To appear in *Science and Golf V: Proceedings of the 2008 World Scientific Congress of Golf*

Assessing Golfer Performance Using Golfmetrics

Mark Broadie, Columbia University

ABSTRACT

The software application Golfmetrics was created to capture and store golfer shot data and to quantify differences in shot patterns between players of different skill levels. Across golfers it is shown, somewhat surprisingly, that longer hitters tend to be straighter than shorter hitters. Individual golfers can be measured relative to a benchmark to assess relative accuracy and to suggest whether to focus on increasing distance or decreasing directional errors. For amateur golfers, distance errors on short game and sand shots are shown to be about three times larger than direction errors. Shot value is a quantitative measure of the quality of each shot in comparison to a scratch golfer. Shot value analysis is a useful way to measure consistency, assess a golfer's relative strengths and weaknesses, and to indicate where practice and improvement are most needed. For amateur golfers a significant contributor to high scores is inconsistency, i.e., a relatively small number of awful shots. This research also quantifies the contributions of each part of the golf game (putting, short game, sand game or long game) to overall scores for golfers of different abilities. The long game is found to be the biggest factor in the difference in scores between pros and amateurs and between low- and high-handicap amateurs.

Keywords: Golf, benchmarking, shot value, shot patterns, statistical analysis

INTRODUCTION

Objective and quantitative analysis of the game of golf is greatly facilitated by golfer shot data collected under real golf conditions. Standard statistics, such as number of putts and greens hit in regulation, have at least two drawbacks. First, most statistics measure the effect of a combination of shots and do not isolate the quality of individual shots. For example, if a golfer misses a green and then chips in, the number of putts recorded will be smaller not because of good putting, but because of an exceptional chip shot. Fewer putts may be an indication of good putting, good chipping, or poor iron play. A second drawback is that most statistics involve counting (e.g., number of fairways hit) and do not distinguish between large and small errors (e.g., whether a fairway is missed by 1 yd or 30 yds). Starting and ending position information of individual shots allows the quality of each stroke to be measured directly and in isolation from other shots. Cochran and Stobbs (1968) pioneered the idea of collecting and analyzing golf shot data. However, they collected a relatively small amount of data, the analysis was done prior to modern computer technology, and their results pertain to golf in another era, when equipment and course conditions were very different than today. Soley (1977) collected and analyzed putting data with similar limitations. The PGA tour's excellent ShotLink™ system is used to record shots of golfers at PGA tournaments. This system contains extensive data, but is limited to the very best professional golfers and does not allow custom analyses to be performed.

Description of the Golfmetrics program and database

Golfmetrics allows golfer, round and shot information to be entered into a computer, stored in a database and analyzed. (Trademark and patent rights for Golfmetrics are claimed.) A map of each hole is created using an accurate rendering of the hole (e.g., a satellite image from Google Earth) and a separate hole editor program. Hole maps can be developed for any course with accurate hole information, but the Golfmetrics program is currently limited to six courses. Using the program's graphical interface, a user can, for example, click near a 150-yd marker and then click in the middle of a greenside bunker to indicate the starting and ending positions of a shot. The program stores this information and then computes the shot distance, the error relative to the hole position, and using the hole map it can determine that the shot started in the fairway and ended in a bunker. This design allows detailed information to be entered relatively quickly, easily and accurately.

The Golfmetrics database currently contains almost 40,000 shots representing about 500 rounds of golf from over 130 golfers on six courses in tournament and casual play primarily during 2005-2007. Golfer ages in the database range from 9 to 70 years and the scores range from 64 to 120. PGA and LPGA tour pros, club professionals, and amateur golfers are included. PGA tour shot information was transferred from ShotLink™ into Golfmetrics. The data were divided into five groups for analysis: Pro1 (PGA tour players scoring in the range 64-71), Pro2 (PGA tour players scoring in the range 72-79), Am1 (low-handicap amateurs scoring in the range 70-83), Am2 (middle-handicap amateurs scoring in the range 84-97) and Am3 (high-handicap amateurs scoring in the range 98-120). The Pro1 and Pro2 groups were sometimes combined into a single Pro group.

