The Mathematical Analysis of Games, Focusing on Variance

Dr. ir. Tom Verhoeff, TU Eindhoven

Last November, I gave a lunch talk on game analysis for W.I.S.V. 'Christiaan Huygens'. If you missed it, you can catch up here.

Games, Mathematics, and Decisions

Games have always been loved by mathematicians. Because of their (usually) well-defined rules, games admit a formal analysis, sometimes with surprising results. It is impossible to provide comprehensive coverage in this article. Instead, I will give an overview and focus on the role of variance, which is often overlooked.

Jörg Bewersdorff wrote a nice book on applying mathematics to the analysis of games [2]. I recommend it to the mathematically inclined reader. It classifies games according to the sources of uncertainty that confront the players, also see Figure 1.

Combinatorial Games generate uncertainty by the many ways in which moves can be combined, in spite of the fact that the players have open access to all information. There are some general bits of mathematical theory for combinatorial games, but most of it is rather ad hoc. A very inspiring reference work for combinatorial games is the monumental and entertaining book known as *Winning Ways* [1].

Figure 1: *Game classification based on source of uncertainty*

In *Games of Chance*, uncertainty comes from stochastic processes, such as dice rolling and card shuffling. The field of probability theory was founded in the 17th century while analyzing games of chance. Probability theory and especially Markov Decision Processes are the tools of choice to attack these games.

Strategic Games arise when players have secrets for each other. In Rock, Paper, Scissors (RPS), both players must simultaneously select an object, not knowing the other's choice. The mathematical field of *Game Theory* addresses such strategic games. There is no Nobel Prize for Mathematics, but mathematicians have obtained several Nobel Prizes for Economics through their work on Game Theory.

Many games are not pure but rather a weighted mixture of the three classes. For instance, poker contains all three elements: cards are shuffled (luck), players cannot see each other's cards (secrets), and choices can be combined in many ways (logic). Poker is hard.

Games involve one or more players making moves with the aim of achieving the game's objective. The rules determine when players may move, what their options are, and what the goal is. The goal could just be winning (binary outcome) or maximizing some profit function (discrete or continuous).

Mathematically interesting questions are: Given the game's state, how do you decide your best move? Which player can force the best outcome from the start? What is the (long-term) optimum result?

A Strategic Coin Game

Let us dive into a simple strategic coin game. The two players, Alice and Bob, simultaneously and secretly each choose 0 (head) or 1 (tail)[1]. After committing their choices, the payoff is determined by the matrix in Figure 2. Alice receives € 2 from Bob when the choices

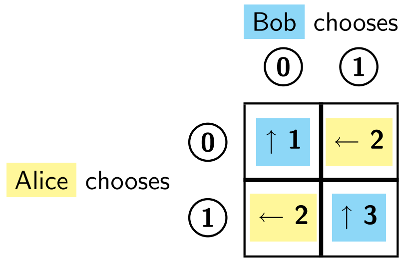

Figure 2: *Payoff matrix for a coin game*

differ. When both chose 0, Bob receives € 1 from Alice, and € 3 when both chose 1. This is a zero-sum game, because one player's profit is the other's loss. The goal is to maximize the total profit under repeated play.

The average profit per move of $(-1+2+2-3)/4 = 0$ seems to make this a fair game. However, Alice can outplay Bob. Alice needs to use a *randomized strategy* to guarantee that she cannot be exploited by Bob. She chooses 1 with probability x. For $x = 50\%$, her expected payoff equals $(-1+2)/2 = +0.5$ if Bob chooses 0, and $(2-3)/2 = -0.5$ if he chooses 1. In that case, Bob would always choose 1 to outplay Alice big time. In general, her expected profit can easily be determined for both Bob's choices. It turns out to be a linear interpolation between the extremes, as shown in Figure 3.

