

A SIMULATION MODEL TO ANALYZE THE IMPACT OF HOLE SIZE ON PUTTING IN GOLF

Matulya Bansal
Mark Broadie

Graduate School of Business,
Columbia University,
New York, NY 10027, USA

ABSTRACT

We develop a model of golfer putting skill and combine it with physics-based putt trajectory and holeout models to study the impact of doubling the radius of the hole on the putting performance of professional and amateur golfers. The putting skill model reflects golfer execution errors, i.e., that golfers cannot hit the ball at exactly their intended velocity and direction. A green reading skill model reflects a golfer's inability to perfectly estimate the slope or contour of the putting surface. The model is calibrated to professional and amateur putting data. Optimal putting strategies are computed using stochastic dynamic programming. Quasi-Monte Carlo and other methods are used to speed up computations. Doubling the hole radius improves the putting performance of both professional and amateur golfers, as expected. However, the improvement for amateur golfers is shown to be relatively larger than for professionals.

1 INTRODUCTION

Gene Sarazen, one of the best golfers in the modern era, believed that putting was too important in golf, and that a larger hole size would make for a better balance between putts and all other golf shots. On his profile in the World Golf Hall of Fame it is written that he "lobbied unsuccessfully to have the hole enlarged from four inches in diameter to eight." Ben Hogan, Johnny Miller, and other prominent golfers have expressed similar views. In a May 2005 article in *Golf Digest* (p.60), Johnny Miller asserted, "The experiment [of enlarging the hole] was tried and promptly **deemed a failure**, because it had the unintended effect of giving an even bigger advantage to the best putters." The decline in the number of people playing golf in recent years has been partly attributed to the time it takes to finish a round and the frustration associated with the process of putting the ball into the hole (Thomas et al. 2008). A larger hole is sometimes suggested as a means to reduce the time to play and increase the enjoyment by amateur golfers.

In this paper, we use simulation to analyze the effect of increasing the radius of the hole. Simulation is ideally suited for this analysis for several reasons. First, there is little available putting data for an enlarged hole. Collecting such data would be time-consuming and expensive, because thousands of putts by golfers of different and known putting abilities would be necessary for reliable results. It would take time for golfers to identify how their optimal putting strategy would be affected by the increased hole size. And finally, it would be nearly impossible to obtain such data under tournament conditions comparable to those facing professionals on a weekly basis. Instead, we use simulation to analyze how increasing the radius of the hole would affect the putting performance of professional and amateur golfers. The simulation model is calibrated to professional and amateur putting data using the current standard hole size and then results are generated with a larger hole size. This approach fully takes into account the physics of a holeout (e.g., that gravity will have a larger effect with a larger hole size) and the change in optimal putting strategies to account for the larger hole size.

We develop models of golfer putting and green reading skill and combine them with physics-based putt trajectory and holeout models for sloped green surfaces in order to determine optimal putting strategies. Equations from Newtonian physics are used to determine the trajectory of a putt on a sloped planar green taking into account friction. Physics principles are also used to determine whether the trajectory of a putt will lead to a holeout or a miss. The golfer skill model has two components. The first is a physical skill model, which includes execution errors in velocity and direction. The second is a green reading model, which reflects errors in the golfer's estimate of the slope of the green. A putting strategy refers to the target velocity and direction chosen by the golfer. In choosing a strategy, a golfer should consider the likelihood of a holeout, how far a putt might finish from the hole in case it misses and the likelihood of making the subsequent putt, which will depend on the hole radius. Another consideration is the slope of

the green, which causes a putt to follow a curved trajectory, referred to as the *break* of the putt. Hitting the ball with a larger initial velocity causes the putt to break less initially, and so it is less affected by green reading error, i.e., errors in the estimate of the green slope. However, a larger initial velocity will lead to misses that tend to be further from the hole, resulting in a greater chance of a three-putt (i.e., taking three putts to holeout). The determination of the optimal strategy involves the solution of a two-dimensional stochastic dynamic program. The model is calibrated to an extensive set of amateur and professional putting data. The calibrated model is used to investigate how the expected number of putts for professional and amateur golfers changes with a doubling of the hole size. We quantify the impact of a larger hole on highly skilled professional and less skilled amateur putters in order to predict which type of putter would benefit more.

Gelman and Nolan (2002), Hoadley (1994) and Tierney and Coop (1999) developed simple putting models that were fit to professional putting data. However, these models are not rich enough to analyze putting with larger hole sizes.

The remainder of the paper is organized as follows. In Section 2, we present models for the ball trajectory, the putting green, and the golfer, and discuss the golfer objective of minimizing the expected number of putts. A numerical algorithm to compute these objectives is given in Section 3. Numerical results, including calibration and computation of optimal putting strategies for professional and the amateur golfers, are presented in Section 4. Concluding remarks are given in Section 5.

2 MODEL

In this section, we describe models for the ball trajectory and determining whether the trajectory leads to a holeout. Then we describe models for the green and the golfer, and discuss the golfer objective.

2.1 Trajectory model

A *trajectory* is the path followed by the ball on a green given an initial velocity and direction. The trajectory model used in this paper is from Vanderbei (2001), who considers the problem of finding the velocity and direction to putt on a green so that the ball comes to rest in the hole. In Vanderbei (2001), the movement of the ball is modeled as sliding on a surface with friction. Perry (2002) develops a slightly more realistic model that considers both sliding and rolling effects, but the extra level of complexity is not necessary for our purposes. We numerically solve the system of differential equations in Vanderbei (2001) to obtain the ball trajectory.

2.2 Holeout model

A holeout, i.e., the ball finishing at the bottom of the hole, can occur in several ways, including the ball falling into the hole, the ball hitting the back of the hole and dropping in, or the ball rolling along the rim and eventually falling into the hole. Holmes (1991), Hubbard and Smith (1999) and Penner (2002) derive equations of motion for a ball interacting with a hole and they determine the maximum velocity that will lead to a holeout, accounting for all of these holeout possibilities, as a function of the distance of the ball from the center of the hole and the hole radius. The current standard hole radius is 2.125 inches. We also consider an enlarged hole with a radius of 4.25 inches.

