Probability of an Ace in Disc Golf

by stevewestdiscgolf.com
April 18, 2011

Data

Ace Race

Charlie Hutchinson of GottaGoGottaThrow ran an Ace Race in November, 2010. He was kind enough to measure the distances from each Ace Race tee, and also to share the results with me. With 240 players throwing 2 discs at each of 28 targets, this represents 13,440 attempts at an ace, from distances of 84 to 144 feet resulting in 79 aces.

In addition to aces, players also recorded any time they hit metal (any part of a Mach Lite, including the legs and flag). 574 throws hit some part of the target (including the 79 aces).

The event was held on a golf course, so almost all the holes were open, and flat. There was little or no wind. The temperature was cool, but not cold enough for most of the players to be wearing bulky clothing.

I found Ratings information for 56 of the 240 players. For these rated players, I have data on 3,080 throws and 169 metal hits.

Minnesota Indoor Putting Championships.

Mike Snelson ran the Minnesota Indoor Putting Championships inside his Fairway Flyerz store. 25 players threw two rounds of two discs at 9 targets, for 900 throws from distances of 29 to 50 feet, resulting in 298 aces.

Of course, the weather indoors was perfect, with no wind. Some of the holes were tricked-out (through the office window, from your knees, etc.)

Model

I modeled the throws using my previously developed Throw Model[1]. The Throw Model produces distributions of distance and accuracy based on player type, rating, and distance to the hole. Using the Throw Model provides a natural way for the probability of an ace to go to 100% at very short distances, and zero at distances beyond a player's range.

Given the proper definition of the target, the model can be adjusted to produce any probability of an ace. So the model will just provide a framework to fit the data to.

There is more data about discs hitting metal, so I used that data to define the method.

[1] See http://stevewestdiscgolf.com/ThrowSimulator.aspx

Target

The target leaves a "disc-shadow" on the ground behind it, and the Throw Model tells me where a disc landed. So, if I define the disc-shadow properly, the Throw Model-generated probability of a disc landing in the disc-shadow will be the same as the probability of a disc hitting metal.

Actually, I can directly compute the probability of a disc passing within half a disc's radius of the width of the target. So all that is left if to solve for the length of the disc-shadow which will cause the Throw Model to generate the probabilities found in the data.

The following graph shows the size of the targets used in the ace race. The red dotted line shows the area where a disc would have hit some part of the target. To simplify the calculations, I used the equivalent-sized green rectangle as the "target".