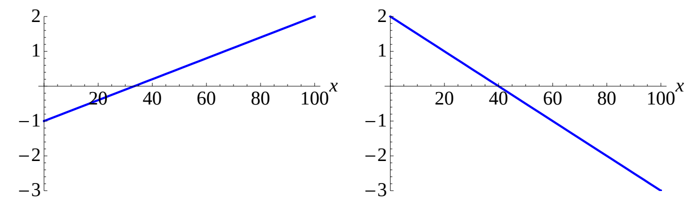

Figure 3: *Expected payoff for Alice if she chooses 1 with probability x (in %), and Bob chooses 0 (left) or 1 (right)*

Of course, Alice does not know what Bob will do. So she assumes the worst, viz. that Bob knows her x. Bob doesn't know Alice' choice, but he can also inspect the graphs and choose whichever provides him the best expectation given x. This is shown in Figure 4 on the left, where Bob maximizes his profit by minimizing Alice' and opting for the lower graph. Fortunately for Alice, there is a small window of opportunity, peeking at $x = 3/8 = 37.5\%$ and bringing her an expected profit of 0.125, or 6.25% of the average stake of € 2 (not a bad return on investment these days :-).

It is worth noting that Alice' expected profit is independent of Bob's strategy, even if he randomizes as well. The graph on the right in Figure 4 shows the expected profit when both players independently randomize their choices. At Alice' optimum of $x = 0.375$,

[1] The Math & CS department of TU/e issued these binary PR coins.

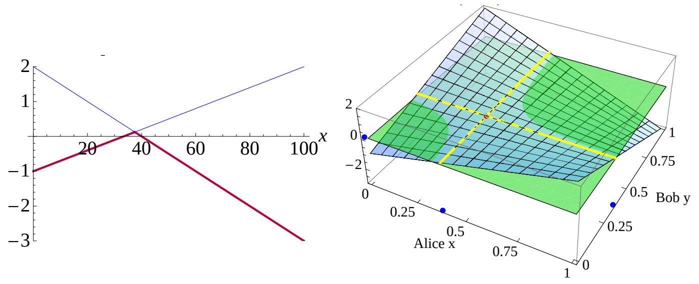

Figure 4: *Expected payoff for Alice if she chooses 1 with probability x, and Bob optimizes his choice (lower graph in red, left), or chooses 1 with probability y (right)*

the expected profit is constant, and similarly for Bob's optimum (yellow lines). This so called Nash equilibrium is immune to the other player's strategy. When players have three or more options, it becomes harder to determine optimal strategies, but John Nash showed that they always exist, even for nonzero-sum games.

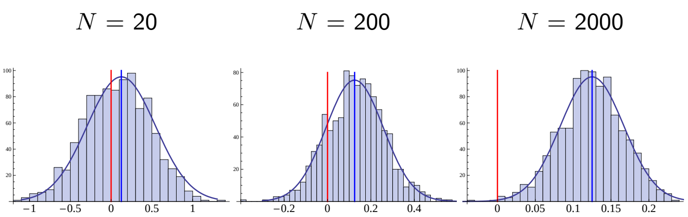

Figure 5: *Distribution of 'profit after a run of N games', based on 1000 simulated runs; blue line at expected profit, red line at 0 profit*

This all looks very nice, and can even be of practical value. But there is a danger lurking: variance. When playing this game 20 times (both players using an optimal randomized strategy), the expected average profit for Alice is indeed $\mu = 0.125$, but with an overwhelming standard deviation of $\sigma = 0.42$. This means that in roughly 38% of the 20-game runs, Alice ends in the negative (Figure 5 left). Playing 200 games reduces the standard deviation by a factor $\sqrt{10}$ to $\sigma = 0.13$, still giving a whopping 17% probability to end in the red (Figure 5 middle). Only by playing 2000 games with $\sigma = 0.04$ is the probability to lose acceptable at 0.1% (Figure 5 right). Bob cannot affect the expected profit if Alice randomizes optimally, but he can affect the variance by deviating from optimal play himself (homework)!

Solitaire Yahtzee: A Game of Chance

A pure game of chance is solitaire Yahtzee, played with five dice and a scorecard with thirteen primary categories (marked with * in Figure 6). Each turn, the player rolls all dice, and may re-roll any subset two more times. At the end of the turn, the roll must be scored in one of the empty primary categories. The actual score is determined by the rules, depending on the roll and the category (not explained here, for instance see [3, 4]).

