

PROBLEM SET #1: REVIEW OF DECISION PRINCIPLES

Maximax Principle (“aim for the best — and ignore everything else”):

- (1) find the maximum payoff in each row, and
- (2) choose the row with the maximum of the maximums.

Every row has a maximum payoff, and one of these maximum payoffs must be the overall maximum payoff, so a player always has a maximax strategy. However, several rows may be tied with the same overall maximum payoff, so a player may not have a unique maximax strategy. (The player’s several maximax strategies probably are not equally good with respect to other decision principles.)

Maximin Principle (“avoid the worst — and ignore everything else”):

- (1) find the minimum payoff in each row (its security level), and then
- (2) choose the row with the maximum of the minimums (the highest security level).

Every row has a minimum payoff, and one of these minimums payoffs must be the maximum of the minimum payoffs, so a player always has a maximin strategy. However, several rows may be tied with the same maximin payoff, so a player may not have a unique maximin strategy. (The player’s several maximin strategies probably are not equally good with respect to other decision principles.)

Maximize Average Payoff (“don’t focus on the best outcome or the worst outcome but on the average outcome”):

- (1) add up all the payoffs in the row,
- (2) divide by the number of contingencies, and
- (3) choose the row with the highest average.

Every row has an average payoff, and one of these average payoffs must be the maximum average payoff, so a player always has a strategy that maximizes average payoff. However, several rows may be tied with the same maximum average, so a player may have several strategies that maximizes average payoff.

Maximize Expected Payoff (i.e., take account of the differing probabilities of contingencies):

- (1) determine (or make a subjective estimate of) the probability p_k of each contingency c_k (where $p_k > 0$ and $\sum p = 1$),
- (2) multiply each payoff by its probability and add these products together to get the expected payoff of each row, and
- (3) choose the row with the highest expected payoff.

Note: average and expected payoff are the same in the event each contingency has equal probability.

For any probability distribution over contingencies, every row has an expected payoff, and one of these average payoffs must be the maximum expected payoff, so a player always has a strategy that maximizes expected payoff. However, several rows may be tied with the same maximum expectation, so a player may have several strategies that maximizes expected payoff.

Dominance (or “Sure Thing”) **Principle**:

Definition: strategy s_k *dominates* strategy s_h if

- (1) s_k gives at least as high payoff as s_h in every contingency, and
- (2) s_k gives a higher payoff than s_h in at least one contingency.

Basic Principle: don't choose a dominated strategy (or, always choose an undominated strategy).

It is quite possible that no strategy dominates another, in which case every strategy is undominated. On the other hand, suppose s_1 is dominated. Then there must be another strategy s_2 that dominates s_1 . Remember that this means that s_2 has at least as great a payoff as s_1 in every contingency and a greater payoff in at least one contingency. Now suppose s_2 is dominated. Then there must be another strategy s_3 that dominates both s_2 and s_1 , because domination (like “greater than”) is a *transitive* relation. If the number of strategies is finite, we must find (at least) one strategy that is undominated.

Definition: a strategy is *dominant* if it dominates every other strategy.

Corollary Principle: always choose a dominant strategy, if you have one.

Two or more strategies may be undominated, in which case none is dominant. A player clearly can have at most one dominant strategy.

After the player has chosen his strategy and discovers what contingency nature has chosen, the player may regret his strategy choice and wish he had chosen a different strategy. Can you identify the condition under which a player would never have reason to regret his strategy choice?

For each contingency that may arise, the player has a strategy that is a *best reply*, i.e., that gives the best payoff in that contingency (column). (Two or more strategies may be “tied” in this respect, in that they all give same [best] payoff in that contingency.) A player has no reason to regret his strategy choice in a given contingency if and only if that strategy is a best reply in that contingency. A player has no reason to regret his strategy choice in any contingency if and only if that strategy is a best reply in every contingency. You can verify that a strategy that is a best reply in every contingency is a dominant strategy. So the condition under which a player would never have reason to regret his strategy choice is when the player has a dominant and uses it.

Which decision principles do you think are most justifiable? Least justifiable?

In games against nature (as opposed to games against other rational players with their own interests):

- (a) Maximax is unduly optimistic (nature isn't benevolent).
- (b) Maximin is unduly pessimistic (nature isn't malevolent [even if it sometimes seems otherwise]).
- (c) Maximize average payoff probably doesn't make sense because we always have some sense as to which contingencies are more and less probable.
- (d) Maximize expected payoffs makes the best sense (provided that we have some

sense as to what the probabilities are and the payoff numbers represent an appropriate kind of cardinal utility)

- (e) It's very hard to see anything wrong with principle that you should never use a dominated strategy (at least in a game against nature). The problem is that it's a very weak principle that usually doesn't tell you what strategy to choose. You can also check that the Dominance Principle (because it is so weak) never conflicts with any of the other Decision Principles, even the ones that don't make much sense.

Preview to think about only (no written answer required). Suppose that "nature" is no longer a disinterested player but rather that nature is "malevolent" and is trying to hold the player's payoff down to a minimum. Under this circumstance, do some decision principles become more justifiable? Less justifiable?

If nature is "malevolent" and is trying to hold your payoff down to a minimum, you are in a *two-player zero-sum game*. Expected Payoff Maximization is no longer directly usable, because your opponent will choose his strategy on the basis of rational principles, not probability. (However, as we will see, these rational principles imply that players should delegate their strategy selection to a chance mechanism in some types of ["non-strictly determined"] zero-sum games.) The Maximin Principle, which is too pessimistic in a game against nature (or non-zero-sum games with other players), is much more justifiable in such a context of total conflict. And the Dominance Principle (for what's worth) still makes sense in a zero-sum situation.