K**   **MAX**

6

9   X   9   6

4 5 6 6 X X 5 4 X X 2 6 X X X 8 6 3

Longer Alpha-Beta Example

$\alpha \geq \beta$!!
Prune!!

α=6
β =+∞

α=6
β=6
kid=D

6 MAX

A 6   B 5   C 6   D 6   MIN

E 6   F 6   G 5   H ?   I 6   J ?   K ?   MAX

9   X   9   6

4 5 6 6 X X 5 4 X X 2 6 X X X X X X

Alpha-Beta Example #2

return node value, $\alpha=6$
MAX updates α, $\beta =+\infty$
no change to α
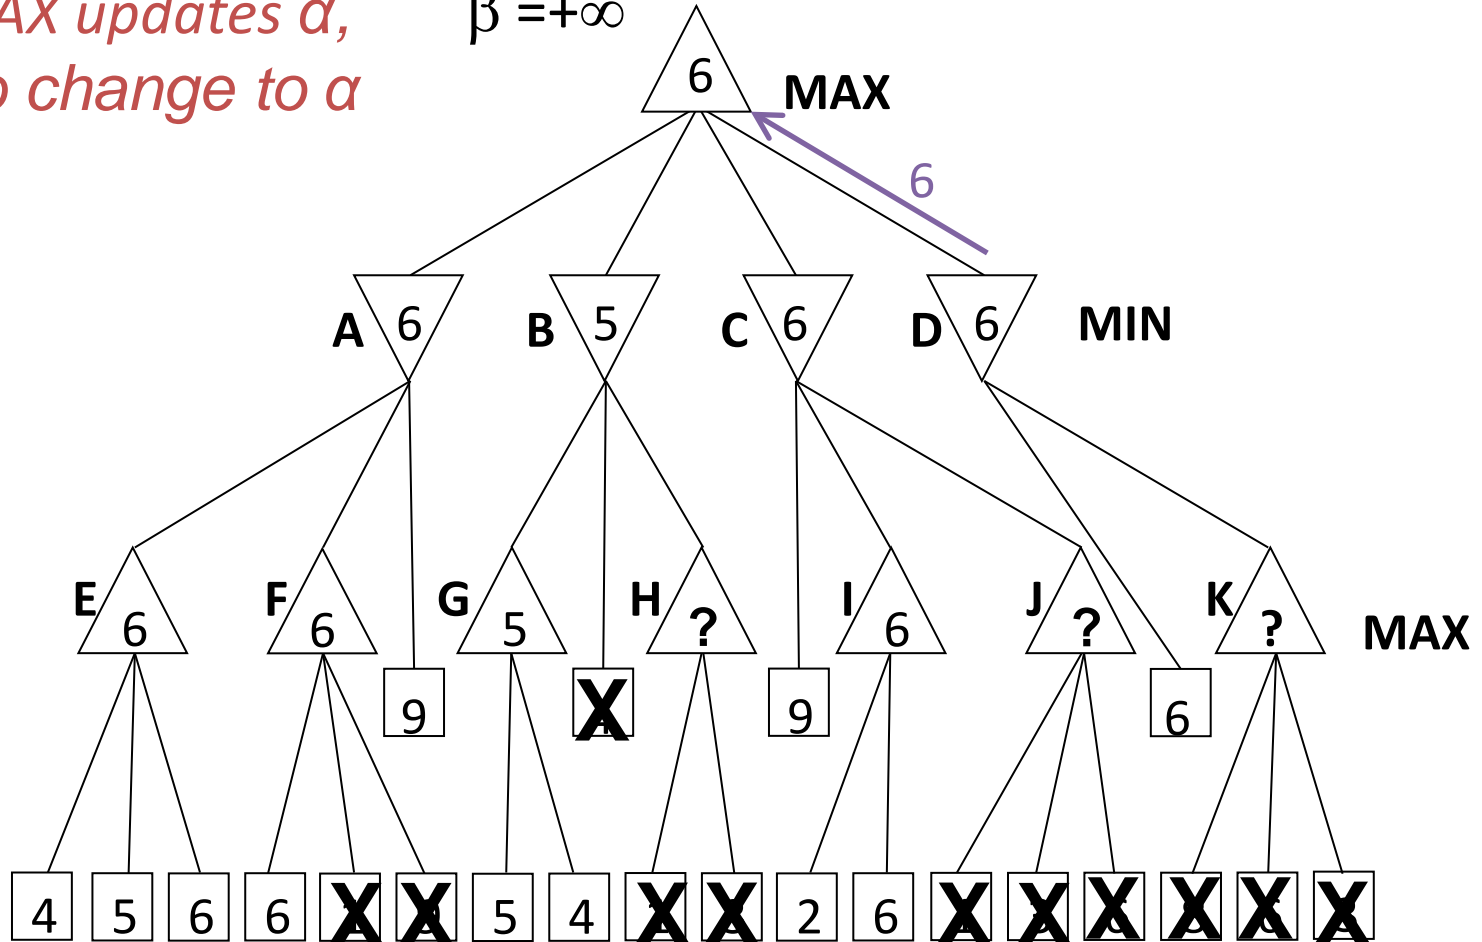
Alpha-Beta Example #2

MAX moves to A,
and expects to get 6

MAX's move

A 6	B 5	C 6	D 6

MAX

MIN

E 6   F 6   G 5   H ?   I 6   J ?   K ?

