

Beauty Contests

CSCI 1440/2440

2025-09-03

1 Historical Background: A Newspaper Contest

Imagine the following scenario: A popular newspaper holds a beauty contest. The newspaper presents to you and the rest of its readers the same set of a hundred photos of people's faces. You are asked to vote for the six faces that you find most attractive. After every reader selects their six faces, the newspaper company tallies up the votes, and the face that was selected by the most participants is deemed the most attractive. The newspaper then gives a prize to any reader who voted for the most attractive face.

Because this scenario involves at least two acting agents whose choices affect each other's outcomes, it is technically a game. How should you play this game, if you want to win the prize?

The naive strategy would be to follow the instructions and simply choose the six faces that are the most appealing to you. But if you expect that most other people will also abide by this strategy, then you should ignore the instructions and instead choose the six faces that you expect will be most appealing to *the other readers*. But, then, if you expect that most people will choose the faces that they expect that most other people will find attractive, then you should go one step further and choose the faces that you expect that most people will expect that most people will find most attractive. This reasoning can recur infinitely.¹

This game was envisioned by John Maynard Keynes in 1936, as analogous to explain price fluctuations in the stock market, as stock prices are based not just on their intrinsic values, but also on how valuable people expect them to become, and thus also on how valuable people expect other people expect them to become, and so on.

¹ As you can see, this newspaper contest is not a good *mechanism* for eliciting people's true attractions. Learning to design mechanisms that incentivize truthful preference reports is a central focus of this course.

2 Formalization: p -Beauty Contest

The **p -beauty contest** was designed to capture the strategic features of Keynes's beauty contest, but is simpler to execute and analyze. In this game, there is a set N of numbers, and a real number $p \in (0, 1]$, which is common knowledge.² Each of the players chooses a number in N , simultaneously. Then, the mean m of their choices is computed. The winner is the player whose number is closest to pm .

Remark 2.1. In the version of this game that we played in class, $N = \{0, 1, \dots, 100\}$, and $p = 2/3$. However, in the standard version of this


² Something is **common knowledge** if every player knows it, and every player knows that every player knows it, and every player knows that every player knows that every player knows it, and so on.

game, N is the interval of *real numbers* in $[0, 100]$. This is convenient because the real-number version has a unique Nash equilibrium in which every player chooses 0, whereas the integer version has additional, awkward equilibria, which are frankly uninteresting in a strategic sense. (For example, there is an equilibrium in which every player plays 1; $2/3$ of the mean is $2/3$ in this case, to which 1 is the closest integer. Depending on the number of players, there can also be equilibria in which players mix between playing 0 and 1.)

3 Empirical Results

This game was studied empirically by Rosemarie Nagel,³ among others. In Nagel's study, between 15 and 18 participants played the same discretized version of the p -beauty contest we played in class. The participants' choices were mostly 33 and below, with clusters around 33 and 22. A histogram of the results is shown below.

³ Rosemarie Nagel. Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5):1313–26, 1995



What is interesting about these results is that people don't even come close to playing an equilibrium. Some people employ other-order reasoning: they simply play the average of what they expect others to play, namely 50. Others are a bit more sophisticated, and employ 1st-order reasoning: they play p times the average of what they expect others to play, namely 33. Still others employ 2nd-order reasoning: they play p times p times the average of what they expect others to play, namely 22. But notably, not a *single person* chose 0! One takeaway message from this story is that Nash equilibrium is not always a very accurate predictor of how people will play a game.

References

- [1] Rosemarie Nagel. Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5):1313–26, 1995.