2.3 Green model

The two main characteristics of our model for the putting green are its slope and its speed. In reality, greens are curved surfaces where the slope varies from one point to the next. However, many greens are nearly flat surfaces and the hole is almost always positioned on a flat portion of the green. For these reasons, we assume that the entire green has a fixed slope, i.e., we model the green to be a planar surface in three-dimensional Euclidean space given by $z = ax + by + c$, where a and b are constants. With this green specification, it suffices to denote any point (x, y, z) on the green as (x, y) . We will refer to a planar green with $a = b = 0$ as being *level*. The coefficients a and b are referred to as the *grade*. We report the green slopes $\tan^{-1}(a)$ and $\tan^{-1}(b)$ in degrees. We assume that the center of the hole lies at position $(0, 0)$.

The speed of a putting green is defined to be the distance a golf ball travels on the green when rolled off a stimpmeter onto a level portion of the green (Holmes 1986). The stimpmeter is a device designed to release a golf ball from a length of 30 inches along an inclined plane making an angle of 20° with respect to the green. As shown by Holmes (1986), the initial velocity of a ball rolling off a stimpmeter is 1.83 m/s. If the ball rolls d feet off the stimpmeter, the *green speed* is said to be d feet. Greens with large stimpmeter speeds, e.g., 11 feet or greater, are called *fast greens* while greens with small stimpmeter speeds, e.g., 8 feet or less, are called *slow greens*. The speed of a green is determined by the height, type and grain of the grass on the green, the wetness and hardness of the green, and other physical features which cause friction between the ball and the green. We assume that the entire green has a constant coefficient of friction denoted η . For a level green, the equations of motion in Vanderbei (2001) can be solved to give $d = v^2 / (2\eta g)$, where d is the distance traveled from the initial position, v is the initial velocity of the ball, and g is acceleration due to gravity.

2.4 Holeout region

The *holeout region* is the set of velocity-angle combinations that lead to a holeout for a given initial ball position. Figure 1 shows holeout regions for 5-foot sidehill putts with standard and enlarged hole sizes. The minimum velocity and maximum angles that lead to a holeout are similar in both cases, since these trajectories correspond to paths where the ball’s velocity declines to zero as it reaches the edge of the hole. However, the maximum velocity for a holeout almost doubles from the standard hole size to the enlarged hole size. The larger size of the holeout region illustrates how putting to a larger hole is easier.

2.5 Golfer skill model

We model three different aspects of golfer putting ability: errors in putting the ball with a desired velocity, errors in putting the ball in a desired direction, and errors in estimating the slope of a green. We refer to these errors as velocity error, direction error, and green reading error, respectively.

The trajectory model requires the ball’s initial velocity. On a level green the distance a ball travels is proportional to the square of the initial velocity, so our primitive variable will be \tilde{v}^2 , the ball’s random initial velocity squared. We assume that

$$\tilde{v}^2 \sim \mathcal{N}(\mu_v^2, g(\mu_v)^2), \quad (1)$$

i.e., \tilde{v}^2 is normally distributed with a mean μ_v^2 , where μ_v is the *target velocity* chosen by the golfer, and $\tilde{v} - \mu_v$ is the *velocity error*. The variance of \tilde{v}^2 is denoted $g(\mu_v)^2$. We motivate the functional form for $g(\cdot)$ next. Differences between a ball’s initial velocity and the golfer’s target velocity contribute to distance errors, i.e., the realized length of the putt is different from the target length. Putting data shows that distance errors are roughly proportional to the length of the putt, which implies $g(\mu_v)$ should be roughly proportional to μ_v^2 . However, 20-foot putts on a fast level green will typically have greater distance errors than 20-foot putts on a slow level green. This implies that lower velocities will have slightly higher relative errors, i.e., distance error normalized by length of the putt, than larger velocities. Similarly, shorter putts tend to have slightly larger relative distance errors than longer putts on the same green. These considerations suggest that $g(\mu_v)$ is a convex increasing function of μ_v^2 . We assume that $g(\mu_v)$ is a piecewise-linear convex function given by:

$$g(\mu_v) = \begin{cases} \beta_2 v_\beta^2 - \beta_0(v_\beta^2 - \mu_v^2), & \mu_v^2 \leq v_\beta^2 \\ \beta_2 v_\beta^2 + \beta_1(\mu_v^2 - v_\beta^2), & \mu_v^2 > v_\beta^2 \end{cases} \quad (2)$$

where v_β is termed the *breakpoint velocity*, and β_0 , β_1 and β_2 determine how distance error changes with velocity. We impose $0 \leq \beta_0 \leq \beta_1 \leq \beta_2$ to ensure non-negativity and convexity of $g(\mu_v)$. As a special case, taking $\beta_0 = \beta_1 = \beta_2$ leads to $g(\mu_v) = \beta_2 \mu_v^2$, which implies that relative distance error is constant on level greens. We sometimes denote \tilde{v}^2 by $\tilde{v}(\mu_v)^2$ to emphasize the dependence on the target velocity.

Direction errors occur because golfers are unable to putt the ball in exactly the desired target direction. Given that a golfer chooses a *target angle* of μ_α (measured relative to the ball-hole line), we assume that the ball starts at a random angle $\tilde{\alpha}$ which follows a normal distribution:

$$\tilde{\alpha} \sim \mathcal{N}(\mu_\alpha, \sigma_\alpha^2). \quad (3)$$

We sometimes denote $\tilde{\alpha}$ by $\tilde{\alpha}(\mu_\alpha)$ to emphasize the dependence on the target direction. Smaller values of σ_α correspond to more highly skilled putters. We assume that \tilde{v}^2 and $\tilde{\alpha}$ are independent.

