CSE 40171:
Artificial Intelligence

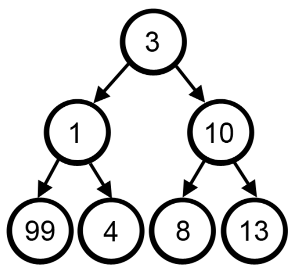

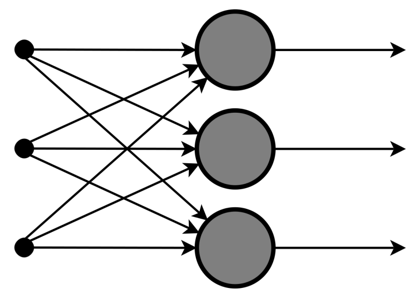

Adversarial Search: Expectimax; Partial Observability

Homework #4 has been released
It is due at 11:59PM on 10/18

Horizon Effect

When the program is facing an opponent's move that causes serious damage and is ultimately unavoidable, but can be temporarily avoided by delaying tactics.

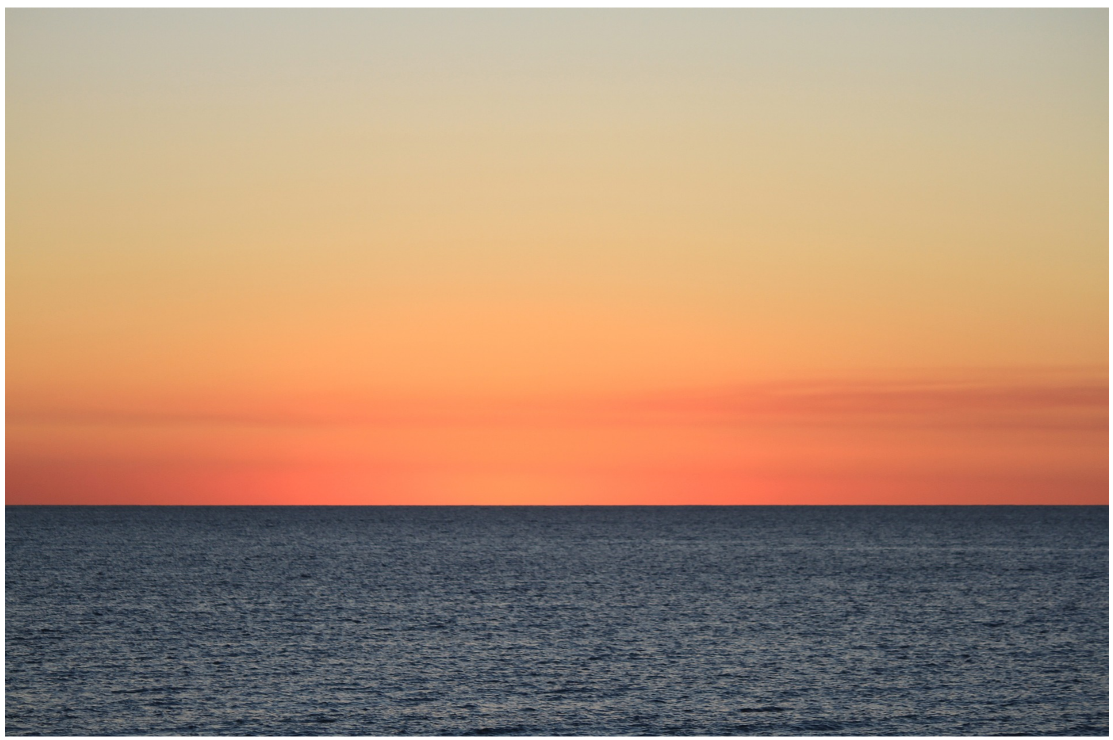

Horizon Effect

Inevitable Loss

Image credit: Russell and Norvig

Horizon Effect

The loss is simply delayed

Image credit: Russell and Norvig

Search vs. Lookup

There are many standard openings and closings in chess

Why bother with search when you can simply use a lookup table?

