

Using Simulation to Improve the Pace of Play

MoonSoo Choi

Harvard Business School

Qi Fu

MSCI, Inc.

Ward Whitt

Columbia University

We develop a flexible and efficient simulation algorithm of a stochastic (queueing network) model of group play over a conventional 18-hole golf course. We apply this algorithm to study ways to improve the pace of play in recreational golf. We consider: (i) using a wave-up rule for par-3 holes to allow two groups to play at the same time, (ii) finding a “minimum cost” tee schedule (the intervals between successive groups starting play on the first hole), and (iii) making the first hole a bottleneck par-3 hole (when they are bottlenecks). The simulation exposes complexities in the wave-up rule, but nevertheless shows that the wave-up rule can increase the daily number of groups that can play each day by about 13%. A simulation optimization is used to select a good constant tee schedule. It shows that the cost of making the tee interval too short tends to be far greater than the cost of making it too long. Finally, when there is a bottleneck hole and it appears first, then the greatest delay occurs there, making congestion developing over the day more evident to both golfers and management.

Keywords: alternative tee schedules, pace of group play in golf, queueing network models, recreational golf, stochastic simulation models, waving up in golf

1. Introduction

1.1. The Problem of Slow Play in Recreational Golf

In recent years there has been concern about slow play in recreational golf, i.e., when the time to play a full round of 18 holes far exceeds the 4-hour standard. In response, (Riccio, 2014a) established the Three/45 Golf Association for golfers as well as owners, managers and designers of golf courses, “dedicated to leading, educating and advocating for a quicker pace of play.” (Riccio, 2014b.) also conducted a supporting data analysis using data from several courses. As a remedy, (Riccio, 2012) proposed analyzing the pace of play by applying industrial engineering methods (mathematical models and computer simulation) that have been applied with great success to improve the efficiency of production facilities, e.g., (Hopp & Spearman, 1996). That approach was illustrated by several deterministic models. Further efforts to address the pace-of-play problem have been reported in (Tiger, Trent, & Haney, 2015) and (Tiger & Ellerbrook, 2016).

Although it may not be evident from watching professional golf tournaments, the play of successive groups over a golf course is actually a quite complicated process in recreational golf, which is our primary concern. The individual golfers often have different experience and skill levels; the groups may have different size, and the golfers may either walk or use golf carts. It is challenging to find models that can provide useful insight considering such complexity. If we model each shot by each golfer, then the model becomes too unwieldy. On the other hand, if we make too much simplification, we may draw unjustified and unreliable conclusions.

To help address this problem, we develop a new flexible and efficient simulation algorithm of group play over a standard golf course and apply it to explore three complex operational issues, which are difficult to study experimentally by trial and error on a golf course. Earlier simulation models of golf were constructed by (Kimes & Schruben, 2002) and (Tiger & Salzer, 2004).

1.2. The Stochastic Model of Group Play

A major contribution of this paper is creating and applying a new simulation algorithm based on the stochastic queueing network model proposed by (Whitt, 2015). That model was used to mathematically determine the capacity (the maximum possible throughput, i.e., the rate at which groups can complete play) on each hole and on the entire course, but a full simulation algorithm for that model on a general course was left for future work.

The model in (Whitt, 2015) differs significantly from previous models. Instead of focusing on the actions of individual golfers, the new high-level model focuses on the times required for the entire group to pass through “stages” of play on each hole. The model primitives are the times required for each group to play each of the stages. In order to represent the inevitable randomness, including the possibility of lost balls, the stage playing times are modeled as stochastic random variables, whose probability distributions are part of the model input.

The stages provide a level of detail intermediate between the actions of individual golfers and the time for the group to play the hole. The stages are important

because multiple groups can play at the same time on some holes, provided they follow appropriate rules of play, which prevent the group in front from getting hit by a ball from the group behind. Typically, two groups can be playing on the most common par-4 hole at the same time, while three groups can be playing on a longer par-5 hole at the same time. On the other hand, the shorter par-3 holes are more elementary, because only one group can play at the same time. (We explain in more detail in Section 3.)