Green reading errors occurs because golfers cannot estimate green slopes perfectly, i.e., a golfer’s estimate of the green slope is different from the actual green slope. Suppose a golfer estimates the green slopes to be $\theta = (\theta_x, \theta_y)$, where θ_x and θ_y are estimates along the x -axis and y -axis, respectively. The actual slope is randomly chosen by nature and its distribution is given by

$$(\tilde{\theta}_x, \tilde{\theta}_y) = (\theta_x, \theta_y) + (\sigma_g Z \cos(2\pi U), \sigma_g Z \sin(2\pi U)), \quad (4)$$

where $Z \sim \mathcal{N}(0, 1)$, $U \sim U[0, 1]$ and Z and U are independent. The green reading skill parameter is σ_g and high values of σ_g imply greater errors in the golfer’s estimates of the green slopes. To motivate equation (4), observe that θ_x and θ_y can be represented as a point in two-dimensional space. Adding $(\sigma_g Z \cos(2\pi U), \sigma_g Z \sin(2\pi U))$ leads to green slopes that are uniformly distributed on a circle centered at (θ_x, θ_y) with radius $\sigma_g |Z|$.

Since we restrict our analysis to planar greens, without loss of generality, we change coordinates so that the golfer’s green slope estimate is zero along the x -axis, i.e., we set $\theta_x = 0$. In other words, the negative y -axis is the downhill direction to the hole, also called the *fall line*. We use the notation $K = (\beta_0, \beta_1, \beta_2, v_\beta, \sigma_\alpha, \sigma_g)$ to denote a golfer’s putting skill parameters.

2.6 Golfer objective

We assume that the golfer’s objective is to minimize the expected number of putts to holeout, defined as follows. Suppose the golfer starts at $I = (x, y)$ and putts until the ball falls in the hole. The golfer’s slope estimates are $(0, \theta_y)$ and the random realized green slope is $(\tilde{\theta}_x, \tilde{\theta}_y)$ defined in equation (4). Suppose the golfer chooses a target velocity μ_v and

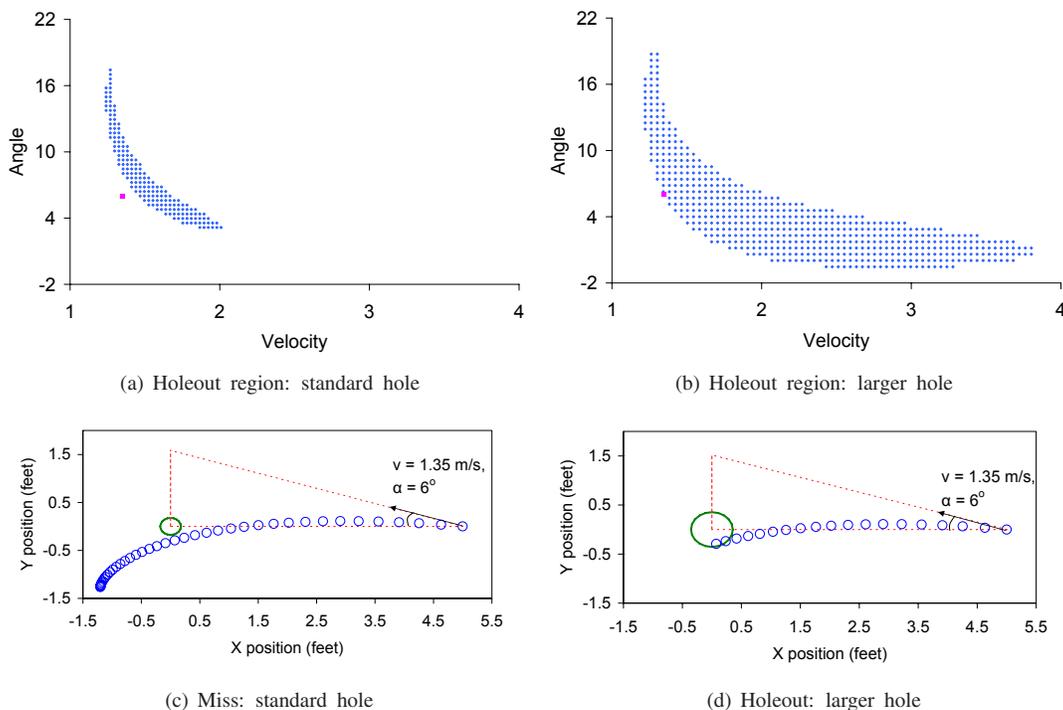


Figure 1: Graphs (a) and (b) show how the holeout region changes for a 5-foot sidehill putt when the hole radius is doubled. The green has a slope of 1.5° along the y -axis, and the green speed is 11 feet ($\eta = 0.0510$). For a sidehill putt, the initial ball position makes an angle of 0° with respect to the x -axis. Each dot on the graph represents a velocity-angle combination that results in a holeout. Velocity is measured in meters per second and the angle refers to the initial direction of the putt relative to the ball-hole line measured in degrees. A zero degree angle means the putt starts directly toward the hole. A positive angle means the ball starts to the right of the hole, and then gravity causes the ball to *break*, i.e., curve, toward the hole. Graphs (c) and (d) show trajectories for $v = 1.35$ m/s and $\alpha = 6^\circ$ (indicated by a square in graphs (a) and (b)). With the standard hole size, the ball misses the hole while the same trajectory results in a holeout with the larger hole.

a target angle μ_α for the putt. The realized velocity $\tilde{v}(\mu_v)$ and the realized angle $\tilde{\alpha}(\mu_\alpha)$ are given by equations (1) and (3), respectively, and the random trajectory of the putt starting at position I is $\tilde{\mathcal{T}}(I, \mu) = \tilde{\mathcal{T}}(I, \tilde{\theta}, \eta, \tilde{v}(\mu_v), \tilde{\alpha}(\mu_\alpha))$, where η is the friction coefficient, and $\mu = (\mu_v, \mu_\alpha)$. The stopping point of trajectory \mathcal{T} will be denoted by $\mathcal{S}(\mathcal{T}) = (\mathcal{S}_x(\mathcal{T}), \mathcal{S}_y(\mathcal{T}))$, where we assume that the hole is covered, so that trajectories for putts that would otherwise lead to holeout do not necessarily end at the hole. The holeout function $h(\mathcal{T})$ maps a trajectory to its outcome: $h(\mathcal{T}) = 1$ if \mathcal{T} leads to a holeout, and 0 otherwise. The expected number of putts objective leads to a dynamic programming problem, because the optimal strategy for the initial putt depends on the strategy of the second putt in the case of a miss on the first putt. In order to begin a policy iteration algorithm, and for interest in its own right, it is useful to consider the simpler objective of maximizing the probability of a one-putt, i.e., a holeout on the first putt. The probability of a one-putt depends on the target velocity and angle μ , the initial position I , the slope estimate θ , the

friction coefficient η , and the golfer skill parameters K :

$$P_1(\mu, I, \theta, \eta, K) = E[h(\tilde{\mathcal{T}}(I, \tilde{\theta}, \eta, \tilde{v}(\mu), \tilde{\alpha}(\mu)))]. \quad (5)$$

