Skunk Redux

DAVID KONG
Queen's University
Kingston, Ontario, Canada
x.kong@queensu.ca

PETER D. TAYLOR
Queen's University
Kingston, Ontario, Canada
peter.taylor@queensu.ca

Peter Taylor opens his linear algebra course at Queen's University by having the students play and analyze a simple dice game called Skunk Redux. This is a variation of a common game known as Skunk or Pig. The dialogue below is an account of what happened last fall, when David, one of Peter's students, asked some intriguing questions which prompted the two of them to wrestle with an unexpected problem.

Skunk is a dice game played in elementary classrooms to illustrate the fundamentals of probability [**1**]. Players are given a table with the letters SKUNK across the top like this:

S	K	U	N	K

Each column is used to record the results from one of the five identical rounds. Several players play simultaneously. The objective is to have the highest cumulative *payoff* (the sum of the payoffs from the five rounds) at the end of the game. This is how points are earned in each round:

1. At the beginning of the round, you stand.
2. Two dice are thrown.
3. If at least one 1 appears, the round is over and you have payoff 0. Otherwise you begin with a *score* equal to the total showing on the dice.
4. If you wish, you may sit down. If you do, your payoff is your score.
5. Otherwise, the dice are thrown again.
6. If at least one 1 appears, the round is over and your payoff for the round is 0.
7. Otherwise you add to your score the total showing on the dice. This gives you a new, larger score.
8. Go back to 4.

Eventually a 1 appears and the round is over.

For example, for the sequence of rolls (2, 5), (4, 2), (6, 1), if you sit after the first roll you get payoff 7, if you sit after the second roll you get payoff 13, and if you stay standing for the third roll you get payoff 0.

Extensive work has been done on variations of this game, most notably a 2-person game where an optimal strategy must take into account the opponent's score and strate-

Math. Mag. **85** (2012) 267–276. doi:10.4169/math.mag.85.4.267.

gies [**2**]. For example, the player in second place would likely use a riskier strategy than the player in first place.

This paper is concerned with a 1-person game—one person and a single round. Our analysis focuses on optimizing the expected payoff for this single round—hence, "Skunk Redux."

The first day of class

PETER: I open my linear algebra course with this game because it creates a fun environment, generates a lively discussion, and encapsulates many of the important concepts in the course—strategy, probability, movement between states, taking an average, and so on.

After playing something like 10 rounds in a row, I have the students average their 10 payoffs. This serves as an estimate of "average payoff per game." Naturally there is a tendency to see who receives the highest payoff, or more precisely, once we start talking about strategies, what *strategy* receives the highest payoff.

I want to get a good class discussion going about the different types of strategies. First of all, what is a strategy? It is a rule that tells you whether to sit or stand in any situation. Any situation in this game can be specified in terms of two variables: the number of times the dice have been rolled, and the current score.

I find that students have differing opinions on how to make use of these two variables. Some strategies are highly intuitive and students sit when they "feel" the time has come. Some sit after a certain number of rolls while others pay attention only to the score (e.g., "sit when I get above 25"). Still others use a mixture (e.g., "sit after 25 or after the third roll, whichever comes first").

DAVID: Surely the number of rolls is irrelevant and should not be a factor in any optimal strategy. The rolls are independent events! The only quantity of relevance is the current score.

PETER: David is, of course, correct. But this issue always generates a fascinating and surprising debate. A number of students will argue quite vociferously that if the dice have been thrown, say, ten times without showing a 1, the chances are increased that a 1 will appear on the next roll.

Moving on, we restrict attention to strategies that take account only of the current score. Such a strategy must specify, for each possible score s, whether you should stand or sit.

DAVID: Let's begin by defining s as your current score. If you decide to sit, your payoff will be s. If you decide to stand, your score will be either better or worse. If, on average, your new score is greater than s, you should stand for the next roll; if it is less, you should sit.

