Don’t try for the triple 20

Where to aim if you are bad at darts

Significance not only increases your understanding of statistics, it also increases your score at darts. Where you should aim depends on just how bad a player you are. **Ryan Tibshirani** shares the story of how his own lack of skill called for a better darts strategy. Now you can learn from his mistakes, and rack up more points the next time you play!

Don’t aim at the triple 20 – that is, unless you are a professional. This tip might be known to a committed darts player or a “student” of the game, but a true amateur would never know it. I fall into the latter category, and started playing darts in my second year of graduate school because my roommate, Andy Price, was a darts aficionado. Andy was also a graduate student, and we often threw darts at night as a study break.

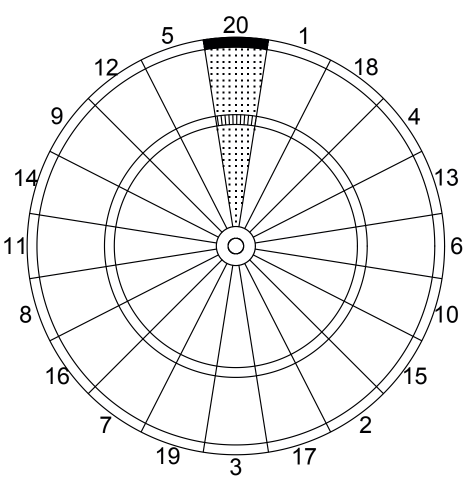

Figure 1. A standard dartboard. The small circle at the centre of the board is called the “double bulls-eye” or just the “bulls-eye”, and is worth 50 points. The surrounding ring is called the “single bulls-eye”, and is worth 25 points. Beyond these two regions, the board is divided into 20 pie-slices, with point values ranging from 1 to 20. There are two rings, a “double” and “triple” ring, that span these pie-slices, and they multiply the score by a factor of 2 and 3. For example, the dotted region above is called the “single 20” and is worth 20 points; the solid region is called the “double 20” and is worth 40; the striped region is called the “triple 20” and is worth 60

First, some basic darts rules: games are played by throwing sharp metal missiles, the “darts”, at a circular target, the “dartboard”. The board is divided into many regions, and a player receives a different score depending on where his or her dart lands. Figure 1 shows the layout and scoring system. In this article, it is assumed that players are interested in achieving the highest score possible with each throw. Common games will often have a more complicated objective in mind (in “501”, for example, players count down from 501 points to zero and are penalised for going negative), but still involve this principle of high score.

Now, my roommate Andy was always quite a bit better than me – if his throws landed with sniper-like accuracy, then mine looked more like buckshot flying out of a hunting rifle. One night, we decided to record our average scores over 100 consecutive dart throws. Andy’s average was 15.9 points, and mine was 11.7. I already thought that my average score seemed low, but Andy, amused with this result, delivered the ultimate insult: he claimed that my average was lower than that for a dart throw distributed *uniformly at random* over the board! Though I immediately disputed this ridiculous-sounding claim, I was hesitant to actually check it by computing the expected score in question. It wasn’t so much that the calculation itself was difficult – this was just a matter of looking up the dimensions of the standard dartboard and working out the areas of the different regions. It was rather an issue of pride. As a statistician, I interpreted the uniform distribution as a complete lack of skill or strategy on behalf of the darts player, the least favourable distribution (in the literal, not statistical, sense) that one could think of!

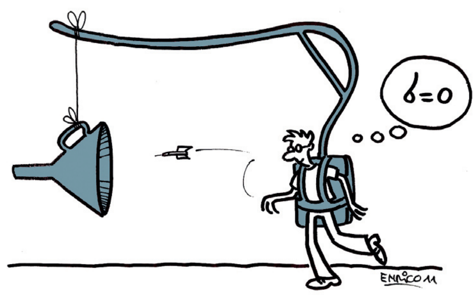

You can imagine my dismay when I eventually calculated the average score for a dart throw distributed uniformly over the board, and discovered that Andy was right: the average is 12.8 points, higher than my own (empirical) average of 11.7. Besides acknowledging my poorly tuned motor skills, how could I explain this phenomenon? I could think of two reasons:

First, in the uniform model, dart throws never miss the board entirely (a score of zero), but in