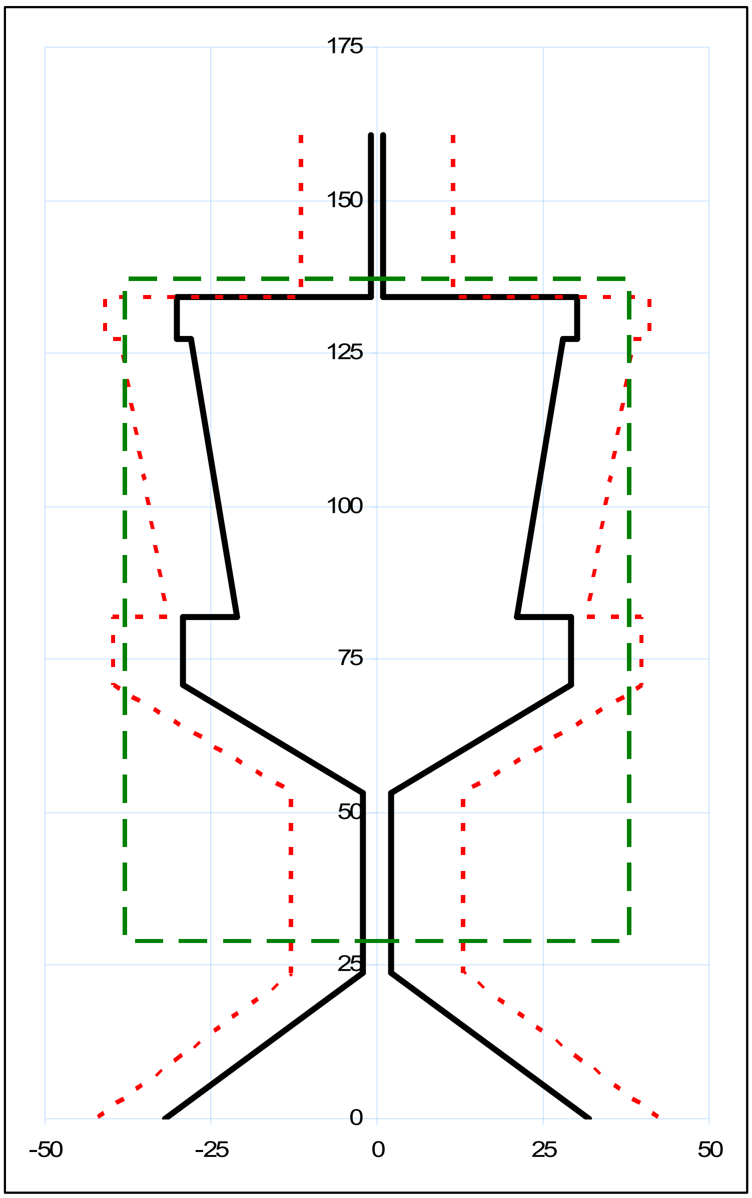

To calibrate the model I set a "trial" maximum distance behind the target. That distance represents the disc-shadow of the top edge of the green area. For every rated player and every hole, I compute the probability that a disc will 1) land shorter than the trial maximum distance, and 2) land farther than 21% of the way from the target to the trial maximum distance (the bottom line of the green rectangle), and 3) pass within 38 centimeters right or left of the target.

If the average computed probability across all rated players and holes is bigger than the observed probability of hitting metal, then the maximum distance is reduced until the probabilities match.

Give this percentage-reproducing virtual landing area, I can then directly compute the probability of hitting metal for any distances and player ratings.

The virtual landing area that calibrates the Throw Model to match observed number of metal hits is: pass within 38 centimeters left or right of the center of the pole, and land between 6 and 29 feet past the target.

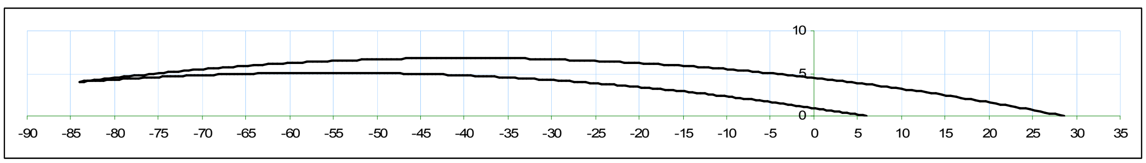

View from the side.

This seems entirely reasonable. From work I did with Chuck Kennedy, I already knew that a disc that could have hit the chains will land about 10 to 25 feet behind the target.

Remarkably, the Throw Model did not need any short-range adjustments to produce this result. Sometimes everything works out.

Just for a rough reasonableness check, here are the theoretical probabilities compared to the observed, by length of hole:

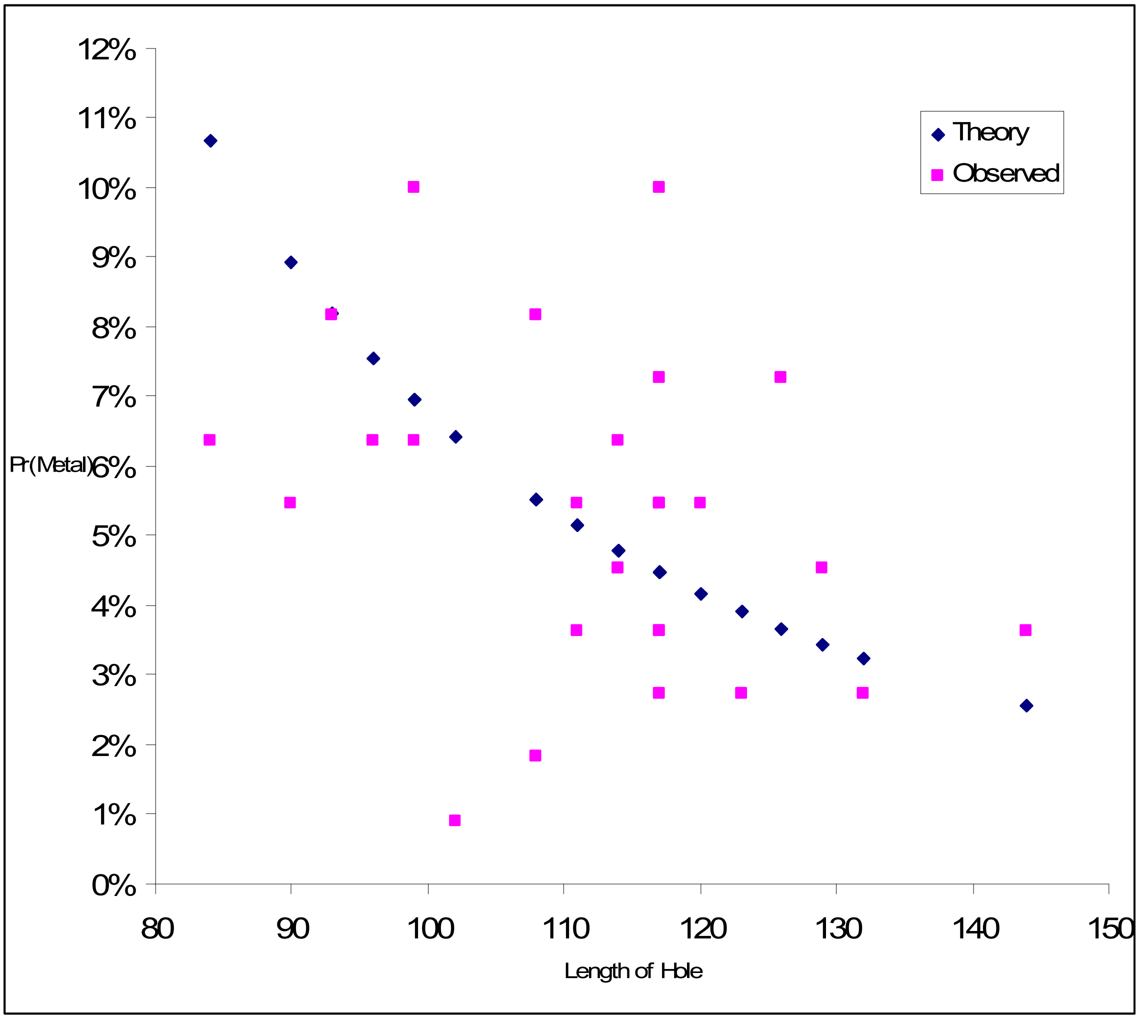

And by Player Rating:

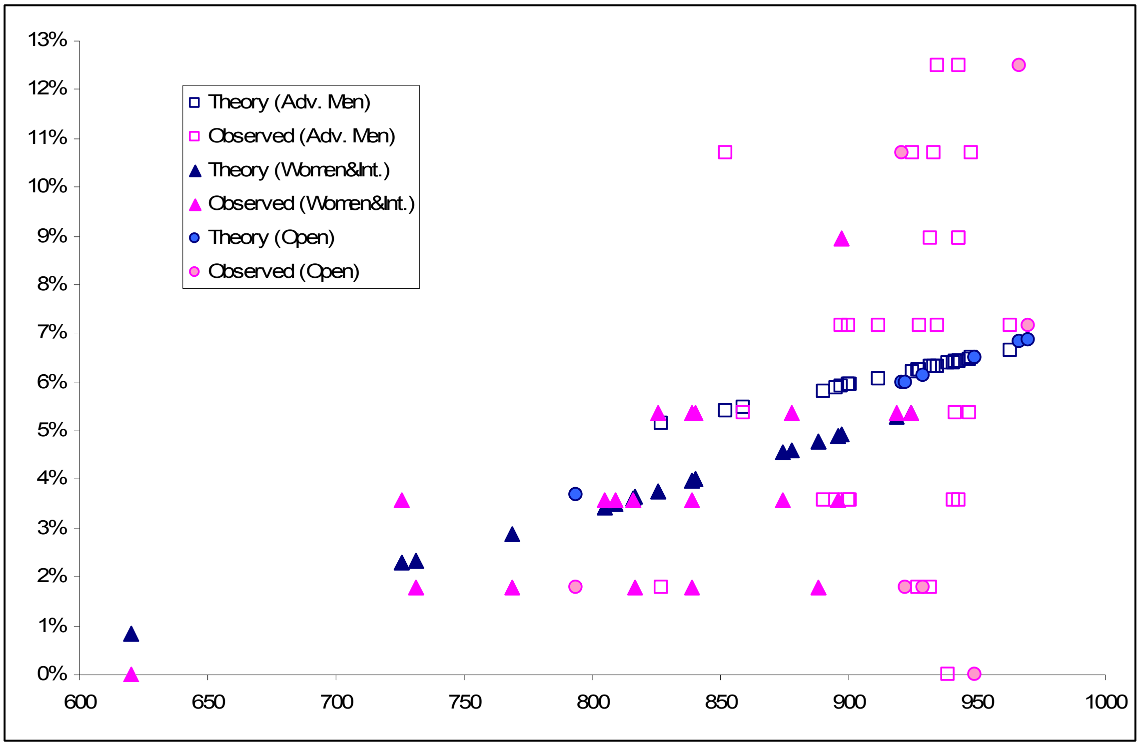

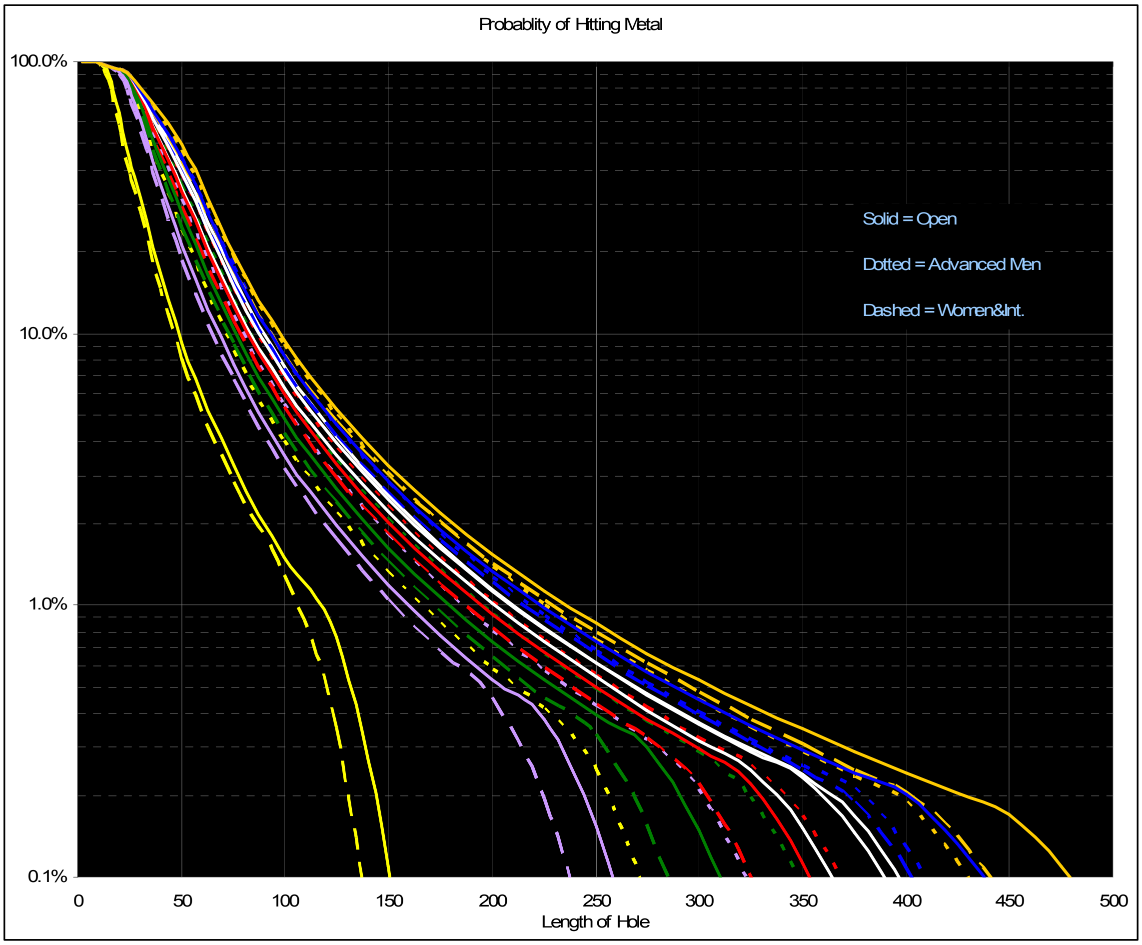

This graph shows the calculated probability of hitting metal for various Ratings, Divisions, and Hole Lengths. The formula seems to behave well. All players are almost certain to hit metal from 10 feet away. Also, for every player there is a hole length at which the probability of hitting metal falls off sharply.

For example, a 1000-rated (Gold) Open player would hit metal on a 450 foot hole about once every 600 attempts, but only one time in 900 on a 475 foot hole.

Metal vs. Ace

Having found a formula for the probability of hitting metal, there are two ways to extend that to the probability of hitting an Ace. The simpler method is to recognize that 13.7% of the throws that hit metal also went in for an Ace. So, one could just multiply the probability of hitting metal by 0.137.

However, that assumes that the probability of an ace - given that metal was hit - does not vary by player type. This would not seem to be a bad assumption; if a throw was good enough to hit metal, what does it matter who threw it?

However, the data supports the assumption that a better player will have a higher probability of hitting an ace, given that a throw was good enough to hit metal: Pr(Ace|Metal).

For the Rated Players Pr(Metal) = 5.3%, and the Pr(Ace) = 0.9%, so the Pr(Ace|Metal) = 17.2%.

For the Unrated Players Pr(Metal) = 4.0%, and the Pr(Ace) = 0.5%, so the Pr(Ace|Metal) = 12.4%.

So it seems that the better players have a better chance of getting an ace even if you only look at the throws that are "close".