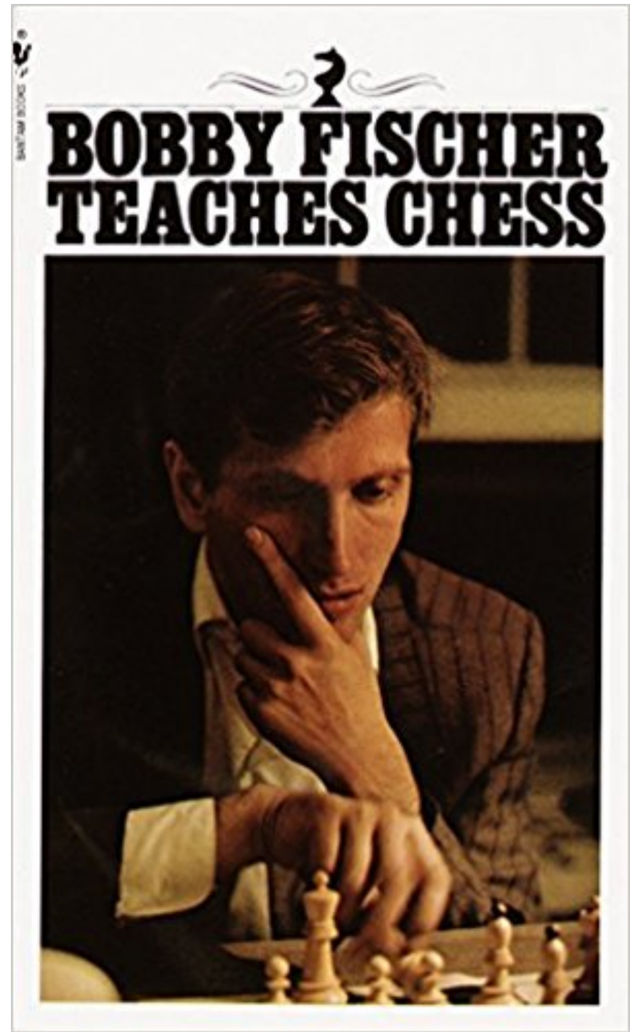

Search vs. Lookup

Computers are particularly good at the endgame

Example: king, bishop, and knight vs. king

462 ways a king can be placed without being adjacent

62 empty squares for the bishop, 61 for the knight, and 2 players to move next

$462 \times 62 \times 61 \times 2 = 3{,}494{,}568$ possible positions

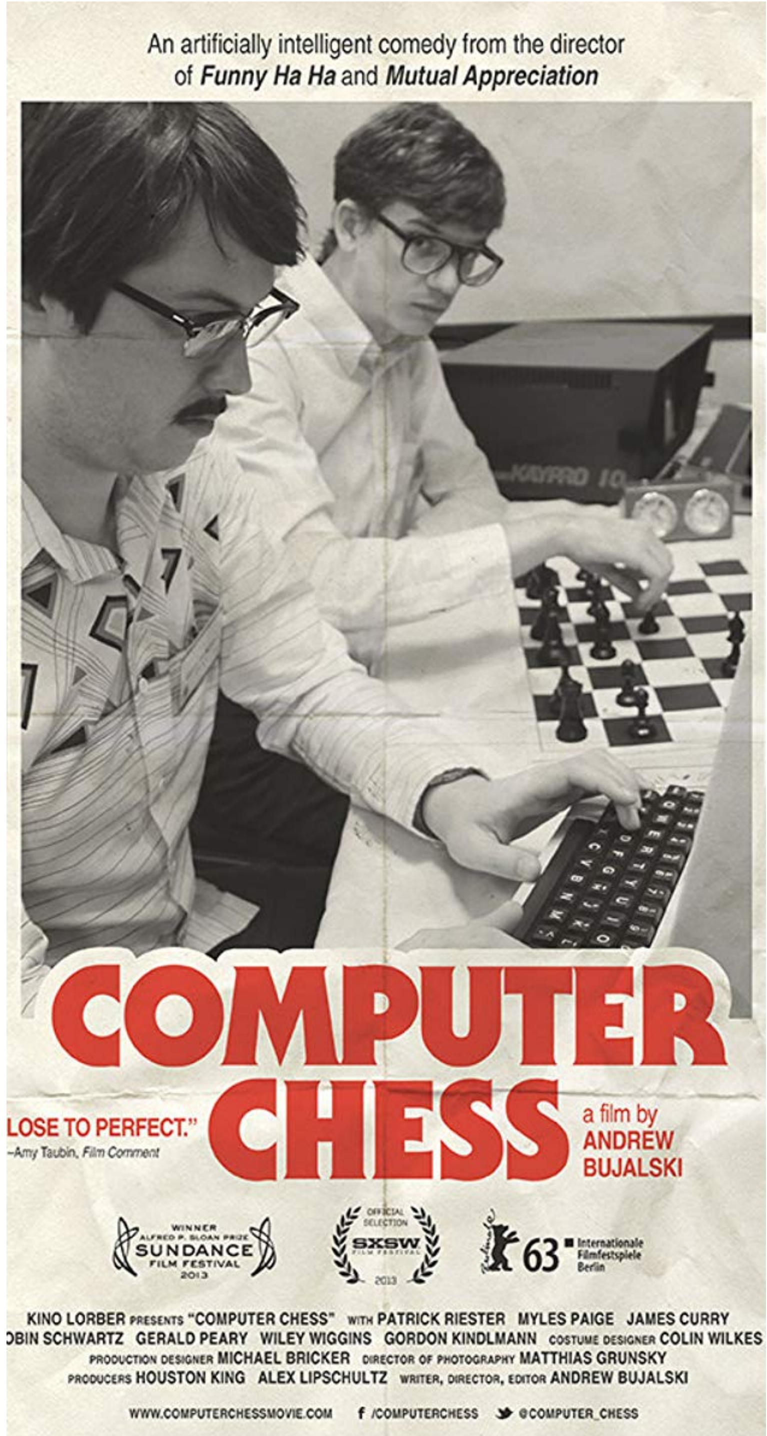
An artificially intelligent comedy from the director
of *Funny Ha Ha* and *Mutual Appreciation*
COMPUTER
CHESS
LOSE TO PERFECT."
–Amy Taubin, *Film Comment*
a film by
ANDREW
BUJALSKI
WINNER
ALFRED P. SLOAN PRIZE
SUNDANCE
FILM FESTIVAL
2013
OFFICIAL
SELECTION
SXSW
FILM FESTIVAL
2013
63
Internationale
Filmfestspiele
Berlin
KINO LORBER PRESENTS "COMPUTER CHESS" WITH PATRICK RIESTER MYLES PAIGE JAMES CURRY
OBIN SCHWARTZ GERALD PEARY WILEY WIGGINS GORDON KINDLMANN COSTUME DESIGNER COLIN WILKES
PRODUCTION DESIGNER MICHAEL BRICKER DIRECTOR OF PHOTOGRAPHY MATTHIAS GRUNSKY
PRODUCERS HOUSTON KING ALEX LIPSCHULTZ WRITER, DIRECTOR, EDITOR ANDREW BUJALSKI
WWW.COMPUTERCHESSMOVIE.COM
f /COMPUTERCHESS
@COMPUTER_CHESS

Stochastic Games

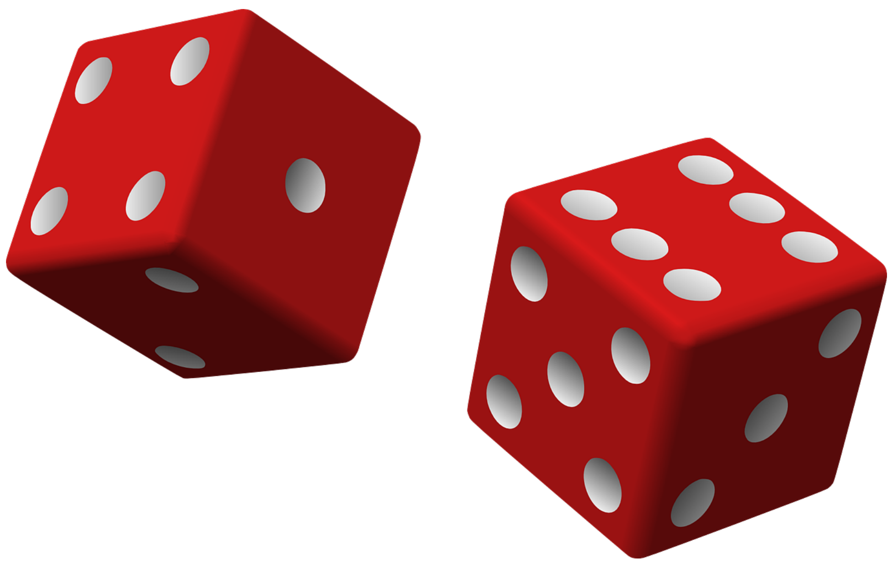

What we've assumed thus far…

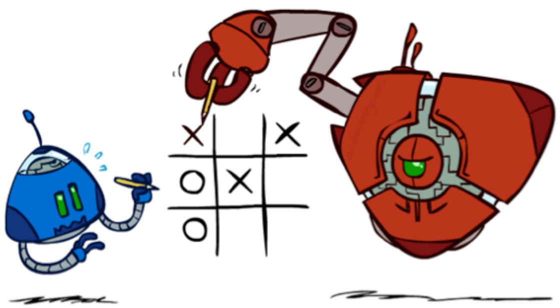

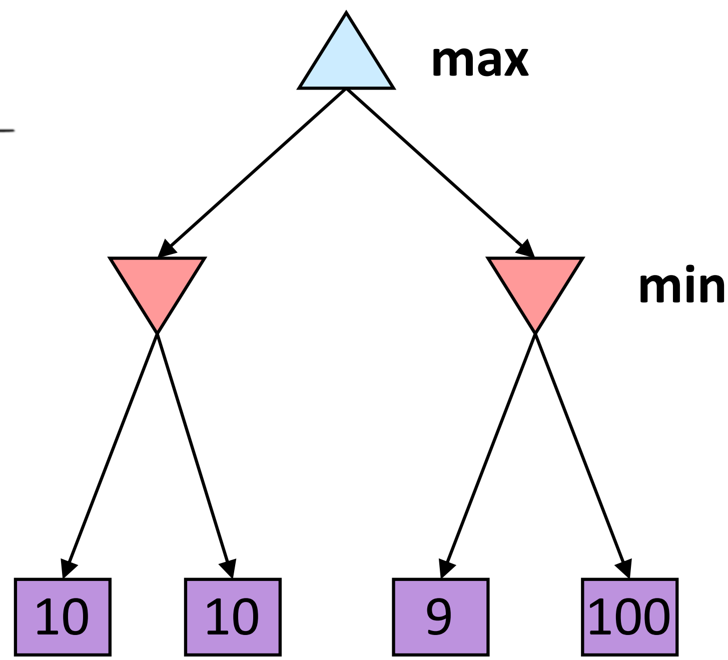

Image credits: Dan Klein and Pieter Abbeel, UC Berkeley CS 188

What if uncertain outcomes are controlled by chance, and not an adversary?