Using these stages, group play over the entire golf course during a day is represented as the flow through a series of 18 queues with precedence constraints. The precedence constraints produce an unconventional queueing network model, e.g., compared to the queueing models in (Hopp & Spearman, 1996) and the large literature on queueing networks, e.g., (Serfozo, 1999) and (Walrand, 1988). Given all the stage playing times, the play of multiple groups over the course is defined by recursive formulas. Thus, given stage playing times over the course for all groups, the starting times and finishing times for all groups can be readily computed by a computer algorithm. However, for random stage playing times, we need to apply simulation to estimate the expected starting and finishing times of all groups. In this paper, we develop a flexible and efficient simulation algorithm of the stochastic queueing network model exploiting the stage structure.

1.3. Aims of the Paper

The main purpose of this paper is to develop a simulation tool and carry out studies with that tool that can help improve the pace of play in recreational golf. We also use the simulation to validate the conclusions of (Whitt, 2015) about hole capacities and the advantage of having a balanced course (where the capacities of the individual holes are all approximately equal).

To improve the pace of play on any given golf course, we recommend collecting data on the playing times of all the groups over multiple days. That leads to the question: What data should we collect? The ideal would be to record the time and location of each shot by each golfer. Then we can aggregate this data into the stage playing times for each group, and obtain the data required to estimate all the stage playing times. Of course, to fit the model we use, it would suffice to have only the more parsimonious data of the stage playing times of each group. However, it is not sufficient to record only the group start and finish times on each hole, because that ignores the interaction among successive groups.

For such a data analysis, the model from (Whitt, 2015) and our simulation algorithm make two important contributions: (i) to identify the relevant data and (ii) to provide a framework for incorporating the data in order to understand the performance implications. The model is chosen so that the stage playing times are likely to be independent of the operational choices. Hence, one set of stage-playing-time data should serve to analyze many different operational scenarios.

In this paper, we show that the simulation algorithm also can be used to investigate a broad array of operational questions, applicable to many golf courses, without using specific course data, provided that representative stage playing time distributions are used. Our parameters are stochastic generalizations of the deterministic parameters in (Riccio, 2012).

1.4. Organization

In Section 2 we describe the three operational problems we consider in more detail. In Section 3 we review the stochastic model of group play. In Section 4 we briefly describe our simulation algorithm (providing more details in an Online Companion). In Section 5 we report results of simulation experiments comparing alternative course designs with and without the wave-up rule. In Section 6 we present the results of our simulations to choose an optimal tee interval subject to constraints. In Section 7 we investigate the possible advantage of an uneven tee schedule, using a shorter interval and then shifting to a longer one. In Section 8 we discuss the advantages of placing a bottleneck par-3 hole first. Finally, in Section 9 we highlight key numerical results and draw conclusions. Additional material appears in the Online Companion, which is available from the authors' web pages.

2. Three Operational Problems

In addition to investigating the validity of the main conclusions in (Whitt, 2015) about capacity, loading and balance for general golf courses, we apply the new simulation algorithm to study three possible ways to improve the pace of play: (i) introducing a wave-up rule for the bottleneck par-3 holes, (ii) optimizing over alternative tee schedules, i.e., the intervals between successive groups starting play on the first hole, and (iii) placing a bottleneck hole first on the course.

2.1. Golf Course Design and the Wave-Up Rule

Golf courses typically have 18 holes of different length (the distance from the initial shot from the “tee” to the final shot i.e., the last “putt” on the green). Golf is typically played in groups of individual golfers, with 4 being a common group size. The par value, i.e., the target number of strokes to play the hole, typically ranges from 3 to 5, with the values increasing with the length and difficulty of the hole. The design of golf courses varies, but there usually are 8–12 of the average-length par-4 (P4) holes with the remaining 6–10 holes either the shorter par-3 (P3) holes or the longer par-5 (P5) holes.

Counter to initial intuition, the shortest P3 holes tend to be the bottleneck holes (i.e., have lower capacity), where groups of golfers experience the greatest delay (before starting to play). That can be explained by the different numbers of groups that are allowed to play on each hole at the same time. The shorter P3 hole is usually a bottleneck, because only one group can play on it at any given time.

In an effort to increase the pace of play, some golf courses have adopted a special wave-up rule to use on P3 holes. The wave-up rule allows two groups to play at the same time there too, while still maintaining the order determined by their arrival; we call this a P3WU hole. The wave-up rule stipulates that, after all members of the group have hit their tee shots and walked up to their balls near the green, they should wait before hitting their next shots and clearing the green until the following group hits its tee shots. However, each group only waits after it gets to its balls near the green if the following group (1) has already arrived and (2) is

ready to play. The waiting group near the green can watch the subsequent tee shots of the next group to avoid danger. If the following group has not yet arrived at the hole, then the current group completes its play on the hole. The following group then cannot start play on the hole until after the current group completes play and departs. The wave-up rule is intended to reduce the expected time between successive groups clearing the green, and thus increase the capacity of P3 holes and improve the pace of play. (Whitt, 2015) showed that the wave-up rule increases the capacity of a P3 hole in the model.

