Stay in Command: Optimal Play for Two Person Generala

Joseph Heled

November 5, 2019

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. All four games 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). The player rolls all five dice and then may re-roll some of them. She may then re-roll a subset again, but after the third roll must score the dice by matching it 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. The player may *waive* a combination, i.e. pick a combination which does not match the dice, for a score of zero. For example, a player might waive Generala with 1-1-1-1-5 if Ones are no longer available. Each combination can be scored exactly once. The player with most points wins.

There are various scoring bonuses like scoring an extra 5 or 10 points when hitting a combination on the first roll, but here I shall focus on the two person "Vanilla" Generala for one simple practical reason: size.

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). There are 216 (6^3) ways to score those three combinations, but only 31 possible total sums, since most totals can be obtained in several different ways. A 3, for example, can be obtained by three 1's, one 1 and one 2, or one 3. When it comes to game strategy, the particular composition of the score does not matter: what matters is the total score so far, and which combinations are still available. So, a position can be encoded as a quadruplet: (X-combinations, X-score, Y-combination, Y-score), which gives $4,719,060,648$ ($\approx 5G$) Generala positions. This number can be greatly reduced by the simple observation that the actual scores themselves don't 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 coded as a triplet: (X-combinations, Y-combination, 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. 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.

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)

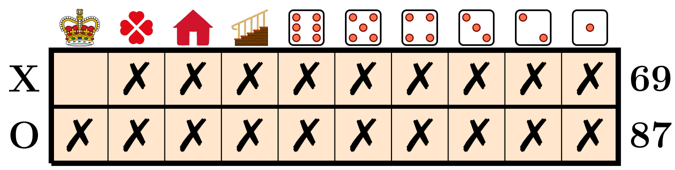

For simplicity, the points of previously scored combinations are not given, only their total sum is shown at the right. Here X must get a Generala on her last turn to win. Since the probability for getting a Generala is $p = {}^{2783176}/_{6^{10}} = 0.046$, the equity is $p \cdot 1 + (1 - p) \cdot -1 = -0.908$. 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)

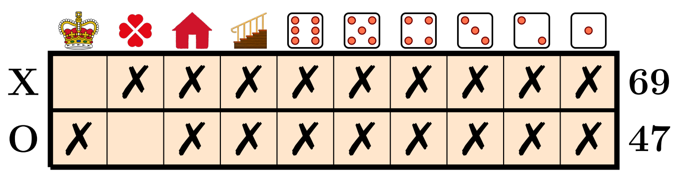

Here X needs either to get a Generala or hope O fails to get Four-of-a-Kind on her last turn. The probability of getting FOAK on your last turn is $q = {}^{17583176}/_{6^{10}} = 0.291$, and so X's equity is $p \cdot 1 + (1 - p)((1 - q) \cdot 1 + q \cdot -1) = 0.445$. After computing the equities of all 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 left out some important details. To compute the equity you need to 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 the combination maximizing the equity. This is done in a similar way, by tracking backwards action by action, starting with the equity after the third roll for every dice combination, then working out the equity after the second roll (again for each dice combination), then computing the equity after the first roll, and finally taking the average of all rolls equities to get the position equity.

It would be extremely tedious and unhelpful to work thorough a complete computation, 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 a three-sided dice – 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, X has Twos'. This equity of this position is -0.056, so we put 0.056 in the table (if O wins x on average, then X loses x 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 columns are 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	–
after 2^{nd} roll	0.056	-0.380	0.628	-0.380	-0.628	-0.726	✗	✗	✗	✗
before 2^{nd} roll										
after 1^{st} roll										

Now to the second row, before the second roll. If X keeps the two dice as-is, rolling none, she will get the equity as given in the row above. If she keeps a 1, she gets 1-1 one third of the time (equity 0.056), 1-2 a third of the time (equity -0.380) and 1-3 third of the time (equity -0.380). The equity is then the average of those, or -0.235. Similarly for keeping a 2 and a 3. When rolling both we need to average the nine possible rolls: 1-1,1-2,2-1,1-3,. . . ,3-2,3-3.

	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}	0.056	-0.380	0.628	-0.380	-0.628	-0.726	-0.235	-0.127	-0.578	-0.313
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 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	1	2	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	−0.235	−0.127	−0.578	−0.313
af. 1^{st}	0.056	−0.127	0.628	−0.235	−0.127	−0.313	✗	✗	✗	✗

Now we can compute the position equity, the average equity of all possible rolls, which is $-0.289 = (0.056 + 2 \cdot -0.127 + 0.628 + 2 \cdot -0.235 + \cdots)/_9$.

This smaller example should make it clear how the computation works for Generala and other similar games. 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 question of implementation quality.

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, we can explore only a very small part of this large elephant. Let me start with the most basic question: Is it an elephant or a mouse?

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 which 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 if 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. 1-1-1-1-5 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 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 score more points. You can find the full details in the code, but all in all 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 loses 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. So stick with Generala when playing for money.

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		✗	✗		✗	✗	✗	✗			**116**
Maximus		✗	✗			✗	✗	✗	✗	✗	**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 (−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 and/or an Escalera.

Sometimes Malina’s play seems to defy reason.

Me		✗		✗		✗			✗		**76**
Maximus		✗	✗		✗	✗					**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 and 1’s, 2’s and 3’s to counter those 12 points plus what Maximus might get from his 1’s to 4’s, and obviously he has six more moves to try 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. 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?

