What it Takes to Win on the PGA TOUR

(If Your Name is "Tiger" or if it isn't)[1]

Robert A. Connolly and Richard J. Rendleman, Jr.

December 14, 2010

[1]Robert A. Connolly is Associate Professor, Kenan-Flagler Business School, University of North Carolina, Chapel Hill. Richard J. Rendleman, Jr. is Visiting Professor, Tuck School of Business at Dartmouth and Professor Emeritus, Kenan-Flagler Business School, University of North Carolina, Chapel Hill. The authors wish to acknowledge with gratitude Max Grender-Jones, who provided invaluable assistance in the collection of data. They also wish to acknowledge Kin Lo, Ken Lovell and Tom Wade of the PGA TOUR and Don Chance and Paul Wolfson who provided helpful comments. Please address comments to Robert Connolly (email: robert_connolly@unc.edu; phone: (919) 962-0053) or to Richard J. Rendleman, Jr. (e-mail: richard_rendleman@unc.edu; phone: (919) 962-3188).

What it Takes to Win on the PGA TOUR (If Your Name is "Tiger" or if it isn't)

Abstract

In this study we show what it takes to win on the PGA TOUR for Tiger Woods and other professional golfers as a function of individual player skill, random variation in scoring, strength of field and depth of field. When Woods wins, he wins by scoring 0.71 strokes per round less than other winning players. This difference reflects a) that Woods may play better than other winning players when he wins and b) that Woods tends to play in tournaments with the strongest fields which, in turn, require lower scores to win. To make this assessment, we develop a new and novel simulation-based estimate of relative tournament difficulty – the mean score per round that it takes to win a PGA TOUR event. We also explore the extent to which players could have won tournaments on the PGA TOUR by playing their normal game, with no favorable random variation in scoring. We estimate that Woods is the only player who could have won events on the PGA TOUR over the 2003-2009 seasons by simply playing normal.

KEY WORDS: PGA TOUR; Tiger Woods; Simulation; Strength of field; Tournaments; Skill; Luck; Smoothing spline.

On December 16, 2009, The Associated Press named Tiger Woods "Athlete of the Decade" (Ferguson (2009)). Since becoming a professional golfer in the late summer of 1996, Woods has won 93 tournaments, 71 on the PGA TOUR, including the 1997, 2001, 2002 and 2005 Masters tournaments, 1999, 2000, 2006 and 2007 PGA Championships, 2000, 2002, and 2008 U.S. Open Championships, and 2000, 2005 and 2006 British Open Championships. (From tigerwoods.com/aboutTiger/bio. These four tournaments, considered to be the most prestigious in professional golf, are known as golf's "majors.") Although there is often debate about whether Woods or Jack Nicklaus is the greatest golfer of all time, there is no question that since winning the Masters tournament in 1997, Woods has had no peer among professional golfers of his era.

As the title suggests, we attempt to shed light on what it takes for Tiger and others who play golf at the highest level to win on the PGA TOUR. How much 'luck,' or favorable random variation in scoring, does Tiger need to win? To what extent does the strength of field come into play in determining tournament winners? Can Tiger, or any other player on the PGA TOUR, win by simply playing his 'normal' game?

In Connolly and Rendleman (2008, 2009), we show that it takes approximately 10 strokes of cumulative abnormal favorable random variation in scoring to win a typical four-round PGA TOUR event. Moreover, almost all who finish among the lowest scoring 25 to 30 experience favorable abnormal performance relative to their estimated skill levels. Thus, most who win tournaments on the PGA TOUR must not only experience some degree of good luck themselves, but must perform with a sufficient level of combined skill and luck to overcome the collective luck of the field. (This point is also made by Berry (2001, 2008). In a typical 156-player event, a handful of players will perform far beyond their skill levels due to natural random variation in scoring, making it very difficult for any one individual player to win. Despite the odds against winning, as of the end of the 2009 PGA TOUR season, Woods had won 71 of 239 PGA TOUR events, almost 30%, since turning professional, a remarkably high winning percentage by the standards of professional golf. (We arrived at these numbers by counting the number of victories and events played from the results section of Tiger Woods' bio on www.pgatour.com.) By contrast, Phil Mickelson, generally regarded as the second-best professional golfer in Tiger's era, had won only 9% of his PGA TOUR events since turning pro.

As is well known, with just a few exceptions, Woods has tended to limit his participation on the

PGA TOUR to the 'majors,' tournaments in the World Golf Championships series and tournaments among the remaining events with the strongest fields. Since Tiger is typically competing in strong-field events, the degree of skill and luck required for him to win should be even greater than if he were competing in tournaments against players of lesser quality. Therefore, to properly assess what it takes for Tiger or any others to win, we employ simulation to estimate the relative difficulty all tournaments in our sample, which covers the 2003-2009 PGA TOUR seasons. With the same simulations, we are also able to estimate the probabilities of winning each event for all tournament participants. We highlight the estimated winning probabilities of actual tournament winners as well as those of the top five players in major tournaments. We also estimate the extent to which Woods and other top players could have won tournaments over the 2003-2009 period without experiencing any favorable random variation in scoring. As it turns out, Woods generally needs a little 'luck' to win, but he is the only player who could have actually won tournaments over our sample period by simply playing his normal game.

1. Data

Our data, derived from ShotLink, and provided by the PGA TOUR, covers the 2003-2009 PGA TOUR seasons. It includes 18-hole scores for every player in every stroke play event on the PGA TOUR for years 2003-2009 for a total of 133,645 scores distributed among 1,731 players.

Among these players, 643 recorded only two scores, and the median number of scores recorded was four. Generally, these golfers are one or two-time qualifiers for the U.S. Open, British Open and PGA Championship who, otherwise, would have had little opportunity to participate in PGA TOUR sanctioned events and, clearly, are not representative of those who compete regularly on the TOUR. Therefore, to reduce the influence of non-representative players, we limit the sample to players who recorded 10 or more scores over the 2003-2009 period. The resulting sample consists of 130,122 observations of 18-hole golf scores for 653 PGA TOUR players over 321 stroke-play events.

In other work, we have limited our samples to players who recorded 91 or more scores. We established the 91-score minimum in Connolly-Rendleman (2008) as a compromise between having a sample size sufficiently large to employ Wang's (1998) cubic spline model (which requires 50 to 100 observations) to estimate player-specific skill functions, while maintaining as many established

PGA TOUR players in the sample as possible. The censoring of a sample in this fashion will have a tendency to exclude older players who are ending their careers in the early part of the sample and younger players who are beginning their careers near the end. If player skill tends to vary with age, such a censoring mechanism can create a spurious relationship, where mean skill across all players in the sample appears to be a function of time. (Berry, Reese and Larkey (1999) show that skill among PGA TOUR golfers tends to improve with age up to about age 29 and decline with age starting around age 36. Thus, ages 30-35 tend to represent peak years for professional golfers.) To eliminate any type of age-related sample bias arising from a censored sample, we employ a 10-score minimum, rather than a 91-score minimum, and use simpler linear functions to estimate skill for those who recorded between 10 and 90 scores. In the sample, 354 players recorded 91 or more scores over the 2003-2009 period, and 299 recorded between 10 and 90 scores. A total of 119,060 and 11,062 scores were recorded for the two groups, respectively.

2. Estimating Skill and Random Variation in Scoring

2.1. Basic Scoring Model

We employ a newly modified version of the Connolly and Rendleman (2008) model to estimate skill and random variation in scoring for our sample of PGA TOUR players. We organize the model using the following general structure:

$$\mathbf{s} = \mathbf{P}\, f\left(\bullet\right) + \mathbf{R}b_2 + \mathbf{C}b_3. \tag{1}$$

In (1), $\mathbf{s} = (s_1, ..., s_m)'$ is an $N = 130,122$ vector of 18-hole scores subdivided into player groups, i, with n_i scores per player i and $m = 653$. Within each player group, the scores are ordered sequentially, with $s_i = (s_{i\,1}, ..., s_{i\,n_i})'$ denoting the vector of scores for player i ordered in the chronological sequence $g_i = 1,\ 2, ..., n_i$. We refer to g_i as the sequence of player i's "golf times." The usual error term is part of $f\left(\bullet\right)$.

$\mathbf{P}f\left(\bullet\right)$ captures time variation in skill for each of the m golfers in the sample. $\mathbf{P}$ is a matrix that identifies a specific player associated with each score. $f\left(\bullet\right) = \left(f_1\left(\bullet\right), ..., f_m\left(\bullet\right)\right)'$ is a vector of m player-specific skill functions described in more detail in the next subsection.

We assume that there are two important sources of golf-related random effects, one due to daily round-course interactions, and another related to player-course interactions (both explained further in Subsection 2.3). The $N \times 1,470$ matrix $\mathbf{R}$ identifies round-course interactions associated with each score, defined as the interaction between a regular 18-hole round of play in a specific tournament and the course on which the round is played. The vector of estimated random effects associated with each of the daily round-course interactions is denoted by b_2.

In our model, we identify player-course interactions associated with each score using an $N\times$ matrix, $\mathbf{C}$, containing 653 groups of nested player-course interactions. The vector of nested random player-course effects grouped by player is denoted $b_3 = (b_{3\ 1}, ..., b_{3\ m})'$, with $b_{3\ j} = \left(b_{3\ j\ 1}, ..., b_{3\ j\ q_j}\right)'$, and q_j is the total number of nested player-course interactions associated with player j.

2.2. Player Skill Functions

Our skill function, as applied to individual player i, takes two forms depending upon the number of sample scores recorded by player i, and may be written as follows:

$$\begin{aligned} f_i\left(\bullet\right) &= z_i\left(g_i\right) + \boldsymbol{\theta_i} \\ z_i\left(\bullet\right) &= h_i\left(g_i\right)\ \ for\ n_i \geq 91 \\ &= l_i\left(g_i\right)\ \ for\ 10 \leq n_i \leq 90. \end{aligned} \tag{2}$$

In (2), $h_i\left(g_i\right)$ is Wang's (1998) smoothing spline function applied to player i's golf scores, reduced by estimated random round-course and player-course effects, over his specific golf times $g_i = 1,\ 2, ..., n_i$, for $n_i \geq 91$. (As noted above, g_i counts player i's golf scores in chronological order.) The vector of potentially autocorrelated random errors associated with player i's spline fit is denoted $\boldsymbol{\theta_i}$ with $\boldsymbol{\theta_i} = (\theta_{i\ 1}, \theta_{i\ 2}, ..., \theta_{i\ n_i})' \sim N\left(0, \sigma_i^2 \mathbf{W}_i^{-1}\right)$ and σ_i^2 unknown. In Wang's model, $\mathbf{W}_i^{-1}$ is a covariance matrix whose form depends on specific assumptions about dependencies in the errors, for example first-order autocorrelation for time series, compound symmetry for repeated measures, etc. (See Wang (1998, p. 343) for further detail.) $l_i\left(g_i\right)$, applied to players for whom $10 \leq n_i \leq 90$, is a simple linear function of player i's golf times $g_i = 1,\ 2, ..., n_i$. We note that for the 354 players for whom we estimate skill using Wang's smoothing spline model, 163 of the spline fits turn out to be linear.

For any given player, i, $\mathbf{f} = (f_1, ...f_n)'$ denotes the vector of the player's n sequentially ordered golf scores, reduced by estimated round-course and player-course effects. If $n \geq 91$, $\mathbf{h} = (h(t_1), ...h(t_n))'$ denotes a vector of values from the player's estimated cubic spline function evaluated at points $t_1, ..., t_n$, which represent golf times $g = 1, 2, ..., n$ scaled to the [0, 1] interval. If $10 \leq n \leq 90$, $\mathbf{l} = (l(t_1), ...l(t_n))'$ denotes a vector of values from the player's estimated linear skill function evaluated at points $t_1, ..., t_n$.

In Wang's model, as applied here, for each player, one chooses the cubic spline function $h(t)$, the smoothing parameter, λ, and the first-order autocorrelation coefficient, ϕ, embedded in $\mathbf{W}$ that minimizes $\frac{1}{n}(\mathbf{f} - \mathbf{h})' \mathbf{W} (\mathbf{f} - \mathbf{h}) + \lambda \int_0^1 \left(d^2 h(t) / dt^2\right)^2 dt$. "The parameter λ controls the trade-off between goodness-of-fit and the smoothness of the [spline] estimate" (Wang (1998, p. 342)).

