Skunk redux

This is a simplified version of the game "skunk" that is popular in elementary school math classrooms. The original game is hard to "solve" completely, but this version provides an excellent grade 12 problem.

This is a good game to play in class, because students can develop a pretty good intuition for how to play. I have been using it as the opening lecture of the course because it encapsulates many of the important notions in the course: movement between states, strategy, probability, taking an average, etc.

Here's how the game is played. To start with you stand. At any point in the game you will have a *score*. You begin with score 0. When the game is over for you, you will get a *payoff* determined as follows.

Two dice are thrown.
If at least one 1 appears, the game is over for you and your payoff is 0.
Otherwise you add to your current score the sum of the numbers showing on the dice.
You now decide whether to sit or stay standing.
If you sit, the game is over and your payoff is your current score.
If you stay standing, you go to the beginning and the two dice are thrown again.

For example, for the sequence of rolls:

(2, 5)
(4, 2)
(6, 1)

Those who sit after the first roll get payoff 7. Those who sit after the second roll get payoff 13. Those who stay standing for the third roll get payoff 0.

The problem is to find the strategy which maximizes your average payoff per game.

Notes. A couple of remarks while you are starting to think about the game.

First, I play this game with an entire class, up front rolling the dice. Everybody stands at the beginning and they make individual choices as to when to sit. So as the game progresses, students sit down bit by bit. The game ends when I roll a 1 (and get a collective groan from those still standing) or when everyone has sat down. I might play 10 games in a row and have the students average their 10 payoffs and this will serve as an estimate of their "average payoff per game." So naturally there is a tendency to see "who gets the highest score," or more precisely, once we start talking about strategies, what strategy gets the highest score. Now it's important to point out that there are two different games here that shouldn't be confused. One is the game we are now analyzing, and that's to maximize average payoff per game, and the other is to get the highest average payoff in the class. The second game is what is called an *N*-person game (where *N* is the number of people in the class) and is much more complicated. An optimal strategy in this game must take account of the strategies the others in the class are using, and this presents difficulties. The game we are concerned with here is a "1-person" game—one person and a pair of dice.

It's important to emphasize that the game we are analyzing here assumes that a player is simply trying to maximize his total payoff from playing a long series of games.

Another objective might be to try maximize the number of times he had the highest score in the class. That's quite a different (and much harder) game.

Secondly, it's not clear to many students what I mean by a "strategy." Essentially it's a set of rules which tells you what to do in any situation. Another way to think of it is as a computer program that would play the game. The computer can also roll the dice so the whole game could take place within a computer. One nice thing you could do with such a program is run a large number of games in a short time and get a very good estimate of the average payoff for any particular strategy.

Possible forms of a strategy. First of all, what is a strategy? Well it's a rule which will tell anyone playing the game what to do in any situation, and of course there are only two choices, sit or stay standing. Now what are the factors on which the decision can be made? Well, there are only two things that change from roll to roll, and that's the number of times the dice have been rolled so far, and your current score. So the question is, what account should we take of each of these.

Well that's the question I pose to the class after we've run the game a few times. In fact I go around and collect various strategies that have been used. Some students use a strategy that's not really precise but based on intuition. They come to "feel" that the time to sit has arrived. They generally report that in making this decision, they are influenced by both these factors—their current score and the number of times the dice have been rolled. Other students develop a precise strategy, for example, "sit when the score gets above 20," or "sit after the third roll." Or a hybrid strategy: "stand for at least 3 rolls, and then sit when the score exceeds 20."

The declaration of these various strategies sparks a discussion in the class. A few students assert that the number of rolls is irrelevant and should not be a factor in any optimal strategy. This idea generates a fascinating and often surprising debate. It turns out those few students *are* correct—*the optimal strategy should take account only of the current score*, regardless of how many rolls have been required to get there. The reason for this is that the dice have no memory, and what they've done in the past does not influence the probabilities for the next roll.