Performance measures

The fractional remaining length of a shot, or FRL, is the distance of the endpoint to the target divided by the initial distance to the target. A useful robust measure of error of a group of shots is the median (50th percentile) FRL. The median FRL measure combines distance and direction errors into a single number which measures a player's ball-striking ability. A measure of the distance potential of a golfer is $d_{0.75}$, the 75th percentile shot distance of a group of long tee shots. This percentile measure is less sensitive to very poor shots than average distance. Directional error is measured by the standard deviation of direction, $\sigma(\alpha)$. Let α represent the angle, in degrees, between the start-end line of a shot and the start-target line. A shot with a +4 (deg) error will finish 14 yds to the right of the target on a 200-yd shot and 21 yds to the right on a 300-yd shot (since $\tan(4°)=14/200=21/300$). For a group of shots, the standard deviation of these angles, $\sigma(\alpha)$, is a measure of direction error. Both median FRL and $\sigma(\alpha)$ are useful because they are normalized by the shot distance, i.e., they automatically account for the increase in absolute error that occurs with longer shots. In contrast, standard "greens hit" and "fairways hit" statistics are more sensitive to the size of greens and width of fairways. Shot value, defined next, is useful for measuring the contribution of individual shots to overall scores.

Shot value is a measure of the quality of a shot relative to a scratch golfer's average shot from a given situation. To define shot value, the notion of fractional par is introduced. Fractional par is an estimate of the average number of strokes that a scratch golfer needs to complete a hole, and this value depends on whether the shot starts from the fairway, rough, sand, etc. For example, a 140-yd par-three hole has a fractional par of 3.2, while a 200-yd hole has a fractional par of 3.5. A putt from a distance of 14 ft has a fractional par of 1.8, since a scratch golfer will one-putt about 20% of the time, two-putt about 80% of the time, and rarely three-putt from this distance. The fractional par function is estimated from Golfmetrics data. The shot value, v, is computed from the fractional par at the start of a shot (denoted f_s) and the fractional par at the end of a shot (denoted f_e) according to equation (1):

$$v = f_s - f_e - 1 \quad (1)$$

If the shot value is positive, it means the golfer used one shot to reduce the average number of remaining shots by more than one. For example, suppose on a par-three hole with a fractional par of 3.2, a golfer hit a shot onto the green to 14 ft from the hole. The golfer gained 1.4 in fractional par (the difference between 3.2 and 1.8) but took one shot to do so. The shot value is $v = 3.2 – 1.8 – 1 = 0.4$, meaning it is 0.4 strokes better than a typical scratch golfer shot. If the golfer missed the 18-ft putt by one inch, the shot value of the putt is $v = 1.8 – 1.0 – 1.0 = –0.2$ (since the fractional par of a one-inch putt is 1.0). If a golfer hits a shot out of bounds, applying equation (1) twice gives a shot value of –2, because the golfer will be at the same position (with the same fractional par) two shots later. Landsberger (1994) introduced a related performance standard, but did not apply the standard to analyze a large amount of data.

There are many ways to shoot the same score on a hole, but the shot value calculation shows which shots contributed more (or less) to the overall score. The shot value measure adjusts for the difficulty of each shot, including the distance of each shot to the hole. In particular, a putter will not be rated as a better simply because of starting closer to the hole. Strengths and weaknesses of a golfer's game can be identified by aggregating shot values from many shots.

An overlooked golf performance measure is consistency. Consistent golfers have few very poor shots and few "blowup" holes. Consistency can be measured by tracking exceptional shots – those with very large or very small shot values. An *awful* shot is any shot with a shot value less than –0.8; a *great* shot has a shot value greater than +0.8. A *typical* shot is any shot that is neither awful nor great. Examples of awful shots include any shot hit out of bounds, most shots hit into water, hitting a 60-yd shot into a greenside bunker, hitting a tree on a drive so the ball only travels 70 yds, and missing a 2-ft putt. Examples of great shots include hole outs from off the green, hitting a 90-yd shot to 2 ft from the hole and hitting a 210-yd shot to within 10 ft of the hole. A simple measure of consistency is the number of awful shots hit in a round.

