STAY IN COMMAND: OPTIMAL PLAY FOR TWO PERSON GENERALA

Joseph Heled¹

jheled@gmail.com

1 The game of Generala

Generala is a social multi-player dice game popular in Spanish and Portuguese speaking countries. Generala is similar to Yacht, Yatzy, Kniffel, and other close varieties. The most well known variant, a commercial game marketed as Yahtzee, still sells 50 million games each year.

All five variants above use five dice and the player scores by matching the numbers on the dice with certain combinations. Generala has 10 such combinations: Ones, Twos, Threes, Fours, Fives, Sixes, Escalera (1-2-3-4-5 or 2-3-4-5-6), Full-House (XXXYY), Four-of-a-Kind (XXXXY), and Generala (XXXXX). In a turn, the player rolls all five dice and then may either stand or re-roll some of the dice. She may then re-roll some of the dice a second time, but after the third roll the player must score the five dice by matching them with one unused combination. Generala scores 50 points, Four-of-a-Kind 40, Full House 30 and Escalera 20. The Ones to Sixes score according to the number of dice with that value. For example, 1-1-3-4-4 scores 2 points as Ones, 3 points as Threes and 8 points as Fours. In some versions a player scores 5 or 10 bonus points when hitting a combination on the first roll, but here I will consider only the “Vanilla” version.

The player may *wave* a combination, i.e. pick a combination which does not match the dice, for a score of zero. For example, a player might wave Generala with 1-1-1-1-5 if Ones is no longer available to score. Each combination can be scored exactly once. After completing one turn the dice pass to the player on the left, and this process continues until all players have completed 10 turns and either scored or waived all their combinations. The player with most points wins.

As is true for most games, the origin and early history of Generala is mostly unknown, and it is simply considered a “traditional game”. Dice of various kinds pre-date recorded history and feature in 5,000-year-old games such as Senet and the Royal Game of Ur

¹Independent researcher:orcid id 0000-0002-2403-7997

([2]), so the Generala family of games might be anywhere from hundreds to thousands of years old. The earliest mention in modern literature, however, seems to be the listing of *Yacht*, together with *Crag*, *Everest*, and *Cargo*, in the comprehensive *Complete Book of Games* published in 1938 ([7]). Yahtzee's popularity in the 80's and 90's led to papers, theses and high-school exercises which explored the various strategies for playing Yahtzee ([3], [4], [1], [6], [5]).

The publications usually promised optimal or near-optimal strategies for playing the game, but the game the papers investigated was not the social "group Yahtzee" but "solitaire Yahtzee", the one player game where a single player is simply trying to maximize her Yahtzee score. This "partisan" strategy is not without merit in the multi-player game, but it is not optimal even for two-person Generala. An optimal strategy must take into account the state of all other players, and this is quite a challenge with Yatzy and Yacht for one practical reason: size.

2 The size of Generala

A player can end up with six possible scores for Ones: 0, 1, 2, 3, 4 and 5. Similarly there are six possible scores for Twos (0, 2, 4, 6, 8, 10) and six for Threes (0, 3, 6, 9, 12, 15). Same for Fours, Fives and Sixes. The other four combinations have two possible scores each (0 or 20, 0 or 30 etc.) Each Generala position consists of the score of each combination: (3 in Ones, 6 in Twos, – in Threes, . . . , 0 in Generala), one such list for each player. The '–' signifies that the combination is still available. In this naive representation the number of possible positions is approximately 10^{25} . Luckily, many of those distinct positions are game-wise identical. For the purpose of playing correctly, individual scores do not matter, only their total sum does. It does not matter if you got 3 points from the Ones, Twos and Threes combinations by scoring 3 1's, one 1 and one 2, or one 3. At the end you got 3 points from those three combinations. Ignoring individual scores and considering only the sum leads to a huge reduction. For example, there are 216 (6^3) ways to score Ones, Twos and Threes, but only 31 possible total sums. So, any Generala position can be effectively encoded as a quadruplet: (X-remaining-combinations, X-score, Y-remaining-combinations, Y-score). In this condensed encoding there are 4,719,060,648 ($\approx 5G$) Generala positions. This number is again reduced by the simple observation that the actual scores themselves do not matter. What matters is the difference between them. The score can be 50 vs. 70 or 100 vs. 120, and in both cases X needs to overcome a 20 points deficient. This almost trivial observation means each position can in fact be encoded as a triplet: (X-remaining-combinations, Y-remaining-combinations, XY-score-difference), for a total of 85,647,207 (86M) unique positions. This makes Generala a very rare beast: a popular game which is small enough to be completely solved. Completely solved means it is possible to construct a perfect player, one which makes the best move

