

Quantum Beauty Contest



Philipp Pestlin
D-CHAB, ETH Zürich

PPESTLIN@ETHZ.CH

Tommaso Antonelli
D-PHYS, ETH Zürich

TANTONELLI@ETHZ.CH

Yoel Zimmermann
D-CHAB, ETH Zürich

YZIMMERMANN@ETHZ.CH

Abstract

We introduce a quantum version of the Keynesian Beauty Contest game, based on the superposition of the players' guesses. We further analyze the strategy space of this game by means of an evolutionary algorithm in analogy to level k thinking. Through minor modifications to the game dynamics, we assess various approaches and observe diverse equilibrium outcomes. We identify a particular version of the game that shows a preference for pure states above the state $|0\rangle$ corresponding to the classical Nash equilibrium, dependent on specific parameters. In a projective approach we find a phase transition from an initially unordered to an ordered phase. The time of the phase transition is found to be strongly dependent on the random seed. This approach favors almost pure $|0\rangle$ states at equilibrium. We also show chaotic games where our algorithm can not find dominant strategies as well as an approach that converges but also shows slight deviation from the classical Nash equilibrium as the expected value.

1. Introduction

Historically, the Keynesian Beauty Contest is a game that describes the interaction of rational agents in a context where they are asked to choose the most attractive person among a group of people. The players who picked the most popular face are then rewarded with a prize.

The catch of this game is that the player should not really pick the most attractive person according to their personal opinion, but rather be aware of the opinions of the other agents in order to formulate a decision. This line of thought can be carried recursively, since the player can then formulate their decision based on their expected opinion of the other players' public perception of beauty, and so on.

In the words of John Maynard Keynes himself:

“...it is not a case of choosing those [faces] that, to the best of one's judgment, are really the prettiest, nor even those that aver-

age opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth and higher degrees.” (Keynes, 1936)

Also, as Keynes remarks in this work, the Beauty Contest has strong parallels with the behavior of professional investors in financial markets, in the sense that investors try to be one step ahead of the average investor, who is seen as an ignorant agent likely to be present in the market (Duffy and Nagel, 1997).

All the features that characterize this historical version of the Beauty Contest can be elegantly put in a mathematical framework that formalizes the setup. In the following section, we will first discuss this classical version of the game, which is a standard introductory example in Game Theory courses. We will then devote our attention to the formulation, computational implementation, and analysis of a quantum version of the game.

2. Theory

2.1. Classical Beauty Contest

The classical Beauty Contest game¹ can be mathematically formulated as a set of N players, each of them choosing an integer number in some interval, often from 0 to 100. The mean value \bar{n} of these N numbers is then calculated, and given a certain positive constant p (generally known by all players before the game), the winning number of the contest is given by $w = p\bar{n}$ (rounded to the closest integer in the range). The value p , commonly taken to be $1/2$ or $2/3$, determines a shift of the winning number from the mean guess and therefore represents the motivation for a certain level of recursive reasoning in the player’s choice, analogous to the original Beauty Contest formulation.

As an example of a contractive Beauty Contest with a factor $p = 2/3$, the first histogram in Fig. 1 shows that a common strategy is to choose numbers

close to the value $50 \cdot p^k$ for some positive integer k (i.e., 33, 22, ...). This indicates the different levels of recursion in the decision process, where one starts with the assumption that everybody will play a random number, producing most likely $\bar{n} = 50$, and hence playing a value $50 \cdot p \simeq 33$. The next step is to think that everybody will most likely play 33, hence the guess is $33 \cdot p = 22$. This line of reasoning is called the k -level thinking model, where a player tries to outsmart the other players by playing one level ahead of them. This model predicts that agents will learn about their respective environment in successively played games and seems to fairly accurately predict the popular guesses, at least for small values of k .

Another popular guess in this iteration of the Beauty Contest seems to be 0, which is basically the limit of this k -level model when $k \rightarrow \infty$. In fact, the scenario where every player guesses 0 is the unique Nash equilibrium for the Beauty Contest game (given $0 < p < 1$) (Kennerberg and Volkov, 2019). However, in this case, the assumption of perfect rationality of the other players is a strong one, and it is very naive to imagine that all people will take this k -level model to the extreme consequences, thus practically invalidating the strategy of playing 0.