The wave-up rule actually is somewhat complicated, because a group will wait before completing play (in order to allow the next group to first hit their tee shots) only if that next group is ready to start. Some groups may experience the good fortune of being waved up by the group in front of them but not having to wave up the group behind, because they are not yet ready. That is good for that particular group, but the next group may have to wave up the following group, even though they themselves were not waved up. Therefore, the wave-up is inevitably applied inconsistently, thus introducing some variability in the delays and the flows. In this paper we apply the new simulation algorithm to carefully study the advantage of the wave-up rule for par-3 holes, comparing P3WU holes to conventional P3 holes.

2.2. Alternative Tee Schedules

We also show how the simulation algorithm can be applied to determine an optimal tee schedule, i.e., the interval between successive groups scheduled to start play on the first hole. To achieve that goal, we formulate an optimization problem. In particular, we maximize the number n of groups that can play on the course during each day subject to two constraints. The first constraint requires that the expected time for any group to play the course be less than γ minutes, which we will stipulate as $\gamma = 240$, corresponding to the well-known 4-hour target; the second constraint requires the expected time for all groups to complete play be less than the total time available, τ minutes, which we take to be $\delta = 840$ minutes, corresponding to 14 hours.

Let $V(\tau, n)$ be the expected time for group n to play a full round and $G(\tau, n)$ be the expected time for group n to complete play when the tee interval is τ . Since these are increasing functions of n , we have the general optimization problem:

Maximize n , such that

$$G(\tau, n) \leq \delta = 840 \text{ and } V(\tau, k) \leq \gamma = 240 \text{ for all } k, 1 \leq k \leq n,$$

where n is the number of groups and thus a positive integer, while τ is the tee interval and thus a positive real number. We solve this optimization problem by applying the simulation algorithm to simulate the play of 100 groups on the full 18-hole course for each value of τ in a suitably large set. We perform a large number of independent replications (typically 2000) of the entire experiment to ensure that we obtain good statistical precision.

In Section 7 we also study a two-level tee schedule, starting with shorter intervals and then switching to longer intervals. We deliberately keep the structure simple in this way to produce a realistic candidate that might be considered in practice.

2.3. Making the First Hole a Bottleneck

As observed on p. 97 of (Riccio, 2012), “one of the rules of factory physics is to design the process with the bottleneck near the beginning of the process.” Moreover, (Riccio, 2012) concludes that “if courses were designed with the longest par-3 hole as the rest of the course would likely flow smoothly.”

We conducted simulation experiments to investigate this phenomenon. First, our simulations show that the ordering of the holes matters very little if the course is balanced, i.e., if the capacities of the individual holes, as defined in (Whitt, 2015), are approximately equal. On the other hand, if the course is unbalanced, with P3 holes as the bottlenecks, and if the course is slightly overloaded, then indeed, most of the delay experienced by later groups occurs before starting play.

On the other hand, we also studied the total waiting times over all holes for each group. We found that the total waiting time differs little upon changing the course ordering.

3. The Stochastic Model of Group Play

In this section we review the stochastic model of group play from (Whitt, 2015). In Section 3.1 we specify the stage-playing-time distributions and the recursive formulas for each of the hole types. In Section 3.2 we specify the parameters of these stage-playing-time distributions.

3.1. The Model Primitives: Stage Playing Times

The primitives for the model of group play on a golf course are the stage-playing-time random variables for the stages of each hole. By focusing on the stage playing times, we do not directly model the actions of each individual golfer. On the other hand, the stage playing times provide essential detail not available from group playing times on each hole. The stages are defined so that the stage playing times of any one group over different stages and of different groups can reasonably be regarded as independent random variables. On the other hand, the times for successive groups to play a P4 or P5 hole are necessarily dependent, because these two groups can be playing on this same hole simultaneously.