RESULTS

Across golfers, are longer or shorter hitters straighter? Short-hitting Fred Funk is straighter than the long-hitting John Daly, but does this intuition apply more generally? Figure 1 shows long tee shot patterns of the groups Pro and Am3. Tee shots are hit from $(x,y)=(0,0)$ and the vertical axis ($y=0$) corresponds to the middle of the fairway. Not only are PGA tour players much longer, but they are about twice as straight as the amateurs. The same conclusion holds for other groups of golfers as shown in Table 1. A likely reason is that golfers who hit the ball longer have better swings, make better contact, and are generally better golfers. Figure 2 provides a benchmark relationship across golfers and illustrates that the *longer hitters are straighter* pattern does not necessarily hold when comparing individual golfers. Individual golfers appearing above the benchmark line in Figure 2 are relatively wild for their tee shot distance; golfers below the line are relatively straight hitters who might do well to work on increasing their distance.

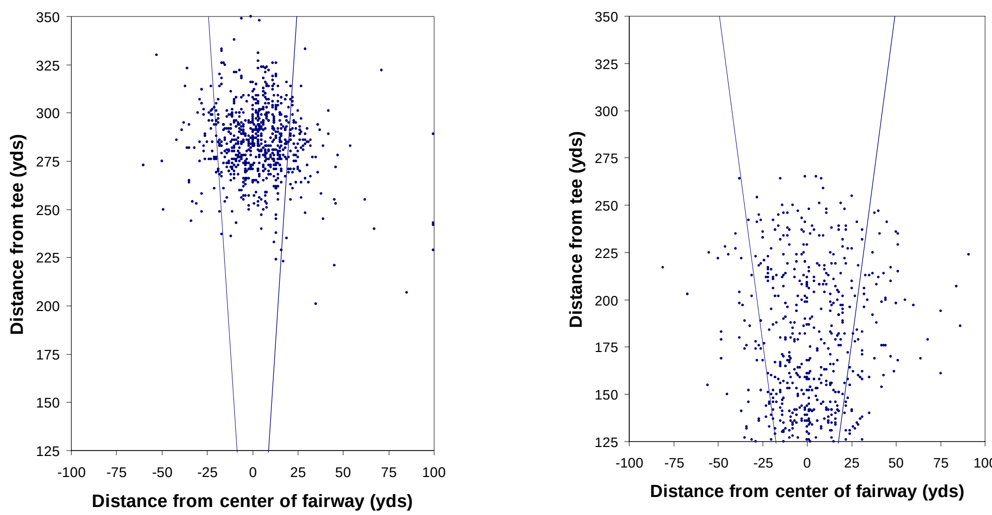

Figure 1 Long tee shot patterns: Pro (left panel, $d_{0.75}$ = 297, $\sigma(\alpha)$ = 4.0) and Am3 (right panel, $d_{0.75}$ = 216, $\sigma(\alpha)$ = 8.1). Solid lines indicate +/- one standard deviation of direction error.

Table 1 Putting, sand game and long tee shot results.

	Putting		**Sand game**		**Long tee shots**	
	50% prob distance	**Avg 2-putt distance**	**Sand save (%)**	**Median FRL (%)**	**75% tee distance**	**Std dev direction**
Pro (64-79)	8.2	30	50	16	297	4.0
Am1 (70-83)	5.8	25	26	30	248	5.4
Am2 (84-97)	5.1	19	17	40	237	6.4
Am3 (97-120)	3.8	12	7	49	216	8.1

Table 1 also gives putting and sand game results. The *50% (probability) distance* is defined as the length of the initial putt for which a golfer has a 50% chance of a one-putt. For PGA players the 50% distance is about 8 ft and it decreases to 4 ft for high handicappers. Pelz (1989) reports a 50% distance of about 6 ft for PGA players – the increase from 6 to 8 ft in the past 20 years is likely due to better and deeper tournament fields and better conditioned greens. The *average two-putt distance* is defined as the length of the initial putt for which a golfer will average two putts, i.e., will have an equal number of one-putts and three-putts. For PGA players the average two-putt distance is 30 ft and it decreases to 12 ft for high handicappers. Figure 3 shows how the average number of putts to hole out increases with the initial putt distance by golfer category. PGA tour players, as a whole, are better putters from every distance than low-handicap amateurs.