As expected, in this first iteration of the Beauty Contest with $p = 2/3$, the winning number is not $w = 0$, but rather $w = 20$. If the same game is repeated with the same group of N people, the consequence is visible: people seem to adapt to the game environment and play on average much lower numbers, as seen in the last two histograms in Fig. 1.

It’s now time to go quantum!

2.2. Quantum Beauty Contest

The game we are considering again involves N players, each of whom will formulate a guess. In this scenario, however, the guess is formulated in terms of quantum states. For this, we construct a $(M + 1)$ -dimensional Hilbert space \mathcal{H}_M , where, for concreteness, we will choose $M = 100$. This space is spanned by the following base kets, in standard Dirac notation:

$$\left\{ |0\rangle, |1\rangle, |2\rangle, \dots |100\rangle \right\} \quad (1)$$

1. Also widely referred to as Guess 2/3 of the average (Ledoux, 1981), p-contest (Kennerberg and Volkov, 2019), and p-guessing game (Vie, 2021)

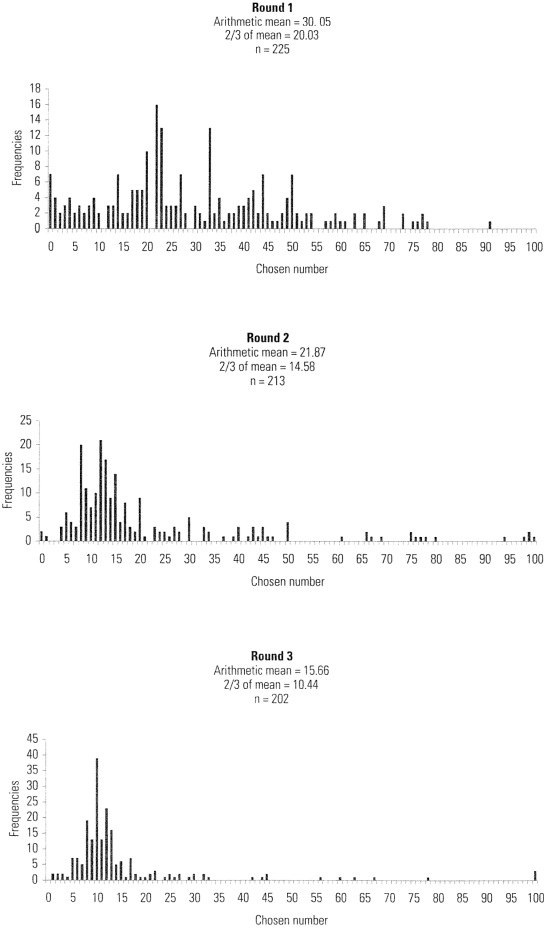


Figure 1: From top to bottom: the first three rounds of the classical Beauty Contest results, with a value of $p = 2/3$ (Diekmann, 2009).

which are the eigenstates of a number operator \hat{N} with the explicit form

$$\hat{N} = \sum_{n=0}^{100} n |n\rangle \langle n| \quad (2)$$

from which we can immediately see the defining property

$$\hat{N} |n\rangle = n |n\rangle. \quad (3)$$

Furthermore, the basis is generally chosen to be normalized:

$$\langle n|m\rangle = \delta_{nm} \quad (4)$$

where δ_{nm} is the Kronecker delta.

In a certain sense, the base kets represent the guesses a player can make in the classical game. Now,

in this variation of the game, each player is asked to give a guess in the form of a normalized ket:

$$|\psi_i\rangle \in \mathcal{H}_M \quad \text{for } i \in \{1, 2, \dots, N\}, \quad (5)$$

i.e.,

$$|\psi_i\rangle = \sum_{n=0}^{100} c_n^{(i)} |n\rangle \quad \text{and} \quad \sum_{n=0}^{100} |c_n^{(i)}|^2 = 1 \quad \forall i. \quad (6)$$

Then the following ket state is constructed:

$$|\psi\rangle = \mathcal{N} \sum_{i=1}^N |\psi_i\rangle \quad (7)$$

where \mathcal{N} is some normalization constant.

The next step is to collapse our state $|\psi\rangle$ into the basis given in Eq. 1, that is to say, we perform a measurement of \hat{N} on $|\psi\rangle$.

