

1999, STEP II, q14 solution

Question: *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.*

Solution (JW):

Do we understand the question. Technical notion is “expected score”. But we don’t know how describe strategies, so it’s 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 than the expected score if we choose to throw the dice a second time. On the other hand, if we have a 4, 5, 6 then this is more than we could expect from a second roll. So our improved strategy is as follows: 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 or 3 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 yield 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 achieve a maximal expected score of 4!

Let’s 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 an 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:

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

If $s = V$ then it doesn’t 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.

Comments, modifications and relationship to content of degree courses offered by Warwick Statistics (JAB):

- This STEP problem is not an easy question for Sixthformers. It requires an idea or an intuition about random sequences. Some of our modules give students the tools to solve these and similar harder problems of the same flavour (see e.g. modifications below); this includes *Introduction to probability* (1st year), *Stochastic Processes* (2nd year) and *Probability theory* (third year).
- This is a toy example for an optimal stopping problem. The latter are very relevant in probabilistic finance, for example the pricing of financial derivatives (e.g. *American option*). Several of our 3rd and 4th year modules (e.g. *Brownian Motion*, *Applied Stochastic Processes*, *Introduction to Mathematical Finance*, *Random Events*, *Dynamic Stochastic Control*) develop theory and applications related to this. Optimal stopping has also applications in many other fields, e.g. clinical trials (medical statistics), house selling, fishing, employment,
- The solution above uses the (implicitly made) assumption that the dice is fair. Where does this assumption come into the calculation?
- Study a modified problem that assumes the dice shows 1 with probability p for some fixed p between 0 and 1. All other numbers come up with probability $(1 - p)/5$. What is the optimal strategy now? This will depend on p . Or, to phrase this in another way, calculate of each of the sets S a condition that makes it the (or a) set producing the optimal result (i.e. maximal expected score).
- Should you modify your strategy based on rolls obtained so far? Here is an (extreme) example: You have rolled the dice 10 times, and it showed 2 each time. You could interpret this as evidence for the hypothesis you are dealing with a loaded dice. At the next roll, it shows 3. Your strategy tells you to keep rolling, but given past observations the next roll is likely to be 2 again. Should you modify your strategy and stop at 3?
- The question asks to maximise the expected score. That means that getting 6 is twice as useful to you as getting 3. But is that really true? Say this was about gaining the corresponding amount of pounds. Economists have observed that in people seem to use utility functions converting the numerical value of money into what it's really worth to them. Typically, they grow slower than linearly, e.g. the frequently used logarithmic utility. Economics modules part of our degree programs and the 2nd year statistics module *Games, Decisions and Behaviour* consider such models.
- Some people may want to be really careful and strongly prioritise to avoid getting only 1. Say in their first roll of the dice they obtain 2. They may accept it, because not stopping mean they may get 1 in their second roll and regretting they did not stop earlier. Anticipated regret could here also be modelling by a utility function (e.g. $u(1) = 1$ and $u(x) = 2$ for $x \geq 2$).

For a review article on optimal stopping see

<http://www.americanscientist.org/issues/pub/knowning-when-to-stop>