9   X   9   6

4   5   6   6   X   X   5   4   X   X   2   6   X   X   X   X   X   X
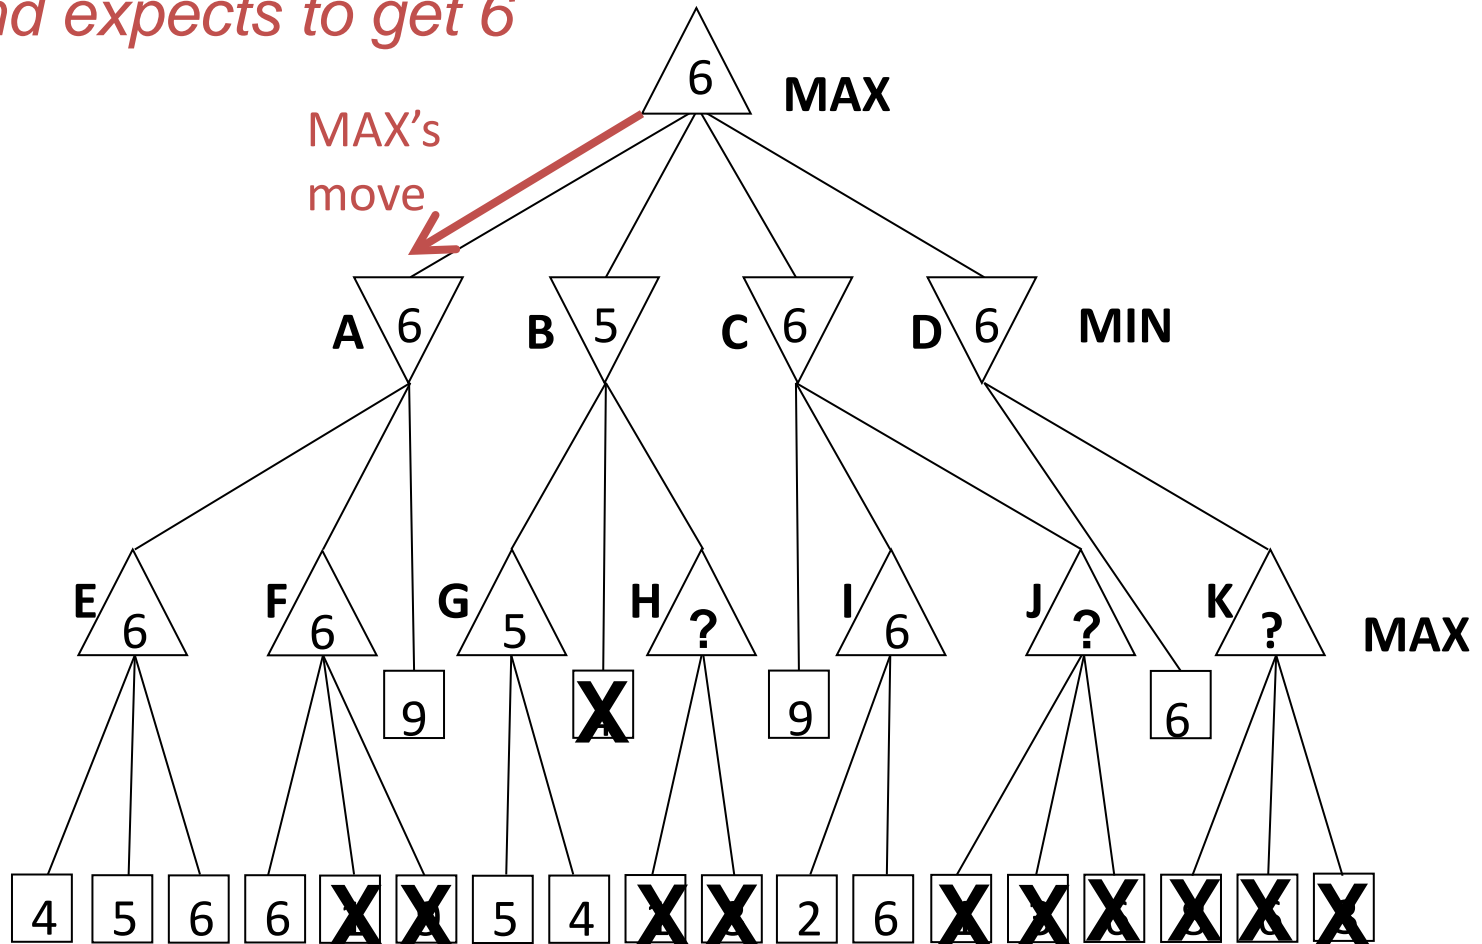
Although we may have changed some internal branch node return values, the final root action and expected outcome are identical to if we had not done alpha-beta pruning. Internal values may change; root values do not.

Nondeterministic games

- Ex: Backgammon
  - Roll dice to determine how far to move (random)
  - Player selects which checkers to move   (strategy)
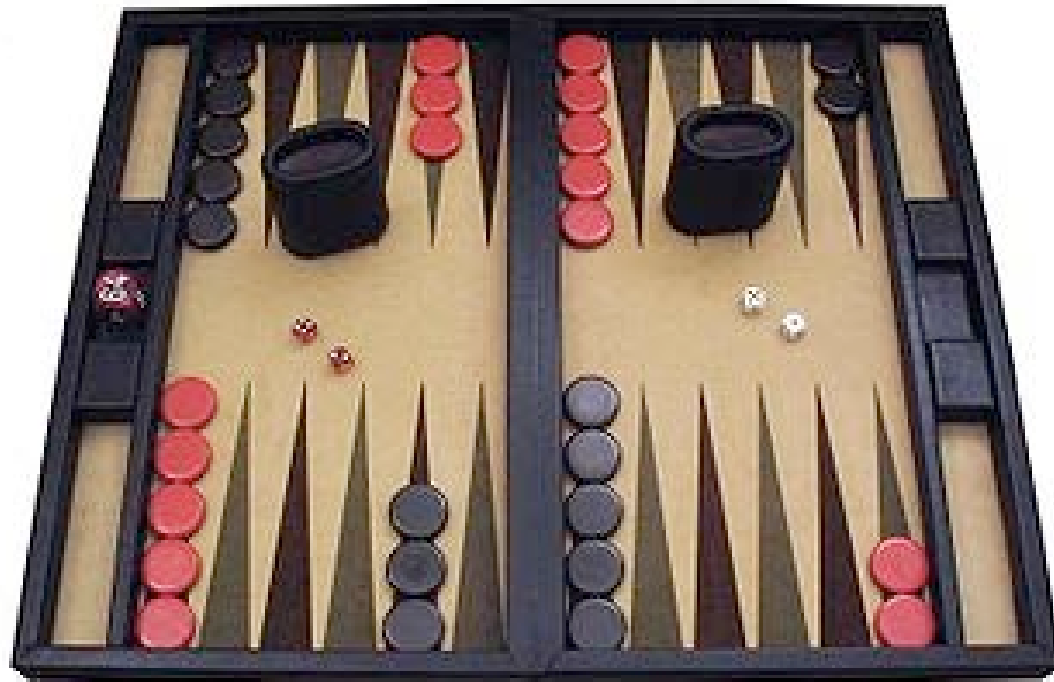
Nondeterministic games