in any possible Generala situation. In comparison, Yacht (with 12 combinations) has 1,557,687,661 (1.5G) positions and Yatzy (15 combinations) has 98,793,323,900 (99G) positions. While those are not impossible sizes for a large server farm, Generala can be solved on a small modern personal computer in less than a day.

3 Solving Generala

A Generala game can end up in a win for X (+1 points for X), a win for O (−1 points for X), or a draw (0 points for X). Every position (assuming best play) has an equity, which is the expected number of points won by the player to move. For example, consider position (A) below.

(A)

X		X	X	X	X	X	X	X	X	69
O	X	X	X	X	X	X	X	X	X	87

For simplicity, the points of previously scored combinations are not given, only their total sum is shown at the right. Here X must get a Generala on her last turn to win. Since the probability for getting a Generala is

$$p = \frac{2783176}{6^{10}} = 0.046$$

the equity is $p \cdot 1 + (1-p) \cdot (-1) = -0.908$. After calculating the equities of all positions one combination away from the end we can compute the equities of positions two combination away from the end:

(B)

X		X	X	X	X	X	X	X	X	69
O	X		X	X	X	X	X	X	X	47

Here X needs either to get a Generala or hope O fails to get Four-of-a-Kind on her last turn. The probability of getting FOAK on your last turn is

$$q = \frac{17583176}{6^{10}} = 0.291$$

and so X's equity is $p \cdot 1 + (1-p)[(1-q) \cdot 1 + q \cdot (-1)] = 0.445$. After computing the equities of all positions two moves from the game end we can back to positions three moves away from the end and so on, all the way back to the start position.

I obviously skipped many important details. I omitted how the probabilities of making Generala or FOAK were obtained, and how equities are computed in more complex situations. In general, to compute an equity you must take the best course of action at each stage, that is, select the best subsets to roll and the best combination to score against, where 'best' means whatever maximizes the equity. This is done in a similar way to the examples above, by tracking backwards action by action, starting with the equity after the third roll for any possible dice roll, then working out the equity after the second roll (again for all possible dice), then computing the equity after the first roll, and finally taking the average of the equity of all possible rolls to determine the final position equity. An illustration how this process works is found in Appendix A.

4 Untangling Generala

The fan finally stopped whining and the smoke around the laptop has cleared. Now what? We are like the blind men and the elephant, only our elephant is 85,647,207 parts strong, and every way we slice it will be incomplete. The 86M expectations lets us easily build a perfect player, which I shall name Malinalxochitl (Malina for short), the goddess of snakes, scorpions and insects of the desert. We can play against Malina and study her moves. We can let her comment on our play and get wiser from that. We can pit her against other strategies and see how well they do, or analyze the effectiveness of general rules of thumb. Here, obviously, I can fit only a very small part of this large elephant. Let me start with the most basic question: Is it an elephant or a mouse?

4.1 Luck and skill in Generala

In other words, is it worth the effort? Perhaps Generala is, as most casual players tend to see it, a game of luck requiring very little skill? Even with Malina there is no clear cut answer. On the one hand, Generala is not a game of pure luck like Snakes and Ladders or a slot machine. Some skill is clearly required; a totally random player – which never re-rolls and randomly selects a combination – will win roughly one game in 4600 when playing against Malina. A better player who keeps the dice but selects the combination yielding the most points will win about one game in 70.

And when it's not pure luck, it is **always only about skill**.