To calculate your average new score, note that with probability $25/36$ (see TABLE 1 below) your score increases by the dice sum and with probability $11/36$ your score drops to 0. Now the average dice sum, given that a 1 does not appear, is 8. (This is nicely seen in TABLE 1 by pairing each entry with its mirror image in the diagonal of 8's.) The average new score from standing is then:

$$\frac{11}{36}(0) + \frac{25}{36}(s + 8).$$

TABLE 1: Addition to the score for each of the 36 possibilities

	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

You should remain standing when this exceeds s, and that happens when

$$25(s+8) > 36s$$
$$s < 200/11 \approx 18.2.$$

Thus you should remain standing as long as $s \leq 18$ and sit when $s \geq 19$.

Peter started rolling the dice on that first day of class. As usual, I did not bring anything to class, not even a calculator, so I had to ballpark it. "How much would I be willing to risk to get an average reward of 8?" Somehow I came up with the number 20, which in hindsight was fairly close to the actual answer. From there, I rigorously abided by my strategy, sitting when the score surpassed that critical value. It took some willpower not to allow my emotions to steer me toward the standard freshman crowd—the eternal optimists who luckily see the world as their oyster, untainted by the rationality I sometimes wish I could do away with. There were times when I would begrudgingly sit from the sidelines while the most risk-friendly participants racked up unimaginable sums. But in the long haul, my strategy paid off.

First day of class and already an interesting (yet accessible) problem. I was truly excited for university. What I did not realize at this point was that I was soon to be led to something even more interesting.

The assignment

PETER: For their first assignment, I usually give the students an extension of the game to analyze. For example, the dice may be replaced by a few coins. One of my favorite (and most demanding) extensions has been the following:

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 round is over with zero payoff, if you are standing and *any* of the dice show a 1. A strategy must now specify, for each score s, whether to remain standing and if so, how many dice to use. Find the optimal strategy.

DAVID: Now that's an enticing problem! Rolling more dice at a time will help you increase your score more quickly, but it also increases the probability of rolling a 1. The key difference between this problem and the simpler one is that now there are two decisions to make for each value of s—whether to remain standing, and if so, how many dice to roll. But I expected the solution not to be much different than before.

PETER: Like David, most students find this problem challenging. Not many manage to come up with a good argument. But there are always a few students who produce the following solution and for some years I have always accepted it as being correct. It is based on the idea that we employed in the solution for the original game, that the correct decision at each step is the one that maximizes the expected new score.

The $A(n)$ strategy. Let $A(n)$ be your expected new score if you stay standing and choose to roll n dice. Note that the probability of not throwing a 1 is $(5/6)^n$ and (as above) the average outcome on a single die is 4. Then:

$$A(n) = (5/6)^n(s + 4n).$$

DAVID: I got the above equation for $A(n)$ without much difficulty. Now the problem was to find the maximum value of $A(n)$. When in doubt, a first-year student differentiates. The result was correct enough but it was ugly with logarithms and decimals. A bit later, I found a much nicer *algebraic* solution. I thought of it as discrete maximization, and it worked beautifully. The idea was that for $A(n)$ to be a maximum at a particular n, it must be at least as great as the neighboring $A(n)$ values, $A(n-1)$ and $A(n+1)$.

$$A(n-1) \le A(n) \ge A(n+1)$$

The first part is:

$$\begin{aligned} A(n-1) &\le A(n) \\ (5/6)^{n-1}(s + 4(n-1)) &\le (5/6)^n(s + 4n) \\ s + 4(n-1) &\le (5/6)(s + 4n) \\ 4n &\le 24 - s \end{aligned}$$

and the same for the second:

$$\begin{aligned} A(n) &\ge A(n+1) \\ (5/6)^n(s + 4n) &\ge (5/6)^{n+1}(s + 4(n+1)) \\ (s + 4n) &\ge (5/6)s + (5/6)4(n+1) \\ 4n &\ge 20 - s \end{aligned}$$

Putting these together, the condition for a maximum $A(n)$ is that

$$20 - s \le 4n \le 24 - s$$

For example, given the score $s = 10$, there is only one integer value ($n = 3$) that satisfies this inequality.