We will therefore obtain the state ket

$$|n\rangle, \quad \text{with } n \in \{0, 1, \dots, 100\} \quad (8)$$

with a probability

$$P_n = |\langle n|\psi\rangle|^2. \quad (9)$$

The winning number is then determined as usual:

$$w = \text{round}(p \cdot n) \quad (10)$$

for some positive p . With this number, we can single out the basis ket $|w\rangle$ and determine the winner(s) of the Quantum Beauty Contest by assigning a payoff to each player using the following formula:

$$\phi_i = |\langle w|\psi_i\rangle|^2 \quad (11)$$

which is basically a simple measure for determining how much of the winning outcome was contained in the player's initial guess.

From a theoretical point of view, the Quantum Beauty Contest we propose does not reduce to its classical analogue for any choice of the players' strategies. This is due to two independent features of the quantum version of the game.

Firstly, the winning number w is produced through a measurement procedure of the state $|\psi\rangle$, hence there is a component of randomness in this game

that is not present in the classical version.

Secondly, even if one repeated the measurement of the same state $|\psi\rangle$ over and over again, thus effectively extracting the mean value of the operator \hat{N} (which will be in fact done in some later computational implementations of the game), in general, there is no strategy of the players that can simulate the classical analogue of the game. This is because the way we decided to construct the state $|\psi\rangle$ is through the sum with equal weights of the guesses $|\psi_i\rangle$. In this game, the only strategy that resembles a classical game is for every player to play a ket in the basis of Eq. 1:

$$|\psi_i\rangle = |n\rangle, \quad \text{for some } n \in \{0, 1, \dots, 100\} \quad (12)$$

In this way, though, if two players decide to play the same number, i.e., $|\psi_i\rangle = |\psi_j\rangle = |n\rangle$ with $i \neq j$, then the probability amplitudes in the state $|\psi\rangle$ for the basis state $|n\rangle$ add up, while the actual probability of extracting the number n goes as the amplitude squared. This is in contrast with the classical game, where every instance of a guess n simply adds up, and the mean is calculated afterwards without squaring the columns of the histogram.

Anyway, apart from these differences, the spirit of the game remains the same, but in this case, the strategy space is much richer than the classical counterpart.

3. Computation

The strategies of the Quantum Beauty Contest reside in a high-dimensional Hilbert space. To efficiently explore and optimize these strategies, an evolutionary algorithm is employed. This algorithm leverages reproduction and mutation techniques, enabling effective sampling of the strategy space to discover optimal strategies. A similar algorithm has been recently utilized for computational analysis of the prisoner's dilemma (Vie, 2021).

3.1. Reproduction

Let $|S_1\rangle$ and $|S_2\rangle$ represent the quantum states of two parent strategies. The reproduction operator

$$R : \mathcal{H} \times \mathcal{H} \rightarrow \mathcal{H}$$

combines these states to generate a new strategy for the offspring:

$$|S_{\text{offspring}}\rangle = R(|S_1\rangle, |S_2\rangle) = \mathcal{N}(|S_1\rangle + |S_2\rangle) \quad (13)$$

This process ensures that the offspring inherits characteristics from both parents while exploring new regions of the strategy space.

3.2. Mutation

The mutation operator introduces variation by perturbing individual strategies after reproduction.

Let's denote the mutation operator as $M(|S\rangle, \pi, \sigma)$, where $|S\rangle$ is some strategy, π is the mutation rate, and σ is the mutation strength. The mutation operator can be defined as follows:

$$M(|S\rangle, \pi, \sigma) = |S'\rangle \quad (14)$$

$$|S'\rangle_i = \begin{cases} |S\rangle_i + \delta & \text{with probability } \pi \\ |S\rangle_i & \text{with probability } 1 - \pi \end{cases} \quad (15)$$

where $|S'\rangle$ represents the mutated strategy, and δ is a random perturbation drawn from a uniform distribution $\mathcal{U}([- \sigma, \sigma])$. If a mutation event occurs (with probability π), the corresponding component of the strategy is perturbed by adding a random value from δ . Therefore, a higher mutation rate leads to more frequent mutations and greater exploration of the strategy space, while a lower mutation rate promotes exploitation of existing elite strategies.

On the other hand, a larger mutation strength results in larger deviations from the original strategy, allowing for more significant exploration of the strategy space. Conversely, a smaller mutation strength limits the extent of perturbations, promoting finer adjustments to the strategies.