Each game involves 38 decisions: which subset to re-roll (26×), which category to score (12×). Note that after the final roll, there is only one empty category left, leaving no choice. A careful analysis reveals that there are slightly over 10^9 choice states in the reduced game graph, whose edges represent roll and choice events.

Since all probabilities are known, it is possible to optimize the decisions for maximizing the expected final score, by modeling Yahtzee as a Markov Decision Process. In 1999, I wrote a computer program to carry out the calculations. Later that year, we made available an on-line advisor [4], who can rate all options in any game state that you present. On that web site, you can also do a Yahtzee proficiency test and have your decisions compared against the optimal strategy. More recently, my programs were applied for [3].

Under optimal play (using official rules), the expected final score is just over 254 points. Human players have a hard time matching that with their long term average. Again, however, variance upsets the picture. I discovered a new method [5] to calculate variances for Markov chains with rewards, which I applied to the optimal Yahtzee strategy. It turns out that the standard deviation in the final score under optimal play is almost 60 points. To guarantee (at 3σ confidence level) an average final score of over 250 points, the optimal strategy needs to play over 2000 games. Figure 6 shows some characteristics of the optimal strategy. It obtains a Yahtzee almost every other game!

Category	**E**	**SD**	**% 0**
* *Aces*	1.88	1.22	10.84
* *Twos*	5.28	2.00	1.80
* *Threes*	8.57	2.71	0.95
* *Fours*	12.16	3.29	0.60
* *Fives*	15.69	3.85	0.50
* *Sixes*	19.19	4.64	0.53
U. S. Bonus	23.84	16.31	31.88
* *Three of a Kind*	21.66	5.62	3.26
* *Four of a Kind*	13.10	11.07	36.34
* *Full House*	22.59	7.38	9.63
* *Small Straight*	29.46	3.99	1.80
* *Large Straight*	32.71	15.44	18.22
* *Yahtzee*	16.87	23.64	66.26
* *Chance*	22.01	2.54	0.00
Extra Y. Bonus	9.58	34.08	91.76
GRAND TOTAL	254.59	59.61	0.00
Yahtzees Rolled	0.46	0.69	63.24
Jokers Applied	0.04	0.19	96.30

Figure 6: *Optimal solitaire Yahtzee: expectation (E), standard deviation (SD), and percentage of zeroes (% 0), per category*

Conclusion

Many real-life situations resemble game playing, in the sense that both involve the need for making good decisions in the presence of constraints and uncertainty. It is not surprising that the same mathematical techniques are helpful in both situations.

Combinatorial games have seen much progress recently due to improved algorithms and computer hardware. For instance, computer go is now making a leap forward through Monte-Carlo methods and UCT (Upper Confidence bounds applied to Trees).

It may come as a surprise that for repeated strategic decisions, it can be optimal to toss a (well-chosen) 'coin'. The so-called Mixed Nash Strategy employs randomization to prevent predictability and hence exploitation. Clever randomization can be profitable.

For repeated tests of fortune, it can be optimal to make a (well-chosen) fixed choice, based on so-called Markov Decision Processes.

Do not, however, underestimate the role of variance. A large standard deviation requires patience (in the form of a counterintuitively large repeat count N) to increase your odds, due to the factor $1/\sqrt{N}$. Larger variance reduces predictability and hinders planning. In fighting traffic congestion problems, it is much better to aim at reducing the variance (improving predictability) than at minimizing the expected travel time (which often increases variance considerably).

References

[1] E.R. Berlekamp, J.H. Conway, R.K. Guy. *Winning Ways for Your Mathematical Plays*. Second Edition, 4 Volumes. A K Peters, 2001–2004.

[2] J. Bewersdorff. *Luck, Logic, and White Lies: The Mathematics of Games*. A K Peters, 2004.

[3] P. Gordon. *Yahtzee Scratch & Play* (2 volumes). Sterling Press, 2007, 2009.

[4] T. Verhoeff. *On-line Optimal Solitaire Yahtzee Advisor and Proficiency Test*. URL: `www.win.tue.nl/~wstomv/yahtzee` (1999).

[5] T. Verhoeff. "Variance in Markov Chains: A Calculational Approach" in *Proceedings of the Eindhoven FASTAR Days 2004* (3-4 Sept. 2004).