We often abbreviate the one-putt probability as $P_1(\mu, I)$.

We optimize a given objective over the set of velocity-angle combinations defined by:

$$\mathcal{U} = \{(\mu_v, \mu_\alpha) | \underline{\mu}_v \leq \mu_v \leq \bar{\mu}_v, \underline{\mu}_\alpha \leq \mu_\alpha \leq \bar{\mu}_\alpha\}, \quad (6)$$

where $\underline{\mu}_v$ and $\bar{\mu}_v$ are the smallest and largest candidate velocities, respectively, and $\underline{\mu}_\alpha$ and $\bar{\mu}_\alpha$ are the smallest and largest candidate angles, respectively, that are considered for optimization. The one-putt probability maximizing velocity and angle are given by:

$$\mu^{(1)}(I) = \arg \max_{\mu \in \mathcal{U}} P_1(\mu, I). \quad (7)$$

The expected number of putts, N , depends on the ball's initial position (I) and the golfer strategy ($\mu(I)$), in addition

to the golfer slope estimates (θ), the friction coefficient (η) and golfer skill parameters (K). The result of the first putt is either a holeout or a second putt which begins from the stopping point of the first putt. This leads to the recursion:

$$N(I, \mu(I)) = E[1 + N(\mathcal{S}, \mu(\mathcal{S}))(1 - h(\widetilde{\mathcal{F}}))], \quad (8)$$

where $\widetilde{\mathcal{F}} = \widetilde{\mathcal{F}}(I, \mu(I))$ and $\mathcal{S} = \mathcal{S}(\widetilde{\mathcal{F}})$. The Bellman equation for the optimal expected number of putts is:

$$N^*(I) = \min_{\mu \in \mathcal{U}} E[1 + N^*(\mathcal{S}(\widetilde{\mathcal{F}}))(1 - h(\widetilde{\mathcal{F}}(I, \mu)))]. \quad (9)$$

Denote the optimal choice of target velocity and angle in equation (9) by $\mu^*(I)$. Using results from Whittle (1983) it can be shown that the standard policy iteration algorithm converges to the optimal policy for this problem.

3 COMPUTATIONAL METHODS

In this section, we show how the optimization problems in equations (7) and (9) are solved to identify the optimal strategies for a given golfer. Both the state and control spaces in equations (7) and (9) are continuous, so we discretize these to proceed with the computation.

3.1 State and control space discretization

The state space $I \subset \mathbb{R}^2$ is continuous, so to solve equations (7) and (9), we discretize I . It is convenient to denote the position of the ball on the green $I = (x, y)$ in polar coordinates as (d, γ) , where $d = \sqrt{x^2 + y^2}$ and $\gamma = \tan^{-1}(y/x)$, $\gamma \in [0, 2\pi)$. We discretize the (d, γ) -space into a finite number of points $I_{ij} = (d_i, \gamma_j)$, $i = 1, \dots, n_d$, $j = 1, \dots, n_\gamma$. Here $d \in (0, \bar{d}]$, where $\bar{d} < \infty$ is the length of the longest putt we consider. We assume that the probability of a one-putt from any point on the green, for any golfer, is strictly positive.

The set of feasible controls \mathcal{U} in equations (7) and (9) is continuous. Since a closed-form solution to the objectives in these equations is not available, we optimize over a discrete subset of $\widehat{\mathcal{U}}$ of \mathcal{U} .

3.2 Probability estimation

To estimate the one-putt probability, $P_1(\mu, I)$, we generate n samples, $(\widehat{\theta}^{(k)}, \widehat{v}^{(k)}(\mu), \widehat{\alpha}^{(k)}(\mu))$, $k = 1, \dots, n$. Then

$$\widehat{P}_1(\mu, I) = \frac{1}{n} \sum_{k=1}^n h(\widetilde{\mathcal{F}}^{(k)}) \quad (10)$$

gives an estimate of $P_1(\mu, I)$, where $\widetilde{\mathcal{F}}^{(k)}$ is short for $\widetilde{\mathcal{F}}^{(k)}(I, \widehat{\theta}^{(k)}, \eta, \widehat{v}^{(k)}(\mu_v), \widehat{\alpha}^{(k)}(\mu_\alpha))$.

To identify $\widehat{\mu}^{(1)}(I)$, the estimate of the velocity-angle combination that maximizes the probability of a one-putt, we perform a grid search over $\widehat{\mathcal{U}}$ to obtain

$$\widehat{\mu}^{(1)}(I) = \arg \min_{\mu \in \widehat{\mathcal{U}}} \widehat{P}_1(\mu, I). \quad (11)$$

3.3 Expected putts estimation

Next we describe how we solve equation (9) to find the strategy that minimizes the expected number of putts. This is an instance of a two-dimensional stochastic shortest path problem discussed, for example, in Whittle (1983). We use policy iteration to solve for the optimal expected number of putts. We first discuss the policy iteration algorithm for the continuous state and control space case, and then show how to implement it after discretizing the state and control space.