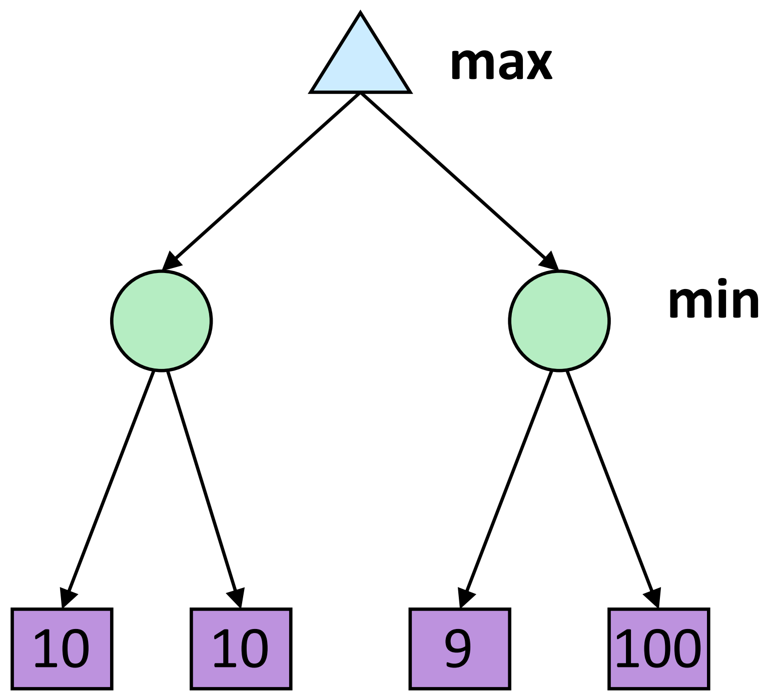

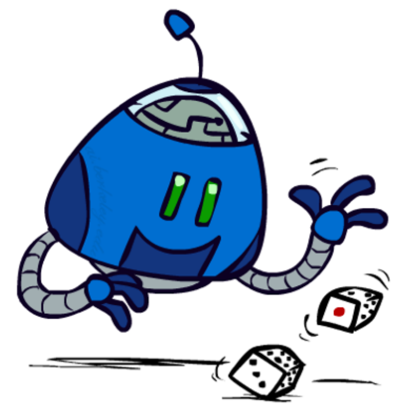

Image credits: Dan Klein and Pieter Abbeel, UC Berkeley CS 188

Expectimax Search

Why wouldn't we know what the result of an action will be?

- Explicit randomness: rolling dice
- Unpredictable opponents: the pacman ghosts respond randomly
- Actions can fail: when moving a robot, wheels might slip

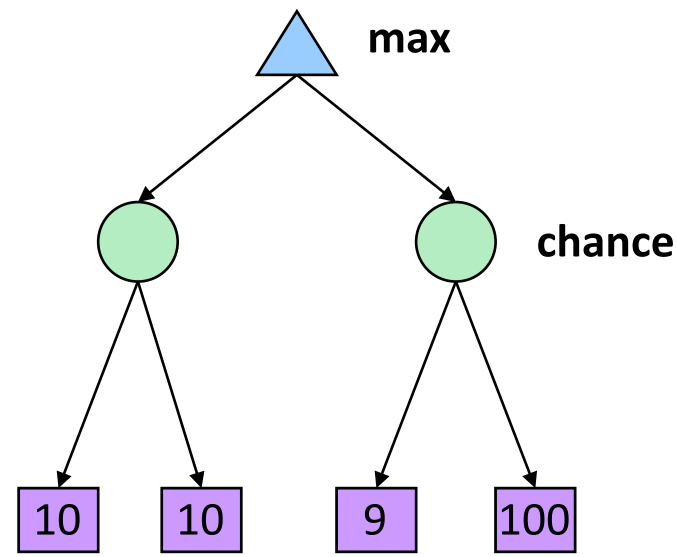

Values should now reflect average-case (**expectimax**) outcomes, not worst-case (minimax) outcomes

Expectimax Search

Expectimax search: compute the average score under optimal play

- Max nodes as in minimax search
- Chance nodes are like min nodes but the outcome is uncertain
- Calculate their expected utilities
- i.e., take weighted average (expectation) of children

