

BETWEEN FEAR AND GREED: THE SIX LOOSES

Joachim Engel

Ludwigsburg University of Education, Germany

engel@ph-ludwigsburg.de

The paper reflects on students' intuitive strategies in a game of chance and contrasts their reasoning with a normative probabilistic point of view. The game involves selecting optimal strategies, outweighing potential gains with small probabilities with more probable losses and provides some insight into students' probabilistic reasoning.

INTRODUCTION

In the dice game „The Six Looses“ a fair die is being tossed repeatedly and the up face is recorded. As long as no 6 occurs, the rolled numbers are added up to a score. The game can be stopped at any time; then the obtained score is the prize (e.g. as money in Euros) won. However, if a 6 shows up, everything is lost and the game is over. If, for example, the die shows the consecutive numbers 4, 2, 3, 3 and then you stop, the gain is 12 Euro. With the sequence 2, 1, 6 you go empty handed, because you didn't stop after the second roll and the 6 in the third roll destroys all your previous gains. Imagine you play the game very often. Which strategy for stopping the game should be employed to get on average over many games a score as high as possible? In other words, what is an optimal strategy to maximize your expected gain? We presented this game to college students and inquired about their choice of strategy. While the students were familiar with basic concepts of probability such as the notion of the expected value, they lacked the formal probabilistic knowledge for a full mathematical analysis of this problem. Therefore, they were challenged to reason informally about risks involved in continuing to roll the die, based on intuition and guts feeling. We present and discuss students' answers followed by a normative presentation of an optimal stopping strategy.

THE DICE GAME

The game “The Six Looses” was presented to 46 second-year students preparing to be secondary school mathematics teachers. Figure 1 informs about the tasks the students were asked to do. After devising a strategy and noting down a rationale for the proposed approach students were asked to actually play the game at least ten times in a row sticking exactly to their recommended strategy. Thus, they collected some empirical evidence regarding the efficacy of their proposed action. Finally, after having played the game, they were asked to reconsider their originally proposed strategy and reason if they stick to the original plan or favor a possible change of strategy.

In previous class meetings students were introduced to the concept of probability and its different notions (classic, frequentist, subjective), learned about computing probabilities in a multi-stage situation via tree diagrams and were introduced to the notion of the expected value as a theoretical average of repeated outcomes that is based on the law of large numbers.

Thus, it can be expected that the participating students were capable of calculating the probability of, say, “no 6” occurring in 1, 2, 3, ... repeated rolls of a die. Below we present a mathematical

analysis of the optimal solution based on the notion of conditional expectation. But as these techniques were not available to the students, their judgment about optimal strategies was mainly based on intuition. Assessing the problem required them to recognize that with an increasing number of rolls the probability of a 6 (and hence losing everything “earned” so far) increases, but also the potential gain in terms of raising the score may get larger. Hence, a decision in a typical risk-benefit trade-off situation had to be made. Their reasoning may shed some light on their probabilistic thinking, in particular their understanding of the independence of consecutive throws of a die (*Does the likelihood of a 6 stays the same or does it increase if a 6 hasn’t shown up in the rolls before?*) and their frequentist interpretation of the results after they have played the game under the proposed strategy several times. On the one hand, the concrete experience of playing the game may provide a valuable background to confirm or reject the originally proposed strategy. On the other hand, as the number of games played was rather low (“at least 10”), students may run the risk of overgeneralization and be subject to the *small sample fallacy* by considering their small sample experience as representative.

Game of dice: A fair die is thrown repeatedly. As long as no 6 occurs, the up face numbers are added up to a score. The game can be stopped at any time. Then the obtained score is your gain (in Euros). However, if a 6 occurs, everything is lost. Which strategy for stopping do you recommend in order to maximize your average gain, provided you are playing the game very often.

1. Note down which strategy you choose in order to maximize your average gain when playing this game very often.
2. Explain, why you think this is a good strategy.
3. Now take a die and play the game at least 10 times following exactly your proposed strategy.
4. After testing your recommended strategy, would you still recommend this strategy? Why or why not?
5. If you favor a change of strategy, what is your preferred choice?