- Chance (random effects) due to dice, card shuffle, …
- Chance nodes: expectation (weighted average) of successors
- Simplified example: coin flips

MAX's move

"Expectiminimax"

MAX

3

Chance

3   -1

0.5   0.5   0.5   0.5

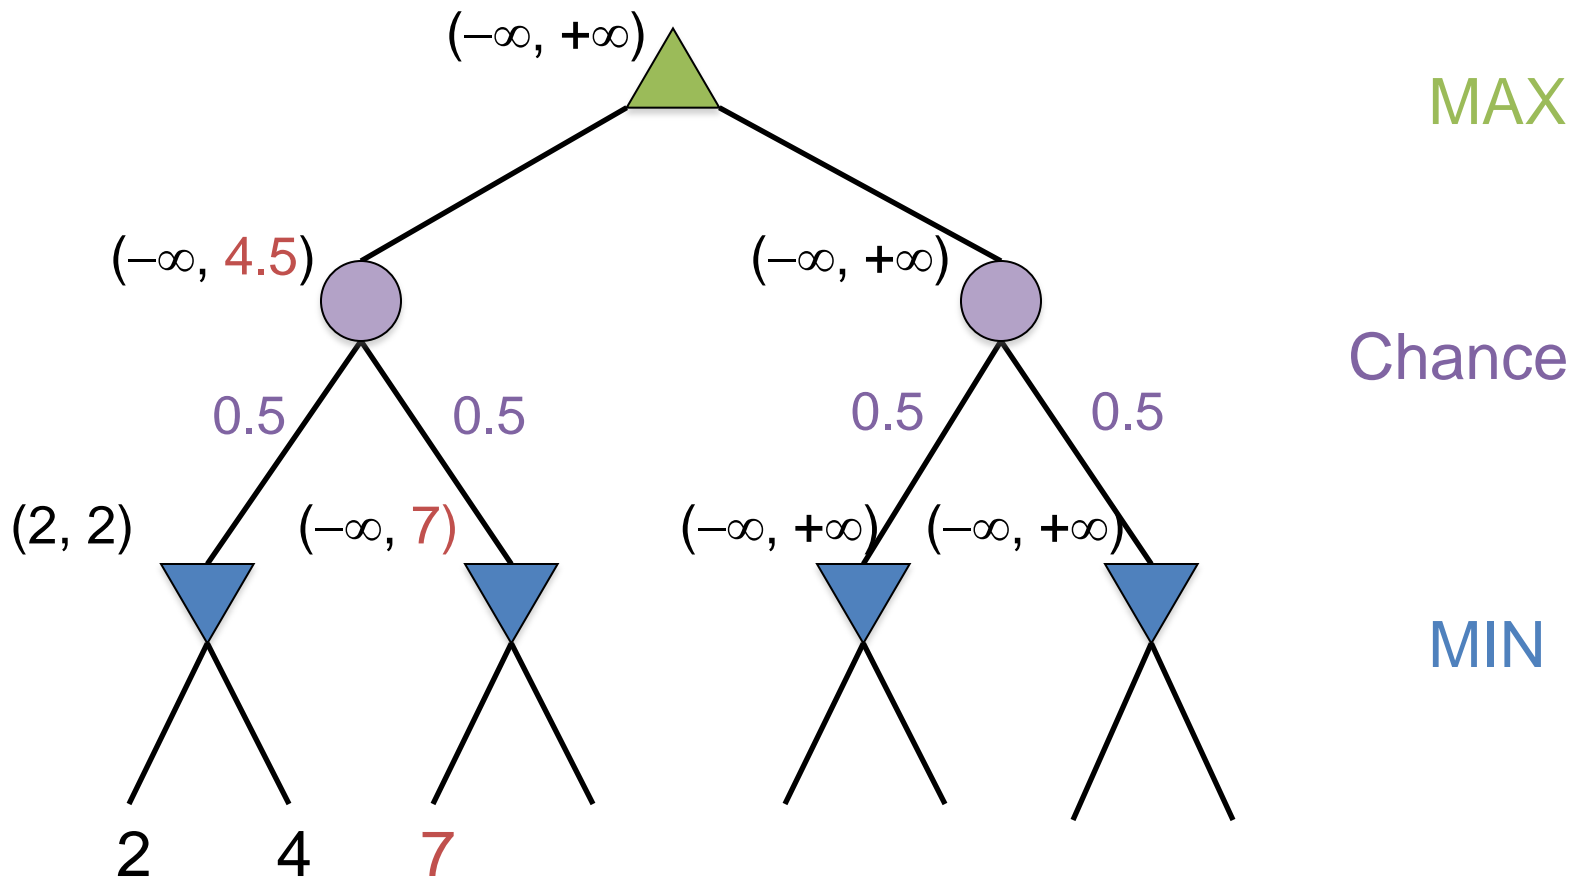2   4   0   -2

MIN

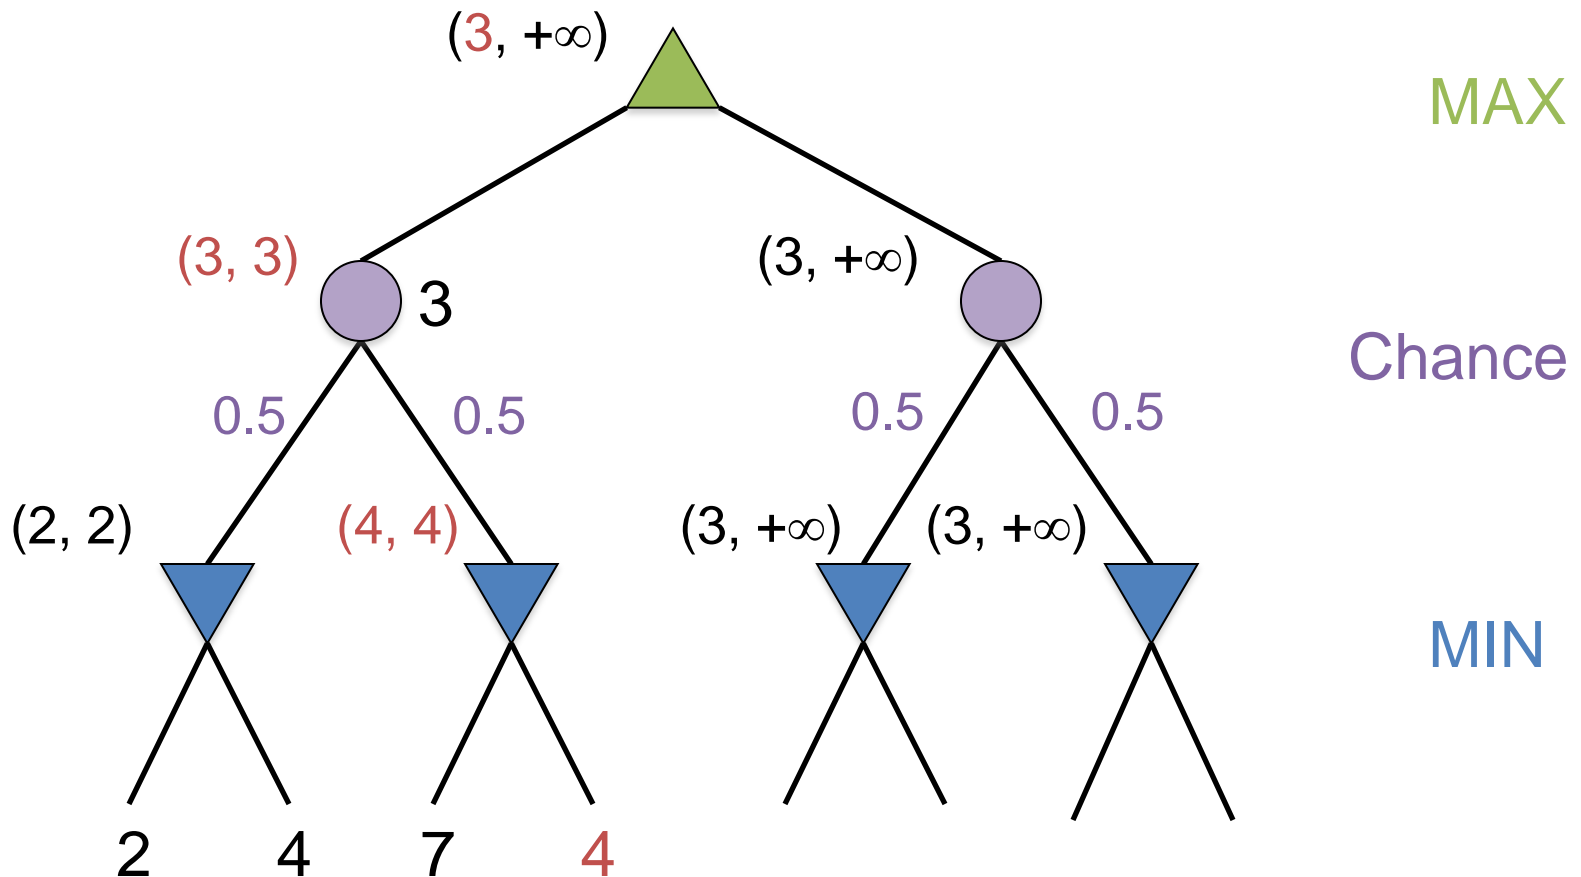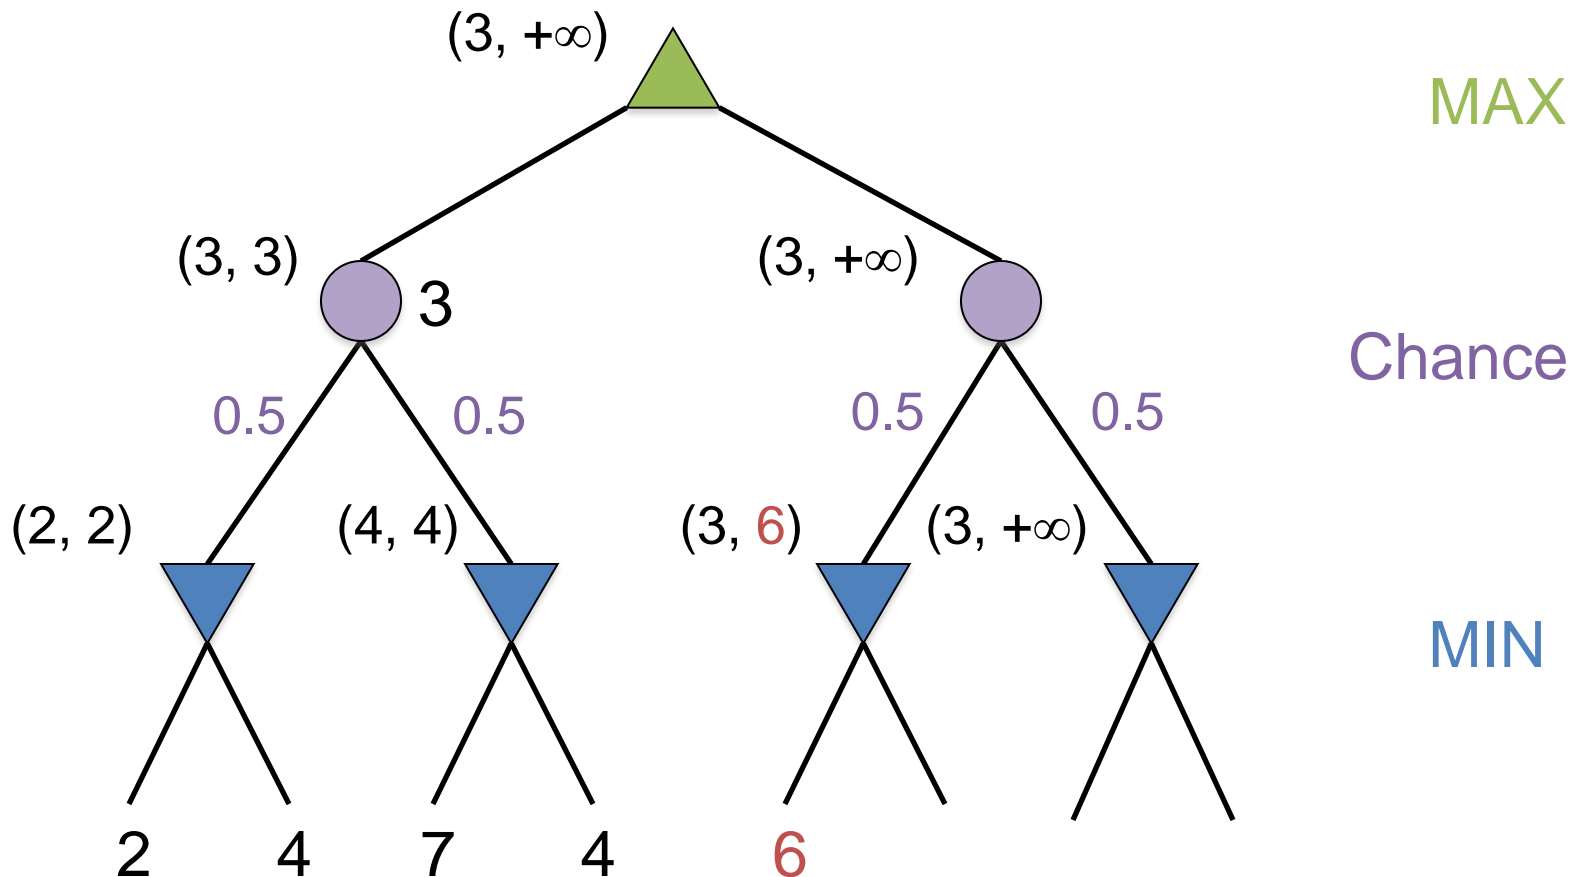2   4   7   4   6   0   5   -2