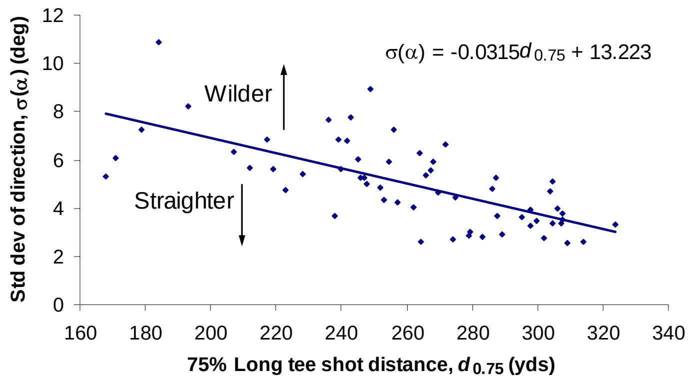

Figure 2 Longer and straighter: directional accuracy versus distance (each point represents an individual golfer).

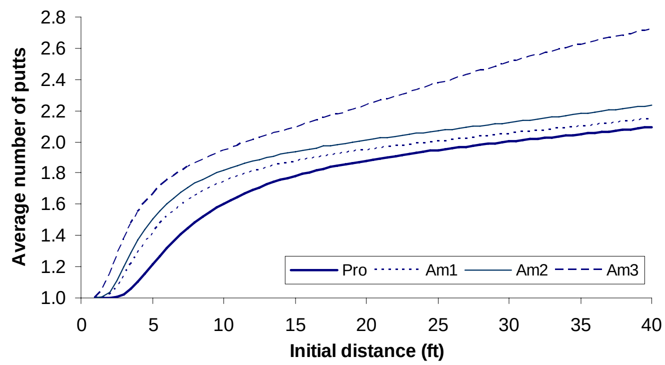

Figure 3 Putting: average number of putts versus initial putt distance by golfer category.

Figure 4 shows short game and greenside sand shot patterns of the middle-handicap group (Am2). All shots are scaled for consistent plotting. The hole is located at (0,0) and each start point is scaled to (0,y), where y = –40 in the left panel and –25 in the right panel. The ellipses indicate regions containing 50% of the shots and are useful for visualizing and quantifying distance errors versus direction errors (see also the "dist/dir" column in Table 2). The ellipses are about *three* times as long as they are wide, indicating that *distance errors are a much larger component of overall error than direction errors for amateurs hitting short shots and sand shots*. Distance errors (caused by poor ball contact, errors in swing speed or the amount of sand taken on sand shots, or other factors) are more important than alignment errors for short shots by amateurs. Figure 4 graphically shows that for the Am2 group, errors from the sand are roughly twice as large as from the rough (40% sand versus 21% rough median FRLs in Tables 1 and 2).

Figure 5 shows 50% ellipses for shots in the 100-150 yd range for the Am1 and Am2 groups. Separate ellipses for shots from the tee, fairway and rough are shown. The hole is located at (0,0) and each start point is scaled to (0,–125). The Am1 shots are closer to the hole; rough shots are seen to be more difficult; and for both groups the distance error for tee shots is less than fairway shots, indicating that there is a measurable advantage to using a tee.