PETER: David's analysis so far is the one I have always accepted, and posted on the website for the class. It says that $4n$ has to be between $20 - s$ and $24 - s$. We can summarize this condition with TABLE 2. When s is a multiple of 4, there are two values of n that give the same average score. [At $s = 20$, the "other" value is $n = 0$, which means *sit.*] And by the way, it can easily be verified directly for $s < 20$, that $A(n) > s$ for the indicated n, signifying that you gain on average by standing.

And then David came up to me after class...

DAVID: I had the solution outlined above and it seemed really elegant (isn't that table beautiful?) but it worried me. Maximizing $A(n)$ only maximizes the score *after* the

TABLE 2: The $A(n)$ strategy: Roll n dice with score s

score s	roll n dice
$s = 0$	$n = 6$
$0 \le s \le 4$	$n = 5$
$4 \le s \le 8$	$n = 4$
$8 \le s \le 12$	$n = 3$
$12 \le s \le 16$	$n = 2$
$16 \le s \le 20$	$n = 1$
$s \ge 20$	$n = 0$

next roll, whereas the objective of the game is to have the highest possible payoff, which is your score *at the moment you sit down.* Do we need to worry about this distinction? It is tempting to think that they lead to the same outcome—if you put yourself ahead in the immediate future, wouldn't that also put you ahead in the long run? But I could see no valid argument for this. I spent an entire night (my first university all-nighter) tangled with this question.

At some point I decided that my only hope was to look for a strategy that outperformed the $A(n)$ strategy. I became interested in the strategy of always using one die because it was the simplest strategy around. I decided to "put it to the test," using EXCEL to compare it with the $A(n)$-strategy I had developed so far. After 50,000 Monte Carlo iterations, the differences were insignificant and inconclusive.

The breakthrough occurred when I looked at the case of $s = 15$. I made a few calculations that put the issue to rest.

A counterexample to the $A(n)$-strategy. Take the case of $s = 15$. The $A(n)$ strategy tells you to use $n = 2$ dice. If double 2s are rolled you stand for one more round using 1 die. Otherwise, you sit. The result is summarized in TABLE 3. The average score is approximately 15.98.

Now compare this with the strategy that uses only one die and stands whenever the score is less than 20 (TABLE 4).

TABLE 3: How the $A(n)$ strategy plays out at $s = 15$

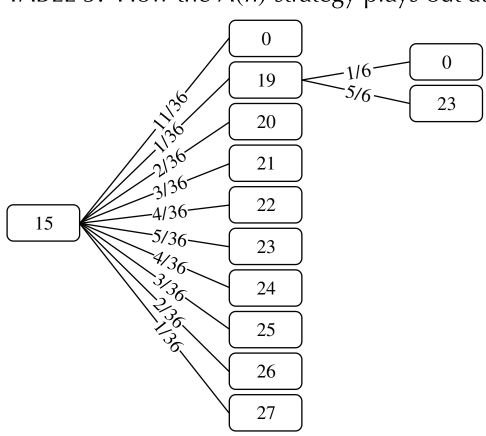

prob.	score $= X$
5/216	23
2/36	20
3/36	21
4/36	22
5/36	23
4/36	24
3/36	25
2/36	26
1/36	27
E(X)	15.97685

