Shut the Box

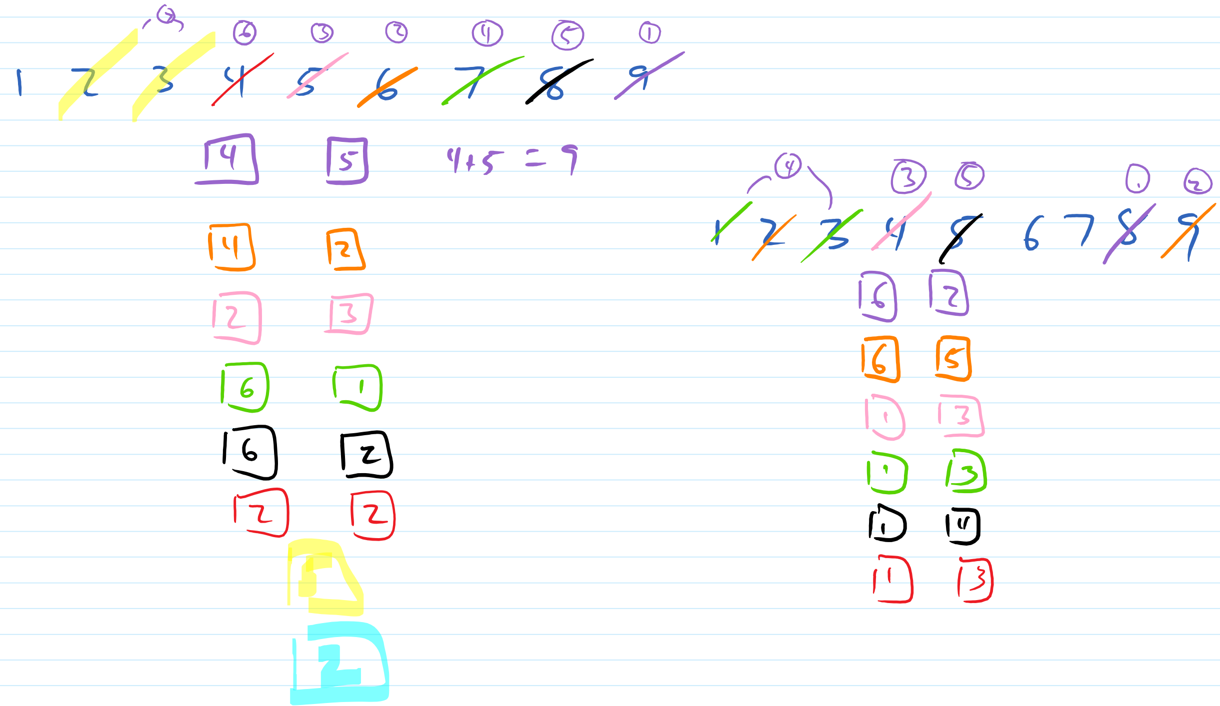

$E_2[(S,t)]$ = expected # wins for P2 with open tiles S, P1's score $= t$

$$= \begin{cases} 0.0 & \text{if } t = 0 \\ 1.0 & \text{if } \text{sum}(S) < t \\ \displaystyle\sum_{\text{rolls } r} P(r) \cdot \begin{cases} 0.0 & \text{if } \forall S' \subseteq S, \text{sum}(S') \neq r \text{ and } \text{sum}(S) > t \\ 0.5 & \text{if } \forall S' \subseteq S, \text{sum}(S') \neq r \text{ and } \text{sum}(S) = t \\ \displaystyle\max_{\substack{S' \subseteq S \\ \text{sum}(S') = r}} E_2[(S - S', t)] & \text{otherwise} \end{cases} & \text{otherwise} \end{cases}$$

$E_1[S]$ = expected wins for P1 with open tiles S

or constant-sum

Two-player Zero-sum, probabilistic finite games → players competing over same resource

value for 1 + value for 2 = 0
= c

For P1 choice position

$$E[pos] = \max E[next(pos, c)]$$

E[pos] = expected # wins for player 1

For P2 choice position

$$E[pos] = \cancel{\max}\ \min\ E[next(pos, c)]$$

For nonterminal random event positions

$$E[pos] = \sum_{\substack{\text{outcome} \\ \sigma}} P(\sigma) \cdot E[next(pos, \sigma)]$$

2-player Yahtzee anchors:

	P1	P2
Y	3	3
other cats	2^{12}	2^{12}
upper subtotal	64	64
score difference	3001 (-1500 … 1500)	

$(3 \cdot 2^{12} \cdot 64)^2 \cdot 3001$

≈ 100 trillion anchors

≈ 3000 yrs @ 1000 anchors/sec

2) compute strategy that maximizes the probability of beating the optimal solitaire player

Pig

2-players, turn-based

On each turn
- roll
- if 1, then turn over
- else add number to turn total
- decide: repeat
  - stop (and add turn total to score)

1^{st} to 100 points wins

$E[x,y] =$ P(next player wins, given that score is x to y)

