

Loss Aversion in Experts: Evidence from the PGA Tour

Steve Lee

2017

Abstract

I study loss aversion in professional golf using a proprietary dataset obtained from a web crawler. I exploit the fact that professional golfers face a “cut” after the second round of a tournament in order to group players into two categories: those who make the cut (and receive prize money) and those who miss the cut (and go home with nothing). Due to this structure, golfers can observe their position after the first round and decide on a strategy. Empirical analysis supports my predictions that 1) players inside the projected cut choose a less risky strategy in the second round than players outside the projected cut; 2) players inside the projected cut after the first round, after controlling for position differences, make the cut more often than players outside of the projected cut; and 3) the magnitude of the effects are smaller for tournaments with more skilled players. These results are consistent with the current loss aversion literature.

1 Introduction

Much existing literature has investigated the role of psychological bias in an individual’s decision making process. While some argue that bias is systematically eliminated in competitive markets (List, 2003), evidence from the PGA Tour supports the presence of loss aversion even in the face of intense competition and expertise (Pope and Schweitzer, 2011).

In this paper, I investigate loss aversion on the PGA Tour by examining second round performance. The PGA Tour is the lead organizer of golf’s most prestigious events, where approximately 150 of the best professional golfers in the world compete (nearly every week) in four-round tournaments. The player with the lowest score after the fourth round is declared the winner. In this analysis, I exploit the fact that less than half of the golfers who start the tournament will actually earn a paycheck. More specifically, after the second round of the tournament (the halfway point), the players with the highest scores are not allowed to continue to the third and fourth rounds and are deemed “cut” from the tournament. All players who are cut from the

tournament receive zero, and all players who successfully “make the cut” receive some reward strictly greater than zero. Thus, after observing the first round results, players are able to anticipate the “cut-line” and decide on an appropriate strategy for the second round of play as to maximize their chance of receiving money.

2 The PGA Tour and the Cut

Founded in 1929, the PGA Tour organizes professional golf events for the best golfers in the world. In general, each event is designed as a four-round tournament with approximately 150 players at the start and less than half allowed to finish. For each round, players complete 18-holes whereby, for each hole, they start at a designated point (said to “tee-off”) and end when they successfully hit their ball into a 108mm wide hole. The goal is to minimize the total number of hits (called “shots” or “strokes”) needed to put the ball in the hole. Prize money is paid out to each player who returns a valid score for all 72 holes in the four-round tournament. Each player’s percent of the total prize money (called “purse”) is decreasing in score, which is to say that the lowest scores receive the highest percentage of the overall purse.

Over the years, professional golf has attracted an increasing amount of attention, and thus financial incentives for players. In 2017, the largest purse was \$12 million with over \$2 million going to the winner alone. However, as mentioned, less than half of the players who start will actually receive any prize money. In fact, only the 70 lowest scoring players, including ties, are allowed to play the last two rounds of the tournament and collect prize money from the purse. The players who are ranked outside the top 70 after the second round (as a result of taking too many strokes to put the ball in the hole), are said to be cut from the tournament. They receive no money. Because of this, there is significant awareness and anticipation of what score will constitute the “cutline”, which, put more precisely, is the highest score a player can have after two rounds and still continue to play the third and fourth rounds. Note that a score exactly equal to the cutline (or “on the cutline”) constitutes “making the cut”, and those players are therefore allowed to continue to the third and fourth rounds and collect a reward.

A popular way of anticipating the cutline involves comparing your score to the score shot by the 70th ranked player after the first round and then doubling them both. While crude, this method is successfully used by players, announcers, and spectators alike, and produces consistently reliable estimates. This paper takes advantage of this simple estimation by comparing the second round performance of players who would be just “inside” and just “outside” of this projected cutline after the first round.

3 Loss Aversion

The loss aversion literature suggests that many individuals may experience greater sensitivity to losses than to gains (Thaler, Tversky, Kahneman, Schwartz, 1997). While some studies suggest that this bias decreases with skill and experience (Haigh and List, 2005), the effects may still persist. Applying this to the PGA Tour, I make my first prediction:

Prediction 1:

Players inside the projected cut after the first round will choose a less risky strategy for the second round than players outside the projected cut.