This principle seems almost obvious to anyone who has worked a bit with probabilities, but surprisingly enough many students have considerable difficulty with it. In fact some students will argue quite vociferously that if the dice have been thrown, say, 10 times and no 1 has appeared, the chances are increased that a 1 will appear on the next roll. They will hold to this even when they admit that the dice have no memory. *The chances of getting at least one 1 in 11 throws are very high,* they will argue, *so if we haven't got a 1 in the first 10 throws, we should have a high chance of getting it in the 11th.* Hmm.

It is in fact very hard to convince these students that they are wrong, at least in the middle of the class. Here's an argument that sometimes helps. Suppose that the game is being played simultaneously in two adjacent classrooms. Suppose that at a certain moment it happens that the current score for those still standing is 13 in both classrooms, but in one this has come about after 2 rolls and in the other it has taken 3 rolls. Is there any reason to believe that the students in one classroom should behave differently from those in the other? I hope that most students will see that the answer is no.

A few students remain convinced that their strategy of "sitting after *k* rolls" is just as good as any other. *Skunk redux* is sufficiently complicated that it's hard to actually calculate average payoff for various strategies. In problem 3 I present a simpler version as a way seeing how this might be done.

Anyway, accepting this, we restrict attention to strategies that take account only of the current score. Such a strategy will then have to specify, for each possible score *s*, whether you should stand or sit. In playing this game with many groups of teachers and students I have seen many ways of tackling this problem, but the most elegant of these is based on the following simple idea.

An idea for finding the optimal strategy. Suppose your score is *s*. Then if you decide to sit, your payoff will be *s*. If you decide to stand, your score will change and the idea is to calculate the average change in your score. Then if this is positive, you stand for the next roll and if this is negative, you sit.

This is a key idea in strategic analysis—the principle of local optimality. If a small change in strategy increases your average payoff, the strategy can't be optimal.

Let's see how this works. How can we calculate the average change if you stand? Well, as a first cut, there are two possibilities, either a 1 comes up or it doesn't. If a 1 comes up, you lose your entire score, so the average change is –*s*. If a 1 doesn't come up, what then is the average change? That's a good little question right there.

There are a couple of ways to make this argument. Here's one. If a 1 doesn't come up, then each die is 2, 3, 4, 5 or 6, with equal probability and the average of these is the middle number which is 4. So the average on both dice is 4+4=8. Another way to make the argument is to take all the possible cases where a 1 doesn't come up and calculate the average of all the outcomes. Now those cases are the 25 unshaded squares in the table at the right and since they are all equally likely, we simply take the straight average of those 25 numbers. Actually there is neat "geometric" way to do that—the antisymmetry around the "diagonal" of 8's, will tell you that the average has to be 8.

	1	2	3	4	5	6
1						
2		4	5	6	7	8
3		5	6	7	8	9
4		6	7	8	9	10
5		7	8	9	10	11
6		8	9	10	11	12

Now to get the overall average from standing, we average those two possibilities, and for that, we need to know their probabilities of occurrence. What's the probability that a 1 comes up? Some students will argue that since it's 1/6 on one die, it will be 1/6 + 1/6 = 1/3 on two dice. But this is not right. (Why not?) In fact, the answer is found in the above table. Of the 36 equally likely possibilities for the two dice, the 11 shaded entries are those that have a 1, and so the probability of getting a 1 is 11/36. Then the probability of not getting a 1 is obtained from the unshaded squares and is 25/36. [One can calculate this result by saying: the probability that each die will not have a 1 is 5/6. so the probability that both dice will not have a 1 will be $(5/6)^2 = 25/36$.]

Let's summarize.

outcome	probability	change in score
a **1**	11/36	–*s*
no **1**.	25/36	+8

Overall average change from standing

$$= \frac{11}{36}(-s) + \frac{25}{36}(8) = \frac{200 - 11s}{36}.$$

We should remain standing when this is positive and that happens when

$$11s < 200.$$

$$s < 200/11 = 18.2.$$

We conclude that you should remain standing as long as your score is ≤ 18. Sit when your score exceeds 18.