The dice might favor the weaker player in the short term, but the stronger player always prevails when enough games are played. That said, it might not matter in practice and Generala may still require very little skill. Because, what is skill? By skill we typically mean performing better than "average", or better than what we consider "normal",

and playing randomly is not considered normal and results in very poor play. We need a “human like” average player to serve as a baseline. There might be many ways to construct this middle-of-the-road Joe Simpleton, but it is important to base him on human-like principles and not on computed equities or expectations. For example, we can base Joe on simple pattern distances. Each dice roll has a natural distance to each combination, which is the number dice needed to complete the combination. 1-1-1-5 is 1-away from Generala and 0-away from FOAK. 1-1-1-2-4 is 1-away from FOAK, 2-away from Ones and Generala, and 2-away from Escalera since a 3-5 together with 1-2-4 will form an Escalera. Joe keeps the dice matching the closest combination still available and rolls the rest. For scoring, again the closest pattern is selected, so 1-1-1-4-5 is scored as Ones (3 points) and not as Fours or Fives, even though those yield more points. You can find the full details in the code, but on the whole it is a very simple-minded strategy and indeed plays very human looking moves. And reasonably well. Joe expected score against Malina is -0.24 , or about 3 wins and 5 losses out of 8 games.

So, Malina shaves about one dollar from Joe every four one-dollar games. For the Chess and Backgammon players among you, this is equivalent to a difference of 85 Elo points. Whether this is high or low is a matter of taste; I would call this medium to medium-low. It is low compared to skill differences common in complex games, yet significantly worse than most slot machines, where you typically lose “only” 10 to 12 cents per dollar. This is not scaled for time, though, as a game of Generala takes much longer than pulling one lever. So, when playing for money, better stick with Generala.

4.2 Seeing the light?

In some games it is easy to see errors in hindsight. Not so in Generala. Here is a position from a game I played against Maximus:

Me		X	X		X	X	X	X			116
Maximus		X	X			X	X	X	X	X	104

I rolled 1-3-3-3-4. I kept the 3's and rolled 3-6. I re-rolled the 6 and ended up with 3-3-3-3-4. Nothing scores, so I waived Escalera. On his next move Maximus equalized with two 6's, and went on to win by making an Escalera on the last move.

Keeping the 3's was an error. Not a large one, only -0.01 , but I should have kept 1-3-4, aiming either at Escalera if I got a 2 or a 5, or at Ones if I got more 1's. Waiving Escalera, on the other hand, was a huge blunder (an equity loss of -0.31). Had I waived Generala I would have been a 3 to 2 favourite to win the game, but after my move it was 4 to 5 against me.

Maximus play, by the way, was also inaccurate. With two 6's he should have waived Generala (0.03 better), hoping to catch me by rolling three or more 6's later on or an Escalera.

Sometimes Malina's play seems to defy reason.

Me		X		X		X			X	76
Maximus		X	X		X	X				114

Here I rolled 1-1-3-3-3 and immediately scored the Full-House. I didn't even stop to think about it. Maximus approves. Malina says it is a -0.04 error, and I should have kept the 3's and rolled on. How can that be? Well, even after scoring the FH I am 12 points behind. I still need to make enough with my 6's, 1's, 2's and 3's to counter the 12 points deficit plus what Maximus might get from his 1's to 4's, and obviously he has six more moves to go for an Escalera as well. In short, the more you study the position, the more attractive Generala becomes. And with three 3's I make that Generala about 9% of the time and still get the Full-House 21% of the time when I fail. At worse I score 9 or 12 points in 3's, and hopefully get a FH or Generala later.

Knowing about the errors is great, but only partially helpful for us humans. To learn and improve we need a story, an explanation, something Malina does not directly provide. For keeping the 1-3-4 instead of the 3-3-3 I came up with the "Escalera or Ones" narrative, but only the hard numbers confirm it is slightly better than aiming for a Generala. You can make you own story why scraping the Escalera was so bad, but will that protect you from making similar mistakes in the future?

4.3 Scaling Escalera

I waived Escalera mostly because my experience-backed-gut-feeling said "Escalera is hard to get". I was wrong and need to work harder on not thinking with my gut¹, but my gut was wrong for a subtle reason. Let me explain.