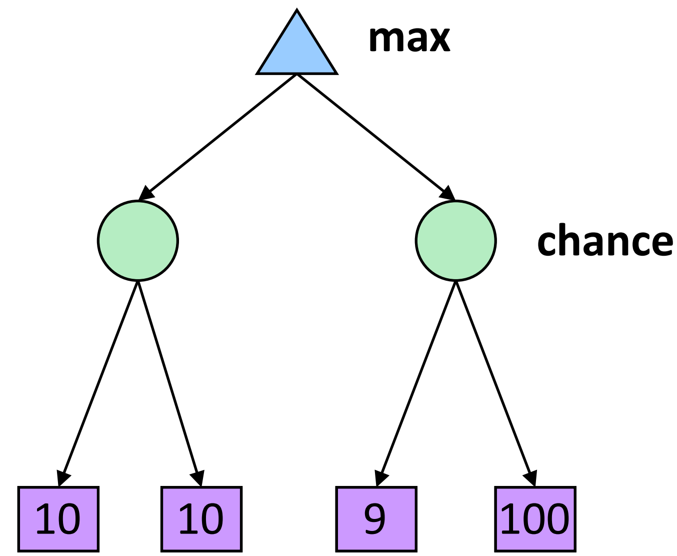

Demo: Minimax + Alpha-Beta Pruning

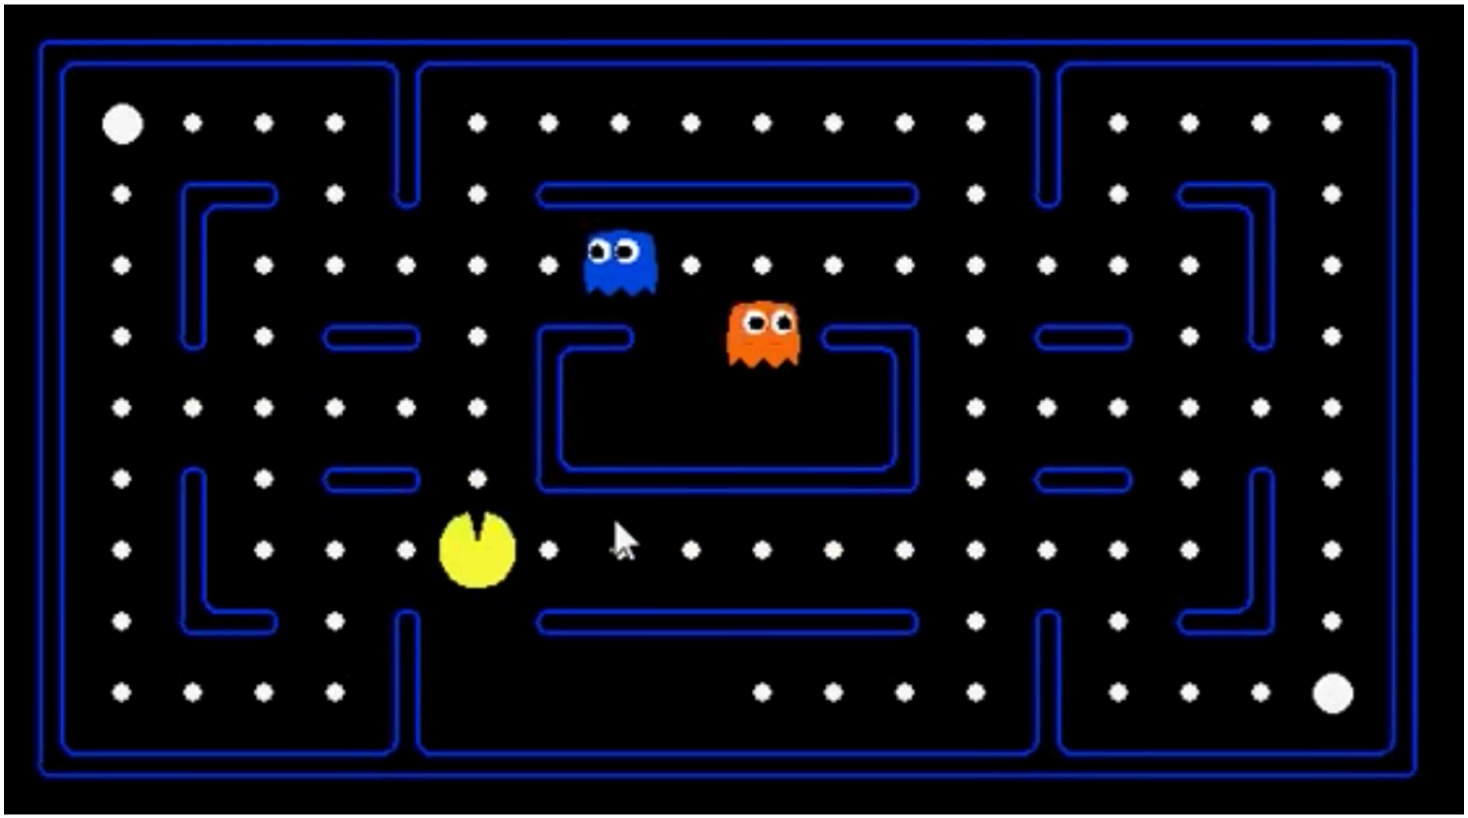

https://www.youtube.com/watch?v=_bEQJKXZ1-U

Demo: Expectimax

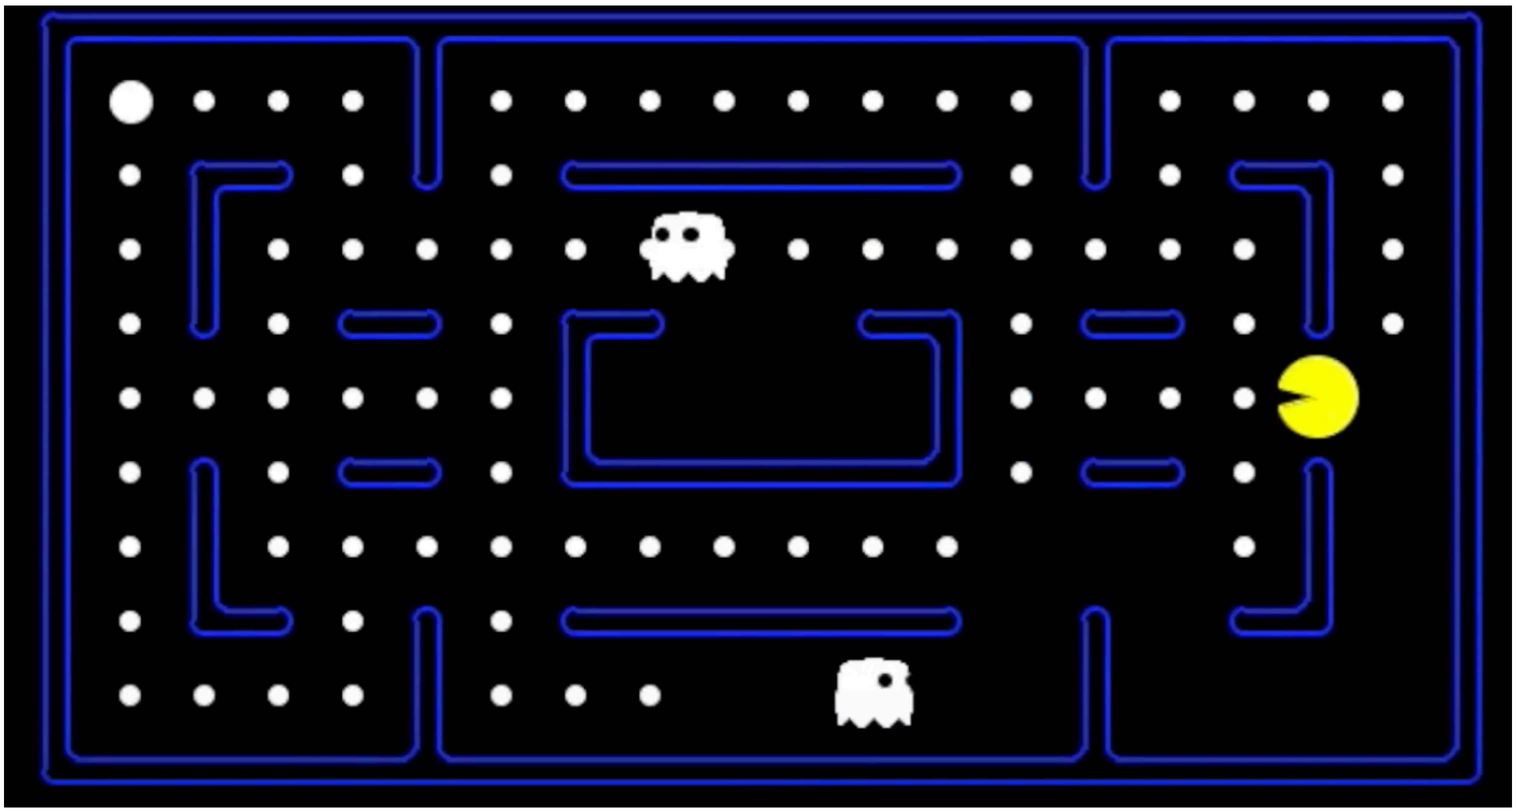

https://www.youtube.com/watch?v=ilxr3IAbpkw

Expectimax Pseudocode

```
def value(state):
    if the state is a terminal state: return the state's utility
    if the next agent is MAX: return max-value(state)
    if the next agent is EXP: return exp-value(state)
```

```
def max-value(state):
    initialize v = -∞
    for each successor of state:
        v = max(v,
            value(successor))
    return v
```

```
def exp-value(state):
    initialize v = 0
    for each successor of state:
        p = probability(successor)
        v += p * value(successor)
    return v
```

Slide credit: Dan Klein and Pieter Abbeel, UC Berkeley CS 188

Expectimax Pseudocode

```
def exp-value(state):
    initialize v = 0
    for each successor of state:
      p = probability(successor)
        v += p * value(successor)
    return v
```