In Equation 3 below, we break $\boldsymbol{\theta_i}$ into two parts, $\boldsymbol{\varphi_i} + \boldsymbol{\eta_i}$, where $\boldsymbol{\varphi_i}$ represents the autocorrelated component of $\boldsymbol{\theta_i}$ and $\boldsymbol{\eta_i}$ is assumed to be white noise.

$$\begin{aligned} \boldsymbol{\theta_i} &= \boldsymbol{\varphi_i} + \boldsymbol{\eta_i}, \; with \\ \boldsymbol{\varphi_i} &= 0 \; for \; 10 \leq n_i \leq 90 \end{aligned} \tag{3}$$

Inasmuch as there are likely to be gaps in calendar time between some adjacent points in a player's golf time, it is unlikely that random errors around individual player spline fits follow higher-order autoregressive processes (i.e., AR(k), $k > 1$). Therefore, we assume that for players with at least 91 scores, each $\boldsymbol{\theta_i}$ follows a player-specific AR(1) processes with first-order autocorrelation coefficient ϕ_i. Otherwise, we assume residual errors are independent.

2.3. Estimated Random Effects

We estimate a time-varying mean skill function for each player, after adjusting the player's 18-hole score by estimated random round-course and player-course effects. We define a daily round-course interaction as the interaction between a specific daily 18-hole round of play in a given tournament and the course on which the round is played. For 283 of 321 tournaments, only one course is used and, therefore, there is only one such interaction per daily round. The remaining tournaments are played on more than one course, generally two courses, but as many as four. For example, the first three rounds of the AT&T Pebble Beach National Pro Am are played on three different courses (usually Pebble Beach, Spyglass and Poppy Hills in the sample) using a rotation that assigns each tournament participant to each of the three courses over the first three days of competition. A cut is made after the third round, and a final round is played the fourth day on a single course. Thus, the Pebble Beach tournament consists of 10 daily round-course interactions - three for each of the first three days of competition and one additional interaction for the fourth and final day.

It should be noted that we do not include specific information about course conditions (e.g., adverse weather as in Brown (2010), pin placements, whether a round is played in the morning or afternoon, etc.) when estimating (1). Nevertheless, if such conditions combine to produce abnormally high or low scores in a given 18-hole round, the effects of these conditions should be reflected in the estimated daily round-course-related random effects. We note that Berry's (2001, 2008) models for predicting player scores employ random effects for daily rounds but do not make a distinction among rounds played on different courses on the same day as we do here. We also note that Broadie (2010) estimates mean player skill while simultaneously estimating random daily round-course effects, as we do in this study.

Continuing the the AT&T example, suppose weather conditions for the first day of the event in year 20XX are relatively benign, and as a result, there is little difference in average scoring on the three courses. On the second day, however, weather conditions are more severe, and accordingly, the Pebble Beach course plays two strokes more difficult than Spyglass, which, in turn, plays one stroke more difficult than Poppy Hills. If this were the case, the six estimated random effects associated with play on the three courses over the first two days of the tournament should reflect these scoring differences.

We treat the effects on golfer scores associated with daily round-course interactions as random, rather than fixed, since the effects can be viewed as random draws from a general population of possible effects rather than as effects arising from specific conditions that could be replicated in a separate sample. Referring again to the AT&T tournament, imagine what the playing conditions might have been like during the second round of the tournament as played on the Pebble Beach course. Most generally, these conditions would have reflected the weather conditions at the time, including wind, rain and temperature, the course setup, and the way the conditions of weather interacted with the course setup. Clearly, these conditions could not be replicated or fixed in a separate sample. Hence, we treat them as random effects.

We also treat player-course effects as random, rather than fixed, since they, too, can be viewed as random draws from a general population of possible interactions between players and courses. Like daily round-course effects, one could not construct a separate sample that replicated the same player-course interactions.

Estimated random round-course effects range from -4.48 to 7.85 strokes per round and by construction, sum to zero. (Inasmuch as we model round-course effects as random effects, rather than fixed effects, we make no claim that the range of these effects is statistically significant. Instead, we present this range as an indication of the variation in scoring associated with the round-course interactions in our sample.) By contrast, estimated random player-course effects are very small, ranging from -0.147 to 0.114.

2.4. Defining Luck

If we substitute $f(\bullet) = z(\bullet) + \boldsymbol{\varphi} + \boldsymbol{\eta}$, (1) can be reexpressed as follows:

$$\mathbf{s} = \mathbf{P}\left(z(\bullet) + \boldsymbol{\varphi} + \boldsymbol{\eta}\right) + \mathbf{R}b_2 + \mathbf{C}b_3 \tag{4}$$

Further, if we subtract the non-random components in (4) from both sides of the equation, we obtain an expression for the random components of scoring.

$$\mathbf{s} - \mathbf{P}\left(z(\bullet) + \boldsymbol{\varphi}\right) = \mathbf{P}\boldsymbol{\eta} + \mathbf{R}b_2 + \mathbf{C}b_3 \tag{5}$$

Equation (5) decomposes random variation in an individual player's score, showing that unusual performance may be due to any of three factors: player-specific effects ($\boldsymbol{\eta}$), round-course effects (b_2), and player-course effects (b_3). The definition of luck turns on an understanding of these three sources of variation in golf scores. As we state in Connolly and Rendleman (2008, pg. 81), "We believe that professional golfers think of luck as sources of variation in scoring outside a player's direct and conscious control. For example, if a player is assigned a relatively easy course rotation in a multiple-course tournament, professionals would say this player had good luck in his course assignments. We can estimate the extent of such luck through the round-course effect. Similarly, if a tournament happens to be played on a course that favors a particular player's style, players might attribute any favorable outcome associated with playing on this particular course to luck, because a player cannot choose the course on which a tournament is played. We estimate the magnitude of this [source of] 'luck' through player-course effects, although they turn out to be very small. Any remaining variation in score, not attributable to round-course and player-course effects, is reflected in the [$\boldsymbol{\eta}$] error."

In our view, some of the $\boldsymbol{\eta}$ error reflects variation due to easily-recognizable influences on scoring that we do not measure directly, for example, fortunate (or unfortunate) bounces of the ball, good and bad lies, relatively favorable or unfavorable weather conditions, imprecision in reading greens and judging effective distances, etc. Of course, some influences may not be nearly as easy for observers to identify, or they may simply represent natural variation in a player's swing, or variation due to judgment or playing conditions. To illustrate, consider a player with an intrinsic skill level that would lead to a 50% chance he will sink an eight-foot putt. (According to Broadie (2010), PGA TOUR golfers one-putt 50% of the time from a distance of eight feet.) If he sinks five such putts in a row, and his intrinsic skill level has not changed, we would say this player experienced good luck (favorable random variation). Although the root cause may be favorable variation in his putting stroke, if the player cannot maintain sufficient control over his putting to sustain this favorable variation, we would call it luck when he sinks five eight-foot putts in a row. On the other hand, if he can sustain the favorable variation, or alternatively, if his rate of success in making eight-foot putts declines, this should be reflected in the model as higher and lower levels of skill, respectively.

We recognize that a portion of what we are characterizing as the purely player-specific random component of scoring may reflect strategic circumstances that might cause a player to take more risk or less risk than 'normal' in his play. If players do attempt to engage in risk-related strategies, such strategies would most likely reflect specific circumstances of competition, which might be detectable in hole-by-hole or shot-by-shot data, but not in the type of round-by-round 18-hole scoring data we employ in this study. As such, we assume implicitly that any effects on a player's variation in scoring arising from conscious decisions to take a higher- or lower-risk approach to his game are part of the residual $\boldsymbol{\eta}$ error and indistinguishable from what would otherwise represent non-strategic random variation.

Among all 653 players in the sample, the standard deviation of individual player $\boldsymbol{\eta}$ residuals ranges from 1.70 to 4.93 strokes per round. Among players who recorded 91 or more scores, Jim Furyk's standard deviation of 2.30 strokes per round is the lowest, and that of David Duval, 3.43 strokes per round, is the highest. Neither of these extremes would surprise those who follow golf, since Furyk is regarded as one of the most consistent players, and over the 2003-2009 period, David Duval became one of the most inconsistent. John Daly, generally regarded as the most inconsistent player on TOUR, has the next-highest standard deviation among those who recorded at least 91 scores, 3.32 strokes per round. Among those in this group, the standard deviations of 93% are greater than that of Tiger Woods, 2.52 strokes per round. Phil Mickelson, generally regarded as an extreme risk taker, has a standard deviation of 2.80 strokes per round, exceeded by only 39% of the players in the same group.

2.5. Hot Hands

We believe that a player's propensity to engage in streaky play can be captured in the estimate of the first-order autocorrelation coefficient associated with his residual $\boldsymbol{\theta}$ errors. In Connolly and Rendleman (2008), we find that 12 of 253 autocorrelation coefficients estimated over the 1998-2001 period are significantly negative at the 5% level and 24 are significantly positive. In a two-tail test, we show that 23 are significantly different from zero. However, using Storey's (2002, 2003) false discovery method, we estimate that only two players show any evidence of significant negative autocorrelation and approximately 14 of the 23 statistically significant positive coefficients are indeed significantly positive. "Thus, there is clearly a tendency for a small number of PGA TOUR participants to experience statistically significant streaky play" (p. 87). (These tests were conducted using the bootstrap, took over a week to complete, and are not repeated in the present study.) We note that approximately 64% of the estimated first-order autocorrelation coefficients are positive among players in our present 2003-2009 sample with 91 or more scores. This compares with 61% in our 1998-2001 sample. The mean values of the coefficient in the 2003-2009 and 1998-2001 samples are 0.0231 and 0.0194, respectively.

In our original study, we also find confirming evidence of streaky play using conventional runs tests and a Markov chain test. Gilden and Wilson (1995) find evidence of streaky play in putting. Clark studies the possibility of hot hands in 18-hole golf scores (2003a, b, 2004a) and in hole-by-hole scores (2004b) but finds little evidence of streaky play among participants on the PGA TOUR, Senior PGA Tour, and the Ladies PGA Tour.

3. What does it Take to Win?

3.1. Neutral and Normal Scoring

Throughout, we assess player performance in terms of neutral scores – scores reduced by estimated round-course and player-course effects. As such, neutral scores provide an estimate of what a player's score would have been after removing the effects of the relative difficulty of the round in which the score was recorded as well as any personal advantage or disadvantage the player might have had when playing the course, therefore neutralizing any effect associated with personal tournament choice. (For example, long courses might favor long hitters. Tight courses might favor those who are the most consistent in controlling their drives. Courses with fast greens might favor certain players over others.)

Mathematically, if we rearrange (4), we obtain neutral player scores as follows:

$$\mathbf{s} - \mathbf{R}b_2 - \mathbf{C}b_3 = \mathbf{P}\left(z\left(\bullet\right) + \boldsymbol{\varphi} + \boldsymbol{\eta}\right). \tag{6}$$

As such, neutral scores reflect scoring estimates from player skill functions, $z\left(\bullet\right)$, autocorrelated components of scoring, $\boldsymbol{\varphi}$, and the purely random components, $\boldsymbol{\eta}$.

We also refer to players playing "normal," where playing normal is defined as recording a score with a zero $\boldsymbol{\eta}$ error. Mathematically, "normal" player scores are defined as follows:

$$\mathbf{s} - \mathbf{P}\boldsymbol{\eta} = \mathbf{P}\left(z\left(\bullet\right) + \boldsymbol{\varphi}\right) + \mathbf{R}b_2 + \mathbf{C}b_3. \tag{7}$$

As such, "normal" scores reflect what players would have been expected to shoot under given playing conditions, taking into account their estimated skill and the potentially autocorrelated components of their scoring.

3.2. Summary Statistics for Tournament Winners

Table 1 provides summary statistics for Woods and other players in tournaments they won over the 2003-2009 period. The table shows that the mean $\boldsymbol{\eta}$ residual among all winning players is -2.51 strokes per round. This value is consistent with the observation made in connection with our original study, covering the 1998-2001 period, that it takes approximately 10 strokes of cumulative 'good luck' to win a typical four-round PGA TOUR event. When Tiger wins, he scores an average of 1.15 strokes per round better than his predicted score. Woods' mean $\boldsymbol{\eta}$ error of -1.15 when winning is significantly greater than that of other winning players (p value < 0.001). Among players who recorded 20 or more winning scores from 2003 to 2009, the player with the next-least negative average winning $\boldsymbol{\eta}$ residual is Vijay Singh (-1.63), followed by Mike Weir (-1.72), Jim Furyk (-1.88), Ernie Els (-2.06) and Phil Mickelson (-2.27).

When Tiger wins, the mean of his neutral winning scores is 0.71 strokes lower than that of other winning players, a difference that is statistically significant at the 0.001 level. This can be interpreted two ways. First, Tiger might simply play better when he wins compared with others who win. Alternatively, or perhaps in addition to playing better when winning, Tiger may choose to play in tournaments that are more difficult to win in terms of their strength and depth of fields. We address the issue of relative tournament difficulty in the next section.