Escalera is indeed harder to roll than the higher scoring combinations of FH and FOAK. The probability of rolling a FH (when you focus solely on getting one) is 36.6%, FOAK is 29.1%, and Escalera is 26.1%. **But only when making the right plays!** For example, what do you keep from 1-2-2-6-6? from 1-1-4-6-6? from 1-2-2-3-6? I used to keep, respectively, 1-2, 4-6 and 1-2-3. Wrong, wrong and wrong. Never keep 1's and 6's, says Malina, unless they are part of a quartet one die away from Escalera. So in the examples above you should keep one 2, one 4, the 2-3, but keep 1-2-4-5 from 1-2-2-4-5

¹Carl Sagan, The Dragon in My Garage

or 2-4-5-6 from 2-4-4-5-6. Two simple directives that are worth between 14 and 20 Elo points! with the right play you can get an Escalera $\frac{1}{4}$ of the time.

4.4 Waiving for fun and profit

I waived Escalera because I had to waive something. Like most humans I dislike wasting resources, and almost never waive a combination when I can score some points. Surprisingly, Malina regularly waives even when points can be made, at a rate of about twice every five games. Naturally, scraping becomes more likely as the game progresses, but Malina might waive at any time. She might even do so on the first turn! If she ends up with a pair of 6's and a bust after the third roll (2-3-4-6-6, 2-3-5-6-6, 2-4-5-6-6 or 3-4-5-6-6), Malina will waive the Ones instead of scoring the two 6's or a single 5, 4, 3, or even a 2! You might think it is unlikely to end up with those rolls at the first turn, but it happens once every 120 games. Or more precisely, it happens at this rate to Malina, not to you, because you misplay your openings. Malina, for example, keeps just the pair of 6's from any X-X-X-6-6 unless it's a FH. In fact, Malina never goes for a FH on the first roll of the first move: she keeps just 2-2 from 1-1-2-2-5 and even 1-1 from 1-1-2-3-4. She will only go for Escalera on the first move when she rolls a combination which is one away, like 2-2-3-4-5, 2-3-3-4-5, where the probability of hitting an Escalera on the next turn is $\frac{1}{3}$, and $\frac{5}{9}$ for getting it in two rolls.

Mind you, scoring the two 6's for 12 point is only a small error (less than 1 Elo point), but remember that Malina's edge over simple-minded Joe is not large to start with, and is made of squeezing every possible bit of advantage out of every action, small as it may be.

Sometimes scraping is easy.

X		X	X	X		X	X	X	X	80
O	X	X	X	X	X	X	X	X	X	99

X is 19 points behind and needs to decide how to score x-x-6-6-6 (where x is not a 6). Scoring 18 points would still require making a Generala to win, so X should waive the Generala and hope to make four or five 6's on her last turn, much more probable than a Generala.

X	X		X	X	X	X		X	X	80
O	X	X	X	X	X	X	X	X	X	88

X is 8 points behind and ended up with x-x-x-4-4. Scoring 8 with Fours would equalize the score, but even a single 1 from O (with three rolls, O roll at least one 1 95.5% of

the time) would require making the FOAK. X should waive FOAK and hope to make enough 4's on her last turn to win. However, if she had 81 points she should score the 8 points! Now a single 1 from O (about 47%) would only tie for O, and she might still make the FOAK. In addition, if she had a a FH still available and not the FOAK, scoring the 8 points and trying for a FH is better than scraping the FH.

X	X		X	X	X	X	X		X	X	92
O		X		X	X	X	X	X	X	X	70

X is ahead by 22 points but it's not over yet. X needs to score a x-x-x-3-3. If she scores 6 points with Threes she will be ahead by 28 points, and still need to make the FOAK if O manages to get a FH in one of her two remaining turns. X has a better chance if she waives the FOAK and goes for three 3's or more the next turn to protect against the FH.

However, if she had 94 points she should score the 6 points, since a FH from O would now only equalize. Somewhat surprisingly, she should also score when she has one point less (91), since now she will need four 3's to win if O makes a FH. Again, she waives only with 92 and 93 points, and scores the two 3's with anything above or below!

X		X	X	X		X	X	X	X	X	80
O	X	X	X	X	X	X	X	X			85

X is 5 points behind and need to score x-x-x-6-6. Scoring 12 points will put her ahead by 7, which is not enough against O 1's and 2's, so she waives the Generala. If she had one more point she should score the 12 points.

X						X	X	X	X	X	30
O				X		X	X	X	X	X	42

X is 12 points behind and needs to score x-x-x-x-6. 6's are her most valuable asset for closing the gap and she should waive Generala. In fact she should waive Generala on anywhere between being 20 points down to 75 up. When down by more than 20 she should waive Escalera, as Generala is more valuable in closing large deficits.

