Quantifying Sorting Ability of Disc Golf Holes

By Steve West Disc Golf, LLC

April 17, 2020

A hole is said to be "lucky" if it assigns scores to players in a way that does not reflect the underlying skill of those players. One aspect of luck is a tendency to add random punishment, resulting in higher scores for some players. Another aspect is how well a disc golf hole allocates those scores to the correct players, called sorting ability. This paper will try to quantify sorting ability.

One way to measure sorting ability is to compare what actually happens to what would happen with known quantities of sorting ability. This paper presents a way to generate known quantities of sorting ability, then compares actual results to results from various levels of sorting ability. The inherent amount of sorting ability in the actual results can be inferred by finding the known quantities of sorting ability which have the same effects.

Synthetic Sorting ability

Synthetic sorting ability is a means of dialing up or down the amount of sorting ability involved. Perfect (100%) sorting ability means that the lowest score on a hole always goes to the best player, the worst score always goes to the worst player, and all the scores in between are allocated in perfect order. Zero sorting ability means that every score has an equal chance of being given to any player

Zero sorting ability is simulated by sorting all the scores of the hole by random values. Every score is assigned a random value between 0 and 1. The scores are then sorted by this random value and applied to the list of simulated players (which are magically already listed in descending order of skill).

Perfect sorting ability is simulated by sorting all the scores of the hole by the value of the score and applying them to the list of simulated players. This is the same as assigning the lowest score a value of zero, the second lowest score a value of 1/(n+1), etc. to assigning the highest score a value of 1. Then the scores are sorted according to those values. In other words, the hole "knows" the skill of each player and perfectly assigns each score according to skill.

The degree of sorting ability can be thought of as one minus the size of the range of random numbers which can be assigned to a score prior to sorting. For zero sorting ability the range has a width of 1; each score can be assigned anything from 0 to 1. For perfect sorting ability, the range has zero width; the "range" is the unique value which represents each score's ranking.

For fifty percent sorting ability, the random number assigned to the lowest score could be anywhere from 0 to .5, the random number assigned to the highest score could be anywhere from .5 to 1, and the random number assigned to the median score could be anywhere from .25 to .75.

For any given level of sorting ability, after the scores are allocated to players according to the random numbers, total scores for each player are added, and the resulting total scores are a single simulated round.

Scoring Spread of Various Levels of Sorting Ability

We will focus on one consequence of various levels of sorting ability: the scoring spread width of the total scores. Scoring spread width is related to the entropy of the distribution of total scores, so it is a measure of the amount of information in the total scores.

Following is a chart of scoring spread width as a function of sorting ability for 1000 simulated rounds on one course.

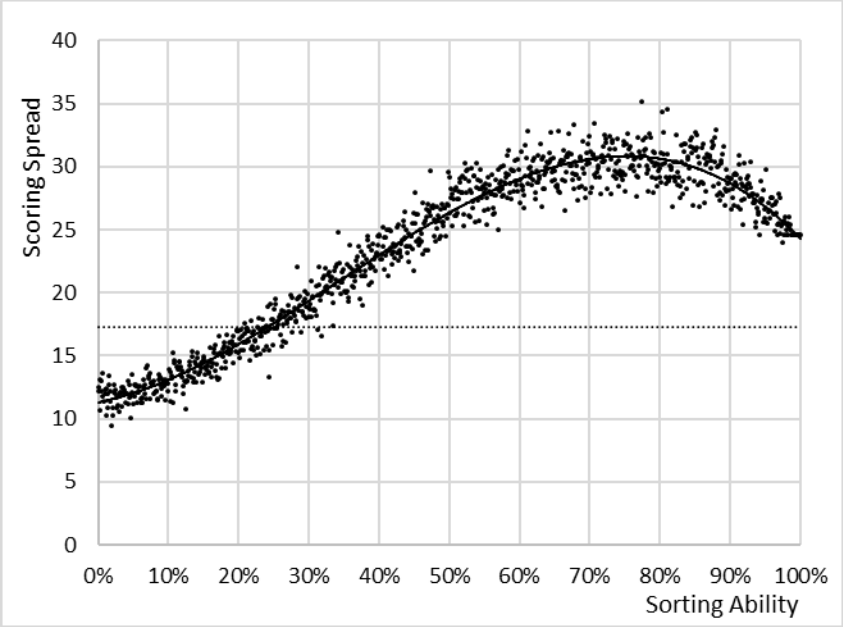

When sorting ability is lowest, the scoring spread is smallest. This is because the amount of information generated by randomly sorting the scores is at a minimum.

When sorting ability is 100% the scoring spread is wide, but not as wide as possible. This is because there aren't enough different scores to break all the ties.

Back-Solving for Sorting Ability given a Scoring Spread

The area we are most interested in is the realistic range of sorting ability; from zero to about 50%. Within that range, more sorting ability results in a wider scoring spread. If we know the scoring spread, we can infer the sorting ability it would have taken to produce that scoring spread.

For the example in the graph, we know the actual scoring spread from this round was 17.28; shown as the horizontal dotted line. Thus, we can be confident the underlying sorting ability was somewhere from 20% to 34% or so. By fitting a cubic to the shape of the function, we can pick a single value for sorting ability where the fitted line is equal to 17.28. In this case, we would say the round exhibited a sorting ability of 24%.

How Courses Generally Perform

The 24% sorting ability in the example is fairly typical. For the 303 rounds (with at least 15 players and including highly rated MPO players) for which I calculated a sorting ability, the range was from 0% to 52% with an average of 22%.

Another way to look at this is that for typical disc golf courses, total scores are the result of a mostly (76%) random assignment of hole scores to players. This should not be a surprise. Each hole gives out only a few scores. These few scores must be allocated across dozens of players. Most of the time, the score given out will be the same for many players.

Let's compare the example of two players who are very different in skill. The best player in a tournament might be expected to average 2.5 on each hole, while the worst player might average 3.5. If the holes all perfectly gave out a 2 or a 3 with 50/50 odds to the best player, a 3 or a 4 with 50/50 odds to the worst player, then about one-fourth of the time, these two players would get the same score on a hole.

If we also factor in the real-world messiness of actual scoring distributions, plus the strong tendency to give out just a single score when the expected average is near an integer, plus the smaller differences between expected scores among all the other players, we can start to see why total round scores are mostly a result of chance.

A Measure of Course Characteristics?

If Scoring Ability calculated in this way is truly a measure of some underlying aspect of a course, it should remain stable from round to round. I had multiple-rounds of data for 86 courses, most with just two rounds, but some with more; one course had 5. The graph below shows the result for one round compared to the result for the next round on that same course.

While the round-to-round results on a course are not totally independent, there is not enough correlation to pronounce this statistic to be a measurement of an underlying characteristic of the course.

The non-independence is certainly partially attributable to these courses hosting the same set of players from round to round and so only partially a result of stability (if any) of the course's ability to sort. Note also that the round-to-round correlation of sorting ability (23%) is lower than the round-to-round correlation of scoring spread width of total scores (61%). This means there is something about the calculation of sorting ability which is scrambling things up.

This potential metric does not measure all aspects of what a good course should do. For example, it does not measure the ability of a course to give out a wide range of scores. The only thing it might measure is the ability of the course to assign hole scores in a way that increases the information content of the total scores.

Therefore, I conclude this is not a good candidate for a way to pass judgement on courses.