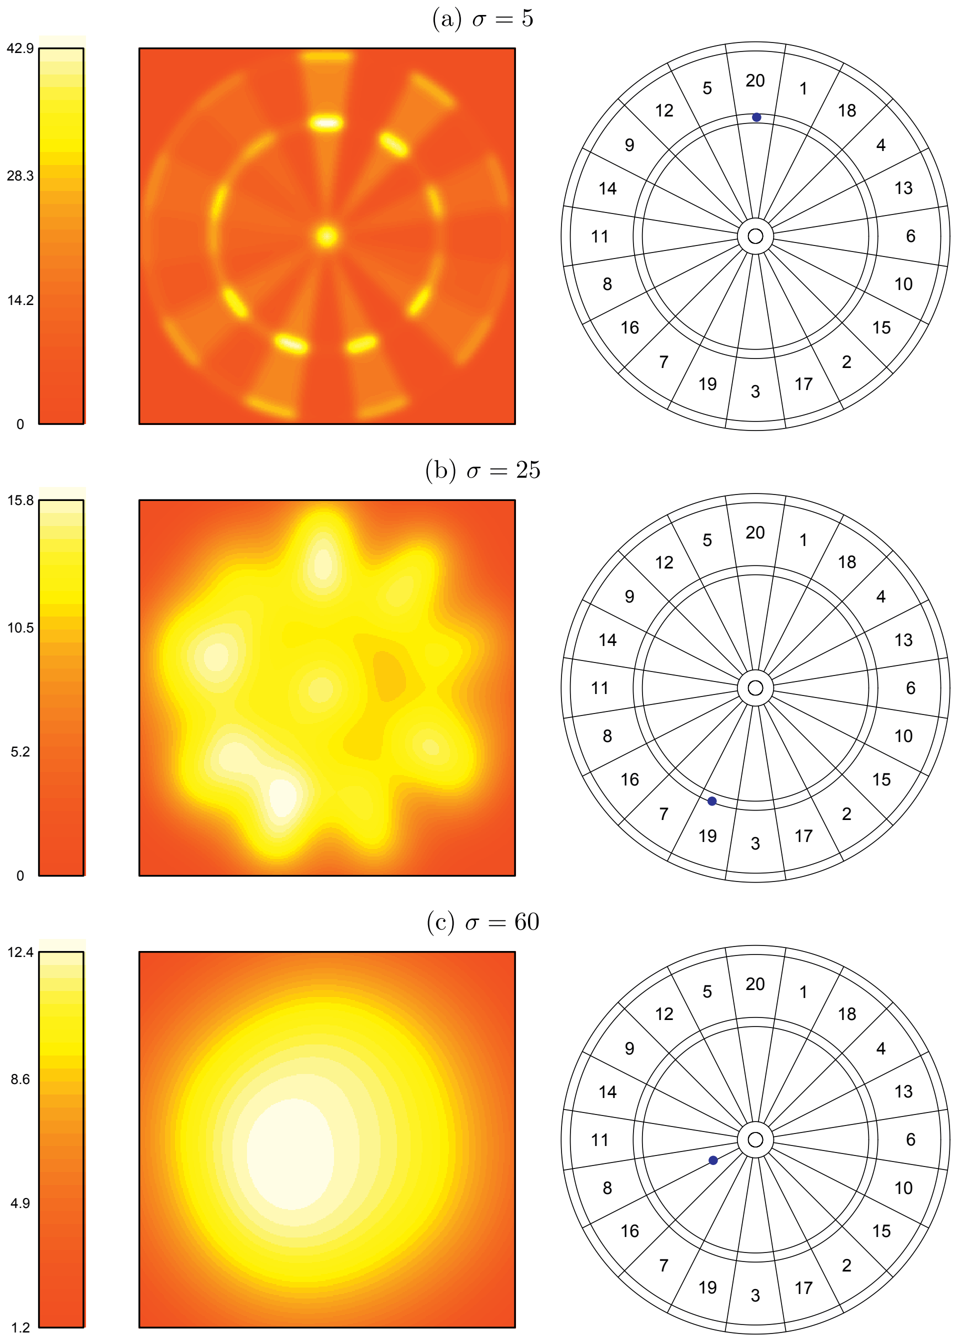

Figure 2: Heatmap of how the expected score varies with the aiming point μ, for (a) σ = 5, (b) σ = 25, and (c) σ = 60 mm. These represent very accurate, moderately accurate, and fairy inaccurate players, respectively. The hotter the colour the more the player can expect to score by aiming at that point. The aiming location μ that maximises the expected score (i.e. the best place to aim) is marked by a blue dot

reality, this sometimes happens (to me, if not to Andy).