Pruning in nondeterministic games

- Can still apply a form of alpha-beta pruning

Pruning in nondeterministic games

- Can still apply a form of alpha-beta pruning

Pruning in nondeterministic games

- Can still apply a form of alpha-beta pruning


MAX

Chance

MIN

Pruning in nondeterministic games

- Can still apply a form of alpha-beta pruning

Pruning in nondeterministic games

- Can still apply a form of alpha-beta pruning

Pruning in nondeterministic games

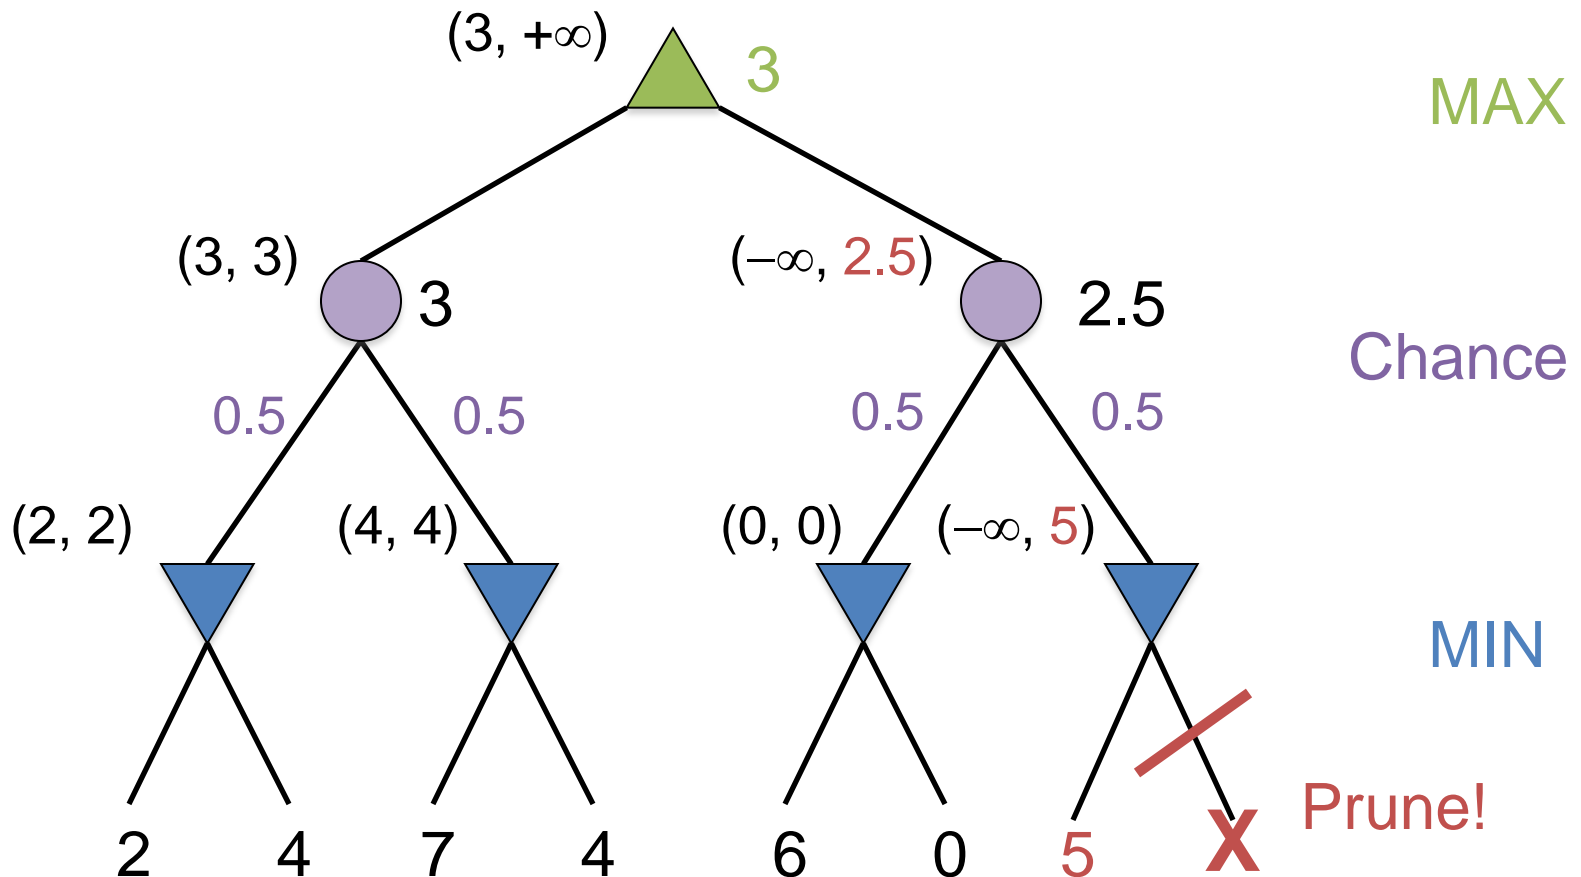- Can still apply a form of alpha-beta pruning