Further, Pope and Schweitzer (2011) provide a conceptual framework for analyzing loss aversion in the context of performance. In general, they propose a model of utility that relies on one's performance, probability of success (as a function of effort), and level of loss aversion such that the player maximizes their overall utility. Using that model in conjunction with the above, I make a second and third prediction:

Prediction 2:

Players inside the projected cutline after the first round will, after controlling for differences in score, make the cut discontinuously more often than players outside of the projected cut.

Prediction 3:

The magnitude of the effects described above will decrease as player skill increases.

4 Data

In order to obtain the data, I wrote a web crawler that recorded scores from an ESPN archive of golf tournament results between the years 2001 and 2017. After removing all tournaments that didn't adhere to the format described above (i.e. some invitation only tournaments that don't have a cut), I obtained a data set with the descriptive statistics as summarized below:

Table 1: Descriptive statistics.

| | mean | sd | min | max |
|---------------|-----------|-----------|---------|------------|
| Year | – | – | 2001 | 2017 |
| Tournament | – | – | 1 | 869 |
| Round 1 | 71.37165 | 3.352653 | 58 | 95 |
| Round 2 | 71.25478 | 3.42269 | 59 | 103 |
| Missed Cut | .491499 | .4999297 | 0 | 1 |
| Round 3 | – | – | 0 | 88 |
| Round 4 | – | – | 0 | 89 |
| Final | – | – | 137 | 316 |
| Projected Cut | 142.2792 | 3.182713 | 136 | 160 |
| Actual Cut | 142.8273 | 3.061517 | 137 | 156 |
| Purse | 3,907,284 | 2,616,664 | 450,000 | 12,000,000 |
| <i>N</i> | 123,632 | | | |

As seen above, this dataset contains the results for 123,632 players over the course of 869 different tournaments. Further, the variable I created for the projected cut (found by doubling the score of the 70th ranked player after the first round) is not statistically different from the actualized cut. This is necessary to support the idea that players can correctly and easily anticipate the cutline, and thus choose an appropriate level of risk for the second round.

Note also that this dataset contains results for some non-PGA Tour events. While still formally organized under the PGA Tour umbrella, events for other “mini” tours (i.e. Web.com Tour) are included. Using the dramatic difference in purse size, we can differentiate between events, which I will use as a proxy representing different skill levels. In this case, “mini” tours all have purses well under \$3,000,000, while all PGA Tour events have purses over that threshold.

5 Empirical Strategy and Results

My first prediction relies on the assumption that a less risky strategy produces outcomes with lower variance. In order to compare similar groups, I restrict the sample to players separated by a single shot - those on the cutline and those one shot outside the cutline (here the projected cutline is taken as the score of the 70th ranked player after the first round). Since players are only separated by a single shot, over the entire sample I assume that, on average, both groups are of the same skill. Under this assumption, any difference in second round outcomes would be primarily explained by psychological effects.

After the first round, players can observe their position relative to the cutline and choose a strategy to play in the second round. If they are loss averse, under the assumptions listed above, we would expect a less risky strategy for players inside the cutline (and thus lower variance) relative to those outside the cutline.

To test, each player is marked “on the cutline” if their first round score equals that of the 70th ranked player and marked “just outside the cutline” if their first round score is one shot higher. After this restriction, I obtain the second round scores for both groups and perform an F-test of equality of variances. The null hypothesis assumes that both normally distributed populations have the same variance. Summarizing the populations below, we obtain a p-value of 0.028, allowing to reject the null. This suggests that as a whole, golfers even just one shot outside of the cutline will play a riskier game than those on the cutline. These effects are magnified as you move away from the cutline (i.e. first round leaders are much less risky than those near the bottom).