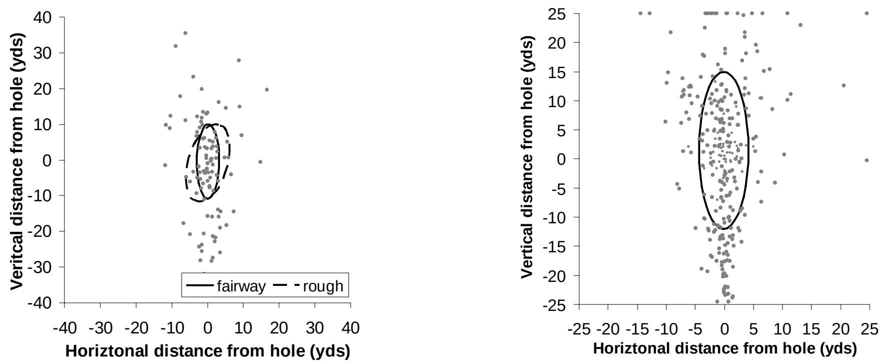

Figure 4 Am2 shot patterns and 50% ellipses: short game 20-60 yds (left panel) and sand game 0-50 yds (right panel).

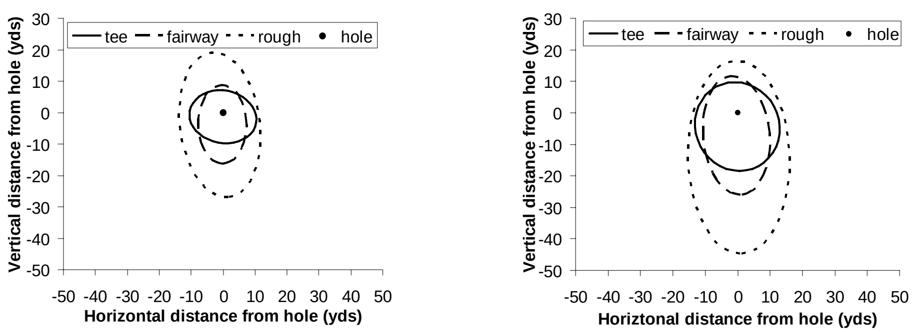

Figure 5 Long game (100-150 yds) 50% ellipses: Am1 (left panel) and Am2 (right panel).

Fairway versus rough comparisons are given in Table 2. The Pro1 (Am1) fairway median FRL of 5.4% (8.7%) from 125 yds means that half of the next shots will start within 20 (33) ft of the hole. Table 2 shows that fewer greens are hit from the rough than the fairway and the shot error is larger from the rough. Also, the rough has a bigger impact on pros than amateurs. For example, from 100-150 yds, the pro's median FRL more than doubles from the fairway to the rough, while the amateur's increases by less than 20%. Here are two possible explanations. First, most of the professional data were collected from the Barclays tournament played at Westchester Country Club in 2006 (with some data from 2005 and 2007). The tournament was played the week before the US Open and the rough was quite heavy. Second, amateurs often have difficulty from tight fairway lies and they occasionally hit better shots when their ball is sitting up in the rough. For all groups of golfers, Table 2 shows that *playing from the rough is more difficult than playing from the fairway, but the rough penalty is greater for professionals than amateurs*.

Table 2 Short game and long game results.

	20-60 yards					100-150 yards			
	Pct on Green		Dist/dir	Median FRL (%)		Pct on Green		Median FRL (%)	
	Frwy	Rough	Frwy	Frwy	Rough	Frwy	Rough	Frwy	Rough
Pro1 (64-71)	96	87	1.5	7.3	12.9	85	63	5.4	11.1
Pro2 (72-79)	95	80	1.6	9.4	15.3	77	46	5.8	12.8
Am1 (70-83)	93	86	3.2	13.8	15.4	63	53	8.7	10.4
Am2 (84-97)	81	72	3.3	16.9	21.0	46	34	12.0	13.5
Am3 (97-120)	75	64	3.1	20.3	25.6	25	25	17.3	18.4

Golfer consistency results are given in Table 3. The middle-handicap group, Am2, has 4.1 awful shots per round, on average. For almost 50 shots in the round, the middle-handicap golfer loses 0.17 shots relative to a scratch golfer for each shot hit, resulting in a total shot value of –8.5. Then with 4.1 awful shots, the golfer loses another 4.7 shots relative to the scratch golfer. Fewer than 8% of the swings produce over 35% of the shots lost relative to a scratch golfer. For the Am1 group, about 4% of the swings produce almost 70% of the shots lost relative to a scratch golfer. The awful shots could come from bad swings or from a strategy that is too risky, e.g., attempting shots with a low probability of success. Regressing the number of awful shots per round (A) on the golfer's score (S) gives a benchmark measure of consistency across golfers:

$$A = 0.24\,S - 17.1 \qquad (2)$$

A golfer who shoots a round of 75 can expect to have one awful shot, while a 95-shooter can expect six awful shots. Individual golfers averaging more awful shots than the benchmark are less consistent and focusing on reducing awful shots may be an easy way to lower scores.