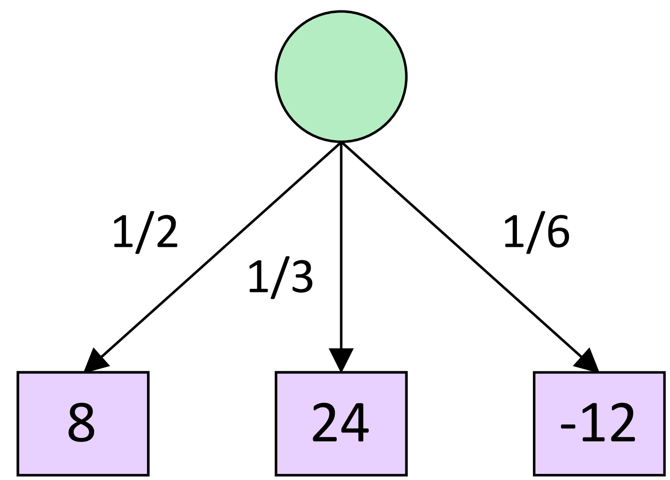

$$v = (1/2)\ (8) + (1/3)\ (24) + (1/6)\ (-12) = 10$$

Expectimax Example

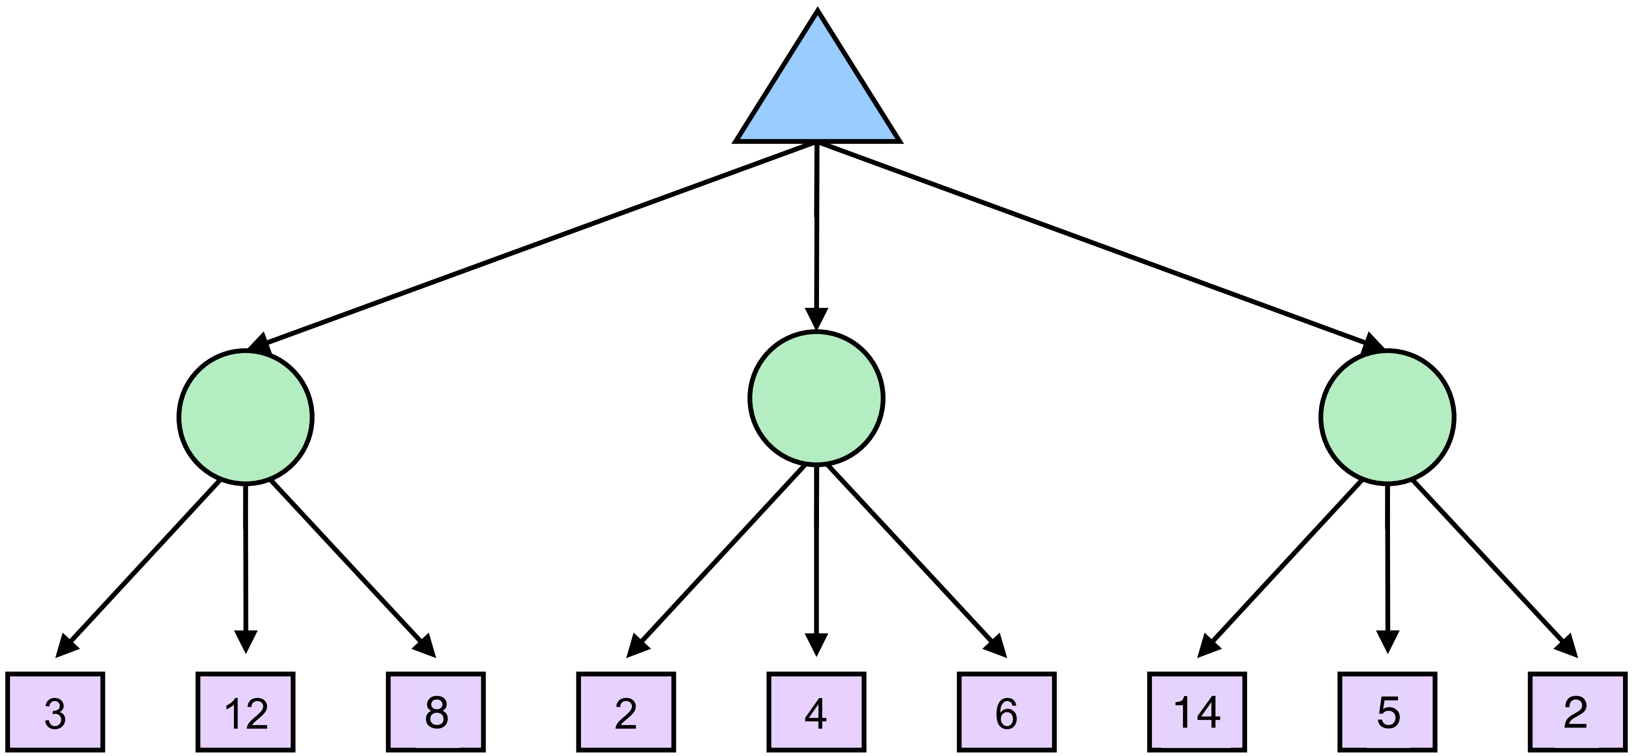

Image credit: Dan Klein and Pieter Abbeel, UC Berkeley CS 188

The Choice of Evaluation Function is Important

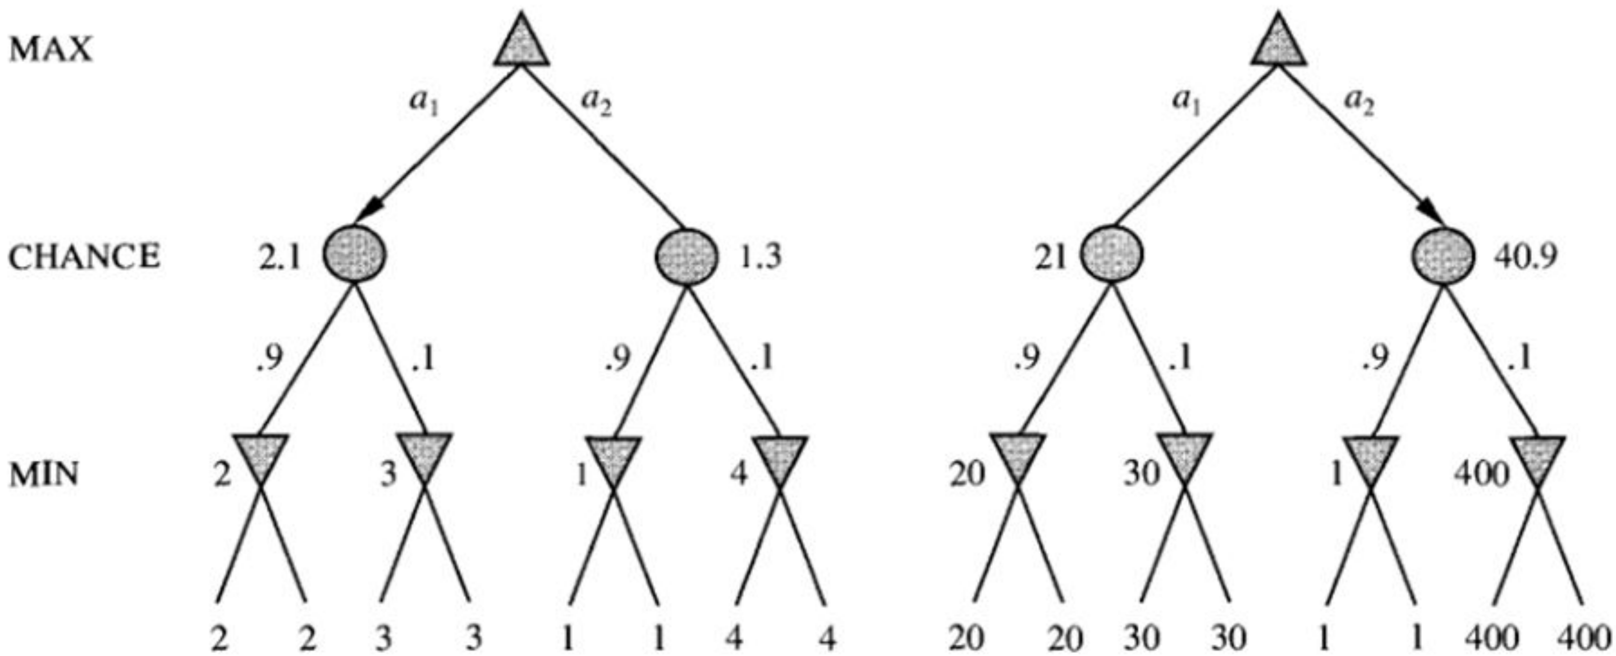

Image credit: Russell and Norvig

Expectimax Pruning?