The one-putt probability maximizing strategy, $\mu^{(1)}(I)$, is the solution of a simple numerical optimization procedure, i.e., one that does not require a recursive dynamic programming algorithm. Furthermore, for short putts, maximizing the one-putt probability is nearly equivalent to minimizing the expected number of putts, since the expected number of putts is approximately $2 - P_1$ when the probability of three or more putts is nearly zero. For these reasons, we use $\mu^{(1)}(\cdot)$ as the initial policy in the policy iteration algorithm.

The expected number of putts starting from the initial position I , and using the policy $\mu^{(p)}(\cdot)$ for the initial and any subsequent putts, is denoted $N^{(p)}(I)$. The number of putts until a holeout occurs is the smallest m for which putt m results in a holeout, so $N^{(p)}(I)$ can be written as

$$N^{(p)}(I) = E[\min\{m = 1, 2, \dots \mid h(\widetilde{\mathcal{F}}(I_m, \mu^{(p)}(I_m))) = 1\}], \quad (12)$$

where $I_1 = I$, and $I_m = \mathcal{S}(\widetilde{\mathcal{F}}(I_{m-1}, \mu^{(p)}(I_{m-1})))$, i.e., the initial position of putt m is the stop point of putt $m - 1$ (if putt $m - 1$ does not end in a holeout). Given a policy $\mu^{(p)}(\cdot)$, equation (12) defines the policy evaluation step. Under our assumption that the probability of a one-putt is strictly positive, $N^{(p)}(\cdot)$ is finite with probability one.

The policy improvement step is:

$$\mu^{(p+1)}(I) = \arg \min_{\mu \in \widehat{\mathcal{U}}} E[1 + N^{(p)}(\mathcal{S}(\widetilde{\mathcal{F}}(I, \mu)))(1 - h(\widetilde{\mathcal{F}}(I, \mu)))]. \quad (13)$$

Equation (13) states that $\mu^{(p+1)}(I)$, the optimal policy given an initial position I , is given by the target velocity-angle combination $\mu \in \widehat{\mathcal{U}}$ that minimizes the expected number of putts to holeout starting from position I , when μ is used for the first putt, and policy $\mu^{(p)}(\cdot)$ is used for subsequent putts, if any. Starting with $p = 1$, we iterate between

equations (12) and (13) until the policy converges, i.e., until $|\mu^{(p)}(I) - \mu^{(p+1)}(I)| < \varepsilon$, for all I , and for some fixed $\varepsilon > 0$.

Since the state space I is continuous, we show how to proceed with the computations in equations (12) and (13) after the state and control spaces are discretized. For each I_{ij} , $i = 1, \dots, n_d$, $j = 1, \dots, n_\gamma$, we solve equation (11) to find $\hat{\mu}^{(1)}(I_{ij})$, the strategy that maximizes the probability of one-putt from I_{ij} . Next we solve equation (12). To estimate $N^{(1)}(I_{ij})$, the objective in equation (12) for $p = 1$, simulate n trials, each trial consisting of a sequence of putts until holeout occurs. Suppose trial k requires $\tilde{m}(k)$ putts, i.e., $h(\tilde{\mathcal{T}}_{u,k}) = 0$, $u = 1, \dots, \tilde{m}(k) - 1$, $h(\tilde{\mathcal{T}}_{\tilde{m}(k),k}) = 1$ and $\tilde{\mathcal{T}}_{u,k}$ denotes the trajectory of putt u for trial k . Then

$$\hat{N}^{(1)}(I_{ij}) = \frac{1}{n} \sum_{k=1}^n \tilde{m}(k) \tag{14}$$

gives an estimate of $N^{(1)}(I_{ij})$, $i = 1, \dots, n_d$, $j = 1, \dots, n_\gamma$. For each simulation trial, the initial position of putt u is the stop point of putt $u - 1$, given that it didn't result in a holeout. The target strategy is the one-putt probability maximizing strategy from the stopping point of putt $u - 1$. Since the stopping point will not, in general, coincide with a grid point, we interpolate to obtain the target velocity-angle strategy $\hat{\mu}^{(1)}$.

To identify $\hat{\mu}^{(p+1)}(I_{ij})$, the policy at iteration $p + 1$, we perform a grid search:

$$\begin{aligned} \hat{\mu}^{(p+1)}(I_{ij}) = \arg \min_{\mu \in \mathcal{W}} \\ E \left[1 + \hat{N}^{(p)}(\mathcal{S}(\tilde{\mathcal{T}}(I_{ij}, \mu)))(1 - h(\tilde{\mathcal{T}}(I_{ij}, \mu))) \right]. \end{aligned} \tag{15}$$

For $p = 1$, we obtain $\hat{\mu}^{(2)}(I_{ij})$, for $i = 1, \dots, n_d$, $j = 1, \dots, n_\gamma$. We repeat this procedure to determine $\hat{N}^{(2)}(I_{ij})$, then we determine $\hat{\mu}^{(3)}$ and so on until the policy converges, i.e., until $|\hat{\mu}^{(p)}(I_{ij}) - \hat{\mu}^{(p+1)}(I_{ij})| < \varepsilon$ for all I_{ij} for some fixed $\varepsilon > 0$.

3.4 Computational speedups

We now discuss some techniques and observations that enable us to considerably speed up the computation of optimal putting strategies.

Quasi-Monte Carlo: For variance reduction, we use the Sobol sequence (Press et al. 2007) to generate samples in equations (1), (3)-(4), (10)-(11) and (14)-(15). Low-discrepancy methods or Quasi-Monte Carlo methods, of which the Sobol sequence is an example, seek to achieve variance reduction by generating samples that are evenly distributed.

We generate a four-dimensional Sobol sequence to estimate the one-putt probability in equation (10). The first two dimensions are used to generate velocity and angle samples using equations (1) and (3), respectively, while the third and the fourth dimensions are used to generate green slopes using equation (4). To estimate the expected number of putts in equation (14), we use a 10-dimensional Sobol sequence, where the third and the fourth dimension are used to generate green slopes using equation (4), and dimensions 1-2 and 5-10 are used to generate velocity and angle realizations for putts 1 through 5, if needed.