The mean round-course effect of 0.85 strokes per round in tournaments that Tiger wins, compared with 0.04 strokes per round when others win, indicates that Woods tends to win on tougher courses and/or under more difficult physical playing conditions. (The difference in mean round-course effects is statistically significant at the 0.001 level.) Although Tiger needs a less favorable player-course effect to win compared with other winning players (p value < 0.001), the magnitude of these effects is too small to be of any practical relevance.

3.3. 12 Top Players

Figure 1 provides plots of neutral scores and predicted neutral scores, including the autocorrelated component, for the 12 most highly-skilled players among those who recorded 91 or more scores in the 653-player sample, where a player's skill level is defined as the average of predicted neutral scores from his estimated skill function, the first number shown to the right of the player's name.

The second number to the right of each player's name is the average value of his $\boldsymbol{\eta}$ residual (or the difference between his neutral score and predicted score, including the autocorrelated component) in tournaments he won. The "+" symbol denotes neutral scores in tournaments the player won.

Note that some of the lines showing predicted neutral scores are jagged, while others are relatively smooth. Deviations from smoothness reflect the autocorrelated components of scoring, which are more pronounced for some players than for others. For example, estimated first order autocorrelation coefficients for Padraig Harrington, David Toms and Vijay Singh, which appear to show the largest autocorrelated scoring components, are 0.117, 0.095 and 0.064, respectively. Autocorrelation coefficients for the remaining players range from -0.04 (Sergio Garcia) to 0.015 (Stewart Cink).

As shown in the figure, Tiger Woods' average predicted score of 68.01 neutral strokes per round is more than a full stroke less than that of the second most highly skilled player, Vijay Singh, with an average predicted score of 69.2 neural strokes per round. (This skill difference of 1.19 strokes per round compares favorably with the 1.1 stroke per round difference as estimated by Broadie (2010) over the 2003-20010 period between the skill level of Woods and that of the next-best player, Jim Furyk. It also compares favorably with the 1.5 stroke per round difference estimated by Berry (2001) in a sample covering 1999, 2000 and a portion of 2001, and a 0.85 stroke difference estimated by Berry (2008) in a 1997-2004 sample.) Moreover, Tiger's average predicted score is almost two strokes per round less than that of Adam Scott, the twelfth-most highly skilled player. This implies that over a typical four-round PGA TOUR event, Tiger would have had more than a four-stroke advantage over Singh and an eight-stroke advantage over Scott. Thus, on average, Scott would have needed eight strokes of favorable random variation in scoring relative to that of Woods to have finished ahead of Woods when Woods was playing his normal game, and Singh would have needed at least four.

Note that predicted neutral scores for Tiger Woods and Luke Donald are generally trending downward over the 2003 sample period, implying that they were improving. By contrast, the predicted neutral scores of several other golfers, most notably Vijay Singh and Ernie Els, are markedly trending upward. Also note that several of the plotted spline fits are close to linear, although a few, most notably those of Jim Furyk, Retief Goosen and Adam Scott, reflect nonlinear tendencies in neutral scoring.

For most players, the great majority of winning neutral scores are below predicted neutral scores. Woods, however, is an exception. 29% of Woods' neutral scores in tournaments he won over the 2003-2009 period are associated with positive $\boldsymbol{\eta}$ residuals, meaning they are higher than predicted, including the predicted autocorrelated component of predicted scores. The percentage of winning scores with positive $\boldsymbol{\eta}$ residuals for Vijay Singh, Jim Furyk and Ernie Els, the next-most highly skilled players, are 18%, 25% and 17%, respectively. Thus, only Woods and, perhaps two or three other top players, could have gotten by with worse-than-normal performance in individual tournament rounds and still have won.

3.4. One-Hit Wonders

Figure 2 provides plots of neutral scores and predicted neutral scores, including the autocorrelated component, for 12 players we define as "one-hit wonders," with neutral scores in tournaments won marked with "+" signs. To compile our list of one-hit wonders, we first created a list of all players who won only one tournament over the 2003-2009 sample period, 71 players total. Using personal knowledge and individual player bios from Wikipedia, we then narrowed the list to 12 players by employing the following criteria:

1. A player must have had only one lifetime PGA TOUR win. Such win could not be in a tournament held opposite a major event or an event in the World Golf Championships series.
2. A player could have no wins on the European Tour.
3. A player could have no more than two additional non-European Tour professional wins. The two wins could not both be in Nationwide Tour events.

As such, these 12 golfers could be considered among the least successful winners on the PGA TOUR.

The players in Figure 2 are presented in the order of the average values of $\boldsymbol{\eta}$ residuals in the tournaments won, starting with Shaun Micheel, who averaged 4.3 strokes per round better than "normal" when he won the 2003 PGA Championship and ending with Martin Laird, who averaged 1.41 strokes better than normal in winning the 2009 Justin Timberlake Shriners Hospitals Open. For the most part, players in this group had to string together a succession of outlying favorable scores to win. Moreover, with the possible exception of Martin Laird, winning among players in this group did not appear to propel them to higher levels of performance. Except for Laird, whose win toward the end of the 2009 season reflected both favorable random variation in scoring and an improvement in skill, nothing other than good luck seems to explain how the "one-hit wonders" were successful in winning their only PGA TOUR events.

4. Relative Tournament Difficulty

4.1. Overview

Suppose the average skill level of participants were the same in tournaments "A" and "B," but "B" included twice as many players. Then, clearly, "B" would be more difficult to win. In fact, "B" could be more difficult to win, even if the average skill level of its participants were not as high as that of "A," because with more players, there is a higher probability that despite being fielded by players of lower average skill, any one player in tournament "B" could get lucky and win (or string together four rounds with very favorable random variation in scoring). Therefore, to win a PGA TOUR event that takes place over anything less than an infinite number of rounds, one must not only overcome the collective skill of the field but also its collective luck, and, obviously, the probability that luck will play a significant role in determining a tournament's winner increases with the number of tournament participants. As a result, some regular large-field PGA TOUR events with relatively weak players could be more difficult to win than the smaller-field select events, such as THE TOUR Championship (the Finals of the FedExCup Playoffs), which includes the 30 players with the greatest number of FedExCup points and The Hyundai Tournament of Champions (previously the Mercedes Championships), limited to winners of TOUR events in the previous year.

During the 2003-2009 period, all tournaments on the PGA TOUR consisted of four scheduled rounds, except for the Bob Hope Chrysler Classic and the 2003 and 2004 Las Vegas Invitationals, all five-round events. Occasionally, tournaments are cut short due to adverse weather conditions. The 2005 Bell South Classic and 2009 AT&T Pebble Beach National Pro-Am were reduced to three rounds, and the 2005 Nissan Open was reduced to two. (Although players in PGA TOUR events that have been reduced to two rounds receive the same prize money they would have otherwise received, they do not receive credit for winning Official Money, which, in turn, can have an effect on their subsequent status on the TOUR.)

In estimating the relative difficulty of winning the various tournaments on the PGA TOUR, consider tournament "C," a four-round event, and "D," a five-round event. If it takes an average score of 66 strokes per round to win both tournaments, "D" would be the more difficult tournament to win, since it would be more difficult for a player with (presumably) a mean score that exceeds 66 to average 66 for five rounds than for four.

We employ Monte Carlo simulation to determine the mean score per round required to win a PGA TOUR event as a function of the number of players participating in the tournament, the number of tournament rounds, the mean skill levels of the tournament participants and their natural random variation in scoring. In this section, our definition of what it takes to win in a given simulation trial is the score of the second-place finisher. For example, if the average score per round of the second-place finisher in a particular simulated tournament competition turns out to be 66.87, a player would need to average better than 66.87 strokes per round to win. The fact that the actual winner in simulated competition might have averaged 66.0 strokes per round is of no consequence, since he could have won by scoring almost 0.87 strokes per round worse (higher).

4.2. Standardizing to Four-Round Equivalent

In estimating relative tournament difficulty, we standardize our estimates of the score required to win each tournament to that of a four-round event. In so doing, if a tournament consists of m rounds of play, the standardized 4-round equivalent score is the 4-round average score per round that yields the same probability of occurring for the winning player in simulated competition (not the second-place finisher) as the winning average m-round score required to win. Obviously, if $m = 4$, there is no adjustment, which is the case for most tournaments.

Assuming $\boldsymbol{\eta}$ errors are normally distributed, the probability that winning player i will average $\tilde{w}_{m,i}$ strokes per round or better over m rounds is given by the standardized normal probability that $\tilde{z}_{m,i} = \left(\tilde{w}_{m,i} - \mu_i\right)/\psi_{m,i}$ or lower, where μ_i is player i's mean score applicable to the m-round tournament, and $\psi_{m,i}$ is the standard deviation of the mean of m scores for player i. Following Zhang (2006), if residual errors about the mean follow an AR(1) process, as we assume in the estimation of (1), the standard deviation of the mean of m scores for player i is given by:

$$\psi_{m,i} = \sigma\left(\eta_i\right)\sqrt{\frac{m - 2\phi_i - m\phi_i^2 + 2\phi_i^{m+1}}{m^2\left(1-\phi_i\right)^2}}, \tag{8}$$

where $\sigma\left(\eta_i\right)$ is the standard deviation of η errors for player i, and ϕ_i is the first-order autocorrelation coefficient associated with player i's $\boldsymbol{\theta}$ errors.

To compute the 4-round average score per round that yields the same probability of occurring for the winning player as the winning average m-round score, we set $\tilde{z}_{4,i} = \tilde{z}_{m,i}$ and solve for $\tilde{w}_{4,i}$, giving

$$\tilde{w}_{4,i} = \mu + (\tilde{w}_{m,i} - \mu) \sqrt{\frac{m^2 \left(4 - 2\phi_i - 4\phi_i^2 + 2\phi_i^5\right)}{16 \left(m - 2\phi_i - m\phi_i^2 + 2\phi_i^{m+1}\right)}}. \tag{9}$$

Note that when $m = 4$, $\tilde{w}_{4,i} = \tilde{w}_{m,i}$, and when $\phi_i = 0$, $\tilde{w}_{4,i} = \mu_i + (\tilde{w}_{m,i} - \mu_i)\sqrt{m/4}$.

4.3. Simulation Methodology

To simulate random variation in scoring for individual player i among the 653 players in the sample, we select a starting (potentially autocorrelated) $\boldsymbol{\theta}$ error at random from the entire distribution of player i's $\boldsymbol{\theta}$ errors estimated over the 2003-2009 period. We then select 1,298 $\boldsymbol{\eta}$ errors, assumed to be white noise, at random (with replacement) from player i's entire distribution of $\boldsymbol{\eta}$ errors. Using the initial $\boldsymbol{\theta}$ error, the vector of 1,298 randomly-selected $\boldsymbol{\eta}$ errors, and player i's estimated first-order autocorrelation coefficient, we compute a sequence of 1,298 random $\boldsymbol{\theta}$ errors. We do not employ the first ten simulated $\boldsymbol{\theta}$ errors in simulated tournament competition in order to give the autocorrelated component of residual scoring errors time to 'burn in.' The last 1,288 estimated $\boldsymbol{\theta}$ errors are those needed to simulate scoring for a player who played in every PGA TOUR event over the 2003-2009 period and made every cut. Obviously, we do not come close to using all 1,288 $\boldsymbol{\theta}$ errors for any player.

We illustrate our simulation methodology through the following example. Assume that player i participates in the first of $k = 1$ to 321 tournaments over the 2003-2009 period, a four-round event with no cut. We determine player i's skill level for tournament $k = 1$, denoted as ts_1, as the mean of the spline or linear model-based estimate of his skill level over the rounds he actually played in the first tournament. His four-round tournament score then becomes $ts_1 \left(\tilde{\theta}_{i,11} + \tilde{\theta}_{i,12} + \tilde{\theta}_{i,13} + \tilde{\theta}_{i,14}\right)$, where $\tilde{\theta}_{i,j}$ is player i's jth simulated $\boldsymbol{\theta}$ residual. Computing the same for all among the sample players who actually participated in the tournament, we obtain simulated four-round scores for each player and determine simulated finishing positions accordingly.