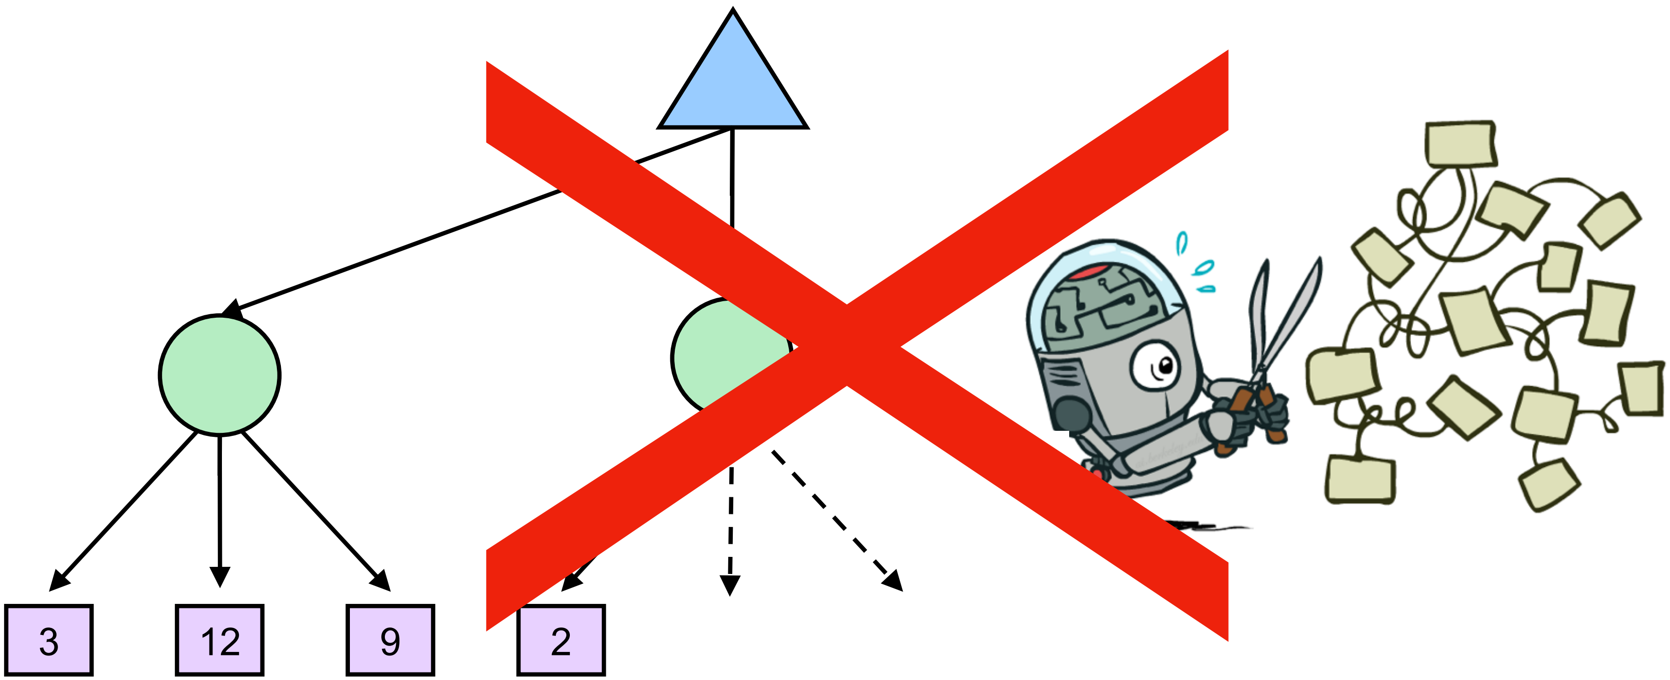

Image credit: Dan Klein and Pieter Abbeel, UC Berkeley CS 188

Depth-Limited Expectimax

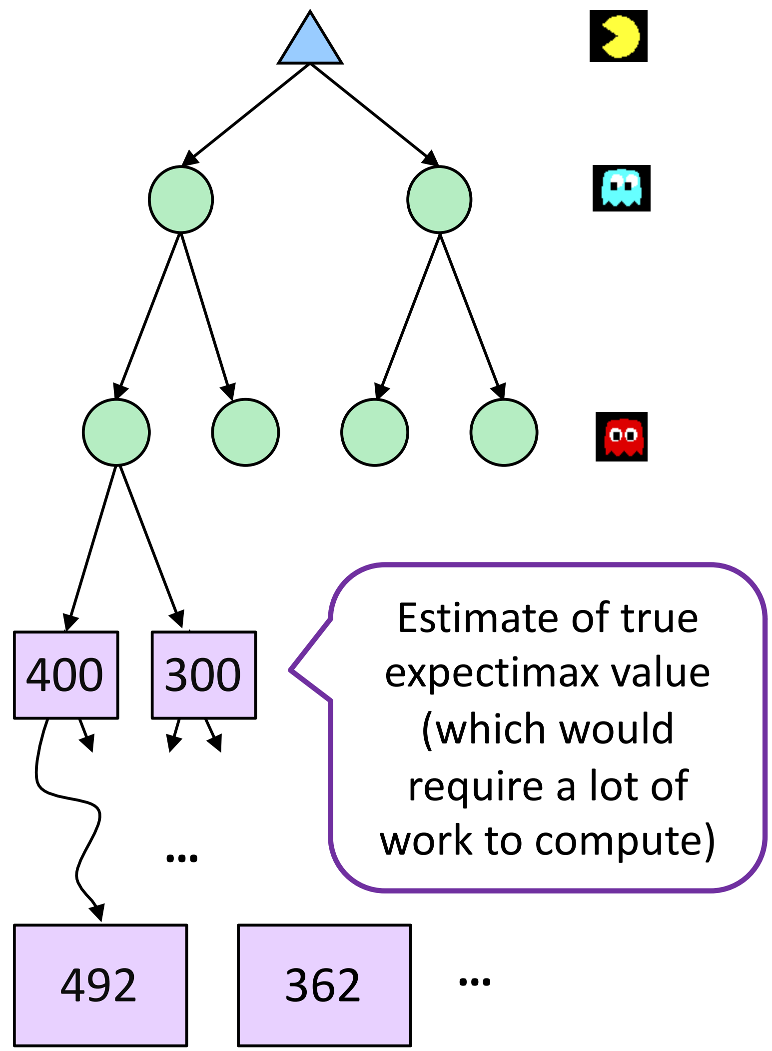

Partially Observable Games

Fog of War

“War is the realm of uncertainty; three quarters of the factors on which action in war is based are wrapped in a fog of greater or lesser uncertainty. A sensitive and discriminating judgment is called for; a skilled intelligence to scent out the truth”.