Reducing the dimensionality of the optimization: The optimization over μ_v and μ_α in equations (11) and (15) can be CPU intensive. All optimal solutions obtained with this two-dimensional optimization procedure were found to possess the property that the ball trajectory at the optimal target velocity-angle combination passes through the center of the hole, a property which makes intuitive sense. It can be shown that this property holds in the special cases of a level green, and for straight uphill and downhill putts. If this property holds in general, then one-dimensional optimization can be used to identify the optimal solution. In particular, suppose that the optimal strategy for the golfer is to target a distance d feet beyond the hole ($d \geq 0$) and that the trajectory corresponding to the optimal velocity-angle combination, (μ_v, μ_α) , passes through the center of the hole. Instead of a two-dimensional search over (μ_v, μ_α) , we perform a one-dimensional search over $d \geq 0$, using a root-finding procedure to solve for the velocity-angle combination $(\mu_v(d), \mu_\alpha(d))$ that leads to a stop point d feet beyond the hole and passes through its center. The computations were done with this one-dimensional procedure, with spot checks for accuracy using the slower two-dimensional search procedure.

Symmetry of the optimal policy: Since we only consider planar greens, putts started on either side of the fall line to the hole will follow symmetric trajectories. Together with the symmetry of the normal distribution used in the putting skill models, this means that we only need to find the optimal solution to equations (11) and (15) for $\gamma \in [-90, 90]$ for any d .

4 NUMERICAL RESULTS

In this section, we present numerical results that allow us to quantify the effect of doubling the radius of the hole. We first describe the putting data and calibration results.

Data: We use amateur and professional golfer data collected under actual playing conditions from regular play and tournaments. PGA TOUR data was collected with their ShotLink™ system. The database contains the start and stop points of approximately 15,000 putts hit by over 100 different golfers. The database is contained in the Golfmet-

rics program which is further described in Broadie (2008). From this data the fraction of one-putts, three-putts, and average number of putts are computed as a function of the initial distance from the hole.

Parameter choices: Public and private courses typically have green speeds in the 7-10 foot stimp meter range. Green speeds at professional tournaments are typically in the 9-13 foot range. For our numerical experiments, we use a green speed of 11 feet ($\eta = 0.0510$) for professional golfers, and a green speed of 9 feet ($\eta = 0.0623$) for amateur golfers. For our numerical computations, we use a constant green slope of 1.5° , which is a typical value found on actual greens.

Computational parameters: We use a time increment of 0.1 seconds to compute the putt trajectories. We use 2^{16} simulation trials to obtain the probability and expected number of putts estimates. We set \bar{d} , the length of the longest putt in our experiments, to be 50 feet. For discretizing the initial putt locations, we use $d = \{1.5, 2, 2.5, 2.75, 3, 3.25, 3.5, 4, 4.25, 4.5, 4.75, 5, 5.25, 5.5, 6, 6.5, 7, 8, 9, 10, 12, 15, 18, 20, 25, 30, 40, 50\}$ and $\gamma = \{0^\circ, 15^\circ, 30^\circ, \dots, 360^\circ\}$. The concentration of gridpoints near the hole allows better interpolation in this region, which is important because a majority of the putts that do not result in a holeout are likely to end near the hole. We use the bicubic spline interpolation implementation from Numerical Recipes in C (Press et al. 2007) for strategy interpolation. With these parameters, errors of approximately 0.005 in the expected number of putts are obtained. Numerical experiments were run on a Pentium 4, 3.2 GHz processor with 1 GB RAM, using the MS Visual C++ compiler.

Calibration: We calibrated the golfer model to professional golfer data and obtained a good match with the parameter values: $\beta_0 = 5.5\%$, $\beta_1 = 6.5\%$, $\beta_2 = 6.5\%$, $v_\beta = 15$ feet, $\sigma_\alpha = 1.0$, and $\sigma_g = 0.15$. For long putts, the relative putt distance error is 6.5% due to velocity errors. The standard deviation of direction error is 1.0° . Calibrating to amateur golfer data (representing golfers with an average score of about 90) we obtained a good fit with the parameter values: $\beta_0 = 6.0\%$, $\beta_1 = 7.8\%$, $\beta_2 = 8.5\%$, $v_\beta = 25$ feet, $\sigma_\alpha = 1.5$, and $\sigma_g = 0.25$. For long putts, the relative putt distance error is 8.5% due to velocity errors. The lower skill level of the amateur golfers is also reflected in the larger standard deviations of direction and green reading parameters. Figure 2 shows the fit for expected number of putts between the model and data as a function of the putt length. For both professional and amateur golfers, the root-mean-squared error in the expected number of putts was about 0.03, which is the same magnitude as the standard errors in the data.

Effect of the hole radius: Table 1 shows the increase in the one-putt probability for professional and amateur golfers with an enlarged hole as a function of the initial putt distance. While one-putt probabilities for both professional and amateur golfers improve significantly, for short initial

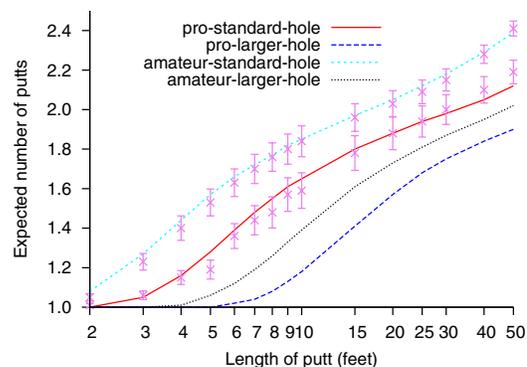


Figure 2: This figure shows how the expected number of putts changes for professional and amateur golfers as the hole radius is changed from 2.125 inches to 4.25 inches. The error bars indicate the confidence intervals for the data used for calibration. The model with the standard hole fits the data well.

putt distances, the relative and absolute improvement for amateur players is much larger. For example, for a 3-foot putt, the one-putt probability for professional golfers increases from 95.4% to 100% with the larger hole, while for amateur golfers probability increases from 76.3% to 98.2%. Table 1 also shows the decrease in the three-putt probability as a function of the initial putt distance. For example, for a 50-foot initial putt distance, the three-putt probability for professionals decreases from 14.6% to 3.2% with the larger hole, while for amateur golfers the probability decreases from 35.8% to 8.9%. The change in expected number of putts for professional and amateur golfers is shown in Figure 2 and Table 1. Professional golfers have less room for improvement compared to amateur golfers.