Pruning in nondeterministic games

- Can still apply a form of alpha-beta pruning

Pruning in nondeterministic games

- Can still apply a form of alpha-beta pruning

Pruning in nondeterministic games

- Can still apply a form of alpha-beta pruning

Partially observable games

- R&N Chapter 5.6 – "The fog of war"
- Background: R&N, Chapter 4.3-4
  - Searching with Nondeterministic Actions/Partial Observations

- Search through Belief States (see Fig. 4.14)
  - Agent's current belief about which states it might be in,
    given the sequence of actions & percepts to that point
- Actions(b) = ?? Union? Intersection?
  - Tricky: an action legal in one state may be illegal in another
  - Is an illegal action a NO-OP? or the end of the world?
- Transition Model:
  - Result(b,a) = { s' : s' = Result(s, a) and s is a state in b }
- Goaltest(b) = every state in b is a goal state

Belief States for Unobservable Vacuum World
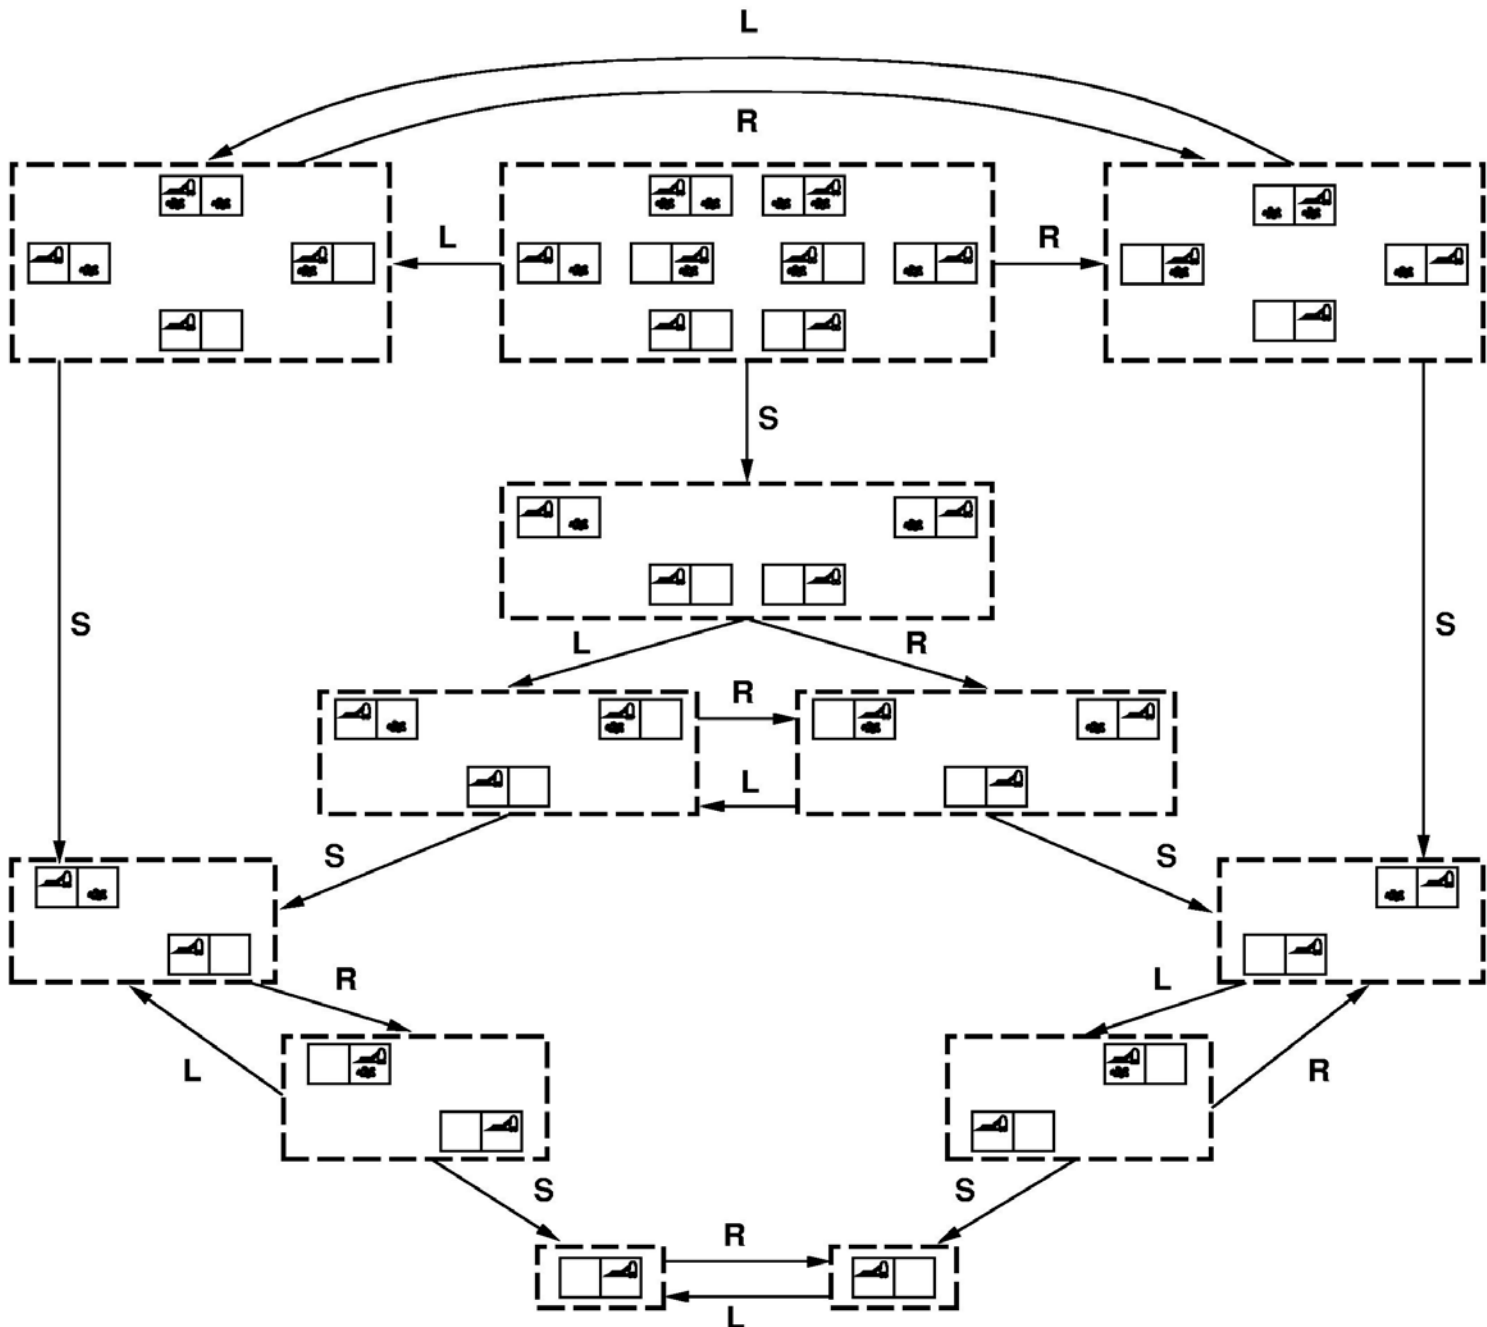
Partially observable games