- Carl von Clausewitz

Example: Battleship

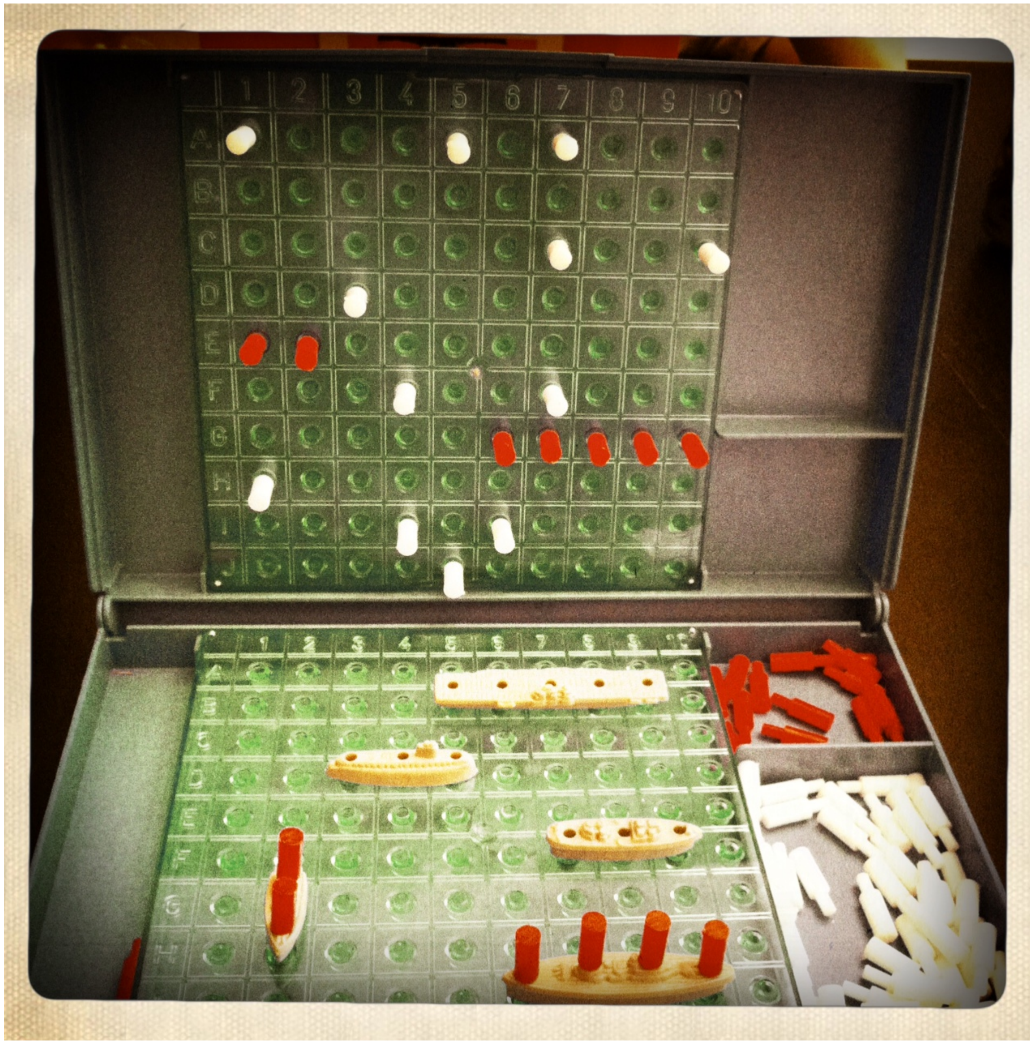

You Sank My Battleship!

Chess Variant: Kriegspiel

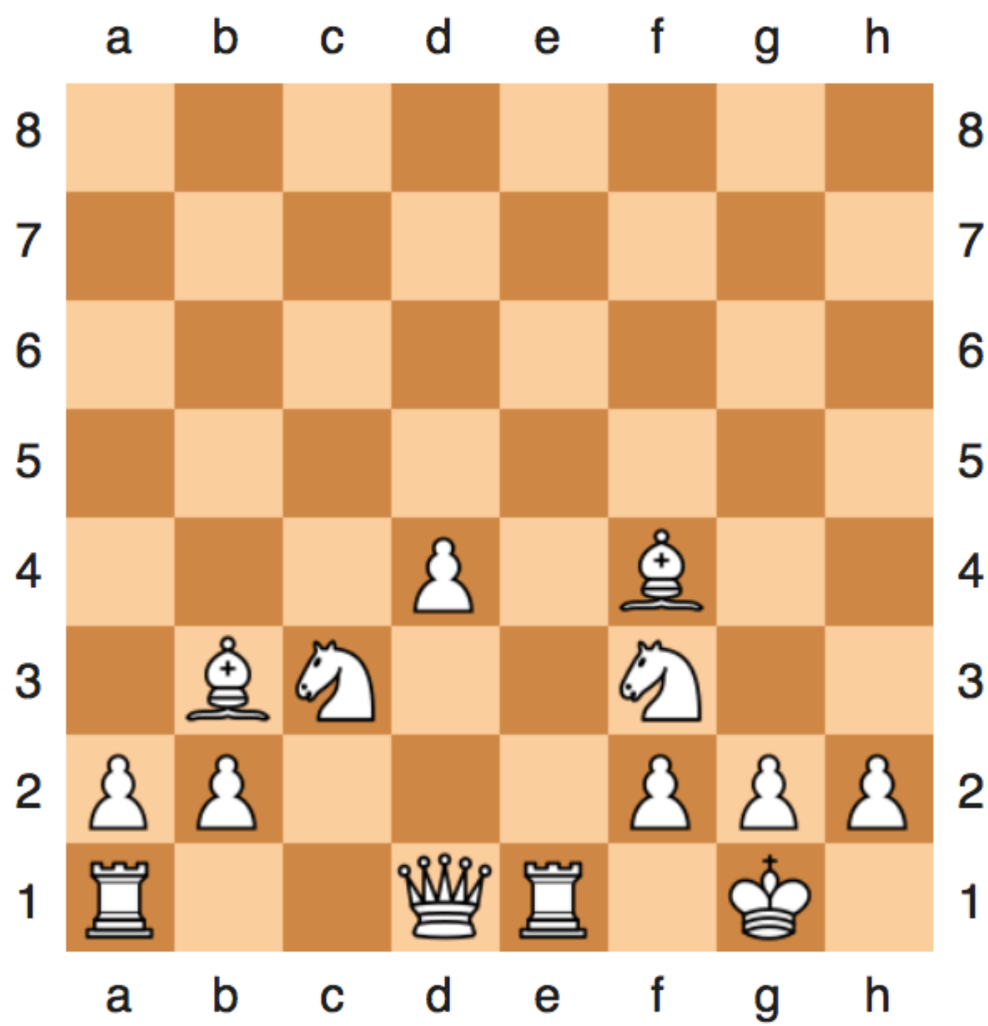

Belief States

Initially, White's belief state is a singleton because Black's pieces haven't moved yet

After White makes a move and Black responds:

- White's belief state contains 20 positions
- Because Back has 20 replies to **any** white move