TABLE 4: How the 1-die strategy plays out at $s = 15$

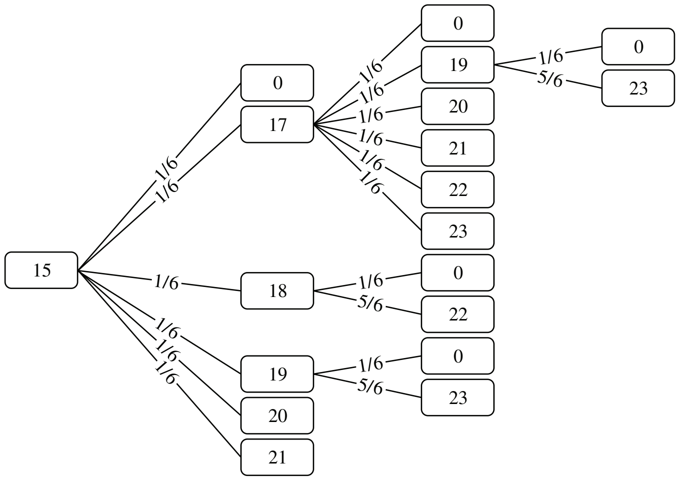

prob.	score $= X$
5/216	23
1/36	20
1/36	21
1/36	22
1/36	23
5/36	22
5/36	23
1/6	20
1/6	21
E(X)	16.00463

The average score is a bit above 16, and higher than was obtained starting with 2 dice. For this particular s-value, the 1-die strategy outperforms the $A(n)$ strategy!

PETER: David's 1-die strategy was a revelation to me and for a time I had a bit of trouble thinking clearly about the situation. The example above of $s = 15$ certainly shows that the $A(n)$ strategy is not optimal. But is the 1-die strategy optimal? Are there situations when it might be better to roll more than 1 die? And suppose that the 1-die strategy *is* optimal. When do we stop? Is $s = 20$ the right place to sit? I was thrown for a bit of a loop and decided to go back to the beginning.

It is surprisingly easy to get confused, particularly when there is more than one question buzzing around. What's needed is to focus on one thing at a time, and hope that it's the right thing to begin with. The next day David came to me with a ridiculously simple argument that nothing could possibly outperform the 1-die strategy.

DAVID: Peter is right—it's so easy to miss simple things. And this is one of them. Suppose your score is s and you are using a strategy that tells you to roll 3 dice. Then you would have exactly the same outcome by standing for the next 3 turns and rolling 1 die each time. The reason for this is that the condition for the game to end with a zero payoff is the same in each case—getting a 1 on any of the three dice. So the 1-die strategy will do just as well as the one you are using. But furthermore, it might even do better because it gives you the option of stopping before the third turn.

PETER: Indeed that's exactly why the 1-die strategy outperformed the $A(n)$ strategy at $s = 15$. If you happen to roll a 6 on your first die (giving you $s = 21$) the 1-die strategy lets you stop and sit down, whereas the $A(n)$ strategy rolls again. Now if you stop, your payoff is 21, but if you roll again, your average score becomes

$$s = \frac{0 + 23 + 24 + 25 + 26 + 27}{6} = \frac{125}{6} \approx 20.833$$

which is less than 21.

DAVID: Always roll one die.

The meeting

After the revelation about a pure 1-die strategy, the final challenge was to determine and prove the critical s-value for when to sit. This appeared obvious enough but a formal proof took quite a while to formulate. We sat down for a final meeting to discuss this.

PETER: So the only question left is, when do you sit?

DAVID: At $s = 20$.

PETER: How do we know?

DAVID: Use the same calculation we made above at $s = 21$. It works for any $s > 20$. Your expected score after one roll will always be less than s.

PETER: Right. It is worth emphasizing that. The $A(1)$ strategy (which is optimal) asks you to compare:

$$s \quad \text{and} \quad \frac{(s+2)+(s+3)+(s+4)+(s+5)+(s+6)}{6}$$

On the left is the payoff if you sit and on the right is your expected new score if you stand. For $s < 20$ the right side is bigger, for $s > 20$ the left side is bigger, and for $s = 20$ they are equal. So the strategy says sit when $s > 20$. But as you pointed out long ago, this only considers the next roll instead of the indefinite future. What we really need on the right is some indication of your payoff at the end of the game, given that you stand and play optimally.