- R&N Chapter 5.6
- Player's current node is a belief state
- Player's move (action) generates child belief state
- Opponent's move is replaced by Percepts(s)
  - Each possible percept leads to the belief state that is consistent with that percept
- Strategy = a move for every possible percept sequence
- Minimax returns the worst state in the belief state

- <u>Many more complications and possibilities!!</u>
  - Opponent may select a move that is not optimal, but instead minimizes the information transmitted, or confuses the opponent
  - May not be reasonable to consider ALL moves; open P-QR3??
- **See R&N, Chapter 5.6, for more info**

The State of Play

- Checkers:
  - Chinook ended 40-year-reign of human world champion Marion Tinsley in 1994.

- Chess:
  - Deep Blue defeated human world champion Garry Kasparov in a six-game match in 1997.

- Othello:
  - human champions refuse to compete against computers: they are too good.

- Go:
  - AlphaGo recently (3/2016) beat 9th dan Lee Sedol
  - b > 300 (!); full game tree has > 10^760 leaf nodes (!!)

- See (e.g.) http://www.cs.ualberta.ca/~games/ for more info

High branching factors

- What can we do when the search tree is too large?
  - Example: Go (b = 50 to 300+ moves per state)
  - Heuristic state evaluation (score a partial game)
- Where does this heuristic come from?
  - Hand designed
  - Machine learning on historical game patterns
  - Monte Carlo methods – play random games

Monte Carlo heuristic scoring

- Idea: play out the game randomly, and use the results as a score
  - Easy to generate & score lots of random games
  - May use 1000s of games for a node
- The basis of Monte Carlo tree search algorithms…

Image from www.mcts.ai
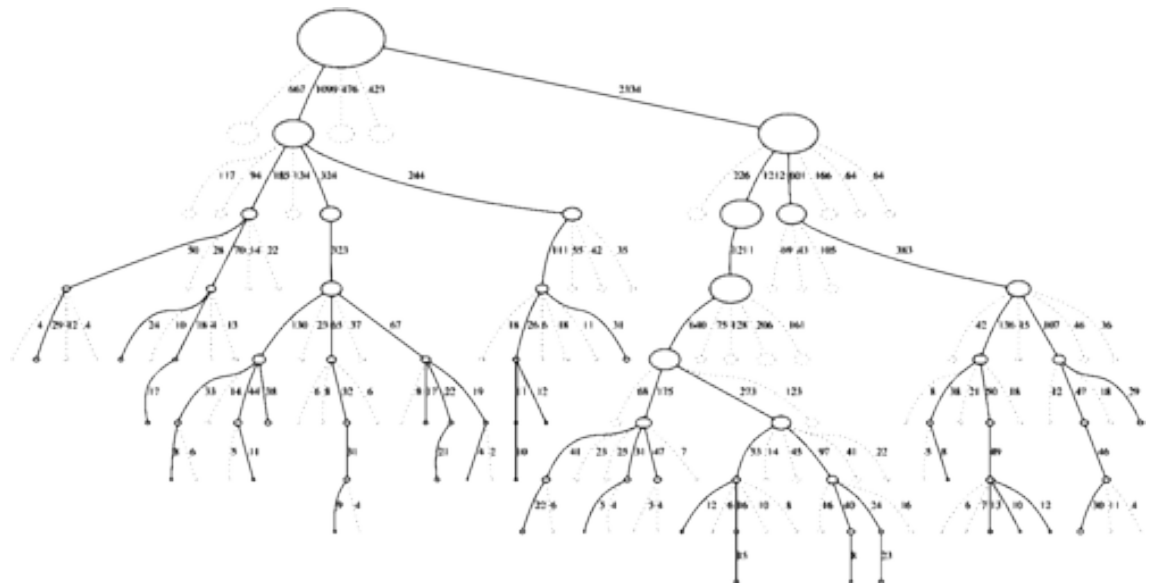
Monte Carlo Tree Search