KRK Endgame

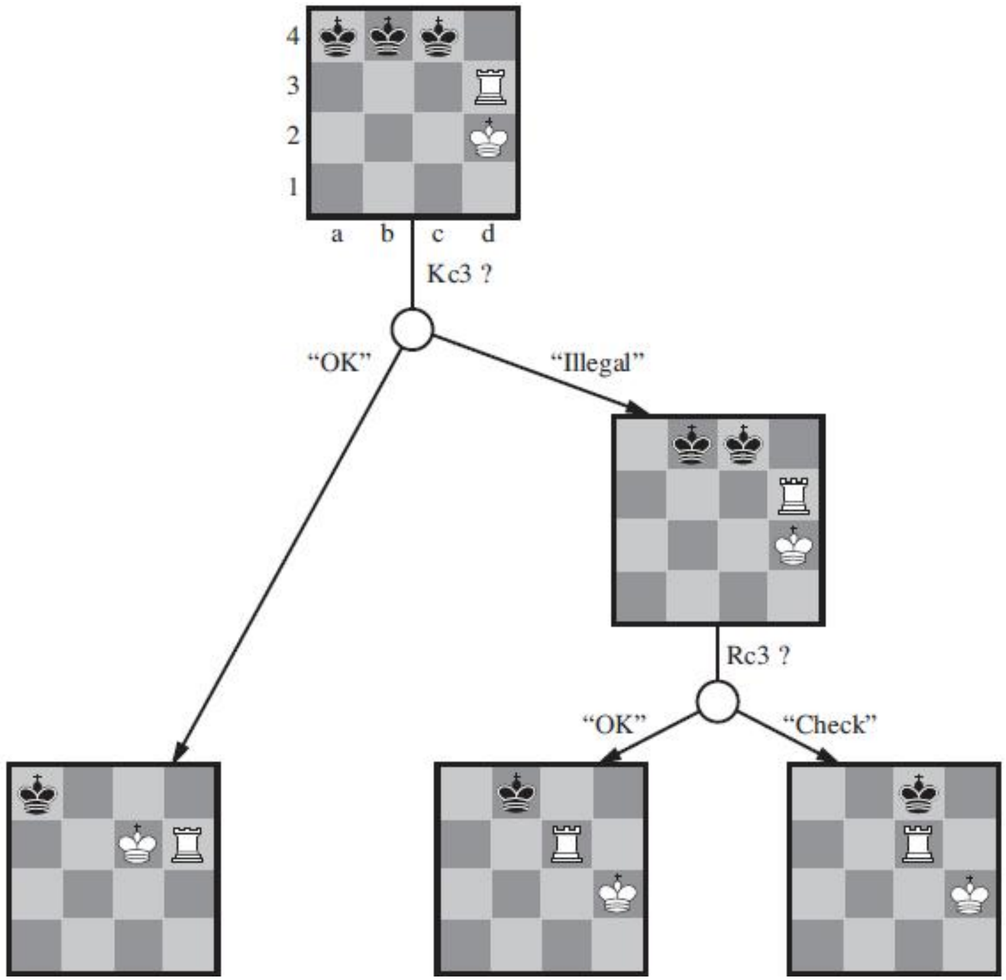

Image credit: Russell and Norvig

Card Games

Naive Assumption: Card Games are Just Like Dice Games

Algorithm: consider all possible deals of the invisible cards; solve each one as if it were a fully observable game

Then choose the move that has the best outcome averaged over all of the deals

Assume that each deal s occurs with probability $P(s)$

Naive Assumption: Card Games are Just Like Dice Games

$$\arg\max_{x} \sum_{s} P(s)\mathrm{MINIMAX}(\mathrm{RESULT}(s, a))$$

Run exact $\mathrm{MINIMAX}$ if computationally feasible

Otherwise run $\mathrm{H\text{-}MINIMAX}$

But the number of deals is very large

Monte Carlo Approximation: instead of adding up all the deals, take random sample of N deals

The probability of deal s appearing in the sample is proportional to $P(s)$:

$$\arg\max_{x} \frac{1}{N} \sum_{i=1}^{N} \text{MINIMAX}(\text{RESULT}(s_i, a))$$

Averaging Over Clairvoyance

Day 1: Road A leads to a heap of gold; Road B leads to a fork. Take the left fork and you'll find a bigger heap of gold, but take the right fork and you'll be run over by a bus.

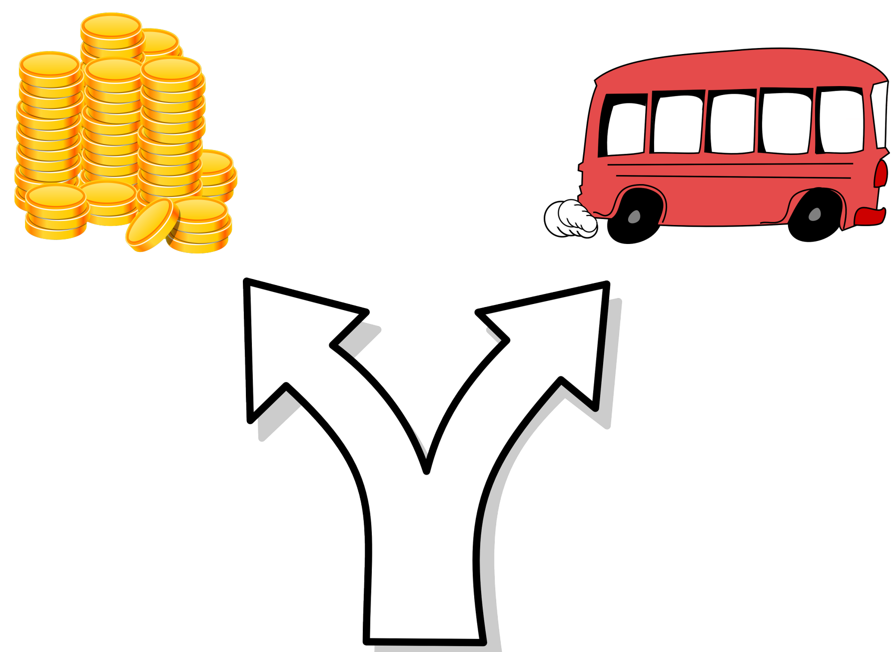

Averaging Over Clairvoyance

Day 2: Road A leads to a heap of gold; Road B leads to a fork. Take the right fork and you'll find a bigger heap of gold, but take the left fork and you'll be run over by a bus.

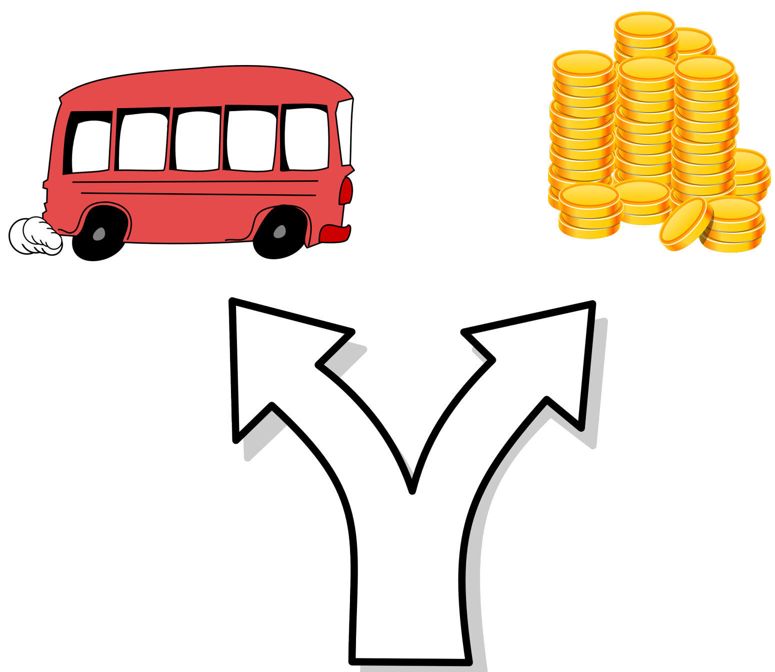

Averaging Over Clairvoyance

Day 3: Road A leads to a heap of gold; Road B leads to a fork. One branch of the fork leads to a bigger heap of gold, but take the wrong fork and you'll be hit by a bus. Unfortunately you don't know which fork is which.

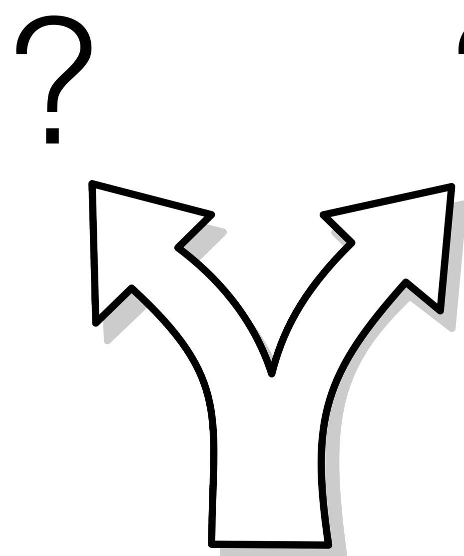

Averaging Over Clairvoyance's Answer

Day 1: B is the right choice

Day 2: B is the right choice

Day 3: B is still the right choice