Assume the second tournament in which player i participates is tournament $k = 3$, a four-round event with a cut after the second round. As with the first tournament, we determine player i's tournament skill level, ts_3, as the mean of the spline or linear model-based estimate of his skill level over the tournament rounds he actually played. Then his score for the first two rounds becomes, $ts_3\left(\tilde{\theta}_{i,15} + \tilde{\theta}_{i,16}\right)$. Computing the same for all tournament participants, we cut the field to the lowest scoring 70 players and ties as of the end of the second round. If player i makes the cut, his total tournament score becomes $ts_3\left(\tilde{\theta}_{i,15} + \tilde{\theta}_{i,16} + \tilde{\theta}_{i,17} + \tilde{\theta}_{i,18}\right)$. Computing the same for all who make the cut, we obtain simulated four-round scores and determine simulated finishing positions accordingly. If player i makes the cut, we use $\boldsymbol{\theta}$ error $\tilde{\theta}_{i,19}$ in connection with his next simulated score. Otherwise, we use $\tilde{\theta}_{i,17}$.

Following this procedure for each tournament, we note the total simulated score of the player who finishes the tournament in second place. We then compute the average per-round score associated with this total. For a given simulation trial, this becomes our estimate of the average neutral score per round required to win the tournament. Using (9), we standardize this average neutral score to a 4-round equivalent. We run 10,000 simulation trials of 2003-2009 PGA TOUR competition and compute the median of 10,000 second-place 4-round standardized simulated scores per round as our measure of relative tournament difficulty. (Berry (2008) uses a similar simulation procedure to estimate winning probabilities for Tiger Woods and other professional golfers over the 1997-2004 period and to estimate the probability that Woods will Jack Nicklaus' record of 18 major tournament wins. Grober (2008) also employs simulation to estimate the likelihood of Woods breaking Byron Nelson's record of winning eleven consecutive tournaments.)

4.4. Estimates of Relative Tournament Difficulty

Tables 2 and 3 list the 10 most difficult and least difficult tournaments, respectively, for each year, 2003-2009. The tables also show the ranking of each tournament based on the average estimated skill level of it participants at the time the tournament was played and indications of tournament type With just a few exceptions, THE PLAYERS Championship and four majors (marked with tournament type "M") occupy the top five positions in every year, and in all but one year, all are among the top 10. (Throughout, we will refer to a tournament in a low-numbered ranking position as being highly ranked and one in a high-numbered position as being low-ranked. This terminology is confusing, but it conforms with the actual usage of the terms.) Although not shown, the Masters is ranked as the 11th most difficult tournament to win in 2009. We note that the largest difference in the score required to win the tournaments in the first and tenth positions is 0.25 strokes per round (2004), which amounts to just one stroke over the course of a four-round event. Thus, winning the 10th place tournament would be equivalent to finishing no worse than a stroke behind the winner of THE PLAYERS Championship or a major. We note that the Bay Hill Invitational (subsequently renamed the Arnold Palmer Invitational), the Memorial Tournament, and the WGC-NEC Invitational (subsequently renamed the WGC-Bridgestone) appear frequently in the top 10, and in two years, the WGC event appears in the top five. (WCG events are marked as tournament type "W.") We also note that since the inception of The FedExCup competition in 2007, each of the first three FedExCup Playoffs events (The Barclays, Deutsche Bank Championship and BMW Championship, all marked as tournament type "F") make the top 10 list at least once. By contrast, THE TOUR Championship, also the Finals of The FedExCup Playoffs since 2007, has occupied positions 17, 31 and 13, since FedExCup competition began (not shown in the table), despite having a field ranked number one on the basis of the mean skill estimate of its participants as of the time of the competition (also not shown) in each year, 2003-2009. The relatively low difficulty ranking reflects that THE TOUR Championship is a small-field event, comprised of 30 players with the highest number of FedExCup points since 2007 and the 30 leading money winners prior to 2007. Thus, it can be more difficult to win a large strong-field event than a tournament, such as THE TOUR Championship, with a small but elite field. We note that in only one instance, the 2009 BMW Championship, is any tournament listed among the top 10 also ranked among the top two on the basis of the mean skill of its participating players.

The tournaments listed among the least difficult in Table 3 should come as no surprise to those who follow professional golf, with the possible exception of the Mercedes Championships ranked 39, 42 and 42 in 2006, 2008 and 2009, respectively, and marked as tournament type "S." Like THE TOUR Championship, The Mercedes Championships (now the Hyundai Tournament of Champions) attracts an elite but small field, consisting of tournament winners from the previous season, but the depth of the field is not sufficient to make this a particularly strong event.

With the exception of The Mercedes Championships, all of the tournaments among the bottom 10 are large-field events of approximately the same size. Many are alternative events held opposite (the same time as) the British Open and events in the World Golf Championships series (tournament type "A") and some are held opposite the Ryder Cup and President Cup (tournament type "R"). As a result, the TOUR's best players are playing elsewhere when these tournaments are being held, and, accordingly, these alternative events attract the weakest fields. Since most of these tournaments are approximately the same size, the mean skill of tournament participants should generally determine relative rankings among the bottom 10. For the most part, tournament difficulty rankings align very closely with tournament skill rankings for tournaments within this group.

We note that the estimated scores required to win the top 10 tournaments in 2003 are all lower than that of the top-ranked tournament in 2004. Despite our efforts to remove any age-related sample biases, this most likely reflects that our skill model produces estimated player skill levels in 2003 that are 0.22 strokes lower than those of 2004. It is entirely possible that players were 0.22 strokes per round better in 2003 than in 2004 and, hence, winning the 10th-ranked tournament in 2003 would have been more difficult than winning the top-ranked tournament in 2004. However, we are reluctant to make this claim and, therefore, do not make any year-to-year comparisons of relative tournament difficulty. Further refinements of our skill model might explain this difference, but we do not address this issue here.

Annual rank orderings of our relative tournament difficulty measure are unaffected by differences in year-to-year estimates of player skill. Therefore, to determine the relative difficulty ordering of the PGA TOUR events played over the 2003-2009 period, we compute the mean relative difficulty ranking for each tournament over this period. Table 4 lists all the PGA TOUR events played over the 2003-2009 period ordered by the mean of their year-to-year relative difficulty rankings. The table also shows the mean of the year-to-year tournament difficulty and field skill measures for each tournament, along with the percentage of the time Tiger Woods participated in the event while played over the 2003-2009 period and an indication of tournament type. We treat pre- and post FedExCup versions of The Barclays, Deutsche Bank Championship, The BMW Championship and TOUR Championship as distinctly different tournaments, since the nature of the competition was much different, especially for the first three playoff events.

THE PLAYERS Championship is the top overall ranked PGA TOUR event, followed by the PGA Championship, U.S. Open, British Open and Masters Tournament (the four majors, marked as tournament type "M"). Next in line are the FedExCup versions of The Barclays and Deutsche Bank Championship (tournament type "F"), the first two events in The FedExCup Playoffs. Note that the FedExCup Playoffs versions of the BMW Championship and TOUR Championship are much further down the list in positions 14 and 23, respectively. Although these two tournaments should have stronger fields than the first two FedExCup Playoffs events, their field sizes are smaller, with 70 and 30 players, respectively, making them less difficult to win than the first two playoff events, with respective fields sizes of 125 and 100.

The two events in the World Golf Championships series (tournament type "W") are both within the top 12 most difficult tournaments to win. The two select small field events, THE TOUR Championship and Mercedes Benz Championship (tournament type "S"), do not rank high on the list, because their small field sizes of approximately 30 players each make winning either of these tournaments less difficult than most larger field events. Eight of the tournaments shown in the last nine positions on the list are large-field events with relatively weak fields held opposite the British Open, events in the World Golf Championships series, the Ryder Cup or Presidents Cup (tournament types "A" and "R").

Note that Tiger Woods played regularly in almost all the events among the top 15, but played in none of the tournaments in positions 38-56. This helps to explain why Woods' mean neutral score in tournaments won is significantly below that of other players – he needs to score better to win, because he tends to participate in the most challenging events.

4.5. Determinants of Relative Tournament Difficulty

Table 5 summarizes the results of an OLS regression of the median score required to win over all 321 tournaments as a function of the mean skill level of tournament participants, the standard deviation of the participant skill levels, the skewness of the skill levels, and the number of tournament participants. Each independent variable is highly significant, and together, they explain a large portion of the variation in our measure of relative tournament difficulty (adjusted $R^2 = 0.9667$). With an estimated regression coefficient of 1.074 associated with mean tournament skill, the median score required to win a given tournament increases by approximately one stroke per unit increase in the mean predicted score of its participants. The score required to win decreases by 0.940 strokes per unit increase in the standard deviation of tournament participant skill and increases by 0.06 strokes per unit increase in the skewness of participant skill. Each unit increase in the number of tournament participants reduces the score required to win by 0.008 strokes. We note that if all of the variables associated with the regression except skewness are expressed in log form, the adjusted

R^2 increases to 0.9891.

4.6. Relationship to OWGR Strength of Field Measure

To help validate our measure of relative tournament difficulty, we obtained end-of-year strength of field data from the PGA TOUR used in the computation of Official World Golf Rankings (OWRG) for years 2003-2009. The OWGR measure does not produce a score required to win, as does our measure, but instead assigns points to each tournament based on the number of participating players ranked in the top 5, 15, 30, 50, 100 and 200, respectively in the World Golf Rankings and in the top 5, 15 and 30, respectively, on the TOUR that conducts the tournament (in our case, the PGA TOUR). Total assigned points range between zero and 1,000 and have no golf-related interpretation other than the tournaments with the highest number of total points are deemed to have had the strongest fields. (See PGA TOUR Communications (2010)).

Since the QWGR measure reflects both the number of golfers in a particular event and the quality of the players, events ranked the highest on the QWGR strength of field measure should be among the most difficult to win. We point out that our measure of relative tournament difficulty and the OWGR strength of field measure are not directly comparable, since they are expressed in entirely different units. Moreover, our measure employs data for the entire 2003-2009 period to estimate the score required to win each tournament over the same time frame, whereas the OWGR method uses data available only at the time the strength of field calculation is made and updates this data on a rolling basis throughout each PGA TOUR season. Nevertheless, we would be concerned if the two measures did not produce similar rankings.

For each year, 2003-2009, we compute Spearman rank correlations and Kendall's tau to estimate the degree of association between the rankings of our measure of tournament difficulty and those of the OWGR's measure of field strength. Spearman rank correlations, which are simply correlations of the two sets of rankings, range between 0.893 and 0.977, and Kendall's tau falls between 0.736 and 0.884. Because of ties in the WRG measure, we are unable to compute significance levels associated with the correlation estimates for any year but 2004. However, for 2004, p-values associated with both rank correlation estimates against a null hypothesis of zero are essentially zero. Thus, the association between the two measures of relative tournament difficulty appears to be sufficiently high to suggest that they produce very similar rankings.

5. Estimates of Winning Probabilities?

5.1. How Difficult is it for Actual Tournament Winners to Win?

In the previous section we describe simulations designed to determine the relative difficulty of the various tournaments on the PGA TOUR. For each of the 10,000 simulation trials, we keep a running count of the number of times each player wins a given tournament. The player's total simulated wins for the tournament, divided by 10,000, provides an estimate of the probability he would have won.

We compiled a list of each of the 321 tournaments showing our estimated probabilities that the actual tournament winners would win, ordered by estimated winning probabilities. Table 6 highlights portions of this list. The top and bottom sections of Table 6 show the ten tournaments for which the probabilities that the actual winner would win are the lowest and highest, respectively. The three middle sections of the table show 10 tournaments each from the middle of the second through fourth quintiles of our estimated probability rankings.

Among those in the top section of the list, which shows the least likely winners, are two notorious unlikely winners of major championship, Ben Curtis, who won the 2003 British Open and Shaun Micheel, who won the 2003 PGA Championship. We estimate the probability associated with Curtis winning as only 0.09%, while the estimated probability of Micheel winning is slightly higher, 0.2%. We also estimate a probability of 0.21% associated with Heath Slocum winning the 2009 Barclays, often mentioned by golf writers and commentators as a very unlikely win.

Interestingly, a number of 'big-name' players are shown in the middle sections of the table, especially in the fourth section, which includes notable winners such as Phil Mickelson, Ernie Els and Jim Furyk. Even in this section, the estimated probabilities of winning are relatively low – on the order of 5%.

Not surprisingly, Tiger Woods is the only player listed at the bottom of the table, where the estimated probabilities of winning are the highest. Among the 10 tournaments shown, Tiger's estimated winning probabilities range from approximately 34% to 52%. Over all tournaments that he won, the mean of Woods' estimated winning probabilities is 29.10%, compared with 13.44%, 9.24%, 7.09% and 6.30% for Vijay Singh, Ernie Els, Jim Furyk, and Phil Mickelson, respectively. (None of these mean probabilities are shown in the table.)