3.3. Algorithm

Here is a quick overview of the algorithm employed:

1. **Initialize:** For every player i , draw random strategy ket $|S_i\rangle \sim \mathcal{U}_{101}([-1, 1]^{101})$ with $\langle S_i | S_i \rangle = 1$
2. **Simulate Game:**
 - a. Draw $|n\rangle$ from $|\langle \psi | \psi \rangle|^2$ as defined in Eq. 7 and determine winner ket $|w\rangle$

- b. Determine fitness of $|S_i\rangle$ with respect to $|w\rangle$ for every strategy

3. Reproduction Cycle

- a. Extract fixed amount of elite strategies into new game
 - b. Select parent strategies $|S_i\rangle, |S_j\rangle$ $i \neq j$ and form new strategy $M(R(|S_i\rangle, |S_j\rangle), \pi, \sigma)$
 - c. Repeat b. until every player has a strategy
4. **Iterate:** Jump to 2. until convergence or fixed number of iterations reached

For the specific implementation of the game we can make different choices for how we determine $|w\rangle$ as well as the fitness. The implementations can be found in the class functions `fitness(self)` and `measurement(self)`, respectively. The full source code can be found in Appendix A.

4. Results and Discussion

For the results we are going to look at the different approaches for a quantum beauty contest, considering variation of parameters, and investigate special behavior while comparing this to the classical results from lab studies. When talking about the expected value of $|\psi\rangle$ or $|S\rangle$ in the following we will only refer to the values where no contraction factor p has been applied yet, if not stated otherwise.

4.1. The Semi Classical Approach

First we look at the approach where the winning state $|w\rangle$ is determined from the expected value of $|\psi\rangle$ and the fitness of $|S_i\rangle$ is determined by its distance of its expected value to $|w\rangle$. This approach is almost identical to the classical game. The main difference are the real values of the expectation from each player. For any simulation these are still rational numbers due to the machines precision of values. An analogue for the classical game could therefore be the formation of alliances of players which then give their average as a guess. In the edge case of an infinite amount of players in an alliance we could still reproduce real numbers as a guess. As this approach is almost isomorphic to the classical game, we expect a similar approach to $|0\rangle$ for $|w\rangle$. While we do observe this as seen in figure 2, we do not quite reach a state of 0 after

many iterations, which is due to our evolution algorithm allowing for relatively strong mutations, even for high states. An analogue to this in the classical game might be a disruptive player who purposefully chooses numbers that can not win ($n \cdot p$ or above) after several rounds, which can also be seen in 1.

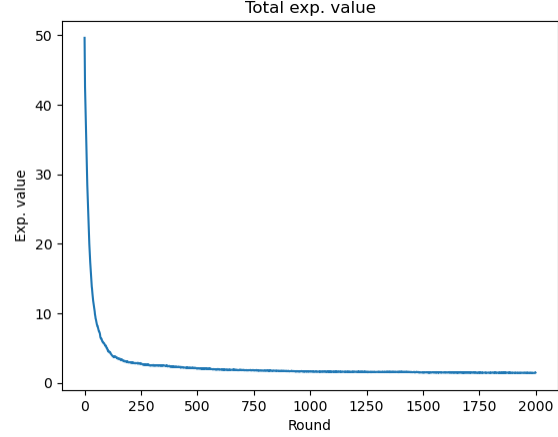


Figure 2: Expected value of $|\psi\rangle$ over 2000 rounds for the semi classical approach.

For this simulation we used 100 players, 20 elite players per round, $\pi = 0.5$, $\sigma = 0.05$ and $p = 0.5$. Using smaller mutation rates such as $\pi = 0.1$ allows for even stronger convergence to 0 and strategies which survive the longest show an almost pure $|0\rangle$ state.

4.2. Collapse with Players' Average

For the next approach we keep the average to determine the players performance, for the player it is therefore still similar to simply providing real (rational) numbers. To determine $|w\rangle$ on the other hand, we use a wave function collapse, i.e. physical measurement of \hat{N} . We therefore have a non zero probability for any state $|w\rangle$ which we can see from small oscillations in the expected value over the course of simulations. If $|\langle w|\hat{N}|w\rangle|^2 \gg |\langle 0|\hat{N}|0\rangle|^2$, only states with expected values larger than 0 make up the elite group of players and therefore increase the expected value far above the noise level for the next rounds as seen in figure 3. Eventually we can still observe an overall convergence towards $|0\rangle$ which is especially clear from looking at the state which was in the elite for the longest time over a simulation of 2000 rounds.

Besides the states $|0\rangle$ and $|1\rangle$, no states have a significant weight which goes beyond the induced mutation. We show these results in figure 4.