DAVID: We need the notion of what a strategy is "worth." If you have score 19, you can expect to increase that on average by staying in the game, so having a score of 19 is actually worth more than 19. However, if you have 21 you can't do any better (in fact, by staying in the game you'll do worse on average), so 21 is only worth 21.

PETER: We could formalize that. Define $v(s)$, the *value* of s, to be the expected payoff for a player who currently has score s and who plays optimally. For example, $v(19) > 19$ and $v(21) = 21$.

DAVID: In fact

$$v(19) = \frac{v(21)+v(22)+v(23)+v(24)+v(25)}{6},$$

and $v(s)$ in general would be

$$v(s) = \max\left(s, \frac{v(s+2)+v(s+3)+v(s+4)+v(s+5)+v(s+6)}{6}\right).$$

The first term represents the payoff if you sit. The second term represents the average payoff if you stand and play optimally. You choose whichever one is greater. If we knew that $v(s) = s$ for big enough s, say for all $s \geq 100$, then we could use the recursive equation to work backwards. We would get $v(99) = 99$, then $v(98) = 98$, and that would keep on working all the way to $v(20) = 20$. The first time s would be less than the expression on the right would be at $s = 19$.

PETER: So what we need to do is to find some large enough s^* for which we can show that $v(s) = s$ for all $s \geq s^*$.

Pretty Black Cat. One way in which Peter creates exercises for the students is to construct variations on what happens when a 1 is rolled. One such variation seems at first quite uninteresting, but in fact it holds the key to a lovely proof of the result we are searching for.

PETER: I've been thinking about a modification called Pretty Black Cat ("PBC") in which you always roll one die, and when a 1 is rolled, the game ends but you do not lose your current score.

DAVID: Not very interesting, of course, because you'd simply always stay in the game.

PETER: Indeed. But the game is so simple that we ought to be able to calculate its $v(s)$ values easily.

DAVID: No doubt. But I'm wondering where this is headed.

PETER: I'm thinking that whatever strategy you choose to use in Skunk, the same strategy used in PBC will give you a payoff that is at least as high. It surely follows that the $v(s)$ values for PBC will always be at least as big as those for Skunk Redux, so PBC's $v(s)$ will give us an upper bound on Skunk's $v(s)$... and that might be useful.

DAVID: Indeed it might. Let's see... in PBC a player with score s would get exactly one more roll with probability $1/6$, exactly two more with probability $(5/6)(1/6)$, exactly three more with probability $(5/6)^2(1/6)$, etc., and the average payoffs would be, s, $s + 4$, $s + 8$, etc. We just have to add a bunch of terms.

PETER: Or perhaps we could try a recursive argument.

DAVID: Yes. I might have thought of that, as it is one of the big themes of the course. Let k be the amount you gain on average by continuing to play. Then, if your next roll is a 1, k is zero, and otherwise, you gain 4 on average and you are able to keep playing so your overall gain is on average $4 + k$. This gives us the recursive equation:

$$k = (1/6)(0) + (5/6)(4 + k)$$

and that solves to give $k = 20$.

PETER: Nicely done. So for Pretty Black Cat, the value of having a score s is $v(s) = s + 20$.

DAVID: We can conclude that for Skunk Redux, $v(s) \leq s + 20$.

PETER: Maybe that will be enough to find a score s for which $v(s) = s$.

DAVID: Let's see. Returning to Skunk Redux, $v(s) = s$ if

$$\frac{v(s+2) + v(s+3) + v(s+4) + v(s+5) + v(s+6)}{6} \leq s$$

and since $v(x) \leq x + 20$, that will hold if

$$\frac{(s+22) + (s+23) + (s+24) + (s+25) + (s+26)}{6} \leq s$$

and that simplifies to $s \geq 120$.

PETER: Wow.

DAVID: We conclude that $v(s) = s$ for every $s \geq 120$.

PETER: That elusive but utterly unsurprising conclusion is just what we need to start the backwards recursion and make all of our deductions legitimate. Finally, we can safely say that 20 is indeed the place to sit.