Waiving is an integral part of Generala. It happens (under optimal play) about 1.2 turns every 10 turns. It is almost always Generala or Ones, and Generala is 3 more times

likely than Ones. Escalera comes at a very distant third, happening less than 1% of the time. As we saw, it might be tricky to choose which combination to waive and especially when to waive voluntarily. Waiving based both on score and the remaining combinations is something unique to Malina and is a big part of her advantage over lesser mortals. Typical Generala strategies, like Joe, consider only the dice and their own remaining combinations.

5 Meet Maximus

I hyped Malina a little, but partisan only-my-side-matters players can be surprisingly strong. At least I was surprised how strong Maximus is. Maximus simply plays to maximize his expected score.

Maximus		X	X			X	X	X	X	X	104
Me		X	X	X	X	X	X	X			116

When Maximus had to score 1-4-5-6-6 against me he informally followed this reasoning. “Scoring the 6’s gives me 12 points and the expected number of points from Escalera+Generala is 10.3, so 22.3 in total. Waiving Escalera gives 0 points, and the expected score of Sixes+Generala is 17.7. Waiving Generala is again 0 and the expected score of Sixes+Escalera is 20.4. 22.3 is the largest, so I score the 6’s”. As you can see, my score or my remaining combinations played no part in this reasoning. The expected score of each subset of combinations is computed in exactly the same way as computing Malina’s game equities, only here the base, what we enter in the first line of the table, is the expected one-sided score and not the expected game score. Since there are only 1023 such subsets this can be done very quickly and in addition those expectations can be easily computed for Yatzy or Yacht or any other game in the family.

6 Maximus strength

How good is Maximus? Usually you need to play the two sides a large number of games to get an accurate estimate. How many? Assume that X wins around 55% of her games. To estimate this probability to (say) three significant digits you want the standard error, the standard deviation of your estimator to be $\frac{0.001}{3}$, since 3 std’s gives you a confidence interval of more than 99%. The standard deviation of running n trials is $\sqrt{\frac{0.55 \cdot 0.45}{n}}$, and so you need to run around 2.2M games. With Generala, however, it is

score	105	104	103	102	101	100
probability	0.013	0.091	0.250	0.344	0.236	0.065

possible to compute the value directly by computing the expected score of each position when played by Malina against Maximus, and Maximus against Malina. It is more work than solving Generala in the first place, but you get a definite answer. And here it is: when going first, Malina expected score against Maximus is 0.033, and 0.045 when going second. Assuming they flip a coin to determine who goes first, Malina wins 3.9 cents per game, or a difference of 14 Elo points.

In short, Maximus is an excellent player, losing just 1 dollar per 25 games. Also note the advantage of going second. The second player has more information at each turn, which for Malina translates to winning 1.2 cents per game more than when she goes first. This is of course true when playing against herself. Assuming perfect play, the second player has an advantage of 1 cent per game.

6.1 Mindful Maximus

Maximus is the best partisan player. Malina is the ultimate player. To operate Maximus needs 1023 floating point values, about 4 kb of computer memory (if using single precision floating point numbers). Malina requires about 268 Mb. I am skipping the details because the exact numbers are not important; what matters is the order of magnitude. Data-wise, the difference between Malina and Maximus is the difference between a Micro and a Major Earthquake. Is there anything in between?

In the position below, Maximus (X) needs to score 1-3-3-3-3.

X		X	X	X	X	X	X	X		55
O	X	X	X	X	X	X	X	X	X	100

The expected score from Twos is 4.2. The expected score from Generala is 2.3, so Maximus waives Generala. This is obviously wrong since a Generala is the only way to close a 45 points gap. How can we make Maximus aware of the score difference?

Here's how. Instead of keeping only the expected score keep the whole distribution. That is, keep the full list of possible scores and the probability of reaching each score. Here, after waiving Twos, X scores 105 with probability of 0.046, and 60 with probability of 0.954. Those are the only scores available for X. The score distribution for O is this:

If X makes the Generala and O scores 5 both will have 105, so the probability of a tie is $(0.046 \cdot 0.013)$. If X score is 105 and O makes anything else, X wins $(0.046 \cdot (1 - 0.013))$.