Models were constructed of each of the basic hole types: P3, P4 and P5, plus the modification P3WU. The model and simulation algorithm allow an option to change the parameters from hole to hole, even if the hole type is the same. However, we assume that the parameters are the same for all holes of the same type. So we have only three sets of parameters, one for each of the three hole types: P3, P4 and P5. The parameters for the P3 hole will depend on whether or not the wave-up rule is being used. Here we first review the model for a P4 hole because it is most common. Then we review the P3, P3WU and P5 models.

3.1.1. The Steps and Stages for a Par-4 Hole

We first describe the steps of group play on a P4 hole. There are five steps, each of which must be completed before the group moves on to the next step. These five steps can be diagrammed as

$$T \rightarrow W1 \rightarrow F \rightarrow W2 \rightarrow G$$

The first step T is the tee shot (one for each member of the group); the second step $W1$ is walking up to the balls on the fairway; the third step F is the fairway shot; the fourth step $W2$ is walking up to the balls on or near the green; the fifth and final step G is clearing the green, which may involve one or more approach shots and one or more shots (putts) on the green for each player in the group.

The rules of play allow two groups to play at the same time on a conventional P4 hole. Two successive groups can be simultaneously playing on the hole, because each group is allowed to hit its initial tee shots after the previous group has hit its fairway shots, and so will be safely out of the way, while each successive group is allowed to hit its fairway shots only after the previous group has cleared the green.

An important part of the modeling approach is to not directly model the performance of these individual steps. Instead, the five steps are aggregated into three stages, which are important to capture the way successive groups interact while playing the hole. In particular, we represent the three stages of a P4 hole as:

$$(T, W1) \rightarrow F \rightarrow (W2, G).$$

Stage 1 is $(T, W1)$; stage 2 is F , and stage 3 is $(W2, G)$. We use this particular aggregation because it is the maximum aggregation permitted by the precedence constraints, which we turn to next.

We now describe the precedence constraints, which follow common conventions in golf. Assuming an empty system initially, the first group can do all the stages, one after another without constraint. However, for subsequent groups, group $n + 1$ cannot start stage 1 until *both* group $n + 1$ arrives at the tee and group n has completed stage 2, i.e., has cleared the fairway. Similarly, group $n + 1$ cannot start on stage 2 until *both* group $n + 1$ is ready to begin there and group n has completed stage 3, i.e., cleared the green. These rules allow two groups to be playing on a P4 hole simultaneously, but under those specified constraints. We may have groups n and $n + 1$ on the course simultaneously for all n . That is, group n may first be on the course at the same time as group $n - 1$ (who is ahead), but then later be on the course at the same time as group $n + 1$ (who is behind). The groups remain in their original order, but successive groups interact on the hole. The group in front can cause extra delay for the one behind.

3.1.2. The Stochastic Model for Group Play on a Par-4 Hole

The description of group play on a P4 hole is the basis for a mathematical model. For that purpose, let $A(n)$ be the scheduled (and assumed actual) arrival time of group n at the tee of this hole on the golf course. Let $S_j(n)$ be the time required for group n to complete stage j ; these are called the stage playing times. The mathematical model data for a P4 hole consists of a sequence of 4-tuples:

$$\{(A(n), S_1(n), S_2(n), S_3(n)) : n \geq 1\}$$

where the four components for each n are nonnegative random variables. All the stage playing times are assumed to be mutually independent. In our simulation we assume that the distribution of $S_j(n)$ is the same for all n on P4 holes, for each stage j separately.

We now define the performance measures for the successive groups playing on the hole. Let $B(n)$ be the time that group n starts playing on this hole, i.e., the instant when one of the group goes into the tee box. Let $T(n)$ be the time that group n completes stage 1, including the tee and the following walk; let $F(n)$ be the time that group n completes stage 2, its shots on the fairway; and let $G(n)$ be the time that group n completes stage 3, and clears the green.

The description above is given a concise mathematical representation as a four-part recursion:

$$\begin{aligned} B(n) &= \max\{A(n), F(n-1)\}, T(n) = B(n) + S1(n), \\ F(n) &= \max\{T(n), G(n-1)\} + S2(n) \text{ and} \\ G(n) &= F(n) + S3(n) \end{aligned}$$

As initial conditions, assuming that the system starts empty, we set

$$F(0) = G(0) = 0.$$

The two maxima capture the two precedence constraints.

The model above extends directly to any number of such single-hole models in series. We simply let the completion time $G(n)$ of group n from one hole be the arrival times $A(n)$ at the next hole.