Epilogue

And thus the four-month journey concludes with the astounding realization that our initial reasoning is flawed. For the $(n = 2)$-dice game we discussed at the beginning, a comparison of the expected immediate gains by sitting and by standing fails to take account of the long-term possibilities. The answer to sit when $s > 200/11$ is correct but requires a more rigorous argument involving $v(s)$.

The reason the 1-die strategy is optimal in an n-dice game, as previously mentioned, is that any gain you can make by rolling n dice can be obtained by rolling 1 die n times. Also, it is important to notice that while the 1-die strategy is optimal, it is not the only optimal strategy. For example, since you will never leave the game with $s < 20$, and 3 dice can only take you to 18, you might just as well throw 4 dice at the very beginning. The same reasoning continues to apply. For example, an optional strategy allows a play of 2 dice for $8 \leq s \leq 14$, and so on.

An interesting problem arises if we *exclude* the option of using 1 die, that is, you can roll any number of dice except 1. In this case your effective choices become sit, stand with 2 dice, or stand with 3 dice. This is because any number $n > 3$ can be written as a linear combination of 2 and 3. The optimal strategy for this game (found with EXCEL) is displayed in TABLE 5. It has an intriguing pattern.

TABLE 5: An optimal strategy when 1 die is forbidden

score s	optimal n
$s = 0$	$n = 3$
$1 \leq s \leq 6$	$n = 2$
$7 \leq s \leq 11$	$n = 3$
$12 \leq s \leq 18$	$n = 2$
$s > 18$	$n = 0$

More generally, suppose there is a given set of available numbers of dice to roll: $\{n_1, n_2, \ldots, n_k, \ldots\}$, where no n_j is a nonnegative-integer linear combination of the other n_i. We invite others to conduct further research on optimal strategies for this and other variations.

REFERENCES

1. D. Brutlag, Choice and chance in life: The game of SKUNK, *Mathematics Teaching in the Middle School* **1**(1) (April 1994) 28–33. http://illuminations.nctm.org/LessonDetail.aspx?id=L248
   In the elementary classroom game, there is an additional twist: if the dice show "double ones," all points from previous rounds are also wiped out. Our version is simpler.
2. T. W. Neller and C. G. M. Presser, Optimal play of the dice game Pig, *The UMAP Journal* **25**(1) (2004) 25–47. http://cs.gettysburg.edu/~tneller/resources/pig/
   Neller's version of the game is called Pig and involves a 2-player race to a score of 100. Otherwise, the rules are the same.

Summary In the simple version of Skunk, a pair of dice is rolled again and again until either you choose to sit or at least one 1 comes up. If you sit, your payoff is the dice sum of all your previous rolls. If at least one 1 comes up while you are still standing, your payoff is zero. Here we look at an extension of the game in which you must also choose the number of dice to use and you can alter this from round to round. The simple game seems to have a straightforward enough optimal strategy but our analysis of the extension reveals flaws in our initial reasoning. Some additional extensions are considered.

DAVID KONG is a 3rd year undergraduate student at the Queen's School of Business and is pursuing a Bachelor of Science in mathematics. He has worked in both the hedge fund and private equity industries (Fore Research, ONCAP Partners) and was invited to the 2010 Canadian Mathematics Olympiad. This has led to an interest in statistics, probability, and modeling, particularly through EXCEL. David enjoys skiing, biking, cooking, and creating latte art.

PETER TAYLOR is a professor in the Department of Mathematics and Statistics at Queen's University, cross-appointed to the Department of Biology and the Faculty of Education. He is a Queen's graduate and has a Harvard Ph.D. His area of research is theoretical evolutionary ecology, particularly the evolution of cooperative behaviour. He is a 3M Teaching Fellow and is Chair of the Education Committee of the Canadian Mathematical Society. Recently he chaired the Task Force at Queen's, charged with writing the university's new Academic Plan.