Homework 5.4

- Recall the following theorem from the "overconfidence versus paranoia" lecture:
- **Theorem**. In a zero-sum perfect-information game, both the paranoid and overconfident models will lead you to choose the same strategy s_1*
- Professor Prune claims that the following game tree is a counterexample to the theorem. Is he right? Why or why not? Can you suggest any alternative interpretations that might resolve the problem?

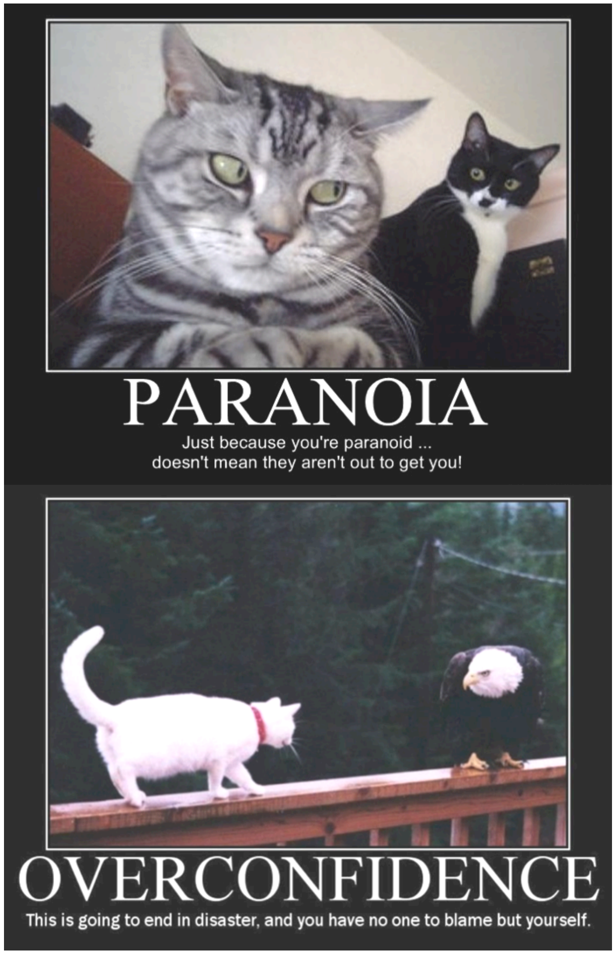

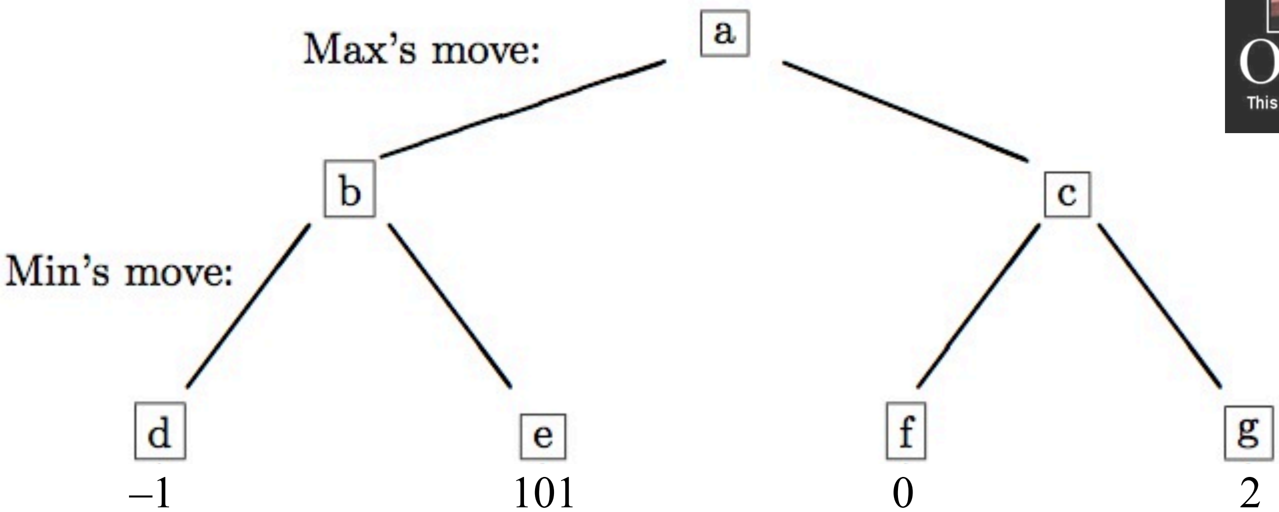

Homework 6.1

In repeated games, which of the following strategies are stationary?