- Should we explore the whole (top of) the tree?
  - Some moves are obviously not good...
  - Should spend time exploring / scoring promising ones
- This is a *multi-armed bandit* (MAB) problem:
- Want to spend our time on good moves
- Which moves have high payout?
  - Hard to tell – random...
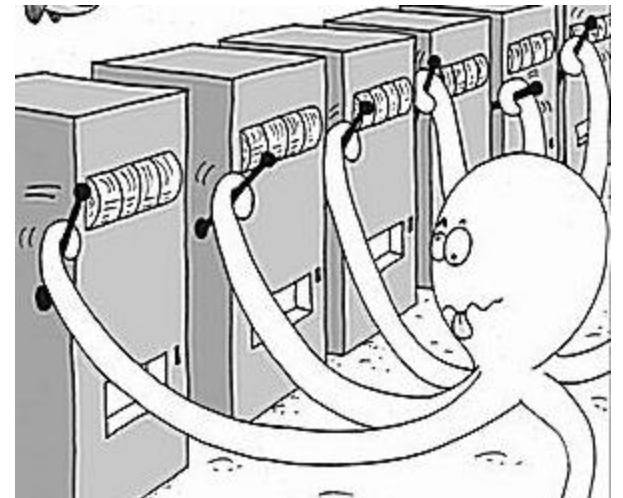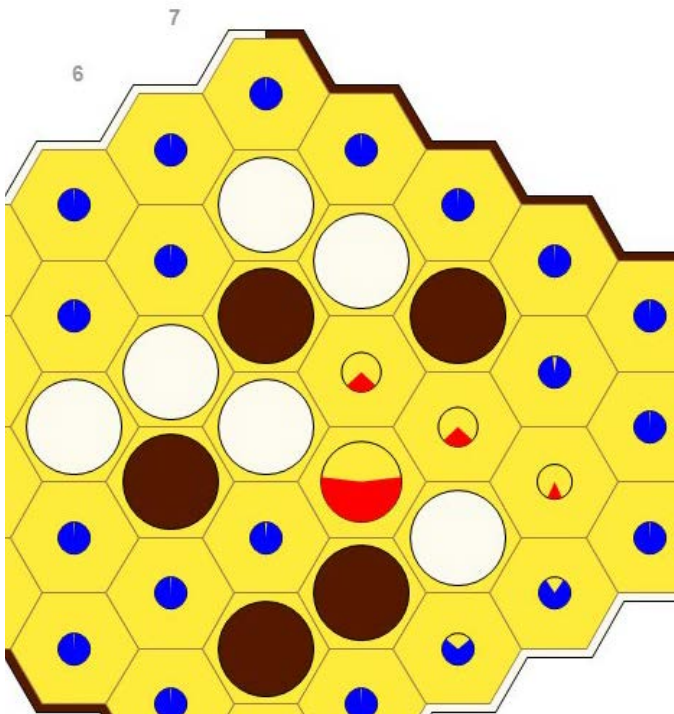- *Explore* vs. *exploit* tradeoff


Image from Microsoft Research

Visualizing MCTS

- At each level of the tree, keep track of
  - Number of times we've explored a path
  - Number of times we won
- Follow winning (from max/min perspective) strategies more often, but also explore others

MCTS

MAB strategy

1/1

1/1

Default / random
strategy

Terminal state
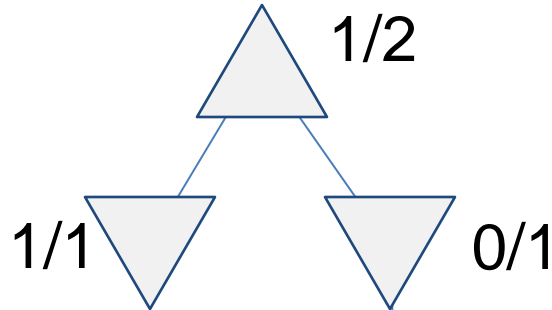
Score consists of
(1) % wins
(2) # times tried
(3) # of steps total

UCT:

$$s(n) = \bar{X}_n \pm \sqrt{\frac{\log n}{t(n)}}$$

MCTS

MAB strategy

1/2

1/1   0/1

Default / random
   strategy

Score consists of
(1) % wins
(2) # times tried
(3) # of steps total

UCT:

$$s(n) = \bar{X}_n \pm \sqrt{\frac{\log n}{t(n)}}$$

Terminal state

MCTS

MAB strategy

1/3

1/2   0/1

0/1

Default / random
   strategy

Terminal state

Score consists of
(1) % wins
(2) # times tried
(3) # of steps total

UCT:

$$s(n) = \bar{X}_n \pm \sqrt{\frac{\log n}{t(n)}}$$

Summary

- Game playing is best modeled as a search problem

- Game trees represent alternate computer/opponent moves

- Evaluation functions estimate the quality of a given board configuration for the Max player.

- Minimax is a procedure which chooses moves by assuming that the opponent will always choose the move which is best for them

- Alpha-Beta is a procedure which can prune large parts of the search tree and allow search to go deeper

- For many well-known games, computer algorithms based on heuristic search match or out-perform human world experts.