

Basic Choice Strategies

The most interesting cases for thinking about the kind of rational choice we're interested in involve some *uncertainty* in how the world is. To start to get a fix on this, let's try to describe principles governing rational choices of this kind.

A choice situation can be represented by a grid. The leftmost column lists the *actions* the agent can choose among. The topmost row lists the possible *states* that the world might be in that could affect the choice-worthiness of the action. For each action-state pair, we can use a number to represent how valued the action-state pair is to the agent. For example:

	<i>Rain</i>	<i>Sunshine</i>
<i>Watch TV</i>	3	2
<i>Go to Beach</i>	1	4

Dominance

The simplest decision rule for such a decision problem relies on the following notions of *dominance*.

Option A ***strongly dominates*** option B if for every state, choosing A leads to better outcomes than choosing B.

Option A ***weakly dominates*** option B if for every state, choosing A leads to outcomes at least as good as choosing B, and at least sometimes better than choosing B.

	<i>History</i>	<i>Sports</i>	<i>Geography</i>	<i>Entertainment</i>	<i>Literature</i>
<i>Alice</i>	8	3	6	4	3
<i>Bert</i>	1	3	2	2	1
<i>Carl</i>	5	2	5	3	2
<i>Dan</i>	2	6	4	3	2
<i>Eve</i>	8	3	7	4	4
<i>Fran</i>	4	4	4	4	4

For example:

Choosing Alice strongly dominates choosing Carl.

Choosing Alice weakly dominates choosing Bert.

Our first rules:

Strong Dominance Rule: Never choose strongly dominated options.

Weak Dominance Rule: Never choose weakly dominated options.

If either version of dominance is going to apply, we have to be a little careful. Consider the following problem.

	<i>Pass</i>	<i>Fail</i>
<i>Study</i>	3	1
<i>Party</i>	4	2

Strong and Weak Dominance here tell you it's *always* irrational to study. What's the problem? The states aren't *independent* of the actions. Some actions make certain states more likely than others. Dominance (and other principles) will only reasonably apply when this isn't so. So we want to be careful to set up our decision tables to keep choices and outcomes independent.

Maximax, Minimax

Dominance reasoning will only get you so far. To get more committal, we need to get more controversial. Consider:

Maximax Rule: Choose only options which for at least one state bring about the best possible outcome.

Maximax tells you to ignore all but the maximum value you could get in each choice, then to choose only from the best of those. Here's another:

Maximin Rule: Choose only options whose worst outcome is at least as good as the worst outcome of every other choice.

Maximin tells you to ignore all but the minimum value you could get in each choice, then to choose only from the best of those.

Ordinal v. Cardinal Utilities

We've been using numbers to represent the extent to which an outcome accords with an agent's values. Let's call these *utilities*. There are two ways of thinking of utilities.

Ordinal Utility: records only the order of an agent's preferences in various outcomes.

Cardinal Utility: records not only the order, but the *relative magnitude*, of an agent's preferences in various outcome.

Suppose you're considering whether to try a drug D to alleviate nausea which, if and only if you have rare condition C, will kill you as a “side effect”. If we are only concerned with ordinal utility we might write:

	<i>You have C</i>	<i>You don't have C</i>
<i>Take D</i>	1	3
<i>Not take D</i>	2	2

But cardinal utilities, which tell us more, might be written as follows

	<i>You have C</i>	<i>You don't have C</i>
<i>Take D</i>	-1000	10
<i>Not take D</i>	-10	-10

Usually in this class we'll work with cardinal utilities. Our final rule makes use of these.

Regret

Let's say

The **regret of choice C given state S** is the utility of the best choice at S minus the utility provided by C at S.

So the regret of *take D* given *you have C* is $-10 - (-1000) = 990$. (that's a lot of regret)

The regret of *take D* given *you don't have C* is $10 - 10 = 0$ (you regret nothing!)

We can actually make a “regret” table which summarizes these results:

(Regret)	<i>You have C</i>	<i>You don't have C</i>
<i>Take D</i>	990	0
<i>Not take D</i>		

Now we can apply rules to this “converted” table.

Minimax Regret Rule: Choose only options whose maximum regret is the least possible.

The Minimax Regret Rule has a funny property. Consider the case Weatherson discusses.

	Sunny	Light Rain	Thunderstorm
Picnic	20	5	0
Baseball	15	2	6
Movies	8	10	9

The regret table looks like this:

	Sunny	Light Rain	Thunderstorm
Picnic	0	5	9
Baseball	5	8	3
Movies	12	0	0

What does Minimax Regret recommend?

But suppose you learn the park is closed, so you can't go picnic. Your new table is:

	Sunny	Light Rain	Thunderstorm
Baseball	15	2	6
Movies	8	10	9

With regret...

	Sunny	Light Rain	Thunderstorm
Baseball	0	8	3
Movies	7	0	0

Now what does it recommend?

What happened? The rule violates an important condition:

Independence of Irrelevant Alternatives: If C is the best choice among options in S, then C will *still* be the best choice if you remove options from S other than C.

To many people this seems like an important constraint on rational choice. Is it?