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 gut[1], 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 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 one quarter of the time.

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 move! 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's unlikely to end up with those rolls, 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 and so on, where the probability of hitting an Escalera on the next turn is $^1/_3$, and $^5/_9$ for getting it in two rolls.

Mind you, scoring the two 6's for 12 point is only a small error (less that 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.

[1]Carl Sagan, The Dragon in My Garage

X		✗	✗	✗		✗	✗	✗	✗	✗	**80**
O	✗	✗	✗	✗	✗	✗	✗	✗	✗		**99**

X is 19 points behind and has 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	✗		✗	✗	✗	✗		✗	✗	✗	**80**
O	✗	✗	✗	✗	✗	✗	✗	✗	✗		**88**

X is 8 points behind and needs to score x-x-x-4-4. Scoring 8 with Fours would equalize the score, but even a single 1 from O (O fails to make a one only 5.5% 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	✗		✗	✗	✗	✗	✗		✗	✗	**92**
O		✗		✗	✗	✗	✗	✗	✗	✗	**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 she will be ahead 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		✗	✗	✗		✗	✗	✗	✗	✗	**80**
O	✗	✗	✗	✗	✗	✗	✗	✗			**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						✗	✗	✗	✗	✗	**30**
O				✗		✗	✗	✗	✗	✗	**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.

Meet Maximus

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

	👑	🍀	🏠	🪜	⚅	⚄	⚃	⚂	⚁	⚀	
Maximus		✗	✗			✗	✗	✗	✗	✗	**104**
Me		✗	✗	✗	✗	✗	✗	✗			**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 the 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.

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 0.001/3, since 3 std's gives you a confidence interval of more than 99%.

The standard deviation of running n trials is $\sqrt{\frac{.55 \cdot .45}{n}}$, and so you need to run around 2.2M games. With Generala, however, it is 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.

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 needs to score 1-3-3-3-3.

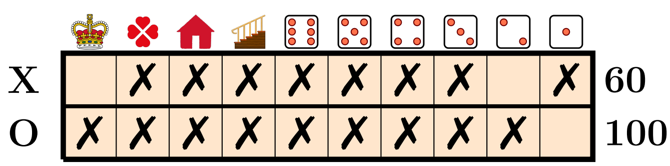

The expected score from Twos is 4.2. The expected score from Generala is 2.3, so Maximus waives Generala. This is obviously wrong here since here making the Generala is the only way to close the 40 points gap. How can we make Maximus aware of the score difference?

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 one of them. Here, after waiving Ones X, scores distribution is 150 with probability of 0.046 and 100 with probability of 0.954. The score distribution for O is this:

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

So X wins with probability 0.046, ties with probability $0.954 \cdot 0.065$, and loses the rest, for an equity of -0.93. After waiving Generala the probability of winning is 0, so the equity is -1. $-0.93 > -1$, so waiving Ones is better.

What we actually did was compute the game equity on the fly **assuming both players are Maximus**. And those equities allow making better decisions than plain Maximus. MM (Mindful Maximus) is only 0.01 (3.5 ELO points) worse than Malina. And the memory usage of MM is 440 kb, two orders of magnitude above Maximus, a light earthquake. This is significant because this approach can be applied to Yatzy, Yacht, or any game in the family, yielding a near optimal player for those games as well. It's not certain that the same holds for the bigger variants, but I think it is quite likely to be the case.

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 person 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 8), using only the distribution of the final scores of each player. Here is a little table showing how much Maximus wins, per game, in a 2,3,4 and 5 players game.

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,

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

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

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

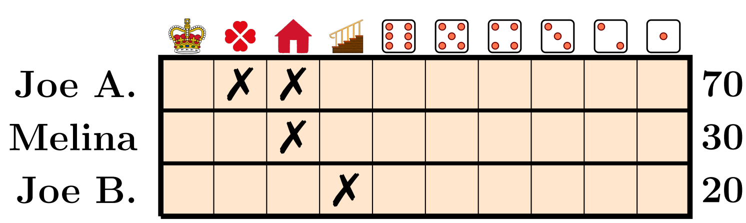

Even though Joe A. currently leads, her equity against him is −0.10 while her equity against Joe B. is −0.12. Malina then plays against Joe B. Informally, making the hard to get Escalera made Joe B. slightly more dangerous than A.

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 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.

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 we can relatively easily 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 hold in Yacht and Yatzy as well, and perhaps even more so, since the difference between Yacht/Yatzy Maximus and 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 several players. All we need to do is combine the final score distribution from each opponent to create one distribution of the highest score, and generate approximate game equity from that. Again, it will not be necessarily accurate, but it does not need to be. It just 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. It is the same principle that lets Backgammon software play so well. To select a move you only need to approximate the equities well enough so 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., his equity comes out as (with a 95% probability) 0.0355 ± 0.005. It is impossible to tell from just 300k games which of them is better, but the difference is very small, and that leads me to 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.

Executive summary

I. Generala, under the right encoding, is small enough to be completely solved.

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

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

IX. The code is available in this GitHub repository.

Generala like games are social in nature and typically played more casually than, say, Chess of 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, if there was such a thing, is much smaller than in games like Chess and Backgammon, but is not insignificant. A professional will still win a 25 point match against an amateur 90% of the time, assuming the amateur plays as well as 'Joe'. 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 made in the course of one game of Generala, and are probably unaware how good or bad their moves 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.