Figure 3 shows how the optimal strategy for the professional golfer changes with the larger hole. The optimal strategy is more aggressive than with a standard hole size, as reflected in the increases in the target distances beyond the hole (i.e., increases in the target velocities) and decreases in the fraction of putts which finish short of the hole. Cochran and Stobbs (1968) performed a small experiment with a larger hole size, and our simulation results are broadly consistent with theirs. They did not test golfers with different putting skills, so their results cannot be used to answer which golfers would benefit more from a larger hole.

To aggregate results from individual putts into an average number of putts per 18-hole round, the expected number results are weighted by the distribution of initial putt lengths. Table 1 shows the effect of doubling the hole radius on putting performance. The expected number of putts per round decreases from 29.3 to 24.3 for professional

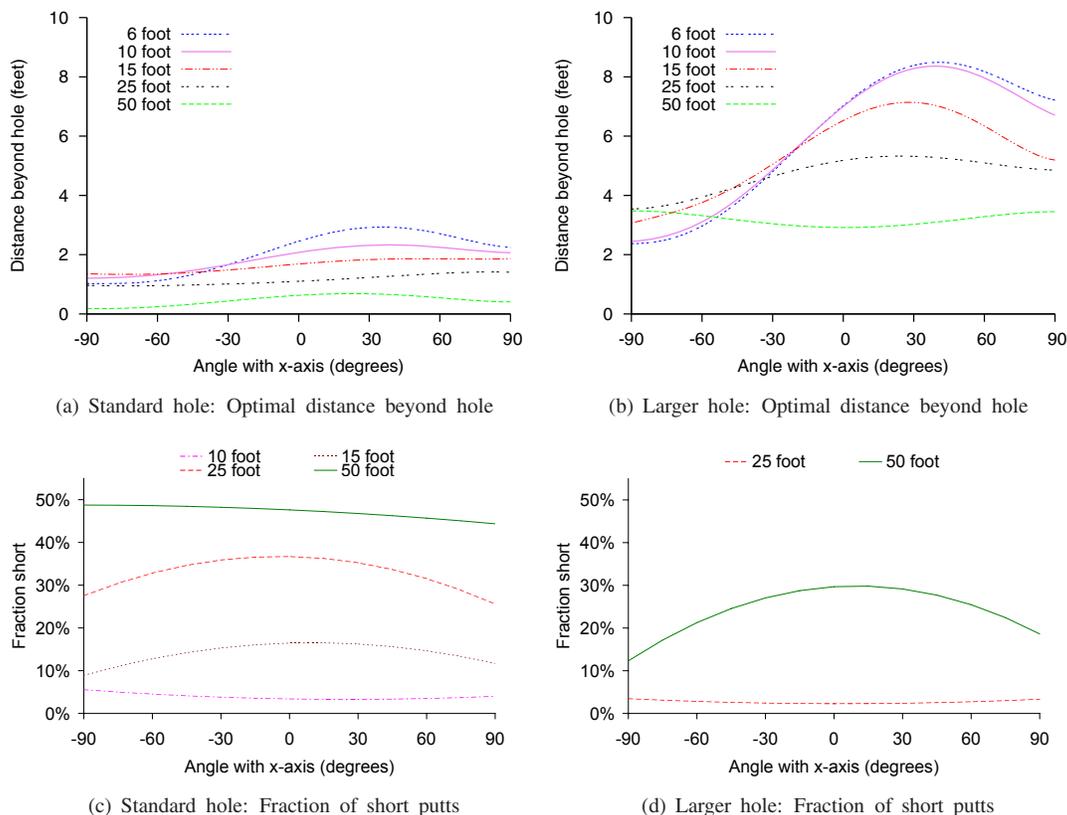


Figure 3: This figure shows how the optimal target distance beyond the hole and the fraction of putts that fall short of the hole vary for a professional golfer as the hole radius is changed from 2.125 inches to 4.25 inches. Sidehill putts have an angle of 0° with the x -axis, downhill putts have a 90° angle, and uphill putts have a -90° angle. Graphs (a) and (b) show that the optimal target distance increases significantly, while graphs (c) and (d) show that the fraction of short putts decreases significantly, together implying that the optimal strategy becomes more aggressive with the larger hole radius. The fraction short values for 10- and 15-foot putts with the larger hole are less than 0.5% and are not shown for clarity.

golfers, and decreases from 33.6 to 26.8 for amateur golfers. Both highly skilled and less skilled golfers benefit from the larger hole size. The improvement by professionals is 5.0 putts per round but amateurs benefit to an even greater extent, with a reduction of 6.8 putts per round. This is in contrast to the Johnny Miller quotation from his May 2005 Golf Digest article which claimed that more highly skilled putters would benefit more. No reference is given for the experiment in which the enlarged hole had the “unintended effect of giving an even bigger advantage to the best putters,” so it is difficult to understand the source of the discrepancy. It is possible that the experiment was not very extensive or carefully done. It is also possible that more highly skilled putters understood that their putting strategy should change with an enlarged hole, while less skilled putters may not have had time to adapt accordingly.

5 CONCLUSION

We developed a model of golfer putting ability and combined it with physics-based putt trajectory and holeout models. The golfer putting model incorporates both physical skill, which reflects the golfer’s ability to putt with a desired target velocity and angle, and green reading skill, which reflects the golfer’s ability to estimate the slope of the green. The model was calibrated to real-world professional and amateur golfer data. Optimal putting strategies were found using stochastic dynamic programming. We analyzed the impact of doubling the hole radius on professional and amateur player putting performance. As expected, doubling the hole radius improves the performance of both professional and amateur golfers. However, the relative performance improvement for amateur golfers is larger.