Recall that I found that the calculated probability of the disc passing within 38 centimeters left or right of the center of the pole, and landing between 6 and 29 feet past the target is the same as the observed probability of hitting metal. Call this the "metal zone".

There is a subset of the "metal zone" where the probability of hitting the smaller zone is the same as the probability of hitting an ace. I chose to define the smaller zone as a certain percentage of the left to right width and short to long distance. Left/right width is of course centered on the middle of the pole. Short/long distance I centered on 15.46 feet = the expected distance beyond the hole for a throw from zero feet, according to the Throw Model. I did the distance this way so that the probability of an ace will approach 100% as the length of the hole approaches zero.

One would expect the width and length of the zone to be about 41.4% (square root of 17.2%) of the width and length of the metal zone. In a nice coincidence, the distance between the top of the rim and the bottom of the chain support is 41.5% of the height of the Metal Zone.

Solving, I found an Ace Zone that is 39.8% of the width and length of the Metal Zone. This is less than 41.4% because of the grouping of throws toward the middle. The Ace Zone is about 15 cm (6 inches) right or left of the center, with the disc landing between 12 and 21 feet behind the hole.

Verification Against Indoor Putting Championship

There were 16 Rated Players in the Minnesota Indoor Putting Championship. All the hole lengths were much shorter than any in the Ace Race. I used the model to compute the number of aces that would be expected of these 16 players.