Figure 1: Tasks given to the students inquiring a written response

FEAR OF LOSING OR GREED FOR MORE: BALANCING RISK AND BENEFIT

Each additional roll of the die offers the possibility of higher gains, but also increases the chance of losing everything that has been gained so far. Therefore, a reasonable strategy has to involve some trade-off between an increasing risk of not gaining anything and of obtaining a higher score. While the probability of “no 6” (and hence maintaining a winning score) in k consecutive throws of a die approaches zero at a rate of $(5/6)^k$, the potential gain increases with each new roll of the die. With a moderate number of, say, 7 rolls, either you win a substantial amount with a rather small probability of $(5/6)^7 = 27.9\%$, or you win nothing with a high probability of 72.1%. On the other hand, if you roll the die only twice, the winning probability is about 69.44%, but the gain will be rather modest. Therefore, the situation can be seen within the framework of risk assessment.

Risk in decision theory is defined as a product of a probability p times a measure for the (dis-)utility of an event (see, e.g. Edwards and Tversky, 1967). In the classical situation, it is about evaluating a loss of a particular resource (money, health, time, energy ...) that occurs with small probability and compare it with consequences of other events that may happen with certain other probabilities. One

perspective at the present case is that you may obtain a score (gain money) with small probabilities as a “positive risk” or benefit and the question is which strategy maximizes your benefit. An alternative perspective is to look at the situation from the viewpoint of consecutive decisions: Assuming, you already achieved a score of $s > 0$ points through previous rolls of the die and you have to decide to stop or continue playing. When you continue, you can increase your score with probability $5/6$ by up to 5 points resulting in an (equally likely) score of $s+1$, $s+2$, ..., or $s+5$. But with probability $1/6$ you’ll loose all the s points you have gained so far.

Hence risk-averse people may decide for very few rolls only, while risk-prone players will try to challenge their luck and opt for many more rolls hoping for a high score. The perspective of sequential decision making highlights the fact that risk is not only about probabilities of loosing a resource, but also it depends equally on the utility of the good that you are loosing.

STUDENTS’ RESPONSES

The strategies proposed by the students can be classified into the following strategies:

Strategy 1: Roll the die r -times

Strategy 2: Roll the die until you have at least a score of s points.

Strategy 3: A mixture of strategy 1 and strategy 2 in the form: Roll the die at least r times.

But if your score is below s , then continue rolling the die one or two more times

Strategy 1: Roll r times (37 students)

The majority of the students chose strategy 1 (“Roll the die r times”) with different recommendations for the number r of rolls. They stuck to that strategy after they played the game ten times or more, some with recommending a change of the number of throws. 21 students recommended only 3 rolls before, 18 students after playing the game. Most students stuck to that strategy, a few of them altered their recommendation about the number of rolls, and a few combined that recommendation with a more flexible rule in the sense of strategy 3. Table 1 presents details. 27 students stuck to the same number of rolls (diagonal in Table 1), while only 10 students changed the number of recommended rolls.

		After rolling the die						
		1	2	3	4	5	6	7
Before rolling the die	1							
	2		2	1				
	3	1		16	3	1		
	4			1	5	1		
	5				1	4		
	6							
	7					1		

Table 1: Students who recommended the „Roll the die r times“ strategy: their choice of the number of rolls before and after they actually played the game at least 10 times.

Some students based their reasoning on computing the probability of “no 6” in several consecutive throws and recognized the (geometrically) decreasing probabilities of “no 6”. However, very few students balanced this observation with the increased potential gain when rolling more often. Hence they solved an easier or a partial problem that they knew how to solve without being aware that this is only a part of the original problem. Some of these answers also could be attributed to the outcome approach (Konold, 1989), because with four or more throws the odds of getting nothing due to a 6 are greater than 1 and increase rapidly with the number of throws, as the following quotes document:

- *the probability of no six still has to be above 50%. This happens, if I roll three times.*
- *Probability of three times “no six” is about 60% $(5/6)^3 = 57.87\%$*
- *the probability of getting a 6 increases rapidly the more often I roll*

Some students based their recommendation just on intuition:

- *If I throw more often, the risk of getting a 6 gets bigger; 4 times appears as reasonable comprise*

The probability of a six is always the same, regardless how often you throw the die. Ignoring this fact, the following arguments are indicative of gambler’s fallacy:

- *Roll 5 times, because on average every sixth roll results in 6*
- *Roll 5 times, because the probability for each side up is 1/6 because of 6 faces. Stop one throw before 6. Therefore roll 5 times*

Several students were discouraged to roll the die 5 times because the actual experiment they performed resulted in many zero-gains or, vice versa, they stuck to the result because of some “lucky” games. No one questioned the random variability inherent in such a small sample size.

Thus, they overgeneralized the experience of playing the game 10 times:

- *With 5 rolls I lost in 6 out of 10 cases. This is more than half, therefore you should throw the die less than 5 times.*
- *Only 3 out of 10 times I won, therefore I’d rather change and roll less frequently*
- *In 9 out of 10 games I won with the “roll 3 times” strategy. Therefore, I would stick to it.*
- *In 9 out of 10 games I won, therefore I keep my strategy [of 3 throws]*

Some students recognized that the decision is about balancing possible gains that increase with the number of throws and the probability of not winning anything:

- *With three rolls there is quite a lot to win, an additional throw would diminish the probability of keeping my gain*
- *A reasonable compromise between high gain and probability of winning*

Some rather idiosyncratic explanations were the following:

- *The die has 5 sides [arguing for 5 throws]*
- *Because of my personality I prefer the safe side; to roll the die three times is too risky (Probability of 57.8% of having a six) [arguing for only 2 throws]*

- *I roll until 1 shows up. Because after the 1 the 6 is most frequent*
- *The probability of “no 6” is 5/6. Hence I roll 5 times*
- *because three is half of six [reasoning for three throws]*

Strategy 2: Rolling until a score of s is achieved (6 students)

While rolling the die until a certain score is obtained is another intuitively reasonable strategy, none of the 4 students who opted for that strategy had striking arguments

- *Roll until a score of at least 15, because $1+2+3+4+5=15$*
- *Roll until 15: risk and probability are always the same. Therefore, it doesn't matter if I go by the number of rolls or the score obtained*
- *Roll until a score of 20, because either I win 20 or I win nothing*
- *If the first few rolls results in a high score, then secure your gain; otherwise continue rolling*

Strategy 3: Mixing 1 and 2 (3 students)

Mixture strategy

- *Roll the die at least 5 times or to a score of 10 points*
- *Roll 3 times except if I got only 1's and 2's, because the loss is small then*
- *At least 3 throws, at most 5; stop if above 10 points*

THE NORMATIVE SOLUTION

Two intuitively plausible strategies are: „roll r times“ and „roll until the score is at least s “.

Probability calculations allow us to determine optimal choices for the parameters r and s under these strategies and to compute the expected gain.

Strategy 1: Roll r times

An easy access is obtained through the formula of the conditional expectation (see, e.g. Feller, 1967, pg 223). Let X denote a discrete random variable assuming values x_1, x_2, \dots, x_k , and A be an event. The conditional expectation of X under the condition A is defined as

$$E(X|A) = \sum_k x_k P(X = x_k | A).$$

Then, provided A_n form a partition of the sample space (i.e. $\bigcup_n A_n = \Omega, A_i \cap A_j = \emptyset$ for $i \neq j$), it holds that

$$E(X) = \sum_n E(X|A_n) \cdot P(A_n)$$

Now let A_r denote the event: “No 6 in r consecutive rolls”, and let the random variable S_r denote the score after r rolls. Then it holds

$$(1) E(S_r) = E(S_r|A_r) \cdot P(A_r) + E(S_r|\bar{A}_r) \cdot P(\bar{A}_r) = E(S_r|A_r) \cdot P(A_r) = 3r \left(\frac{5}{6}\right)^r,$$

because the expected gain in each of r throws is 3 provided no 6 is thrown and $S_r = 0$ otherwise, i.e. $E(S_r|A_r) = 3r$, $E(S_r|\bar{A}_r) = 0$.