47-36 $\xrightarrow{\text{roll } 2}$ 47-36 $\xrightarrow{\text{roll } 3}$ 47-36 (sub = 3) — roll 1 → back to 47-36

cycle → game is infinite
can't compute by plain backwards induction

Make the game finite: introduce a turn limit

as limit → ∞, P(reach limit) → 0
so opt strat w/ limit → opt strat w/o limit

for high enough even turn limit, P(reach turn limit) ≈ 0 and $E[x,y] \approx E[x,y,n]$

$E[x,y,n]$ = expected # wins for P1 given score is x to y w/ n turns left

even n → P1's turn
odd n → P2's turn

$$E[x,y,n] = \begin{cases} 1.0 & \text{if } x \ge T \text{ and } y < T \\ 0.0 & \text{if } y \ge T \text{ and } x < T \\ 0.5 & \text{if } n = 0 \\ E_{xyn}[0] & \text{otherwise} \end{cases}$$

$E[x,y]$ depends on $E[y,x]$,
$E[y,x]$ depends on $E[x,y]$

so $E[x,y]$, $E[y,x]$ are fns of themselves

value iteration: guess values
- perform calculations
- update guesses
- perform calculations
- update guesses
- ⋮

repeat until converged

$E_{xyn}[t]$ = expected wins for P1 when score is x to y, n turns left, turn total = t

$$= \begin{cases} 1.0 & \text{if } n \text{ even and } t+x \ge T \\ 0.0 & \text{if } n \text{ odd and } t+y \ge T \\ \max\Big(\underbrace{E[x+t,y,n-1]}_{\text{value of stopping}}, \underbrace{\tfrac{1}{6}\cdot E[x,y,n-1]}_{\text{roll a 1, lose progress}} + \tfrac{1}{6}\sum_{r=2}^{6} E_{xyn}[t+r]\Big) & \text{if } n \text{ even} \\ \min\Big(E[x,y+t,n-1], \tfrac{1}{6}\cdot E[x,y,n-1] + \tfrac{1}{6}\sum_{r=2}^{6} E_{xyn}[t+r]\Big) & \text{if } n \text{ odd} \end{cases}$$

initialize $E[x,y,0]$ according to base cases of recurrence

$n \leftarrow 0$
repeat

initialize $E[x,y,0]$ according to base cases of recurrence

$n \leftarrow 0$
repeat
for $x \leftarrow 0$ to T
for $y \leftarrow 0$ to T

$$E_{xyn}[t] \leftarrow \begin{cases} 1.0 & \text{for } t \geq T-x \text{ if } n \text{ even} \\ 0.0 & \text{for } t \geq T-y \text{ if } n \text{ odd} \end{cases}$$

for $t = \begin{cases} T-x-1 & \text{if } n \text{ even} \\ T-y-1 & \text{if } n \text{ odd} \end{cases}$ downto 0

$E_{xyn}[t] \longleftarrow$ value from recurrence

$E[x,y,n] \leftarrow E_{xyn}[0]$

$n \leftarrow n+1$

until n odd and $n \geq 3$ and $E[x,y,n-1]$ close enough to $E[x,y,n-3]$ for all x,y

OR (value iteration on each pair of entries separately)

$E[x,y]$ = expected wins for next player when score is $x-y$

could have $E_1[x,y]$ = expected wins for P1 when score is $x-y$, and P_1 is next
$E_2[x,y]$ = expected wins for P1 when score is $x-y$, and P_2 is next
but then $E_2[x,y] = 1 - E_1[y,x]$, so only need E_1
(and just call it E)

$E[T,y] \leftarrow 1.0$ for all y
$E[x,T] \leftarrow 0.0$ for all x

$E[x,y]$, $E[y,x]$ depend on each other, so compute both at same time

for (x,y) in $\{0 \ldots T\} \times \{0 \ldots T\}$ s.t. $x \geq y$ in order of $\downarrow x+y$

E[x,y]
E[y,x]

repeat
$E_{xy}[t] \leftarrow 1.0$ for $t \geq T-x$
$E_{yx}[t] \leftarrow 1.0$ for $t \geq T-y$

account for change of turn

for $t \leftarrow T-x-1$ downto 0

$$E_{xy}[t] \longleftarrow \max\left(1.0 - E[y,x+t],\ \frac{1}{6}\left(1.0 - E[y,x] + \sum_{r=2}^{6} E_{xy}[t+r]\right)\right)$$

for $t \leftarrow T-y-1$ downto 0

$$E_{yx}[t] \longleftarrow \max\left(1.0 - E[x,y+t],\ \frac{1}{6}\left(1.0 - E[x,y] + \sum_{r=2}^{6} E_{yx}[t+r]\right)\right)$$

$E[x,y] \leftarrow E_{xy}[0]$
$E[y,x] \leftarrow E_{yx}[0]$

until $E_{xy}[0]$ and $E_{yx}[0]$ have converged