Table 3 Exceptional shots (excluding putts): total shot value and number of shots per 18-hole round.

	Total shot value per round					Number of shots per round				
	Pro1	Pro2	Am1	Am2	Am3	Pro1	Pro2	Am1	Am2	Am3
Awful shots	-0.1	-0.6	-2.2	-4.7	-10.7	0.1	0.6	1.9	4.1	9.3
Typical	6.7	4.6	-1.6	-8.8	-15.5	36.8	39.6	42.6	49.3	53.8
Great shots	1.6	0.7	0.7	0.3	0.2	1.4	0.7	0.6	0.3	0.2
Total	8.2	4.7	-3.2	-13.2	-25.9	38.3	40.9	45.1	53.6	63.3

Is the biggest difference in scores between a PGA tour player and a low-handicap amateur due to: putting, short game (shots under 100 yds to the target, not including sand shots), sand game (shots under 50 yds to the target starting from the sand), or long game (shots over 100 yds to the target)? This question and related questions can be addressed by analyzing shot values.

Total shot value results are given in Table 4, where *total shot value* is the average shot value (per shot) times the number of shots in a given category in an 18-hole round. For example, suppose a player putts better than a scratch golfer and gains 0.05 in shot value for each of 30 putts in a round. The player will have total putting shot value of 1.5, i.e., will have gained a total of 1.5 shots relative to a scratch golfer in putting. Table 4 shows that the group Pro1 gains 1.4 shots per round from better putting compared to a scratch golfer. The gain is 7.1 for the long game and only about a half-shot for the sand and short games. Comparing the Am1 and Pro1 columns of Table 4 shows that for the Am1 golfers to score as well as Pro1 would require gaining 2.2, 9.3, 1.4 and 0.7 shots per round in the putting, long, short and sand games, respectively. *The biggest contributor to the difference in scores is the long game – shots over 100 yds to the target.* Put another way, if a low-handicap golfer had Tiger Woods do all of the putting, the gain would be about 2.2 shots per round, but having Tiger Woods hit all shots over 100 yds would lower the score by about 9.3 shots per round. If a Pro1 player hit only long tee shots for the Am1 golfer (about 13 or 14 shots per round) the difference would be about 4.3 shots per round.

Table 4 Total shot value per 18-hole round.

	Pro1	**Pro2**	**Am1**	**Am2**	**Am3**
Putting	1.4	-0.4	-0.8	-2.3	-5.1
Long game	7.1	5.0	-2.2	-9.3	-17.5
Short game	0.6	-0.5	-0.8	-3.4	-7.3
Sand game	0.5	0.2	-0.2	-0.6	-1.1
Total	9.6	4.3	-4.0	-15.5	-31.0

The PGA tour player's long game is better because they hit drives longer (about 50 yds) and straighter (about 1.4 deg) and they hit their irons closer to the hole. The PGA tour player's sand game is also better, but there aren't enough sand shots to make a big difference in the overall score. Why isn't putting more important? It is partially because the Am1 golfer has almost 12 putts per round within 2.5 ft that are almost always made. It is interesting to note that the average distance of the initial putt for all five groups is about 17 ft. Even though pros (Pro1) are better putters from every distance than amateurs (Am1), the net effect amounts to 2.2 putts per round – a significant number for sure, but still much less than the long game's 9.3 shots per round.[1]

Comparing any two amateur groups shows that the long game is the most important factor in explaining scoring differences. For example, of the 15-shot difference between Am2 and Am3, more than 50% of the difference (8 shots) is explained by the long game. Comparing the Pro1 and Pro2 groups shows differences of 1.8, 2.1, 1.1 and 0.3 shots per round in the putting, long, short and sand games, respectively. For PGA tour players, putting is relatively more important and is nearly as important as the long game in explaining scoring differences.

