1999, STEP II, q14 solution

You play the following game. You throw a six-sided fair dice repeatedly. You may choose to stop after any throw, except you must stop if you throw a 1. Your score is the number obtained on the last throw. Determine the strategy you should adopt in order to maximise your expected score.

Do we understand the question. Technical notion is"expected score". But we dont know how describe strategies, so difficult to get started.

Can we give an example of a strategy? Just throw the dice once!

If we throw the dice once then our score is equally likely to be any of $\{1, 2, 3, 4, 5, 6\}$, and our expected score is therefore 3.5.

How could we improve on this? Suppose we throw a 2 or 3 then this score is less that the expected score if we choose to thow the dice a second time. On the otherhand if we have a $4, 5, 6$ then this more than we could expect from a second roll. So our improved strategy is if we get a 2 or 3 on the first roll then we should roll again, otherwise stop.

Can we improve it anymore? We can make the same argument again, and if on the second roll we get a 2 or3 we roll a third time and so on.

So we guess a good strategy is continue rolling until we first get a $1, 4, 5, 6$, then stop.

By virtue of *symmetry* the final score following this strategy is equally likely to any of $\{1, 4, 5, 6\}$ so our expected score is $(1 + 4 + 5 + 6)/4 = 4$.

Now lets generalize: Consider a strategy that continues until we first get a score in some given subset $S \subseteq \{1, 2, 3, 4, 5, 6\}$, which must contain 1.

If $S = \{1, 6\}$, the expected score is $(1 + 6)/2 = 3.5$.

If $S = \{1, 5, 6\}$, the expected score is $(1 + 5 + 6)/3 = 4$.

If $S = \{1, 4, 5, 6\}$, the expected score is $(1 + 4 + 5 + 6)/4 = 4$.

If $S = \{1, 3, 4, 5, 6\}$, the expected score is $(1 + 3 + 4 + 5 + 6)/5 = 3.8$.

With all six elements in S the expected score will be 3.5.

So amongst these strategies there are two ways of achieving a maximal expected score of 4!

Lets reflect on our answer. Is it complete? We still need to justify the best strategy overall of being of this form stop when you first get a value in S.

Suppose the optimal strategy has a expected score V. Roll the dice once, and suppose we obtain a score s, we may choose between stopping and accepting s as our score (which we must do if $s = 1$), or continuing. If we continue then the first roll of the dice has no influence on the subsequent rolls, and following the optimal strategy our expected score will be V. So acting optimally in deciding whether to continue we will:

$$\text{Stop if } s \geq V, \text{ or if } s = 1,$$
$$\text{Continue if } s \leq V \text{ and } s \neq 1.$$

If $s = V$ then it doesnt matter if we continue or not. The same argument applies at later throws of the dice, and so we see the optimal strategy is of the form we have considered.