3.1.3. Group Play on a Par-3 Hole Without Wave-Up

In contrast, the conventional P3 hole (without wave-up) is relatively simple, because only one group can be on the course at that hole at any one time. There are three steps for group play on a P3 hole, with or without wave-up:

$$T \rightarrow W \rightarrow G$$

The first step T is hitting shots off the tee; the second step W is walking to the green, possibly including approach shots; and the third step G is putting on the green. For the P3 hole, we identify the stages with steps, but speak of stages, to be consistent with P4.

Given stage playing times $S_j(n)$ for group n on the three stages (indexed by j) as before, the total time for group n to play the hole is $X(n) = S1(n) + S2(n) + S3(n)$.

3.1.4. Group Play on a Par-3 Hole With Wave-Up

The stochastic model of group play on a P3WU hole is considerably more complicated. The wave-up rule stipulates that, after all members of a group have hit their tee shots and walked up to their balls near the green, they should wait before clearing the green until the following group hits its tee shots, provided that the following group has already arrived and is ready to play. If the following group have not yet arrived at the hole, then the current group immediately start stage 3. The following group then cannot start play on the hole until after the current group completes stage 3 and departs.

The wave-up rule makes the formulas for $B(n)$ and $G(n)$ in terms of the other variables somewhat complicated. At a time equal to the larger of the times $W(n)$ and $G(n-1)$, i.e., at the time $\max\{W(n), G(n-1)\}$, group $n-1$ has cleared the

green and group n has completed stage 2, so that group n is ready to play stage 3. However, group $n + 1$ may impose a constraint. At that time, group n can start stage 3 (to play on the green) only if either (i) group $n + 1$ has not yet arrived at the hole and is not ready to tee off or if (ii) group $n - 1$ has completed its tee shots. Otherwise, group n starts stage 3 at time $T(n + 1)$. Finally, if group n has not arrived at the hole when group $n - 1$ is ready to start stage 3, then group $n - 1$ will start stage 3 immediately, and so that group n cannot start to play on the hole until group $n - 1$ has cleared the green and departed, at time $G(n - 1)$. Thus, we introduce the event (set) $E(n)$, defined by

$$E(n) = \{A(n) \leq \max\{W(n-1), G(n-2)\} < T(n)\}$$

And let be the complement of the set $E(n)$. If group n is the last scheduled group, then let $A(n + 1) = \infty$ (or some very large value) so that the event $E(n + 1)$ never occurs.

Thus, the wave-up rule is specified by the following four-part recursion:

$$\begin{aligned} \text{For } n \geq 2, & B(n) = \max\{W(n-1), G(n-2)\} \text{ on } E(n) \\ & \text{and } B(n) = \max\{A(n), G(n-1)\} \text{ on } E(n)^c; \\ \text{for } n \geq 1, & B(1) = A(1), T(n) = B(n) + S1(n), \\ & W(n) = T(n) + S2(n), \\ & G(n) = T(n+1) + S3(n) \text{ on } E(n+1) \text{ and} \\ & G(n) = \max\{W(n), G(n-1)\} + S3(n) \text{ on } E(n+1)^c, \\ & \text{so that} \\ & G(n) = \max\{W(n), G(n-1)\} + S1(n+1) + S3(n) \text{ on } E(n+1) \\ \text{while} \\ & G(n) = \max\{W(n), G(n-1)\} + S3(n) \text{ on } E(n+1)^c \end{aligned}$$

where “on $E(n)$ ” means “on the set $E(n)$ ” or “when the event $E(n)$ holds (is true),” with the event being the complement of the set $E(n)$, i.e., “when the event $E(n)$ does not hold (is false),” as before. As initial conditions, again assuming that the system starts empty, we set

$$W(0) = G(0) = G(-1) = 0.$$

If n is the last group, then instead of the recursion above, we let

$$G(n) = \max\{W(n), G(n-1)\} + S3(n),$$

because the event $E(n + 1)$ cannot occur.

Note that the expression for $B(n)$ involves $G(n - 1)$, because only two groups can be playing on the hole at the same time. Observe that the event $E(n)$ actually simplifies. From the definition of $E(n)$ above,

$$T(n) = B(n) + S1(n) \geq W(n-1) + S1(n) > W(n-1),$$

so that

$$E(n) = \{A(n) \leq \max\{W(n-1), G(n-2)\}\},$$

without the final inequality involving $T(n)$. Note again that care is needed in treating the last group to play; if n is the last group, then $A(n+1)$ is made large, so that $E(n+1)$ never occurs.