Figure 3: Expected value of $|\psi\rangle$ for the collapse - average game.

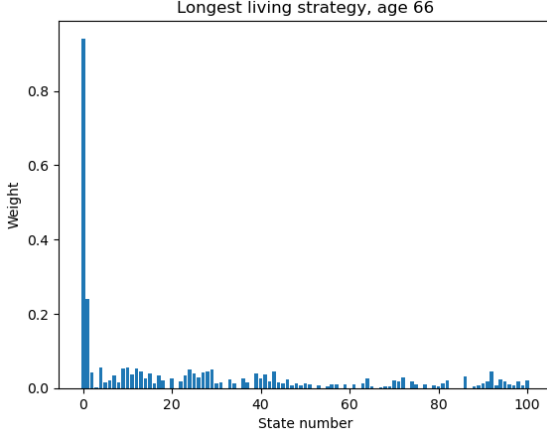


Figure 4: Strategy with most rounds without an update (most rounds inside the elite).

For the results shown we again used 100 players, 20 elite players per round, $\pi = 0.5$, $\sigma = 0.05$ and $p = 0.5$. While a smaller π leads to a decrease in the density of spikes seen in figure 3, a change to $\sigma = 0.1$ gives a non convergent game. We can therefore not find any value as a fix point. The noise is too strong for the convergence rate of the game. Interestingly this value seems to be somewhere in the region of a phase shift as for $\sigma = 0.09$ we can already see much better convergence with a lower bound of ≈ 9 .

4.3. Weight of the Average: A Nonzero Equilibrium?

With this approach we go back to using the average to determine $|w\rangle$ but we now determine the fitness from the weight each player assigned to the respective state $|w\rangle$.

This approach is particularly interesting as the winning players do not necessarily favor states with low averages but are only focused on the single state. This will eventually lead to a slower decay of the expected value. This is also enhanced by the mutation algorithm where the states not equal to $|w\rangle$ are uniformly distributed and only decrease in weight slowly, from the higher weight necessary on $|w\rangle$. Once the decay flattens we do not observe the expected value to reach a value around 0. Instead the expected value stays constant between 8 and 10, therefore giving $|w\rangle$ as $|4\rangle$ or $|5\rangle$. This only happens after some time by chance and shorter simulations often indicate a saturation between 10 and 12.

Additionally, we are able to observe large spikes of the expected value when inside a flat area. This may be caused by large mutations on higher states for multiple states in the elite, leading to a continuous boost of such features over the following rounds. An example for the course of the expected value is given in figure 5.




Figure 5: Expected value of $|\psi\rangle$ for the weight of the average approach.

The reason for this behavior lies in the nonzero probability of higher states. Having a weight of 0.01 on the state $|100\rangle$ increases the players expected

value by 1. As we have positive mutations on 25% of the states as well as the remaining probabilities from the parents on higher states, an average of 0 is, with the parameters as previously used in 4.1 and 4.2, simply not possible.

For the exact point of convergence we were yet able to construct an exact formula but we do observe a linear scaling in σ . While the lower bound for $\sigma = 0.05$ was found to be 5, the lower bound for $\sigma = 0.01$ is determined to be 1. For such low sigma games, we also observe the establishment of dominant players which is again an effect of the mutation algorithm, although the mutations are already quite weak. Once a player finds a strategy with an almost pure 1 state, any mutation on other states leads to the a decrease in the weight of the 1 state and therefore looses against the dominant players. On the other hand, due to the rounding of $|w\rangle$, $|0\rangle$ is always a bad play if all or most players have some other non $|0\rangle$ state. We show the rounds won by each player after the saturation in figure 6, indicating the dominance of a few players.

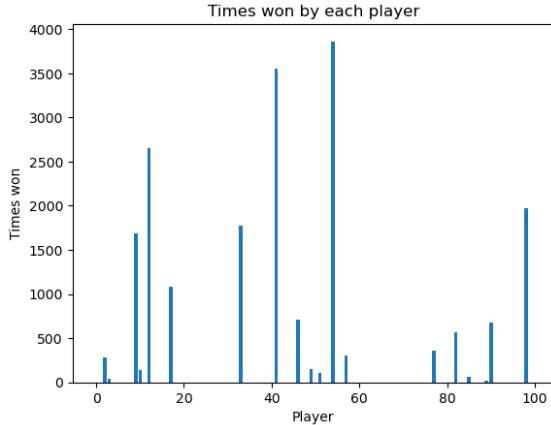


Figure 6: Times won by each player with $\sigma = 0.001$ and after boundary is reached.