.

A few students raise a cheer when this result is established, as their intuitively driven strategies have turned out to be very close to this, or even right on!

Problems.

1. Here's a slight alteration of *Skunk redux*. If a double one appears, then you add 2 to your score and you can stay in the game if you want. Find the optimal strategy.

2. Another version of *Skunk redux* uses a biased coin which comes up heads with probability ¾ and tails with probability ¼. Suppose that if heads comes up you add 1 to your score but if tails comes up the game is over and your payoff is zero. Find the optimal strategy. [Something unexpected happens here which necessitates some extra calculation.]

3. Another version of *Skunk redux* uses a single 6-sided die with two faces labeled 0, two faces labeled 1, and two faces labeled 2. [So that 0, 1 and 2 all come up with probability 1/3.] If you are standing when the die is rolled, then if a 1 or 2 comes up you add that to your current score, but if a 0 comes up the game is over for you and your score is 0. If you leave the game before a 0 is thrown, your payoff is your score at that point.
(a) Calculate the optimal strategy.
(b) Calculate the average payoff obtained using the strategy in (a). [The game is simple enough that you can simply make a list of all possible trajectories and take the weighted average of their payoffs. It's useful to lump together (without listing) all trajectories which give you a zero payoff.]
(c) Calculate the average payoff for the strategy "sit after k rolls" for the values $k = 2$, 3 and 4, and compare with the answer to (b).

4. (a) Consider the version of *Skunk redux* that uses a single standard single 6-sided die. Suppose that the bad outcome is 1, as before, and otherwise you get to add to your score the outcome of the die. How long should you remain standing?
(b) As a variation of (a), suppose that if you decide to stand you must also hold up a number k of fingers on one hand, 1, 2, 3, 4 or 5. Then you get payoff 0 if the outcome of the die is $\leq k$. But if the outcome of the die is $> k$, you get to add to your score k times the outcome of the die. Thus, if you hold up just 1 finger, you are playing the game of (a), but you have to option of more fingers, and you can change this as the game progresses. What is the optimal strategy now?

5. Here's an absolutely lovely elaboration of *Skunk redux.* Suppose that, before each roll, you are able to specify the number of dice that are to be rolled, and you can change this number from roll to roll based on your score. As before, the game is over for you, with zero payoff, if *any* of the dice show a 1. A strategy now must specify, for each score s, whether to remain standing and if so, how many dice to use. Find the optimal strategy.

6. Suppose you are playing skunk but you are not told the outcomes of the dice each time. You have to decide when to sit based solely on the number of rolls so far. When should you sit?

7. Consider the following game. There are six states numbered 1, 2, 3, 4, 5 and 0. Individuals start in state 1 and move according to the probabilities given in the diagram below. That is, on your first move from state 1 you go to state 2 with probability p_1 and to state 0 with probability $1 - p_1$. The same applies to states $i = 2$, 3 and 4, but possibly with different probabilities p_i. *Except* in any of these states, 1, 2, 3 or 4, you can decide to withdraw from the game. If are in state i when you withdraw you win \$$i$ (and the game is over). If you get to state 5 you automatically win \$5 and the game is over. If you ever get to state 0, you win nothing and the game is over.

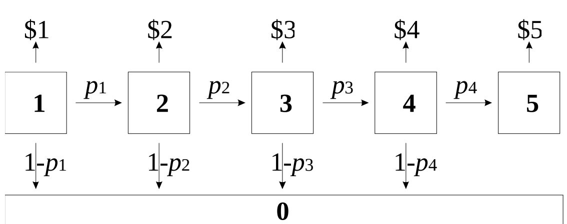

(a) Find a strategy that maximizes your average winnings starting in state 1, and calculate this average.
(b) Specify the optimal strategy and the average winnings for someone starting in states 2, 3 or 4. Does your solution for (a) provide you with an easy way to do this?