Centipede Game

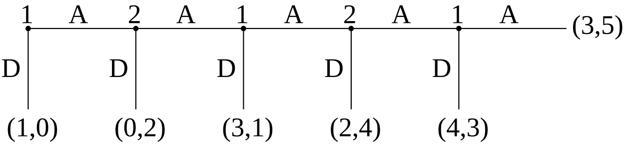

- Play this as a fun game...

Computing Subgame Perfect Equilibria

Idea: Identify the equilibria in the bottom-most trees, and adopt these as one moves up the tree

function BACKWARDINDUCTION (node h) **returns** $u(h)$
if $h \in Z$ **then**
  return $u(h)$
$best\_util \leftarrow -\infty$
forall $a \in \chi(h)$ **do**
  $util_at_child \leftarrow$BACKWARDINDUCTION$(\sigma(h, a))$
  if $util\_at\_child_{\rho(h)} > best\_util_{\rho(h)}$ **then**
    $best\_util \leftarrow util\_at\_child$
return $best\_util$

- $util\_at\_child$ is a vector denoting the utility for each player
- the procedure doesn't return an equilibrium strategy, but rather labels each node with a vector of real numbers.
  - This labeling can be seen as an extension of the game's utility function to the non-terminal nodes
  - The equilibrium strategies: take the best action at each node.

Computing Subgame Perfect Equilibria

Idea: Identify the equilibria in the bottom-most trees, and adopt these as one moves up the tree

function BACKWARDINDUCTION (node h) **returns** $u(h)$
if $h \in Z$ **then**
  return $u(h)$
$best\_util \leftarrow -\infty$
forall $a \in \chi(h)$ **do**
  $util\_at\_child \leftarrow$ BACKWARDINDUCTION$(\sigma(h, a))$
  if $util\_at\_child_{\rho(h)} > best\_util_{\rho(h)}$ **then**
    $best\_util \leftarrow util\_at\_child$
return $best\_util$

- For zero-sum games, BACKWARDINDUCTION has another name: the minimax algorithm.
  - Here it's enough to store one number per node.
  - It's possible to speed things up by pruning nodes that will never be reached in play: "alpha-beta pruning".

Backward Induction

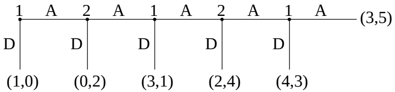

- What happens when we use this procedure on Centipede?
  - In the only equilibrium, player 1 goes down in the first move.
  - However, this outcome is Pareto-dominated by all but one other outcome.
- Two considerations:
  - practical: human subjects don't go down right away
  - theoretical: what should you do as player 2 if player 1 doesn't go down?
    - SPE analysis says to go down. However, that same analysis says that P1 would already have gone down. How do you update your beliefs upon observation of a measure zero event?
    - but if player 1 knows that you'll do something else, it is rational for him not to go down anymore... a paradox
    - there's a whole literature on this question