5.2. How Likely are Major Tournament Wins?

Using the results from the same simulations, we estimate the probabilities associated with each of the five top players (Tiger Woods, Vijay Singh, Jim Furyk, Ernie Els, and Phil Mickelson) winning the four majors and THE PLAYERS Championship over the 2003-2009 period. Table 7 summarizes the estimated winning probabilities. Except during 2003 and 2004, when the estimated probabilities for Tiger Woods and Vijay Singh are of the same order of magnitude, approximately 12%, Woods' estimated winning probabilities are far greater than those of the other four players. (The estimated probabilities of Woods winning in 2003 and 2004 are very close to the 13.5% probability that Woods would win a tournament in 2004 comprised of the top 144 players as estimated by Berry (2008).) Woods' estimated winning rate of 16% to 36% over the last five years, 2005 to 2009, is consistent with his actual winning rate of 27% over the same period. We note that Woods' estimated winning rates for 2009 are more than 10 times those of any among the other four players.

Comparing the skill profiles of each player, as shown in Figure 1, to the estimated winning probabilities, as shown in Table 7, one can see that Woods' estimated winning probabilities are increasing over the 2005-2009 period while his skill is improving, while those of Singh and Els are decreasing as their skill levels are deteriorating. The irregular pattern of estimated winning probabilities for Furyk reflects the irregular pattern in his predicted scoring estimates, while the relatively flat pattern of estimated winning probabilities for Mickelson reflects his relatively flat estimated skill function.

6. Tournament Success when Playing Normal

In this section we explore the extent to which players could have won tournaments by playing "normal," where, again, playing normal is defined as recording a score with a zero $\boldsymbol{\eta}$ error. To compute such a score, we simply take a player's actual score and subtract the $\boldsymbol{\eta}$ error associated with the same score. For example, if a player shoots a 67 in a round in which his estimated $\boldsymbol{\eta}$ error is -2.53 strokes, his "normal" score would be $67 - (-2.53) = 69.53$. In this case, we would recompute the player's place in tournament competition, assuming that he recorded a score of 69.53, with scores in his other rounds adjusted in similar fashion, while using actual scores for all other players in the competition. If the player misses a cut, we do not compute the score he would have recorded by playing normal, since we have no $\boldsymbol{\eta}$ errors for rounds in which he did not participate.

Figure 3 shows neutral scores around Tiger Woods' predicted neutral score, including the autocorrelated component, in terms of tournaments that he actually won and lost and those he could have won or lost by playing normal, with neutral scores broken into the following four categories:

- Scores in tournaments that Tiger lost and would have lost if he had played normal, marked with a small dot.
- Scores in tournaments that Tiger won and would have also won if he had played normal, marked with a "+" sign.
- Scores in tournaments that Tiger won but would have lost if he had played normal, marked with an asterisk.
- Scores in tournaments that Tiger lost but would have won if he had played normal, marked with a triangle.

Throughout, we assume, implicitly, that Tiger's playing normal would not have affected the scores of other players against whom he was competing.

We estimate that Tiger Woods could have won 13 tournaments over the 2003-2009 period by playing normal, but none of the other top five players shown in Figure 1, and by assumption, no one among the remaining 648, could have won a single tournament by playing to his norm. Two of the 13 tournaments that Tiger could have won are tournaments in which he did not finish in first place but could have finished in first by playing his normal game. Except for our estimate that Woods could have won some tournaments by playing normal, these results are consistent with those of Berry (2008), based on a 1997-2004 sample, who states "Woods is so much better than the other players that he is the major determination on whether he wins the tournament. If he plays well above his ability, by random variation, he wins the tournament. This corresponds to him winning tournaments by huge margins, as he has done. When he plays poorly ... other players will beat him. If he plays a little better than average, then he may be beat by another player, and may finish second." The results are also consistent with those of Connolly and Rendleman (2009), which would have shown Woods winning five tournaments he actually lost over the 1998-2001 period if corrected for a slight programming error.

7. Conclusions

In this paper we shed light on what it takes for Tiger Woods and others who play golf at the highest level to win on the PGA TOUR. How much 'luck' does Tiger need to win? To what extent does relative tournament difficulty come into play in determining tournament winners? How likely (or unlikely) were the wins of actual tournament winners? Can Tiger, or any other players on the PGA TOUR, win by simply playing their 'normal' games?

Our results make it clear that measured performance outcomes in sports, business and other settings reflect both the skill and luck components of individuals and the collective skill and luck among those against whom individuals are competing. In many cases, it is not enough to be the most highly skilled if luck, or random variation in performance, among competitors plays a large role in determining performance outcomes. Even Tiger Woods is not sufficiently skilled to win most tournaments in which he competes without a little luck of his own or help from his competitors.

Using a sample that covers the 2003-2009 PGA TOUR seasons, we find that on average, the winner of a PGA TOUR event experiences approximately 2.5 strokes per round of favorable random variation in scoring. However, when Tiger Woods wins, he wins by having experienced approximately 1.15 strokes per round of favorable random variation in his score. Among players who recorded 20 or more winning scores from 2003 to 2009, the player with the next smallest favorable random variation is Vijay Singh (-1.63), followed by Mike Weir (-1.72), Jim Furyk (-1.88), Ernie Els (-2.06) and Phil Mickelson (-2.27).

We estimate that when Tiger wins, he wins by scoring approximately 0.71 strokes per round lower than other winning players, after adjusting for the relative difficultly of each round and the extent to which a course may have played favorably or unfavorably for the player. This difference appears to reflect two things. First, Tiger may simply play better when he wins compared with others who win. Second, Tiger tends to play in tournaments that are more difficult to win, which, in turn, require lower scores to win compared with tournaments in which he does not participate.

To make this assessment, we develop a new and novel estimate of relative tournament difficulty. Using Monte Carlo simulation, we estimate the mean score per round that it takes to win a PGA TOUR event as a function of the number of players participating in the tournament, the number of tournament rounds, the mean skill levels of the tournament participants and their natural random variation in scoring. This measure not only sheds light on what it takes for Woods to win, but it also provides a means to compare the relative difficulty of tournaments played on the PGA TOUR. Generally, we find THE PLAYERS Championship to be the most difficult tournament to win, followed by the four majors. The first three events in the FedExCup Playoffs also rank very high in tournament difficulty as well as the tournaments in the World Golf Championships series. Despite being fielded by the TOUR's most elite players, the two small select-field events, the Mercedes Championships and THE TOUR Championship, are not among the most difficult to win. The reason – in larger-field, less selective events, it is more likely that any given player could string together four successive rounds with favorable random variation in scoring (i.e., good luck). Therefore, to win such a tournament, a player must play with sufficient skill and luck himself to overcome the greater potential in a large-field event for any one player or group of players to 'go low.'

We also estimate the probabilities that tournament winners would have won the tournaments they actually won. Except for Tiger Woods, whose estimated winning probabilities in tournaments won are on the order of 30%, the probability of winning among winning players is relatively low.

Finally, we explore the extent to which players could have won tournaments by playing "normal," where normal is defined as recording a score with no favorable or unfavorable player-specific random variation in scoring. We estimate that Tiger Woods could have won 13 tournaments over the 2003-2009 period by playing normal. Two of the 13 tournaments that Tiger could have won are tournaments he actually lost. All other players would have needed some favorable random variation in scoring to win.

To re-state the obvious, Tiger Woods has had no peer during his era as a professional golfer. In this paper, we are able to illustrate his dominance in not-so-obvious ways and shed light on what it takes to win, for both Tiger himself and others on the TOUR.

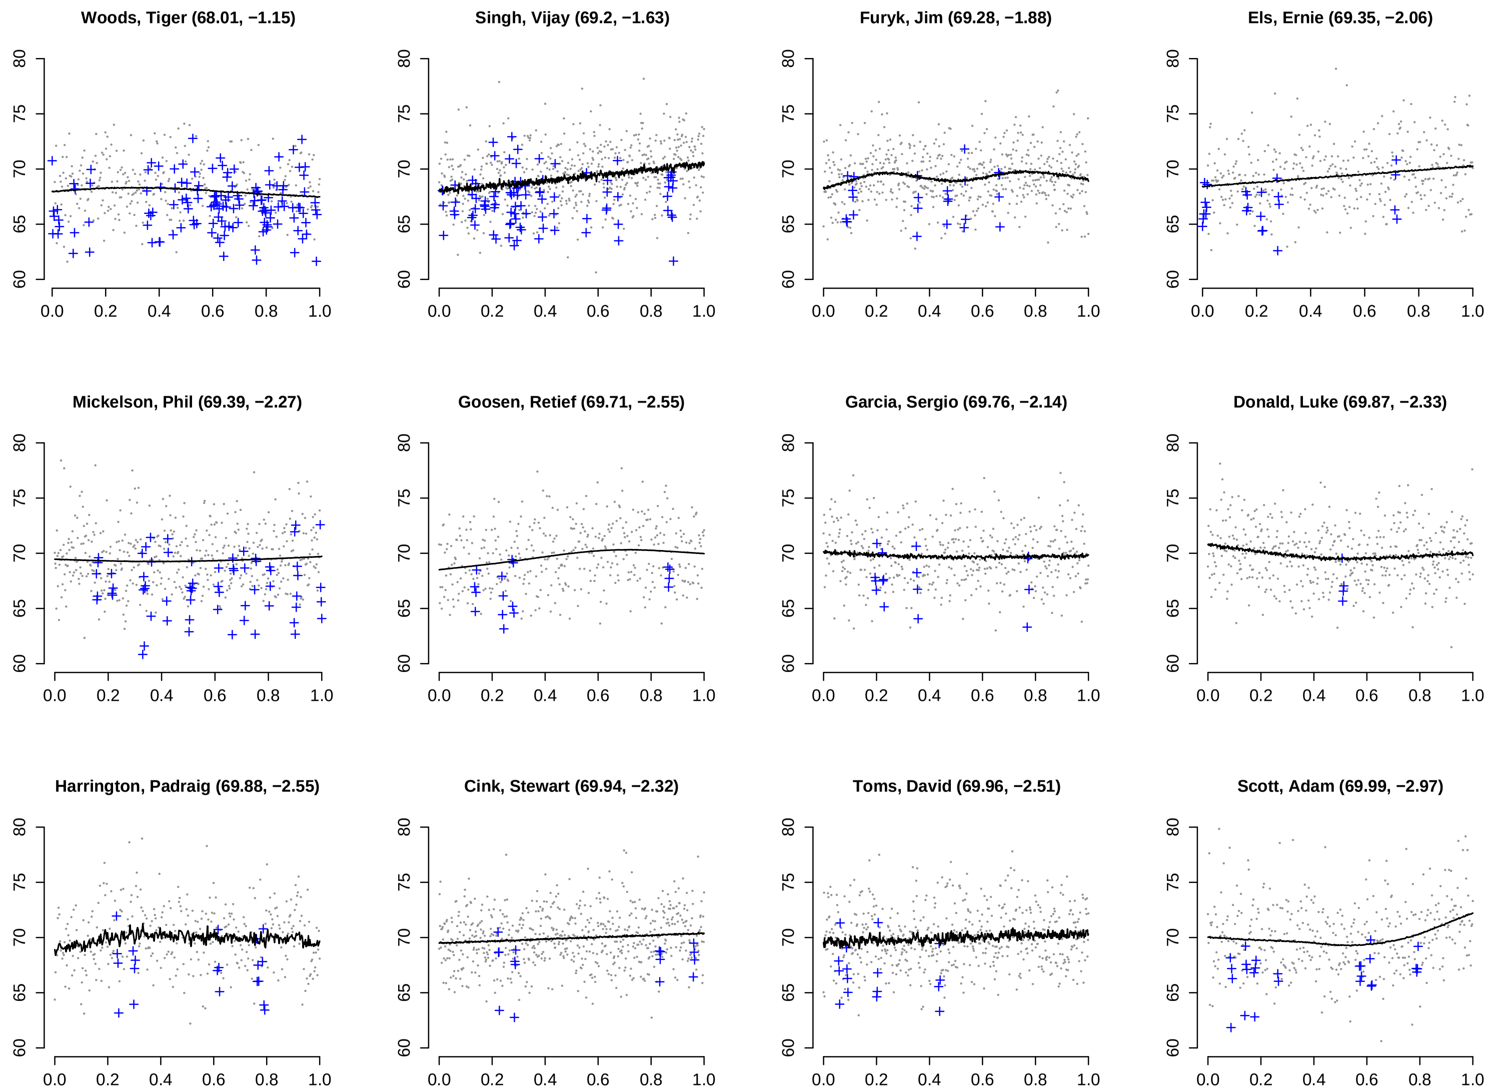

Figure 1: Round-Course/Player-Course-Neutral Scores, 2003-2009, for 12 Most Highly-Skilled Players.