The model predicted 182 aces (out of 576 attempts). Actually, there were 187. If the true expected number of aces is 182, then one would see 187 or more aces in over 30% of the trials.

By Hole length:

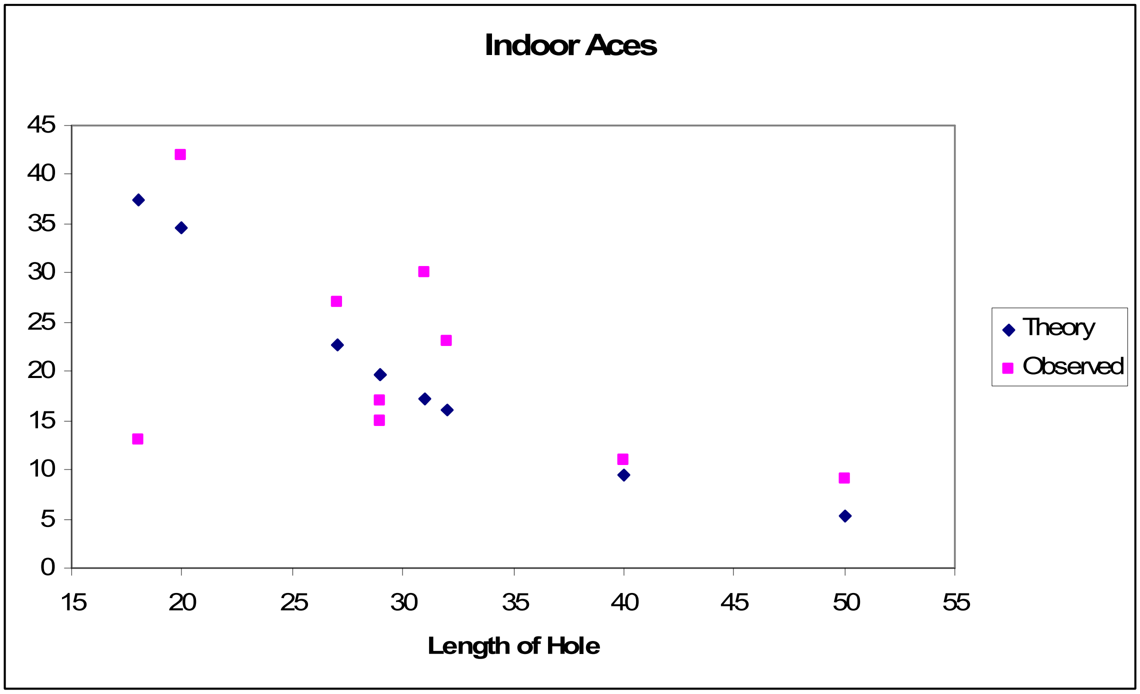

By Player Rating:

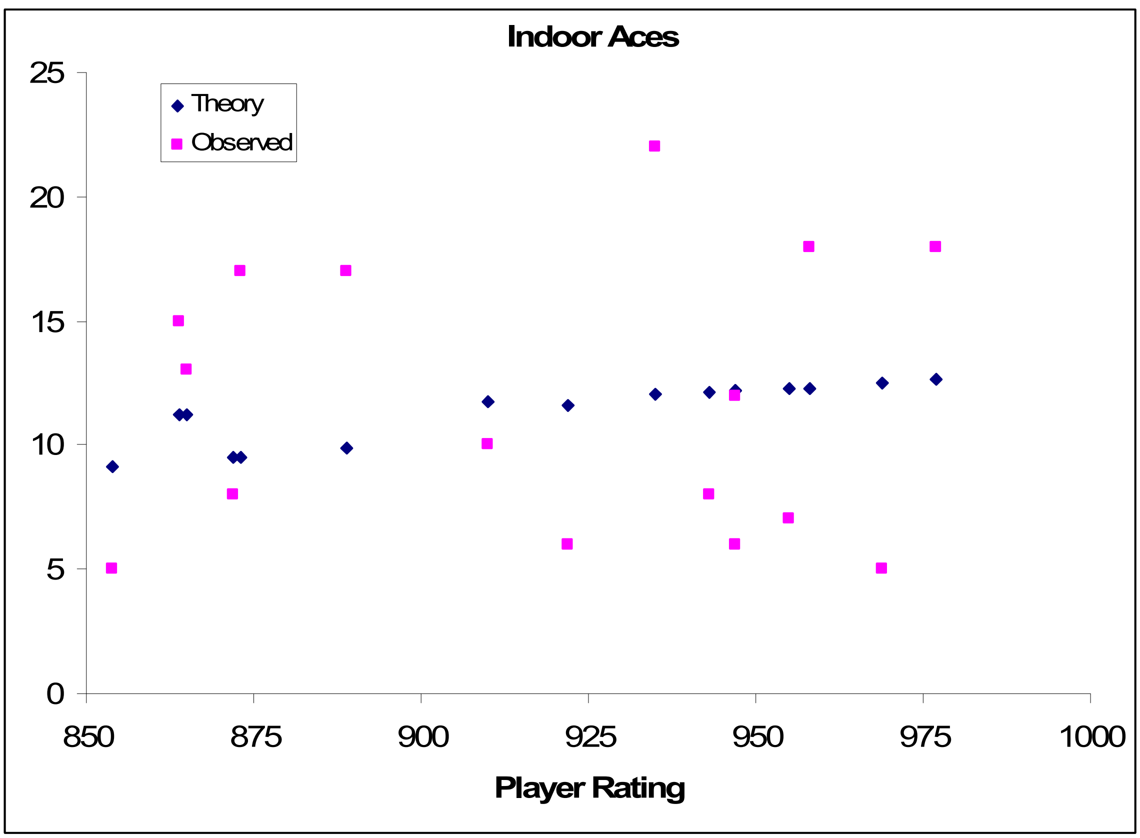

Comparison to PDGA Division Descriptions

The Professional Disc Golf Association publishes examples of the skill levels expected of various divisions. Part of the description is a measure of how well they putt. The following table compares the PDGA descriptions to the model.

Division	PDGA Description	Rating Used	Model
Advanced	From 25 to 30 feet, make 5 to 7	915	3.0 to 4.2
Advanced Master	From 25 to 30 feet, make 5 to 7	875	2.9 to 4.1
Advance Grandmaster	From 25 to 30 feet, make 4 to 6	825	2.6 to 3.8
Advanced Women	From 25 to 30 feet, make 4 to 6	800	1.7 to 2.6
Intermediate	From 20 feet, make 5 to 7	875	4.9
Intermediate Women	From 20 feet, make 3 to 5	750	5.0
Recreational	From 20 feet, make 3 to 5	875	5.4
Top Open	From 25 to 30 feet, make 7 to 9	1000	3.5 to 4.7
Average Open	From 25 to 30 feet, make 6 to 8	950	3.1 to 4.4
Top Open Women	From 25 to 30 feet, make 6 to 8	925	2.6 to 3.8
Average Open Women	From 25 to 30 feet, make 5 to 6	900	2.4 to 3.5
Pro Master	From 25 to 30 feet, make 6 to 8	925	3.1 to 4.3
Pro Grandmaster	From 25 to 30 feet, make 5 to 8	900	3.0 to 4.2

Not a close match. To resolve the difference, I looked at the data I had for Open players at short distances.

The graph on the next page shows the results. The box represents the PDGA description. The fitted line excludes the 18-foot hole, which required the players to shoot from their knees.

The PDGA does not match with the model or the data, so there is no indication that the model needs to be adjusted. The model fits well with the short-range data, even though no short-range data was input into the model. Perhaps the PDGA descriptions are a bit optimistic.