- ***AllC***: always cooperate
- ***AllD*** (the *Hawk* strategy): always defect
- ***Grim***: cooperate until the other agent defects, then defect forever
- ***Tit-for-Tat (TFT)***: cooperate on the first move. On the nth move, repeat the other agent $(n–1)^{\text{th}}$ move
- ***Random***: Randomly intersperse cooperation and defection
- ***Tester***: defect on move 1. If the other agent retaliates, play TFT. Otherwise, randomly intersperse cooperation and defection
- ***Tit-For-Two-Tats (TFTT)***: Like Tit-for-Tat, except that it retaliates only if the other agent defects twice in a row
- ***Generous Tit-For-Tat (GTFT)***: Like Tit-for-Tat, but with a small probability $c>0$ of cooperation on the n'th move if the other agent defected on the $(n–1)^{\text{th}}$ move
- ***Pavlov (Win-Stay, Lose-Shift)***: Repeats its previous move if it earns 3 or 5 points in the previous iteration, and reverses its previous move if it earns 0 or 1 points in the previous iteration

Homework 6.2

- Recall that DBS's opponent model is as follows:
  - A set of rules of the following form

    if our last move was m and their last move was m'
    then P[their next move will be C]

    - Four rules: one for each of (C,C), (C,D), (D,C), and (D,D)
  - For example, TFT can be described as
    - (C,C) $\Rightarrow$ 1, (C, D) $\Rightarrow$ 1, (D, C) $\Rightarrow$ 0, (D, D) $\Rightarrow$ 0

(a) Which of the strategies in Homework 6.1 can be represented by DBS's opponent model?

(b) Are there any stationary strategies that DBS's opponent model *cannot* represent? Why or why not?

(b) Are there any stationary strategies that DBS's opponent model *cannot* represent? Why or why not?

Homework 6.3

- Let S be the following set of interaction traces for the iterated battle of the sexes with 6 iterations.
  (a) What are the compatible subsets of S?
  (b) For each compatible subset, what is the composite trace and its expected utility if we assume agents a_2, a_4, a_6, and a_8 are equally likely?

Trace 1		Trace 2		Trace 3		Trace 4	
a_1	a_2	a_3	a_4	a_5	a_6	a_7	a_8
G	T	G	T	G	T	G	T
T	G	T	T	G	T	G	T
G	T	T	T	T	G	G	G
T	G	G	T	T	G	T	G
G	T	G	T	G	T	T	G
T	G	G	T	G	T	T	G

Homework 6.4

(a) What is the expectiminimax value of the following game tree?