3.1.5. The Stochastic Model for Group Play on a Par-5 Hole

A P5 hole is considerably more complicated than a P4 hole, primarily because three groups can play at the same time instead of the two groups. We model the P5 hole to have seven steps (instead of five steps in P4 holes), and we group the steps into five stages (instead of the three stages in P4 and P3 holes). In particular, we represent the stages of play on a P5 hole as:

$$(T, W1) \rightarrow F1 \rightarrow W2 \rightarrow F2 \rightarrow (W3, G).$$

The following is the six-part recursion for the P5 hole:

$$\begin{aligned} B(n) &= \max\{A(n), F1(n-1)\}, \quad T(n) = B(n) + S1(n), \\ F1(n) &= \max\{T(n), F2(n-1)\} + S2(n), \\ W2(n) &= F1(n) + S3(n), \\ F2(n) &= \max\{W2(n), G(n-1)\} + S4(n) \text{ and} \\ G(n) &= F2(n) + S5(n). \end{aligned}$$

A more detailed description of the mathematical model of the P5 hole appears in Section 6 of (Whitt, 2015). The capacity of each hole type is also derived there.

3.2. The Stage-Playing-Time Distributions

The simulation algorithm allows for the distributions of the stage playing times to be user-defined, but for our analyses, we use concrete models that follow Section 4 of (Whitt, 2015). For all stages of all holes, we assume that the stage playing times $S_j(n)$ are mutually independent random variables with a symmetric triangular ($Tri(m, a)$) distribution, but we also use the modification to allow for a lost ball in the first stage of each hole. The parameter pair (m, a) may depend on both the stage and the hole type.

The triangular $Tri(m, a)$ distribution has a symmetric continuous piecewise-linear density on the interval $[m-a, m+a]$, with a peak at m and the value 0 at the end points $m-a$ and $m+a$, assuming $0 \leq a \leq m$. The two-parameter model provides a convenient characterization of the central tendency or mean via m and the spread or variability via a . The $Tri(m, a)$ distribution has variance $\frac{a^2}{6}$. We let $a = 1.5$ in all cases, but we automatically reduce a to m if the model has $a > m$, then making the density continuous, piecewise-linear and symmetric on $[0, 2m]$ (which has variance $\frac{m^2}{3}$).

In particular, for all P3, P4 and P5 holes, we initially let the mean values of the triangular distributions stage playing times (for the three, three and five stages, respectively) come from the parameter vectors (3.50, 2.00, 2.67), (4.00, 2.00, 4.00) and (4.00, 2.00, 2.00, 1.33, 4.00), respectively. For example, for the second stage of a P4 hole, the density of $S4(n)$ is symmetric on $[0.5, 3.5]$, having mean 2.00, whereas for the fourth stage of a P5 hole, the density of $S4(n)$ is symmetric on $[0.00, 2.67]$, having mean 1.33.

To model the possibility of a lost ball on the first stage, on each hole we let the first stage have a fixed large value of $L = 8$ with probability $p = .05$ and have the original stage playing time otherwise, with probability $1 - p = 0.95$. Thus, including the possibility of a lost ball, the sums of the mean stage playing times on P3, P4 and P5 holes, respectively, are $8.167 + 0.225 = 8.392$, $10.000 + 0.200 = 10.200$ and $13.333 + 0.200 = 13.533$.

The *limiting cycle time* is the average interval between successive departures on a fully loaded hole, i.e., where there always are new groups ready to tee off at the earliest opportunity. The reciprocal of the limiting cycle time is the *capacity of the hole*. From Theorem 1 and Section 4.3 of (Whitt, 2015), we can calculate the limiting cycle time in (6) there is 6.533 for a P4 hole. We directly see that the limiting cycle time for a P3 hole is 8.392, the sum calculated above. We applied simulation to deduce the corresponding limiting mean cycle times are 6.504 and 6.433 on fully loaded P3WU and P5 holes, respectively.

Then, in order to produce a more balanced course, with all holes having limiting cycle time approximately the same as the 6.533 value for P4 holes, we applied simulation to adjust all the parameters for the P3WU and P5 holes by scaling up the means of the triangular distributions. For P3WU, we multiplied (3.50, 2.00, 2.67) by $c = 1.00438$ to get the adjusted mean vector (3.515, 2.009, 2.682), which yielded a limiting cycle time of 6.529. For P5, we multiplied (4.00, 2.00, 2.00, 1.33, 4.00) by $c = 1.0177$ to get the adjusted mean vector (4.071, 2.036, 2.036, 1.357, 4.071), which yielded a limiting cycle time of 6.531. (The lost ball parameters were not adjusted.)