Otherwise O wins. So X's equity will be $0.046 \cdot (1 - 0.013) + (1 - 0.046) \cdot (-1) = -0.906$. After waiving Generala X probability of winning is 0, so his equity is -1 . $-0.906 > -1$, so waiving Twos is better.

What we actually did was compute the game equity on the fly **assuming both players are Maximus**. That is, the game decisions of Mindful Maximus (MM) are based on actual game equities, but only if those games were played by two copies of Maximus. Maximus is not perfect, but he is a very strong player and so those equities are excellent approximations of the true equities. As a result MM is only 0.01 (3.5 Elo points) worse than Malina, and uses 440kb of computer memory, two orders of magnitude above Plain Maximus, a light earthquake. MM is a significant improvement because the approach can be applied to Yatzy, Yacht, or any game in the family, yielding, I believe, a near optimal player for those games as well. It's not certain that MM is an almost optimal player for the bigger variants, but the tiny gap between Malina and MM in the two-person Generala makes it quite likely.

7 Generala in the wild

Generala is a social game typically enjoyed by more than two people. Sadly, even the three player Generala is 3 orders of magnitude bigger than the two player game, with almost 100G positions. However, while three person Generala and up, like Yacht and Yatzy, is too big to be effectively solved, the techniques developed for two player Generala are still useful in multi-player games.

First, partisan players like Maximus are oblivious to the number of players or their state, so can play in any configuration. It is even easy to compare their relative strength, employing the same method used to compute Maximus strength (see page 61), using only the distribution of the final scores of each player. Here is a little table showing how much Maximus wins, per game, against 1, 2, 3 and 4 'Joe's.

no. of players	2	3	4	5
Maximus PPG	0.24	0.31	0.38	0.41

Maximus indeed wins more (per game) as the number of players increase, but this is slightly misleading. Since a four for person game takes twice as long as a two person game, the amounts should be scaled to reflect how much is won per time unit. If the two person game is 1 time unit, we get this table instead,

So Maximus actually wins less per hour, and if playing for money he should seek games with less participants. This pattern seems to hold for all players.

Malina can adapt to handle multi-player Generala. She can't play optimally against

no. of players	2	3	4	5
Maximus PPTU	0.24	0.21	0.19	0.17

all players, but she can pick the most “dangerous” opponents and play optimally against her. The most dangerous opponent is the one most likely to win against her in a two-person game, that is, the player with the lowest pairwise equity. For example, here Malina plays second in a 3-person game.

Joe A.	X	X								70
Melina			X							30
Joe B.				X						20

Even though Joe A. currently leads, her equity against him is -0.10 while her equity against Joe B. is -0.12 . Malina then plays against Joe B. Informally, making the hard-to-get Escalera made Joe B. slightly more dangerous than A.

Since 3-person Generala is not solved we can’t compute directly how good Malina is, and have to resort to running matches and tallying the results. This is a slow process and I run only 300k games between Malina, Joe A. and Joe B., alternating between all six players orders. Those games indicate that with 95% probability Malina’s equity is 0.359 ± 0.005 . That is, I can’t pinpoint the exact value but almost certain she is doing better than Maximus (0.31). A difference of 0.04 seems reasonable since it is the exact same difference between Malina and Maximus in 2-person Generala.

7.1 Global Mindfulness?

Malina play in multi-player Generala is not optimal, but probably quite close. This is of course only a personal opinion, but here is another reason why I think this is the case.

Yacht and Yatzy (and even Crag) are too big to be solved, but it is relatively easy to construct a Maximus and a Mindful Maximus for their team. Given how close Malina and Mindful Maximus are in Generala, it is reasonable to assume this also holds in Yacht and Yatzy, and perhaps even more so, since the difference between the Yacht/Yatzy Maximus and their MM is smaller than the difference in Generala. Can we then adapt MM to play in multi-player Generala, Yacht and Yatzy? quite easily. Remember, MM plays by “faking” game equities on-the-fly, assuming the other player is Maximus. It is easy to do that for multiple opponents. We just need to combine the individual final score distribution from each opponent into one distribution of the highest score, and

generate approximate game equity from that. Again, it will not be the real game equity, but it does not need to be. It only needs to preserve the equities order, that is, if move X is better than move Y it suffices that the approximate equity of X is greater than the approximate equity of Y. Backgammon software plays well for exactly the same reason. To select a move you only need to approximate the equities well enough so that the best move comes on top, or at least one of the very best moves ends on top.