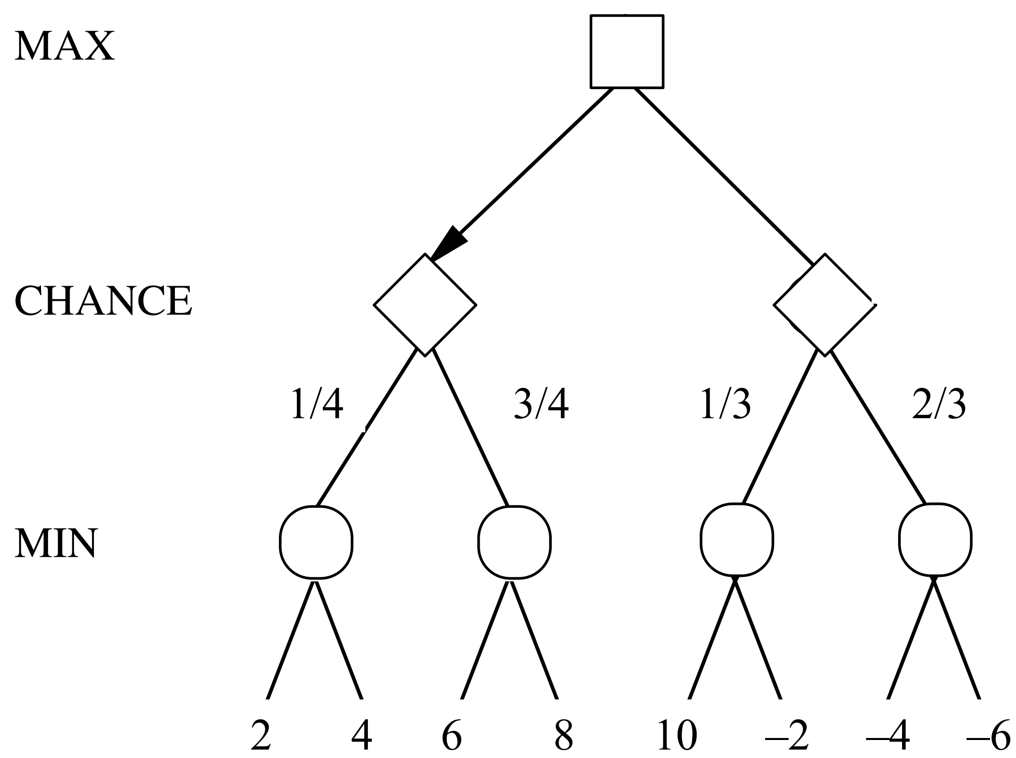

Homework 6.5

(b) Do we need to know *z* in order to know the expectiminimax value of the game tree? Why or why not?

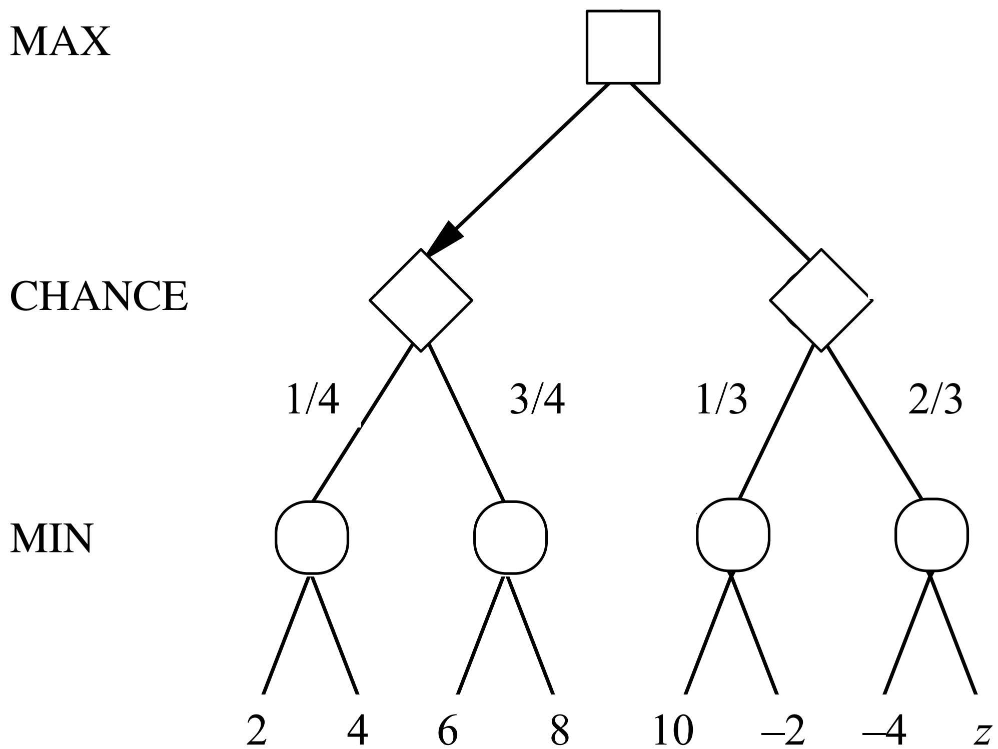

Homework 6.5

(c) Do we need to know *y* and *z* in order to know the expectiminimax value of the game tree? Why or why not?