Second, I was aiming at the triple 20, because this is the highest scoring region of the board. The same logic is probably used by most beginners. Referring back to Figure 1, the 20 pie-slice is adjacent to the 5 and the 1, which are low scores. Therefore, for an inaccurate player like myself, the triple 20 may not be the best place to aim.

The second reason is interesting, and begs the question: where *is* the best place to aim? Based on what we have just discussed, this should depend on a player's level of skill. Perhaps the most natural model assumes that a player hits his or her intended aiming location on average, but individual throws are diverted by random error. If we let the origin 0 correspond to the centre of the board, we can express this idea mathematically as

$$Z = \mu + \varepsilon, \quad \varepsilon \sim N(0, \sigma^2 I),$$

where Z denotes the two-dimensional position of a player's dart throw, μ is the location at which the player is aiming, and ε is a Gaussian error with mean zero and spherical covariance $\sigma^2 I$. Non-mathematicians can skip that sentence and need not worry. All it means is that every dart thrown by the player will land at his aiming target plus some error, and this error depends on how generally inaccurate he is. This inaccuracy is measured by σ, which is the standard deviation of his throws. A smaller σ means a more accurate player who can land the darts closer to where he wants them.

Aiming at the triple 20 would cost a poor player more than two points each throw

Now the question can be phrased as follows: if you know your level of inaccuracy σ, what aiming location maximises your expected score?

Figure 2 uses heatmaps to display the expected score as a function of aiming-point for three examples of player inaccuracy. Here σ is given in millimetres; for reference, the dartboard's radius is 170 mm. The hotter the colour, the more you can expect to score, on average, by aiming at that point. As we anticipated, the heatmaps look quite different depending on the level of accuracy. For a player with near-perfect accuracy, σ = 5 mm, the best place to aim is indeed the triple 20; by aiming there the player can achieve an average of 42.9 points per throw. Meanwhile, a player with decent accuracy, σ = 25 mm, should aim at the triple 19. This earns him or her an average of 15.8 points per throw, and aiming at the triple 20 gives only 14.9, nearly a full point lower. For a beginner, σ = 60 mm, aiming somewhat lower than and to the left of the board's centre gives the highest expected score, 12.4 points per throw. On average, aiming at the triple 20 would cost this player more than two points each throw, yielding only 10.2 points.

Expected score heatmaps, like those in Figure 2, are interesting but not very useful until we

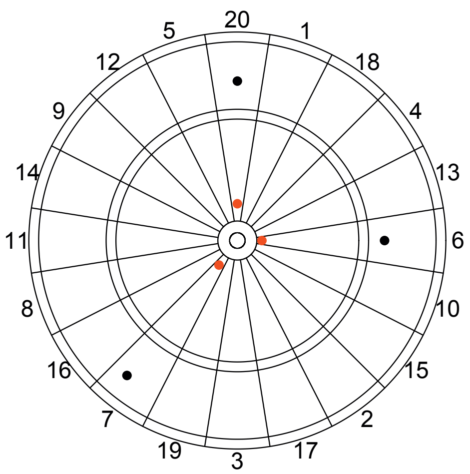

Figure 3: Dart throws from two players aiming at the centre of the board: an inaccurate player (in black) and an accurate one (in red). Both players would record the same set of scores: 20, 6, 7

have a way of estimating a player's σ. This would allow us to produce a customised heatmap that instructs the player where to aim to optimise his or her average score. To estimate the inaccuracy parameter σ, suppose that we had the player throw 100 darts aimed at the centre of the board (so that μ = 0). If we measured the positions of these throws, then we could use standard statistical methods to find σ. This would be quite easy to compute given the simple normal model for dart throws: it is the distance from the target within which 39% of his shots land. While this would certainly give an accurate estimate, no player would want to get out a ruler and measure the distances between 100 of his throws and the centre of the dartboard. This is a pretty dull task! On the other hand, it is easy to record the scores of each throw. But can we get a decent estimate of σ from these?