When running 300k Generala games between Mindful Maximus, Joe A. and Joe B., MM's equity comes out as (with a 95% probability) 0.0355 ± 0.005 . It is impossible to tell from just 300k games if MM did better than Malina or not, but in any case the difference between them is very small, and that makes me believe that Mindful Yacht and Mindful Yatzy are almost optimal players as well. And even if not, they are probably stronger than any human player.

8 Executive summary

- I. Generala, under the right encoding, is small enough to be completely solved. That is, a perfect Generala player is (relatively) easy to construct. The process is straightforward mathematically since the game graph of those games is a simple tree, and so equities can be solved by working backwards.
- II. Yacht, Yatzy (and even Crag) are too big for comfort (1.5G, 99G and 4.8G positions vs. 86M for Generala).
- III. A simple minded opponent (Joe) loses to Malina (the goddess of Generala) about 1 Dollar per four games. This is equivalent to a difference of 85 Elo rating points. Most real humans are probably better than that, but it is unclear how humans fair in general.
- IV. Maximus, a partisan player only interested in maximizing his own score, loses to Malina about 4 cents per game (14 Elo points difference).
- V. Mindful Maximus, a player which generates approximate equities on-the-fly by assuming all players are simple maximizers is near optimal, only 3.5 Elo points behind Malina (−1 cent per game).
- VI. Malina can be adapted to play in a multi-player game by picking the most “dangerous” opponent and playing against her as if it was a two person game. Using this simple trick Malina keeps her 4 cent edge over Maximus in multi-player games.
- VII. Mindful Maximus can also be adapted to play multi-player Generala, by combining the individual score distribution from each opponent into one distribution of the

highest score of all opponents. Mindful Maximus and Malina seem to be of equal strength in multi-player Generala.

VIII. Playing against the most “dangerous” opponent (when you have an optimal two-player strategy) and combining several opponent score distributions into one distribution (effectively treating the group of opponents as a single opponent) are generic ideas that might be useful in other multi-player games. It will be interesting to see if these ideas are also effective in more complex games.

IX. It is relatively easy to build Mindful Maximus for Yacht and Yatzy, creating extremely strong players which can also annotate human play.

X. The code is available in this [GitHub repository](#).

Generala-like games are social in nature and typically played more casually than, say, Chess or Backgammon. On the other hand, people are combative by nature and usually play to win. This short informal account tries to shed some light on the possible scope of skill in Generala-like games, and how much advantage is theoretically possible. The conclusion seems to be that the skill gap between an average player and a top notch professional, assuming there was such a thing, is much smaller than in games like Chess and Backgammon, but is not insignificant. A professional would still win a 25 point match against an amateur 90% of the time, assuming the amateur plays somewhere near the ‘Joe’ level. Most Generala errors can be subtle, at least for casual players. Casual players don’t lose by making big blunders, they make small errors over the course of the 29 decisions taken during one game of Generala, and are probably unaware how good or bad their actions are. If you are one of them, now you have the chance to play against Malina, analyze your games and get better at gaining social mojo.

Image via <http://www.peakpx.com>.

A Computing equities

It would be extremely tedious and unhelpful to work thorough a complete computation of one full Generala position, so let's use a stripped down version of Crag instead. Crag uses only three dice and two rolls per turn, and we will further use three-sided die – faces 1, 2 and 3. Assume that O has only Ones left, X Ones and Twos, and is trailing by 3 points. What is her equity? We shall work it out by filling the table below.