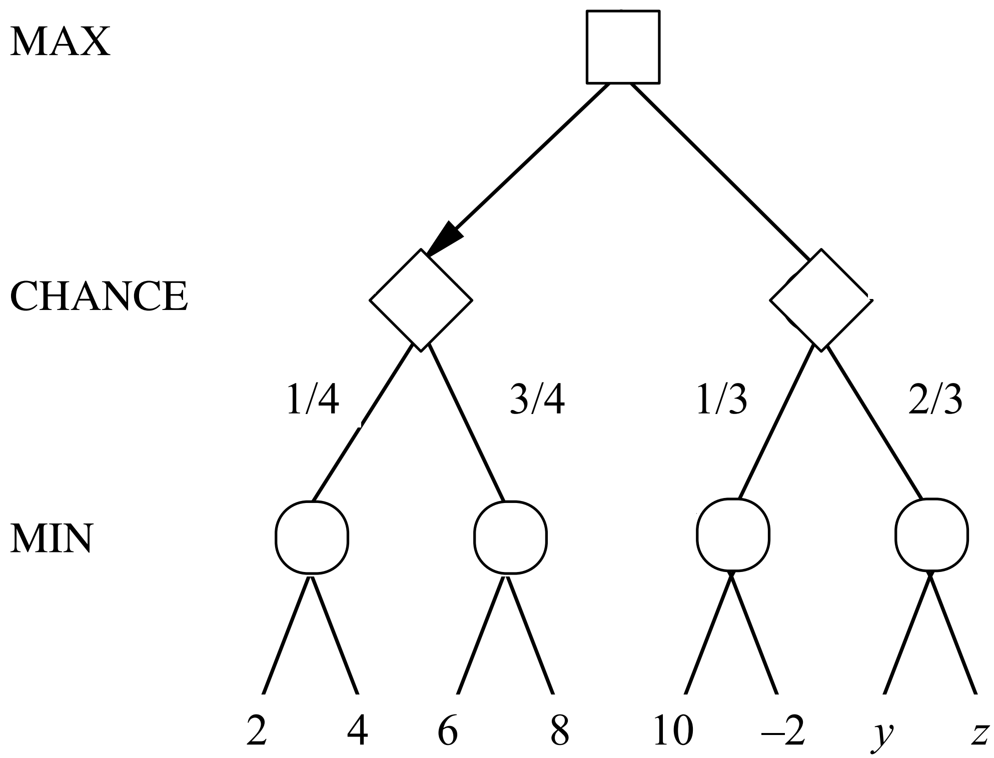

Homework 6.5

(d) Suppose we know in advance that every terminal node has a value in the range [–10, +10]. For what values of x (if any) will y and z have no effect on the expectiminimax value of the tree?

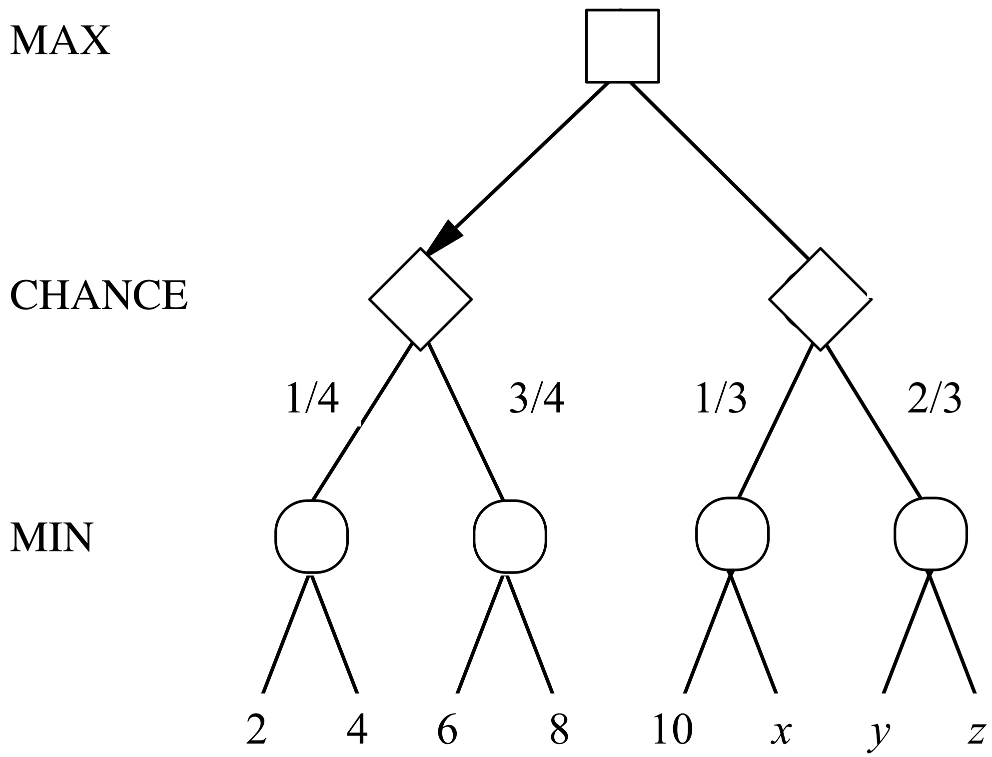

Homework 6.6

- Consider the double lottery game with the imitate-the-better dynamic.

  Suppose agent *a*'s strategy is *R* and agent *b*'s strategy is *RwS*. If we compare the two agents' payoffs in order to decide which of them reproduces, then what are their probabilities of reproducing?