4.4. The Projective Game

In this approach we are using the projection of each player onto a state which is again determined by the collapse of $|\psi\rangle$ as already described in section 4.2. In other words, the more likely a player is to receive the respective state upon measurement, the better its fitness. This game is quite complex for the player. For single rounds or rounds before a clear convergence a player needs to have a good balance

in their states. Distributing all probabilities evenly might not be enough to reach the elite. Putting all the weight on one card on the other hand kicks the player out of the elite for any other state. The second extreme is very much the same as betting all chips on one number at a roulette table but with a less significant reward if the state is actually found.

With the simulation we do find an area of uncertainty at the start of each game. With the two previous approaches we were more or less able to observe some kind of exponential decay in the expected value. With this approach on the other hand such a decay is only observed after many rounds were already played. Before that the game shows intervals of slower decay and even sections where the expected value may increase again. This behavior makes the decay already very much reliant on the seed for the random numbers which we were able to observe from decays starting after a wide range of rounds played. For our parameters this was generally still below 1000 rounds. Eventually the expected value still converges towards $|0\rangle$ where we again observe a similar behavior to the approach described in 4.2 with occasional spikes from non $|0\rangle$ measurements. By looking at the states $|w\rangle$ we can clearly see a phase transition from an unordered phase with many different winners to an ordered phase with only solitary non $|0\rangle$ measurements. The strategy surviving the most rounds even shows an almost pure $|0\rangle$ state with no weights above 0.1. Examples are shown in figures 7 - 9.

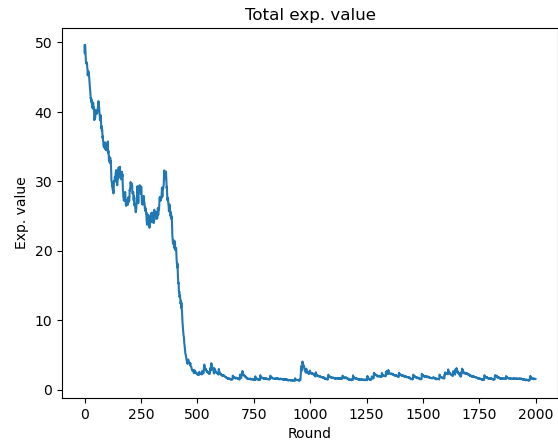


Figure 7: expected value of $|\psi\rangle$ for the projection approach.




Figure 8: Longest living state for the projection approach.




Figure 9: Values of $|w\rangle$ over 2000 rounds for the projection approach.

For these results we used `np.random.seed(42)` and `np.random.seed(42)` with the first run of the simulation after setting the seed. The parameters are kept at the same values as before with 100 players, 20 elite players per round, $\pi = 0.5$, $\sigma = 0.05$ and $p = 0.5$. We again also try to bring up σ . Around $\sigma = 0.1$ we do not observe total chaos, as the longest living state is still dominated by $|0\rangle$ and the expected value is not distributed across all possible values. Still for $\sigma = 0.102$ we do see some form of oscillations over the course of 30.000 rounds where we observe intervals with expected value around 15 as well as 40. At this point we do not have an explanation for this phenomena, especially why there still seems to

be some order with two fix points and the system does not stabilize around 30. What we also observe is a far more uniform distribution across all possible states in $|w\rangle$ when the expected value is around 40 and far less measurements for states other than $|0\rangle$ when around 10. This gives the indication that the game switches between relatively ordered and almost random states and large enough deviations from the mutation change the phase of the game. An example of this effect is shown in figure 10.




Figure 10: expected value of $|\psi\rangle$ for a simulation with on the projection game $\sigma = 0.102$.

4.5. Player Strategy Collapse

Until now we have not looked at approaches where we collapsed each players strategy in order to determine the fitness score. The reason for this is simple: Using our previous approach of keeping the elite wave functions after a measurement not only seems to counter the idea of collapsing a wave function, it also leads to chaotic behavior where the expected value does not clearly converge to any point. This is regardless of determining $|w\rangle$ from the total expected value or also from a collapse of $|\psi\rangle$. Starting with almost uniform states the randomness which governs onto which state each player collapses is so large that good features (population of lower states) do not get the chance to become relevant. An example is shown if figure 11.