Shot value results can be broken down into subcategories. For example, the high-handicap Am3 group loses 5.1 putts per round compared to a scratch golfer: 2.4 come from putts in the 3-6 ft range, 1.9 from the 7-21 ft range and 0.7 from putts over 22 ft. From Figure 3 it might appear that the Am3 group has particular difficulty with lag putting. However, shot value analysis indicates that the Am3 group loses over three times as many strokes from putts in the 3-6 ft range compared to putts over 22 ft. The data show that these putts are missed, in part, because too many are left short and not enough break is played (60% of the misses are on the low side of the hole). These results will vary for other individuals and groups, but they illustrate how shot value analysis can be used to assess golfer performance in various categories and subcategories.

[1] The difference between pros and amateurs may be understated because of difficult PGA tournament course set ups. The shot value computation accounts for length differences, but does not account for differences in green speeds, rough height and other course features. For example, the –0.5 short game total shot value for Pro2 players in Table 4 is likely due to the difficult rough at a pro tournament, not because their short game is worse than a scratch golfer's short game. Comparisons in Table 4 between pro groups or between amateur groups do not have this drawback.

CONCLUSIONS AND APPLICATIONS

The Golfmetrics program was developed to record and analyze detailed golfer shot data. Analysis of the data reveals interesting patterns in golfer performance. Longer hitters tend to be straighter. A benchmark relationship between distance and directional accuracy across golfers is provided in Figure 2. The benchmark can be used to determine whether an individual golfer is relatively wild or straight for a given long tee shot distance. The benchmark may prove useful for a golfer to decide whether to focus on increasing distance or decreasing directional errors. The penalty for hitting from the rough versus the fairway is relatively bigger for pros than amateurs. For amateur golfers, distance errors on short game and sand shots are about three times larger than direction errors, and instruction or practice which focuses on reducing distance errors is a beneficial approach to lower scores. An accurate assessment of a golfer's ability is important in determining course strategy (e.g., how far right to aim when there is out of bounds to the left), but this analysis will be presented elsewhere.

The shot value measure to assess the quality of individual shots was introduced. Shot value can be used to identify exceptionally good and bad shots and measure golfer consistency. For amateur golfers a significant contributor to high scores is a relatively small number of awful shots. Equation (2) relates awful shots to score and gives a benchmark to measure golfer consistency. An inconsistent golfer may find reducing the number of awful shots an easy path to lower scores.

Aggregating shot values enables performance assessment of an individual golfer or a group of golfers. An individual golfer can use shot value analysis to identify strengths and weaknesses in various aspects of their game and to indicate where practice and improvement are most needed. Players who score lower tend to be better at all aspects of the game. However, the long game was found to be the biggest factor in the difference in scores between pros and amateurs and between low- and high-handicap amateurs.

ACKNOWLEDGEMENTS

Special thanks to Lou Lipnickey for exceptional programming on the Golfmetrics project. Additional thanks go to Weiwei Deng, Bin Li, Kivanc Anar, Don Devendorf and Kevin Zhang for programming support, Christopher Broadie for assistance with data input, and the many golfers who helped with the data collection. Development of Golfmetrics was supported in part by a grant from the United States Golf Association.

REFERENCES

Cochran, A. & Stobbs, J. (1968). *The search for the perfect swing*. Philadelphia: J.B. Lippincott Company.

Landsberger, L.M. (1994). A unified golf stroke value scale for quantitative stroke-by-stroke assessment. In A. J. Cochran and M. R. Farrally (Eds.), *Science and golf II: Proceedings of the 1994 World Scientific Congress of Golf* (pp. 216-221). London: E & FN Spon.

Pelz, D. (1989). *Putt like the pros*. New York: Harper & Row.

Soley, C. (1977). *How well should you putt? A search for a putting standard*. Soley Golf Bureau.