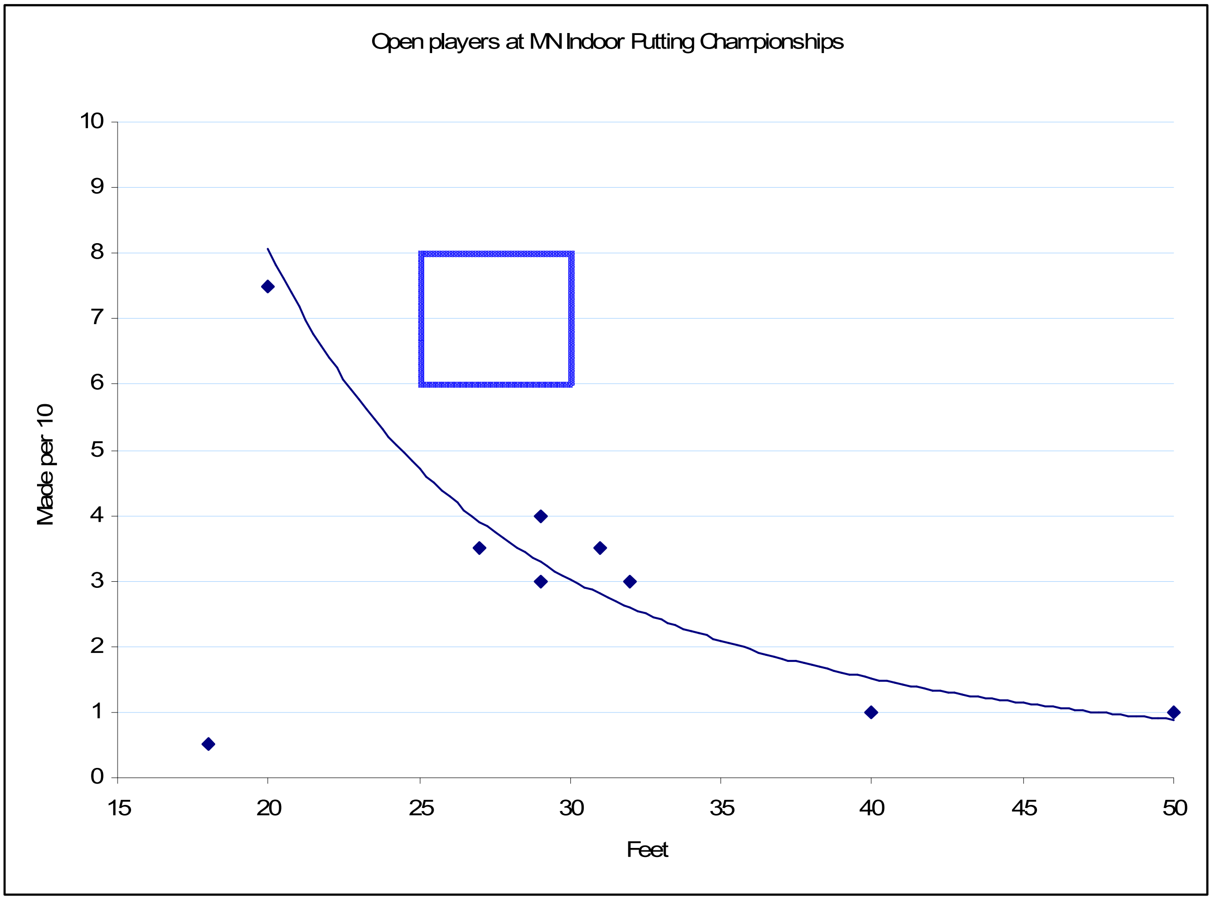
Open players at MN Indoor Putting Championships
Made per 10
10
9
8
7
6
5
4
3
2
1
0
15
20
25
30
35
40
45
50
Feet

Conclusion

A number of things, mentioned above, give me confidence that my Throw Model, with a couple of new parameters for "where the target is" can accurately calculate the probability of hitting metal or getting an ace, for a player of a given rating and division, and a hole of a given length.

I put a table of the computed probabilities here:
http://stevewestdiscgolf.com/ThrowSimulator.aspx