Table 1: Effect of doubling the hole radius

This table summarizes the effect of doubling the hole radius from 2.125 inches to 4.25 inches. Distributions of the number of initial putts in an 18-hole round are given in the columns n_p and n_a for the professional and amateur golfer, respectively.

d	Professional							Amateur						
	n_p	Standard hole			Larger hole			n_a	Standard hole			Larger hole		
		P_1	P_3	N^*	P_1	P_3	N^*		P_1	P_3	N^*	P_1	P_3	N^*
2	1.00	99.8%	0.0%	1.00	100.0%	0.0%	1.00	1.15	93.2%	1.1%	1.08	99.6%	0.0%	1.00
3	1.51	95.4%	0.3%	1.05	100.0%	0.0%	1.00	0.46	76.3%	3.1%	1.27	98.2%	0.0%	1.00
4	0.78	84.8%	0.8%	1.16	99.9%	0.0%	1.00	0.86	60.4%	4.1%	1.44	93.7%	0.1%	1.01
5	0.78	73.0%	1.2%	1.28	98.9%	0.0%	1.00	0.86	48.1%	4.3%	1.57	85.7%	0.4%	1.06
6	0.78	62.3%	1.3%	1.39	95.7%	0.1%	1.01	0.86	39.4%	4.4%	1.66	76.4%	0.9%	1.12
7	0.69	53.8%	1.5%	1.48	91.2%	0.2%	1.04	0.89	33.2%	4.9%	1.72	68.5%	1.4%	1.19
8	0.69	46.8%	1.6%	1.55	85.4%	0.4%	1.08	0.89	28.2%	5.0%	1.77	60.5%	1.8%	1.26
9	0.69	41.0%	1.6%	1.61	79.2%	0.8%	1.13	0.89	24.3%	5.0%	1.82	53.2%	1.9%	1.33
10	1.59	36.4%	1.8%	1.65	72.9%	1.0%	1.18	1.76	21.1%	5.3%	1.85	47.1%	2.2%	1.39
15	2.11	22.0%	2.1%	1.80	48.6%	1.5%	1.41	2.62	12.3%	7.7%	1.97	29.2%	2.8%	1.61
20	1.53	14.6%	2.7%	1.88	33.5%	1.8%	1.57	1.75	8.1%	11.1%	2.05	20.2%	3.3%	1.73
25	1.48	10.3%	3.9%	1.94	24.3%	1.9%	1.68	1.25	5.8%	15.2%	2.12	14.8%	3.4%	1.81
30	1.75	7.6%	5.4%	1.98	18.4%	2.0%	1.75	1.47	4.6%	19.7%	2.18	11.9%	4.3%	1.87
40	1.31	4.7%	9.6%	2.05	11.8%	2.5%	1.84	1.28	3.0%	28.2%	2.29	8.3%	5.8%	1.95
50	1.05	3.3%	14.6%	2.12	8.5%	3.2%	1.90	0.93	2.2%	35.8%	2.39	6.3%	8.9%	2.02
Average putts per round		29.31			24.31				33.59			26.81		

REFERENCES

Broadie, M. 2008. Assessing golfer performance using golf-metrics. In *Science and Golf V: Proceedings of the World Scientific Congress of Golf*, ed. D. Crews and R. Lutz, 253–262. Mesa, Arizona: Energy in Motion Inc.

Cochran, A., and J. Stobbs. 1968. *Search for the perfect swing: The proven scientific approach to fundamentally improving your game*. Chicago, Illinois: Triumph Books.

Gelman, A., and D. Nolan. 2002. A probability model for golf putting. In *Teaching Statistics*, Volume 24, 93–95.

Hoadley, B. 1994. How to improve your putting score without improving. In *Science and Golf II: Proceedings of the World Scientific Congress of Golf*, ed. A. J. Cochran and M. R. Farrally, 186–192. London: E & FN Spon.

Holmes, B. W. 1986, October. Dialogue concerning the stimpmeter. *The Physics Teacher* 24(7):401–404.

Holmes, B. W. 1991. Putting: How a golf ball and hole interact. *American Journal of Physics* 59(2):129–136.

Hubbard, M., and T. Smith. 1999. Dynamics of golf ball-hole interactions: Rolling around the rim. *Transactions of the ASME* 121:88–95.

Penner, A. R. 2002. The physics of putting. *Canadian Journal of Physics* 80(2):83–96.

Perry, S. K. 2002. The proof is in the putting. *The Physics Teacher* 40(7):411–414.

Press, W. H., S. A. Teuklosky, W. T. Vetterling, and B. P. Flannery. 2007. *Numerical recipes: The art of scientific computing*. 3rd ed. Cambridge, UK: Cambridge University Press.

Thomas, F., B. Christina, V. Melvin, and E. Alpenfels. 2008. Growing the game: A survey report. In *Science and Golf V: Proceedings of the World Scientific Congress of Golf*, ed. D. Crews and R. Lutz, 245–252. Mesa, Arizona: Energy in Motion Inc.

Tierney, D. E., and R. H. Coop. 1999. A bivariate probability model for putting efficiency. In *Science and Golf III: Proceedings of the 1998 World Scientific Congress of Golf*, ed. A. J. Cochran and M. R. Farrally, 385–394. UK: Human Kinetics.

Vanderbei, R. J. 2001. A case study in trajectory optimization: Putting on an uneven green. In *SIAG/OPT Views-and-News*, Volume 12(1), 6–14.

Whittle, P. 1983. *Optimization over time*. Chichester, UK: Wiley Series in Probability and Mathematical Statistics.

AUTHOR BIOGRAPHIES

MATULYA BANSAL is a doctoral student in the Decision, Risk and Operations division at Columbia Business School. His research interests are in the areas of financial engineering, revenue management, and computer networks. His email address is <mbansal11@gsb.columbia.edu>.

MARK BROADIE is the Carson Family Professor of Business at the Graduate School of Business, Columbia University. His research focuses on issues in financial engineering, with a particular focus on the design and analysis of efficient Monte Carlo methods for pricing and risk management. His email address is <mb2@columbia.edu> and his webpage is <www.columbia.edu/~mb2>.