In the individual player plots, the first number shown to the right of the player's name is the average of predicted scores from his estimated skill function. The second number is the average value of the player's $\boldsymbol{\eta}$ residuals (or the difference between his neutral scores and predicted scores after taking account of potential autocorrelation) in tournaments he won. "+" symbols denote neutral scores in tournaments the player won. The line represents the player's predicted neutral score, including the autocorrelated component. The X-axis reflects the numerical sequencing of a players's scores, scaled to the [0, 1] interval.

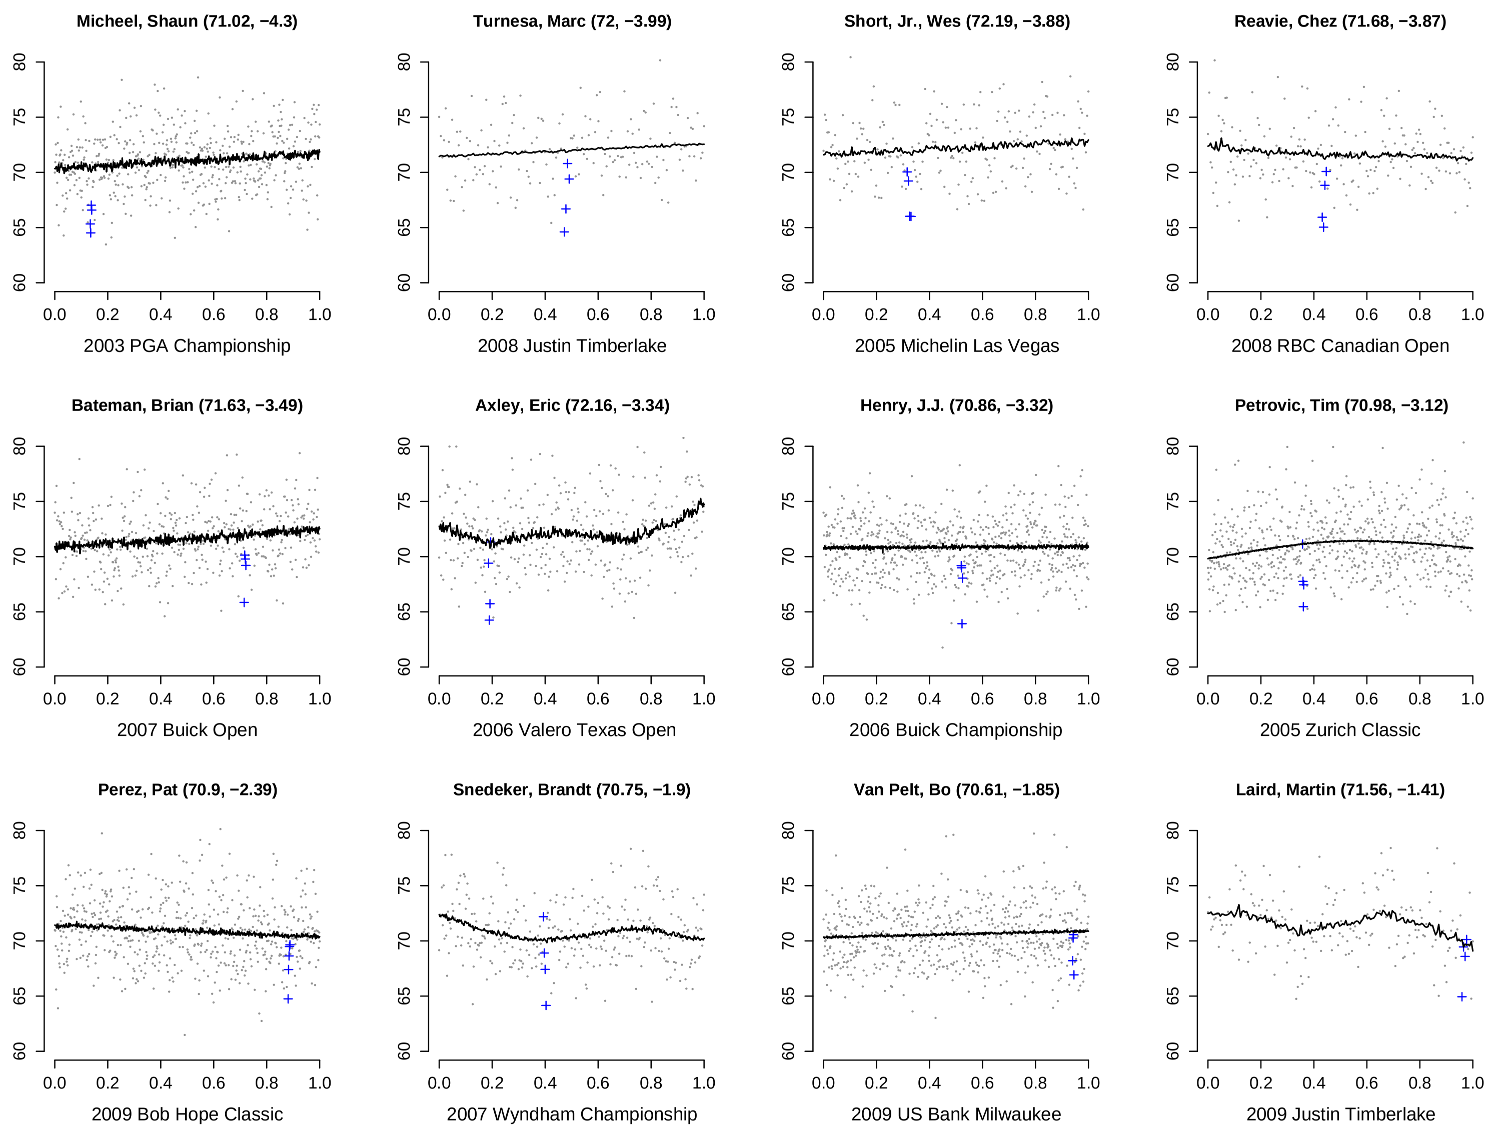

Figure 2: Round-Course/Player-Course-Neutral Scores, 2003-2009, for "One-Hit Wonders".

In the individual player plots, the first number shown to the right of the player's name is the average of predicted scores from his estimated skill function. The second number is the average value of the player's $\boldsymbol{\eta}$ residuals (or the difference between his neutral scores and predicted scores after taking account of potential autocorrelation) in tournaments he won. "+" symbols denote neutral scores in tournaments the player won. The line represents the player's predicted neutral score, including the autocorrelated component. The X-axis reflects the numerical sequencing of a players's scores, scaled to the [0, 1] interval.

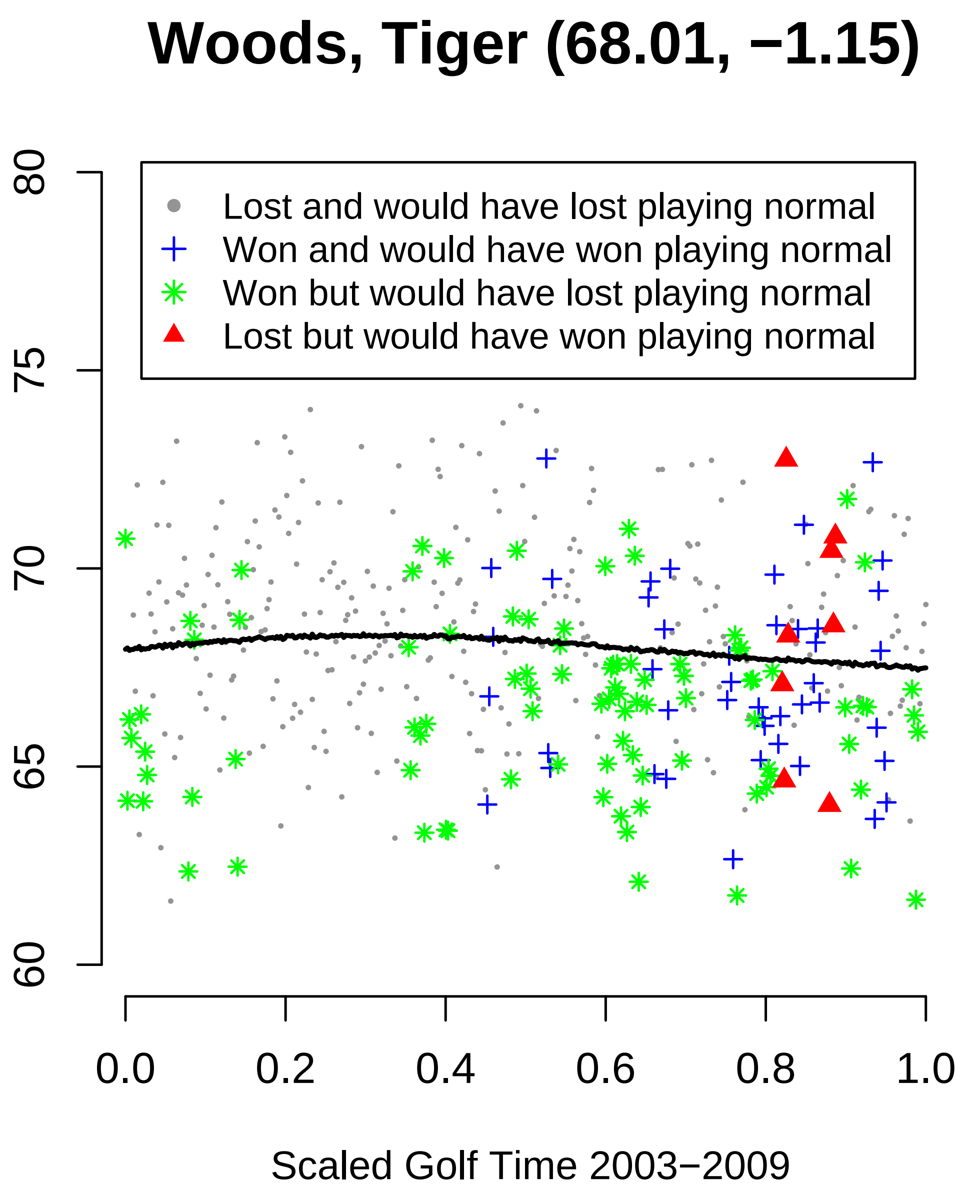

Figure 3: Round-Course/Player-Course-Neutral Scores for Tiger Woods, 2003-2009.

The first number shown to the right of Woods' name is the average of predicted scores from his estimated skill function. The second number is the average value of his η residuals (or the difference between his neutral scores and predicted scores after taking account of potential autocorrelation) in tournaments he won. The line represents Woods' predicted neutral score, including the autocorrelated component. "Scaled Golf Time" reflects the numerical sequencing of Woods' scores, scaled to the [0, 1] interval.

Table 1: Summary Statistics for Winning Players, 2003-2009

Scoring Measure	Players	Mean
η residual when winning	Tiger	-1.148*
η residual when winning	All	-2.507
η residual when winning	Others	-2.668
Raw score when winning	Tiger	67.581
Raw score when winning	Others	67.484
Neutral score when winning	Tiger	66.735*
Neutral score when winning	Others	67.445
Round-course effect when winning	Tiger	0.846*
Round-course effect when winning	Others	0.039
Player-course effect when winning	Tiger	-0.024*
Player-course effect when winning	Others	-0.042

*Difference between Tiger's mean and the mean for all other players is statistically significant at the 0.001 level. We note that although estimated round-course effects sum to zero, and the values in this table reflect all tournament rounds, they do not reflect all round-course interactions, since some tournament rounds are played on multiple courses. Therefore, it is entirely possible that the mean of the round-course effect associated with Tiger Woods' wins and that associated with all others in the sample can both be of the same sign.