In order to find the value of r maximizing this expression, consider the quotient

$$\frac{E(S_{r+1})}{E(S_r)} = \frac{r+1}{r} \cdot \frac{5}{6}$$

This quotient is exactly greater (smaller) than 1, if $r < 5$ (resp. $r > 5$).

Furthermore we have $E(S_5) = 15 \cdot \left(\frac{5}{6}\right)^5 = 18 \cdot \left(\frac{5}{6}\right)^6 = E(S_6) = 6,028$.

When rolling the die 5 times, we can expect an average score of 6.028, we'll end up with a score of 0 with a probability of $1 - (5/6)^5 = 59.81\%$. In about 6 out of 10 games we will gain nothing, while the score in the other games will be distributed between 5 and 25.

Strategy 2: Roll until a score of at least s is obtained

Assume, you already have a score of s and roll once more. Let Y_s denote your gain after the next roll. Y_s assumes the values $s+1, s+2, \dots, s+5$ and 0, all with probability $1/6$, therefore we have

$$(1) \quad E(Y_s) = \frac{1}{6} \sum_{j=1}^5 (s+j) = \frac{5s+15}{6}.$$

To continue rolling the die pays off if and only if $E(Y_s) > s$ which is equivalent to $s < 15$. Hence the optimal rule is to stop rolling when a minimum score of 15 is achieved.

Because of $E(Y_s) = s \Leftrightarrow s = 15$ waiting till 16 will not deteriorate your score. Therefore, the best choice within Strategy 2 is “continue playing until your score is 15 (or more); otherwise stop and cash in your gain”. The formal derivation of the expected gain under this rule is beyond the scope of this paper. Henze (2011) suggests a recursive procedure resulting the Table 2.

s	E(Y_s)	s	E(Y_s)
8	5.318	14	6.130
9	5.550	15	6.154
10	5.744	16	6.154
11	5.896	17	6.133
12	6.003	18	6.094
13	6.081	19	6.040

Table 2: Expected gain under strategy: “stop when the score collected is at least s ”.

Furthermore, the probability of gaining nothing at all under strategy 2 is 62.22% (Henze, 2013), hence even slightly higher than under strategy 1.

Roters (1998) could prove that the Strategy 2 to stop when obtaining a score of at least 15 is not only superior to Strategy 1, but it is the overall optimal strategy for maximizing the expected values of the game.

CONCLUSIONS

While almost all students' chose a reasonable strategy, most students followed an outcome approach (Konold, 1989) by focusing on the high probability of not gaining anything with too many throws. They did not take into account the increase in potential gain that is connected. In addition,

the students' reasoning revealed many typical fallacies and preconceptions (Batanero et al., 2016) most notably gambler's fallacy, small sample fallacy and the outcome approach.

REFERENCES

- Batanero, C., Chernoff, E., Engel, J., Lee, H., & Sanchez, E. (2016). Research on Teaching and Learning Probability. ICME-13. Topical Survey series. New York: Springer.
- Edwards, W., & Tversky, A. (Eds.) (1967). *Decision Making: Selected Readings*. Penguin Books.
- Feller, W. (1967). *An Introduction to Probability Theory and Its Applications*. 3rd Edition, New York: Wiley.
- Henze, N. (2013). Die Sechs verliert. Stochastik in der Schule. *Stochastik in der Schule* 31 (2), 2–5
- Konold, C. (1989). Informal conceptions of probability. *Cognition and Instruction*, 6, 59-98
- Martignon, L., Kuntze, S. (2015). Good Models and Good Representations are a Support for Learners' Risk Assessment. *The Mathematical Enthusiast*, 12 (1, 2 & 3), 157 - 167
- Roters, M. (1998). Optimal stopping in a dice game. *Journal of Applied Probability*. 35, 229–235.