Table 2: Second round scores, grouped by first round position.

| | N | mean | sd |
|--------------------------|-------|----------|----------|
| On the cutline | 17932 | 71.0614 | 3.290207 |
| One shot outside cutline | 15525 | 71.22254 | 3.346826 |

Next, I test my second and third prediction using the following regression discontinuity estimator:

$$MissedCut = \beta_0 + \beta_1 Rank + \beta_2 Rank^2 + \beta_3 InsideCut + \epsilon \quad (1)$$

In this setup, the tournament sample is restricted to only measure players within two shots (above and below) the projected cutline after the first round in order to maintain a plausibly similar sample while still gathering rank specific trends. Note also that the variable “Rank” represents their position after the first round, including ties. In other words, rank is a measure how many players scored lower than you, and every player with the same score will receive the same rank. This helps reduce noise associated with using score as the primary control variable since scores fluctuate significantly depending on golf course difficulty.

Thus, this design tries to capture any psychological effects associated with being inside the cut - namely a loss aversion associated with the expectation of receiving prize money. In order to test my third prediction, that more skilled players are less averse to loss, I use purse size as a proxy for skill where higher skill corresponds to purses greater than \$3 million (i.e. the PGA Tour) and lower skill corresponds with purses less than \$3 million dollars (i.e. not the PGA Tour). Results from a simple linear probability model (with robust standard errors) are shown below:

Table 3: Regression Results

| | Full Sample | Higher Skill | Lower Skill |
|---------------------------|-----------------------|----------------------|-----------------------|
| <i>InsideProjectedCut</i> | -0.0253*** (-3.33) | -0.0233* (-2.42) | -0.0353** (-2.71) |
| <i>Rank</i> | 0.00638*** (11.72) | 0.00590*** (8.65) | 0.00801*** (8.46) |
| <i>Rank</i> ² | 0.00000297 (0.78) | 0.00000765 (1.58) | -0.0000123 (-1.93) |
| Constant | 0.0811*** (3.77) | 0.0678* (2.55) | 0.0957* (2.50) |
| <i>N</i> | 62540 | 39105 | 21148 |

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

We observe that players inside the cutline are, on average, about 2.5% more likely to make the cut than players just outside the cutline after controlling for their positional advantage. This advantage is attenuated for the most skilled subset.

6 Conclusion

In this paper, I set out to test whether professional golfers show signs of loss aversion. While a more comprehensive approach is needed for robustness, I find evidence that doesn't contradict my predictions that: 1) players inside the projected cut choose a less risky strategy in the second round than players outside the projected cut; 2) players inside the projected cut after the first round will, after controlling for differences in position, make the cut more often than players outside of the projected cut; and 3) the magnitude of the effects will be smaller for tournaments with more skilled players.

These predictions are based on existing loss aversion literature along with studies that explore loss aversion in markets of varying skill and experience. In this case, one explanation is that professional golfers, after observing their first round position, can choose a strategy for the second round. Since missing the cut returns zero prize money, players just outside the cutline and players just inside the cutline will, on average, make their choice with different expectations of prize money. Thus, players who expect to receive a positive reward for making the cut (those on and inside the projected cutline) play less risky, while players who expect to receive zero reward for missing the cut (those outside the cutline) play more risky. These results appear to be attenuated as skill increases.

Under the conceptual framework in Pope and Schweitzer (2011), loss averse golfers will exert more effort when inside (or on) the projected cutline and

thus perform better on average than those outside the cutline. The narrative here is that players outside of the cutline may exert suboptimal effort in the second round under the premise that they already expect zero in prize money.

Future work in this area would require a more refined dataset that allows for golf course, player, and time fixed effects. Alternatively, incorporating a player's professional ranking (that accounts for every tournament played in the previous two years) would help remove skill based differences. More research would also need to be done in order to try and separate out competing psychological biases like differences in confidence level. While incomplete, the preliminary results in this paper are consistent with the current literature on loss aversion.

References

- [1] Pope and Schweitzer. *Is Tiger Woods Loss Averse? Persistent Bias in the Face of Experience, Competition, and High Stakes*. American Economic Review, 2011.
- [2] List. *Does Market Experience Eliminate Market Anomalies?*. The Quarterly Journal of Economics, 2003.
- [3] Thaler, Tversky, Kahneman, and Schwartz. *The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test*. The Quarterly Journal of Economics, 1997.
- [4] Haigh and List. *Do Professional Traders Exhibit Myopic LossAversion? An Experimental Analysis* The Journal of Finance, 2005