Table 2: Most Difficult Tournaments to Win, 2003-2009, Rankings Based on Median Simulated Second-Place Finish

	Tournament	Median Simulated Second Place Finish Score	Rank (Yr.)	Mean Skill Score	Rank (Yr.)	n Players	Tourn Type
2003	THE PLAYERS Championship	66.85	1	70.55	5	143	
2003	Masters Tournament	66.89	2	70.63	7	87	M
2003	U.S. Open Championship	66.92	3	70.76	15	139	M
2003	British Open Championship	66.93	4	70.63	8	124	M
2003	PGA Championship	66.97	5	70.68	11	136	M
2003	Nissan Open	66.98	6	70.67	10	139	
2003	Phoenix Open	67.04	7	70.66	9	128	
2003	the Memorial Tournament	67.05	8	70.57	6	102	
2003	WGC-NEC Invitational	67.05	9	70.39	4	81	W
2003	Bay Hill Invitational	67.09	10	70.74	14	118	
2004	THE PLAYERS Championship	67.15	1	70.70	4	147	
2004	Masters Tournament	67.24	2	70.75	6	88	M
2004	British Open Championship	67.24	3	70.89	10	119	M
2004	U.S. Open Championship	67.26	4	70.94	14	138	M
2004	PGA Championship	67.27	5	70.90	11	136	M
2004	Bay Hill Invitational	67.28	6	70.82	9	117	
2004	Nissan Open	67.34	7	70.91	12	141	
2004	the Memorial Tournament	67.36	8	70.76	7	105	
2004	WGC-NEC Invitational	67.38	9	70.54	3	74	W
2004	FBR Open	67.40	10	70.81	8	129	
2005	THE PLAYERS Championship	67.33	1	70.86	5	146	
2005	PGA Championship	67.37	2	71.05	12	138	M
2005	U.S. Open Championship	67.40	3	71.20	22	141	M
2005	British Open Championship	67.40	4	71.14	20	124	M
2005	Masters Tournament	67.40	5	70.90	7	87	M
2005	WGC-NEC Invitational	67.48	6	70.58	3	69	W
2005	Bay Hill Invitational	67.48	7	70.97	8	117	
2005	Ford Championship at Doral	67.49	8	71.05	13	139	
2005	EDS Byron Nelson	67.52	9	71.14	17	148	
2005	the Memorial Tournament	67.52	10	70.88	6	105	
2006	THE PLAYERS Championship	67.36	1	70.92	5	144	
2006	PGA Championship	67.37	2	71.10	10	146	M
2006	U.S. Open Championship	67.39	3	71.14	11	135	M
2006	British Open Championship	67.41	4	71.06	8	126	M
2006	Masters Tournament	67.43	5	70.93	6	85	M
2006	WGC-Bridgestone Invitational	67.45	6	70.66	3	76	W
2006	Bay Hill Invitational	67.52	7	71.06	9	118	
2006	Ford Championship at Doral	67.52	8	71.18	14	139	
2006	Nissan Open	67.53	9	71.16	12	142	
2006	Cialis Western Open	67.55	10	71.34	21	153	
2007	PGA Championship	67.39	1	71.09	15	141	M
2007	THE PLAYERS Championship	67.41	2	70.92	8	142	
2007	Wachovia Championship	67.43	3	71.08	12	150	
2007	U.S. Open Championship	67.45	4	71.11	17	130	M
2007	Masters Tournament	67.47	5	71.02	10	91	M
2007	Deutsche Bank (FedEx)	67.47	6	70.78	5	115	F
2007	British Open Championship	67.47	7	71.15	18	123	M
2007	WGC-Bridgestone Invitational	67.52	8	70.80	6	83	
2007	WGC-CA Championship	67.54	9	70.77	4	72	W
2007	Arnold Palmer Invitational	67.57	10	71.08	14	118	
2008	U.S. Open Championship	67.46	1	71.10	11	127	M
2008	Masters Tournament	67.48	2	71.12	13	94	M
2008	British Open Championship	67.56	3	71.23	20	118	M
2008	WGC-CA Championship	67.57	4	70.92	7	77	W
2008	PGA Championship	67.57	5	71.14	15	141	M
2008	THE PLAYERS Championship	67.58	6	70.93	8	141	
2008	The Barclays (FedEx)	67.59	7	70.91	6	135	F
2008	Arnold Palmer Invitational	67.61	8	71.10	12	119	
2008	Buick Invitational	67.62	9	71.25	22	150	
2008	Northern Trust Open	67.63	10	71.09	10	140	
2009	British Open Championship	67.16	1	70.98	10	124	M
2009	PGA Championship	67.17	2	70.97	9	143	M
2009	The Barclays (FedEx)	67.29	3	70.72	6	124	F
2009	WGC-Bridgestone Invitational	67.30	4	70.58	4	76	W
2009	Deutsche Bank (FedEx)	67.31	5	70.55	3	99	F
2009	U.S. Open Championship	67.32	6	71.04	13	121	M
2009	THE PLAYERS Championship	67.33	7	70.84	7	143	
2009	BMW (FedEx)	67.35	8	70.33	2	69	F
2009	WGC-CA Championship	67.40	9	70.63	5	78	W
2009	the Memorial Tournament	67.40	10	70.99	11	117	

Tournament type codes: M = major; W = WGC event; F = FedExCup Playoffs event.
Tournament difficulty estimates based on 10,000 simulation trials.

Table 3: Least Difficult Tournaments to Win, 2003-2009, Rankings Based on Median Simulated Second-Place Finish

		Median Simulated Second Place Finish		Mean Skill			Tourn
	Tournament	Score	Rank (Yr.)	Score	Rank (Yr.)	n Players	Type
2003	FedEx St. Jude Classic	67.71	37	71.27	35	145	
2003	Chrysler Classic of Tucson	67.72	38	71.26	34	137	A
2003	Bell Canadian Open	67.73	39	71.61	43	147	
2003	Valero Texas Open	67.74	40	71.27	36	139	
2003	Greater Hartford Open	67.79	41	71.46	40	148	
2003	84 Lumber Classic	67.90	42	71.46	41	140	
2003	Greater Milwaukee Open	67.91	43	71.55	42	144	
2003	Reno-Tahoe Open	68.14	44	71.73	44	129	A
2003	Southern Farm Bureau Classic	68.17	45	71.74	45	126	A
2003	B.C. Open	68.35	46	71.99	46	125	A
2004	John Deere Classic	67.96	37	71.68	39	148	
2004	Booz Allen Classic	67.98	38	71.55	35	151	
2004	Buick Championship	68.07	39	71.56	36	148	
2004	FedEx St. Jude Classic	68.09	40	71.69	40	146	
2004	U.S. Bank Championship	68.12	41	71.69	41	152	
2004	Valero Texas Open	68.17	42	71.85	44	142	R
2004	Chrysler Classic of Tucson	68.21	43	71.70	42	139	A
2004	Southern Farm Bureau Classic	68.36	44	71.85	45	129	A
2004	Reno-Tahoe Open	68.39	45	71.83	43	130	A
2004	B.C. Open	68.68	46	72.27	46	119	A
2005	Chrysler Classic of Greensboro	68.03	37	71.44	29	130	
2005	FedEx St. Jude Classic	68.09	38	71.67	39	147	
2005	Buick Championship	68.10	39	71.56	36	148	
2005	Valero Texas Open	68.18	40	71.69	40	141	R
2005	U.S. Bank Championship	68.22	41	71.76	43	148	
2005	John Deere Classic	68.22	42	71.73	42	146	
2005	Chrysler Classic of Tucson	68.27	43	71.77	44	138	A
2005	Southern Farm Bureau Classic	68.33	44	71.72	41	132	A
2005	Reno-Tahoe Open	68.41	45	71.81	45	130	A
2005	B.C. Open	68.61	46	72.18	46	125	A
2006	Booz Allen Classic	68.19	37	71.67	38	150	
2006	U.S. Bank Championship	68.22	38	71.78	39	148	
2006	Mercedes Championships	68.24	39	70.60	2	28	S
2006	FedEx St. Jude Classic	68.25	40	71.80	40	146	
2006	John Deere Classic	68.28	41	71.84	41	148	
2006	Valero Texas Open	68.34	42	72.03	44	144	R
2006	Chrysler Classic of Tucson	68.35	43	71.84	42	138	A
2006	Southern Farm Bureau Classic	68.47	44	71.99	43	130	A
2006	Reno-Tahoe Open	68.53	45	72.10	45	131	A
2006	B.C. Open	68.81	46	72.56	46	116	A
2007	John Deere Classic	68.24	37	71.85	40	150	
2007	Turning Stone Resort Championship	68.26	38	71.62	33	138	
2007	Frys.com Open	68.30	39	71.67	35	143	
2007	Wyndham Championship	68.33	40	71.85	41	152	
2007	U.S. Bank Championship	68.38	41	71.78	39	129	A
2007	Ginn sur Mer Classic at Tesoro	68.38	42	71.87	43	130	
2007	Valero Texas Open	68.40	43	71.87	42	131	
2007	Viking Classic	68.49	44	71.91	44	131	R
2007	Mayakoba Golf Classic	68.50	45	72.07	46	134	A
2007	Reno-Tahoe Open	68.51	46	71.97	45	131	A
2008	Children's Miracle Network	68.16	38	71.56	30	127	
2008	Turning Stone Resort Championship	68.24	39	71.63	34	130	
2008	AT&T Classic	68.28	40	71.93	43	149	
2008	Ginn sur Mer Classic	68.29	41	71.70	39	128	
2008	Mercedes-Benz Championship	68.31	42	70.66	3	31	S
2008	Viking Classic	68.43	43	72.04	44	141	R
2008	U.S. Bank Championship	68.44	44	71.79	41	126	A
2008	Legends Reno-Tahoe Open	68.64	45	72.43	47	128	A
2008	Mayakoba Golf Classic	68.67	46	72.06	45	127	A
2008	Puerto Rico Open	68.70	47	72.17	46	124	A
2009	Canadian Open	68.01	35	71.67	35	140	
2009	Wyndham Championship	68.04	36	71.61	33	150	
2009	Turning Stone Resort Championship	68.05	37	71.55	30	127	
2009	Children's Miracle Network Classic	68.12	38	71.68	37	127	
2009	Valero Texas Open	68.24	39	72.09	43	150	
2009	Legends Reno-Tahoe Open	68.32	40	71.95	42	130	A
2009	Mayakoba Golf Classic	68.35	41	71.77	40	126	A
2009	Mercedes-Benz Championship	68.39	42	70.92	8	33	S
2009	U.S. Bank Championship	68.46	43	71.90	41	125	A
2009	Puerto Rico Open	68.63	44	72.14	44	116	A

Tournament type codes: S = select small-field event; A = alternative event held opposite the British Open or WGC event; R = event held opposite Ryder Cup or Presidents Cup. Tournament difficulty estimates based on 10,000 simulation trials.

Table 4: Tournaments Ordered by Mean Annual Difficulty Rank

Overall Rank	Tournament	Mean Annual Difficulty Rank	Mean Difficulty	Mean Skill	% Tiger Played	Tournament Type
1	THE PLAYERS Championship	2.71	67.29	70.82	85.7	
2	PGA Championship	3.14	67.30	70.99	85.7	M
3	U.S. Open Championship	3.43	67.31	71.04	100.0	M
4	British Open Championship	3.71	67.31	71.01	85.7	M
5	Masters Tournament	4.57	67.33	70.91	100.0	M
6	The Barclays (FedEx)	7.33	67.49	70.84	33.3	F
7	Deutsche Bank (FedEx)	7.33	67.47	70.71	66.7	F
8	WGC-Bridgestone Invitational	7.86	67.42	70.63	85.7	W
9	Arnold Palmer Invitational	8.86	67.45	70.98	100.0	
10	the Memorial Tournament	10.57	67.47	70.90	71.4	
11	Northern Trust Open	10.71	67.48	71.01	57.1	
12	WGC-CA Championship	10.71	67.47	70.69	100.0	W
13	Quail Hollow Championship	11.29	67.47	71.11	57.1	
14	BMW (FedEx)	12.33	67.58	70.50	66.7	F
15	Ford Championship at Doral	14.25	67.48	71.08	50.0	
16	FBR Open	14.71	67.59	71.00	0.0	
17	Buick Invitational	16.29	67.60	71.28	85.7	
18	Cialis Western Open	16.75	67.52	71.25	100.0	
19	AT&T National	17.67	67.74	71.14	66.7	
20	Barclays Classic	18.50	67.55	71.22	25.0	
21	Sony Open in Hawaii	18.71	67.66	71.23	0.0	
22	HP Byron Nelson Championship	19.00	67.68	71.27	28.6	
23	TOUR Championship (FedEx)	20.33	67.82	70.22	66.7	F, S
24	Transitions Championship	21.43	67.73	71.17	0.0	
25	Verizon Heritage	22.00	67.74	71.14	0.0	
26	AT&T Pebble Beach	22.14	67.74	71.62	0.0	
27	Crowne Plaza Invitational	23.29	67.74	71.06	0.0	
28	THE TOUR Championship	24.75	67.69	70.08	75.0	S
29	Shell Houston Open	25.14	67.78	71.38	0.0	
30	The Bob Hope Classic	25.71	67.81	71.18	0.0	
31	The Honda Classic	25.71	67.81	71.27	0.0	
32	Deutsche Bank Championship	27.25	67.72	71.45	100.0	
33	Buick Open	27.71	67.85	71.54	71.4	
34	Booz Allen Classic	30.00	67.83	71.39	0.0	
35	Children's Miracle Network Classic	30.14	67.89	71.43	28.6	
36	Zurich Classic of New Orleans	30.43	67.89	71.48	0.0	
37	Mercedes-Benz Championship	30.71	67.94	70.51	28.6	S
38	Frys.com Open	31.00	68.04	71.53	0.0	
39	Justin Timberlake	31.86	67.97	71.45	0.0	
40	AT&T Classic	33.00	67.94	71.51	0.0	
41	Travelers Championship	33.43	68.00	71.54	0.0	
42	St. Jude Classic	33.71	68.01	71.63	0.0	
43	84 LUMBER Classic	33.75	67.90	71.46	0.0	
44	Canadian Open	33.86	67.98	71.64	0.0	
45	Wyndham Championship	36.00	68.04	71.57	0.0	
46	John Deere Classic	37.14	68.07	71.67	0.0	
47	Turning Stone Resort Championship	38.00	68.19	71.60	0.0	
48	Valero Texas Open	40.43	68.18	71.78	0.0	R (04-06)
49	Ginn sur Mer Classic	41.50	68.34	71.79	0.0	
50	U.S. Bank Championship	41.57	68.25	71.75	0.0	A (07-09)
51	Chrysler Classic of Tucson	41.75	68.14	71.64	0.0	A
52	Viking Classic	44.00	68.37	71.87	0.0	A (03-06) R (07-08)
53	Mayakoba Golf Classic	44.00	68.51	71.97	0.0	A
54	Legends Reno-Tahoe Open	44.29	68.42	71.98	0.0	A
55	Puerto Rico Open	45.50	68.67	72.15	0.0	A
56	B.C. Open	46.00	68.61	72.25	0.0	A