Changing the approach to setting the players strategy to its collapsed state before the next round as well as mutations does also not lead to any

convergence in the expected value. Another possible change is to change the order in the sense that we first collapse each players and then draw from the average. But even with this approach we do not see any convergence in the data again also trying out slight variations in the parameters.

This brings us to the conclusion that the confusion for the players is simply too large and for any described approach they are missing the incentive to develop a 'good' strategy in the sense of convergence.

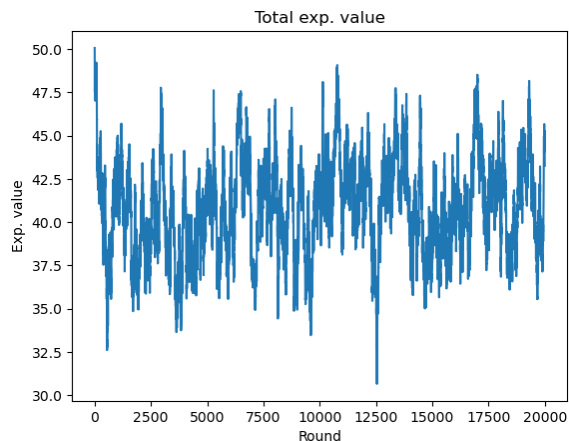


Figure 11: Example of a non convergent game based on wave function collapse of both $|\psi\rangle$ and all $|S\rangle$.

5. Conclusion and outlook

With the Quantum Beauty Contest we were able to apply a formalism known from the theory of quantum mechanics onto the classical game of the Keynesian beauty contest. In particular, we looked at the case of a contractive factor $p = 1/2$ but other cases could also be easily investigated in the future.

With different approaches on how to determine the fitness of the strategy of a specific player we were able to observe and describe several phenomena in an evolutionary setting which are not predicted by the classical game. With phenomena such as a phase transition on the projective game approach we are able to draw connections to human behavior, in this particular example the transition from confusion to clarity on the interactions within the game. Still it is uncertain if these simulations could ever be closely reproduced by humans, considering the relatively

simple mutation algorithm we used throughout the simulations. Here one could perform further investigations by limiting the amount of vector components one strategy can contain, changing the adaptation from parent strategies or limiting the precision. These might be able to closer represent a human game where people would probably only use simple fractions and not choose 101 states for many rounds.

By using the principle of wave function collapse we were also able to bring some instability to the game in its equilibrium position forcing older strategies to adapt to different conditions. While strategies still show dominant $|0\rangle$ states, using the player average does not allow them to fully neglect other states. Only by using the projection operation a player is able to largely focus on this state when at the Nash equilibrium.

With the approach of taking the weight of the average we showed another interesting result as the game did converge but always to a state higher than $|0\rangle$. With humans it might therefore be interesting if we can reproduce such a behavior and especially to which points humans would converge. Here using smaller values for the contraction factor such as $p = 1/3$ might also be interesting to investigate in the future as this currently sets a hard boundary at $|1\rangle$ and makes a measurement of $|0\rangle$ impossible.

One last thing which was not considered yet by us is a variation in initial conditions. While we chose the initial states at random, leaving us with a relative uniform distribution in states, the classical game shows that most players tend to avoid numbers that are impossible to be determined as the average already on the first move. Manipulating the initial states might therefore also be a point for investigation in the future.

References

- Andreas Diekmann. *Diekmann, Andreas, 2009. Rational Choice, Evolution and the “Beauty Contest”. In Mohamed Cherkaoui and Peter Hamilton, Eds., Raymond Boudon – A Life in Sociology. Oxford: Bardwell Press. 01 2009.*
- John Duffy and Rosemarie Nagel. On the robustness of behaviour in experimental ‘beauty contest’ games. *The Economic Journal*, 107(445):1684–1700, 1997.
- Philip Kennerberg and Stanislav Volkov. Convergence in the p -contest, 2019.
- John Maynard Keynes. *The general theory of employment, interest and money*. 1936.
- Alain Ledoux. Concours résultats complets. les victimes se sont plu à jouer le 14 d’atout. *Jeux & Stratégie*, 2(10):10–11, 1981.
- Aymeric Vie. Evolutionary strategies with analogy partitions in p -guessing games, 2021.

Appendix A. Supplementary Material

The source code for the simulation can be found on <https://github.com/yzimmermann/QBC>.