	1-1	1-2	2-2	1-3	2-3	3-3	1	2	3	–
after 2 nd roll										
before 2 nd roll										
after 1 st roll										

The rows specify the stage of the turn, while the columns list the dice kept at this stage. The first row, after the second roll and before picking a combination to score

against, is filled using previously computed equities. With a 1-1 dice, the position becomes: 'O (to play) has Ones and leads by 1 point, X has Twos'. The equity of this position is -0.056 , and we enter its negation in the table (if O wins m on average, then X loses m on average). With a dice of 1-2 we need to consult two equities. After marking Ones, X is 2 behind and still has Twos, equity -0.380 . After choosing Twos, X is 1 behind and has Ones, equity -0.628 . Losing -0.380 is better than losing -0.628 , so that what goes into the 1-2 column. The rest of the first line is filled using the same logic. Remember, I simply stated the equities of future positions, but they are computed in exactly the same way. Turtles all the way down.

	1-1	1-2	2-2	1-3	2-3	3-3	1	2	3	–
af. 2 nd	0.056	-0.380	0.628	-0.380	-0.628	-0.726	✗	✗	✗	✗
bf. 2 nd										
af. 1 st										

Now to the second row, before the second roll. Here we fill the row right to left, starting with the last column, keeping none. If X keeps none of the dice from the second role the equity will be the average of the nine possible rolls: 1-1, 1-2, 2-1, 1-3, $\dots$, 3-2, 3-3. All those values are available in the first row. If X keeps a single 3, the equity will be the average of the equity of 1-3, 2-3 and 3-3 from the first row. Same for 2 and 1. For 3-3, X can either keep both, keep a single 3 or keep none. All those have already been computed, but here we take the **maximum** of those values, since X will take the action with the best (largest) equity. Same for keeping a 2 or a 1. When keeping both (first six columns) we simply copy over the value from the previous line.

	1-1	1-2	2-2	1-3	2-3	3-3
af. 2 nd	0.056	-0.380	0.628	-0.380	-0.628	-0.726
bf. 2 nd	0.056	-0.380	0.628	-0.380	-0.628	-0.726
af. 1 st						

On to the last row. After rolling 1-1 we can re-roll both (equity -0.313), keep one 1 (equity -0.235) or keep both (equity 0.056). The best (largest) equity is 0.056, so this is

	1	2	3	–
af. 2 nd	X	X	X	X
bf. 2 nd	–0.235	–0.127	–0.578	–0.313
af. 1 st				

the equity after rolling 1-1. After rolling 1-2 we need to compare the entries in the 1-2, 1, 2, and – columns and take the largest. Same for the other columns.

	1-1	1-2	2-2	1-3	2-3	3-3
af. 2 nd	0.056	–0.380	0.628	–0.380	–0.628	–0.726
bf. 2 nd	0.056	–0.380	0.628	–0.380	–0.628	–0.726
af. 1 st	0.056	–0.127	0.628	–0.235	–0.127	–0.313

	1	2	3	–
af. 2 nd	X	X	X	X
bf. 2 nd	–0.235	–0.127	–0.578	–0.313
af. 1 st	X	X	X	X

The position equity, the average equity of all possible rolls, is

$$-0.289 = \frac{0.056 + 2 \cdot -0.127 + 0.628 + 2 \cdot -0.235 + \dots}{9}$$

This simplified example should make it clear how the computation works for Generala and other games in the family. The table has more rows and columns but the principle is exactly the same. Each row is filled using the values in the previous row, requiring only maximums and averages. The math is trivial, and doing it fast is just a matter of implementation quality.

References

- [1] CJF Cremers. “How best to beat high scores in Yahtzee: A caching structure for evaluating large recurrent functions”. MA thesis. the Netherlands: Technische Universiteit Eindhoven, 2002.
- [2] Irving L Finkel. “On the rules for the Royal Game of Ur”. In: *Ancient Board Games in Perspective* (2007), pp. 16–32.
- [3] James Glenn. “An optimal strategy for Yahtzee”. In: *Loyola College in Maryland, Tech. Rep. CS-TR-0002* (2006).
- [4] James R Glenn. “Computer strategies for solitaire Yahtzee”. In: *2007 IEEE Symposium on Computational Intelligence and Games*. IEEE. 2007, pp. 132–139.
- [5] Marcus Larsson and Andreas Sjöberg. “Optimal Yatzy Strategy”. In: *Kungliga Tekniska Högskolan* (2012).
- [6] Tom Verhoeff. *Optimal solitaire Yahtzee strategies*. 2000.
- [7] Clement Wood and Gloria Goddard. *The Complete Book of Games*. Blue Ribbon Books, 1938.