The "difficulty" of a given tournament is the median simulated neutral score per round of the second place finisher. The "% Tiger Played" column gives the percentage of times in which a tournament was played over the 2003-2009 period that Tiger Woods participated. Many of the tournaments listed in this table were subject to name changes over the 2003-2009 period. The tournament names shown are the most recent names as of the end of the 2009 PGA TOUR season. Tournament type codes: M = major; W = WGC event; F = FedExCup Playoffs event; S = select small-field event; A = alternative event held opposite the British Open or WGC event; R = event held opposite Ryder Cup or Presidents Cup. (XX-YY) denotes that the designation applies only in years 20XX through 20YY, and, therefore, there is at least one other year when the designation does not apply. Tournament difficulty estimates based on 10,000 simulation trials.

Table 5: OLS Regression to Predict Median Score Required to Win

	Coef	Std error	t-stat
Intercept	-6.9132	0.8258	-8.372
Mean skill	1.0739	0.0118	90.942
Standard deviation of skill	-0.9401	0.0284	-33.080
Skill skewness	0.0601	0.0065	9.314
Number of players	-0.0080	0.0002	-46.973

Adjusted $R^2 = 0.9667$. Number of observations = 321. The median score required to win is estimated from simulations using 10,000 trials per tournament over the 2003-2009 PGA TOUR seasons.

Table 6: Estimated Probabilities of Winning for Tournament Winners

Rank	Tournament		Probability of Winning	Winner
1	2004	Michelin Championship	0.0000	Stolz, Andre
2	2007	Buick Open	0.0003	Bateman, Brian
3	2003	Ford Championship at Doral	0.0006	Hoch, Scott
4	2003	British Open Championship	0.0009	Curtis, Ben
5	2004	Wachovia Championship	0.0009	Sindelar, Joey
6	2008	Justin Timberlake	0.0010	Turnesa, Marc
7	2005	Michelin Championship	0.0015	Short, Jr., Wes
8	2008	Shell Houston Open	0.0016	Wagner, Johnson
9	2003	PGA Championship	0.0020	Micheel, Shaun
10	2009	The Barclays	0.0021	Slocum, Heath
92	2008	Mercedes-Benz Championship	0.0132	Chopra, Daniel
93	2009	Wyndham Championship	0.0137	Moore, Ryan
94	2003	84 Lumber Classic	0.0139	Lewis, J.L.
95	2007	Ginn sur Mer Classic at Tesoro	0.0142	Chopra, Daniel
96	2006	John Deere Classic	0.0146	Senden, John
97	2007	the Memorial Tournament	0.0150	Choi, K.J.
98	2007	AT&T National	0.0150	Choi, K.J.
99	2007	Verizon Heritage	0.0151	Weekley, Boo
100	2003	HP Classic of New Orleans	0.0153	Flesch, Steve
101	2007	Mayakoba Golf Classic	0.0153	Funk, Fred
156	2005	Mercedes Championships	0.0249	Appleby, Stuart
157	2004	MCI Heritage	0.0256	Cink, Stewart
158	2004	U.S. Bank Championship	0.0256	Franco, Carlos
159	2003	The Honda Classic	0.0257	Leonard, Justin
160	2006	Chrysler Championship	0.0257	Choi, K.J.
161	2003	Southern Farm Bureau Classic	0.0262	Huston, John
162	2007	The Barclays	0.0266	Stricker, Steve
163	2003	Shell Houston Open	0.0267	Couples, Fred
164	2008	the Memorial Tournament	0.0267	Perry, Kenny
165	2009	Transitions Championship	0.0268	Goosen, Retief
220	2007	Deutsche Bank Championship	0.0486	Mickelson, Phil
221	2008	The Honda Classic	0.0487	Els, Ernie
222	2006	Mercedes Championships	0.0488	Appleby, Stuart
223	2009	Deutsche Bank Championship	0.0493	Stricker, Steve
224	2006	Masters Tournament	0.0498	Mickelson, Phil
225	2008	John Deere Classic	0.0506	Perry, Kenny
226	2005	Chrysler Classic of Tucson	0.0507	Ogilvy, Geoff
227	2005	PGA Championship	0.0513	Mickelson, Phil
228	2009	Northern Trust Open	0.0529	Mickelson, Phil
229	2003	Buick Open	0.0558	Furyk, Jim
312	2007	Buick Invitational	0.3449	Woods, Tiger
313	2009	BMW Championship	0.3506	Woods, Tiger
314	2006	Deutsche Bank Championship	0.3558	Woods, Tiger
315	2007	BMW Championship	0.3628	Woods, Tiger
316	2008	Arnold Palmer Invitational	0.3778	Woods, Tiger
317	2008	Buick Invitational	0.3836	Woods, Tiger
318	2009	Arnold Palmer Invitational	0.4110	Woods, Tiger
319	2007	THE TOUR Championship	0.4141	Woods, Tiger
320	2009	AT&T National	0.4507	Woods, Tiger
321	2009	Buick Open	0.5171	Woods, Tiger

The top and bottom sections represent the ten tournaments for which the probabilities that the actual winner would win are the lowest and highest, respectively. The three middle sections are from the middle of the second through fourth quintiles of probability rankings. Estimates based on 10,000 simulation trials.

Table 7: Estimated Probabilities of Winning the Four Majors and THE PLAYERS Championship for the Five Top Players, 2003-2009

Tournament		Woods	Singh	Furyk	Els	Mickelson
2003	THE PLAYERS Championship	0.1160	0.1156	0.0429	DNP	DNP
2003	Masters Tournament	0.1172	0.1203	0.0468	0.0685	0.0193
2003	U.S. Open	0.1198	0.1186	0.0285	0.0657	0.0206
2003	British Open	0.1181	0.1109	0.0215	0.0622	0.0174
2003	PGA Championship	0.1127	0.1133	0.0181	0.0654	0.0235
2004	THE PLAYERS Championship	0.1280	0.1017	DNP	0.0702	0.0301
2004	Masters Tournament	0.1353	0.1171	DNP	0.0896	0.0375
2004	U.S. Open	0.1431	0.1088	0.0122	0.0703	0.0435
2004	British Open	0.1383	0.1013	0.0126	0.0645	0.0462
2004	PGA Championship	0.1463	0.0979	0.0119	0.0693	0.0437
2005	THE PLAYERS Championship	0.1585	0.0813	0.0138	0.0638	0.0487
2005	Masters Tournament	0.1699	0.0849	0.0194	0.0663	0.0530
2005	U.S. Open	0.1818	0.0767	0.0231	0.0560	0.0502
2005	British Open	0.1819	0.0670	0.0291	0.0545	0.0554
2005	PGA Championship	0.1840	0.0630	0.0318	DNP	0.0513
2006	THE PLAYERS Championship	0.2041	0.0493	0.0437	0.0466	0.0499
2006	Masters Tournament	0.2203	0.0549	0.0564	0.0458	0.0498
2006	U.S. Open	0.2265	0.0433	0.0558	0.0402	0.0473
2006	British Open	0.2347	0.0416	0.0590	0.0374	0.0506
2006	PGA Championship	0.2285	0.0353	0.0501	0.0374	0.0462
2007	THE PLAYERS Championship	0.2721	0.0241	0.0308	0.0281	0.0453
2007	Masters Tournament	0.2777	0.0282	0.0343	0.0356	0.0484
2007	U.S. Open	0.2948	0.0255	0.0227	0.0263	0.0500
2007	British Open	0.3186	0.0236	0.0205	0.0275	0.0439
2007	PGA Championship	0.2932	0.0204	0.0182	0.0232	0.0440
2008	THE PLAYERS Championship	DNP	0.0194	0.0160	0.0269	0.0573
2008	Masters Tournament	0.3589	0.0180	0.0130	0.0204	0.0409
2008	U.S. Open	0.3429	0.0132	0.0117	0.0216	0.0369
2008	British Open	DNP	0.0169	0.0167	0.0221	0.0446
2008	PGA Championship	DNP	0.0172	0.0192	0.0217	0.0495
2009	THE PLAYERS Championship	0.3254	0.0079	0.0177	0.0086	0.0237
2009	Masters Tournament	0.3397	0.0101	0.0200	0.0116	0.0299
2009	U.S. Open	0.3270	0.0068	0.0187	0.0076	0.0249
2009	British Open	0.2775	0.0041	0.0169	0.0049	DNP
2009	PGA Championship	0.3075	0.0034	0.0219	0.0048	0.0168

DNP = did not participate. Estimates based on 10,000 simulation trials.

References

Berry, S. M. 2001. How Ferocious is Tiger? *Chance* 14(3) 51-56.

Berry, S. M. 2008. Is Tiger Woods a Winner? Working paper.

Berry, S. M., C. S. Reese and P. D. Larkey. 1999. Bridging Different Eras in Sports. *Journal of the American Statistical Association* 94(3) 661-676.

Broadie, M. 2010. Assessing Golfer Performance on the PGA TOUR. Working paper, Columbia University.

Brown, J. 2010. Quitters Never Win: The (Adverse) Incentive Effects of Competing with Superstars. Working paper, Northwestern University.

Clark, R. D., III. (2003a). Streakiness Among Professional Golfers: Fact or Fiction? *International Journal of Sports Psychology* 34 63-79.

Clark, R. D. III. (2003b). An Analysis of Streaky Performance on the LPGA Tour. *Perceptual and Motor Skills*, 97 365-370.

Clark, R. D. III. (2004a), On the Independence of Golf Scores for Professional Golfers. *Perceptual and Motor Skills* 98, 675-681.

Clark, R. D. III. (2004b). Do Professional Golfers Streak? A Hole-to-Hole Analysis. in *Proceeding of the Statistics in Sports Section*, American Statistical Association 3207-3214.

Connolly, R. A. and R. J. Rendleman, Jr. 2008. Skill, Luck and Streaky Play on the PGA TOUR. *The Journal of The American Statistical Association* 103(1) 74-88.

Connolly, R. A. and R. J. Rendleman, Jr. 2009. Dominance, Intimidation and 'Choking' on the PGA TOUR. *Journal of Quantitative Analysis in Sports* 5(3).

Ferguson, D. 2009. Tiger Woods Voted AP's Top Athlete of the Decade. *Associated Press* December 16, 2009.

Gilden, D. and S Wilson. 1995. Streaks in Skilled Performance. *Psychonomic Bulletin and Review* 2(2) 260-265.

Grober, R. D. 2008. PGA Tour Scores as a Gaussian Random Variable. Working paper, Yale Univeristy.

PGA TOUR Communications. 2010. Official World Golf Rankings. January 19, 2010.

Storey, J. D. (2002). A Direct Approach to False Discovery Rates. *Journal of the Royal Statistical Society, Ser. B* 64(3) 479-498.

Storey, J. D. (2003). The Positive False Discovery Rate: A Bayesian Interpretation and the q-Value. *The Annals of Statistics* 31(6) 2013-2035.

Wang, Y. 1998. Smoothing Spline Models with Correlated Random Errors. *The Journal of The American Statistical Association* 93(1) 341-348.

Zhang, N.F. 2006. Calculation of the Uncertainty of the Mean of Autocorrelated Measurements. *Metrologia* 43(4) S276-S281.