Initially, the answer may appear to be "no", because most scores seem to carry relatively little information about the dart throw's position. Figure 3 illustrates this point by showing the positions of three throws from two players each, marked by black and red dots. The black player's throws landed much farther from the board's centre, indicating that he is probably less accurate (given that he was aiming at the centre), but both players would record the same scores: 20, 6, and 7.

The story is different if we collect more scores. Figure 4 shows the positions of 100 throws from the same two players. The black player records several scores that uniquely correspond to the triple and double regions (e.g. triple 16 = 48 points, and double 13 = 26 points), which are far away from the centre. The red player does not hit any triples or doubles, but records many more scores of 25 and 50, which uniquely correspond to the bulls-eye regions. So based on only their scores, we would follow this line of thought to properly conclude that the red player is more accurate than his black opponent. Further, it seems possible to estimate a player's σ from his or her scores, provided that they throw enough darts.

Throw 50 darts, record the scores and find out your deviation at darts

© iStockphoto.com/Alex Hinds

This last intuition is right. The mathematically inclined may like to know that there is an algorithm, an implementation of the expectation–maximisation algorithm, that can effectively estimate a player's standard deviation σ from 100 scores of throw aimed at the board's centre. Actually, even 50 scores will give a fairly good estimate. The details, as well as more complex models for dart throws, can be found in a separate paper in *RSS Journal Series A* published to coincide with this article. But help is at hand – even if you don't understand the maths, all you have to do is throw the darts! Throw them 50 times, record the scores, and go to `http://stat.stanford.edu/~ryantibs/darts/` and enter your scores on the Java applet you will find there. The website will then work out your σ, and print out your own personalised heatmap, adjusted to show exactly where a player of your calibre should aim. Use it in the pub, bar, common-room or wherever you play darts, and even if you do not win at least you should do better than before.

Andy and I could not resist running our own scores through the algorithm to estimate our values of σ. Andy's estimate was 27, and his heatmap instructs him to aim at the triple 19, close to the border that it shares with the triple 7. My estimate was 65, and my heatmap tells me to aim pretty much at the centre (so that I avoid missing the board completely). Andy and I do not live together anymore, but we still compete at various games and sports on occasion. Next time Andy trounces me, I won't feel too bad, because it just might reveal another opportunity to do some statistics!

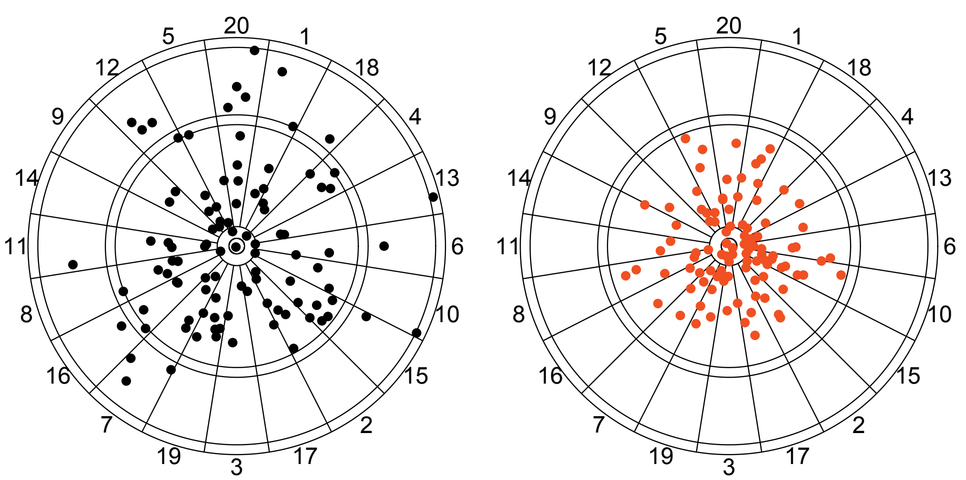

Figure 4: More dart throws from the same two players as in Figure 3. Now their scores are quite different, and allow us to discriminate between their levels of accuracy

References

1. Tibshirani, R. J., Price, A. and Taylor, J. (2011) A statistician plays darts. *Journal of the Royal Statistical Society, Series A*, **174**(1), 213–226.

Ryan Tibshirani is completing his final year of his PhD at Stanford University.