With these adjustments, we have a balanced course with P3WU holes, having approximate course capacity equal to the capacity of a fully loaded P4 hole, $1/6.533$. However, with P3 holes, the P3 holes are bottlenecks. Hence, with P3 holes we have an unbalanced course, having approximate course capacity equal to the capacity of a fully loaded P3 hole, $1/8.392$. With the detailed model specified in this section, we see that using the wave-up rule increases the course capacity by a factor of 1.28. Equivalently, the critical tee interval is reduced by a factor of 0.77. However, that is at the expense of the greater complexity of a P3WU hole. In the rest of this paper we apply simulation to evaluate the actual impact of changing from P3 holes to P3WU holes.

To provide perspective, we also consider scaled P3 holes, referred to as SP3 holes. In the SP3 holes, we reduce the means of the stage playing times (3.50, 2.00, 2.67) in the P3 model by the factor $6.533/8.167$ (approximately 0.8) to produce the stage playing time mean vector (2.800, 1.600, 2.136), but we do not adjust the lost ball parameters. This SP3 model has limiting cycle time 6.793. Hence, this SP3 model is still slightly a bottleneck, being about 6% larger than 6.533. Nevertheless, simulation experiments show that the SP3 courses are slightly more efficient than the P3WU courses, both being much more efficient than the P3 courses.

4. The Simulation Study

We now give an overview of our simulation algorithm and the model parameters. We provide more details about the simulation algorithm in the Online Companion.

4.1. A Flexible Simulation Algorithm

As usually is the case with simulation, there is an important question about how much detail to include in the simulation model. We show that the high-level model proposed in (Whitt, 2015), not including the actions of individual golfers, makes simulation experiments feasible. Moreover, experience indicates that the level of detail in that model is appropriate for operations management issues such as the ones we study (Hopp & Spearman, 1996).

In order to serve as a useful tool for a variety of simulation studies, the simulation algorithm is designed to be flexible. The user can choose various parameters, including (i) the stage playing time distributions and parameters for each hole, (ii) the group arrival times at the first hole, (iii) the hole sequence, (iv) the number of groups playing per day, (v) the type of the P3 holes (P3, P3WU or SP3), and (vi) the number of independent simulation replications.

The simulation algorithm is also designed to be efficient for large-scale analysis, i.e., for achieving good statistical precision for the average performance of all groups over a full day, for a variety of course designs. The high-level model of group play is very important for achieving simulation experiments that can actually be conducted, but nevertheless the simulations are challenging. For example, the course design with 12 P4 holes, three P3 holes and three P5 holes, has $(12 \times 3) + (3 \times 3) + (3 \times 5) = 60$ stage playing times for each group. We allow for up to 100 groups playing on the course each day. Since the performance of the wave-up rule depends on following groups, we simulate 102 groups. That leads to 6,120 stage playing times for one day of golf.

Since the model is a stochastic model, we require multiple (independent) replications of our simulation. We consistently used 2,000 replications. Table 3 shows that the half-widths of the confidence intervals for the estimates of the total waiting times are about 1% of the mean itself. The statistical precision is far less for the individual holes; see Tables 4–6 of the Online Companion. Together, that requires generating $6,120 \times 2,000 = 12.24 \times$ playing times to produce statistically reliable performance estimates of group play for one day and one golf course design.

The simulation experiments become much larger when we study alternative course designs. Hence, the design alternatives must be chosen with care. First, there are $18!/(3!12!3!) = 371,280$ distinct orderings of the three P3, 12 P4 and three P5 holes, but experience indicates that the relevant number of arrangements should be much less, e.g., about 20–30. Second, there are three types of P3 holes (P3, P3WU and SP3), assuming that we treat all the same. Third, we consider a range of 30 different tee intervals. (We use a larger range when we consider two-level tee intervals.) As is, the number of course designs becomes $20 \times 3 \times 30 = 1,800$. That requires generating

$$(12.2 \times 10^6) \times (1.8 \times 10^3) = 22 \times 10^9$$

stage playing times. With MATLAB, time measurements indicate that *a full simulation required about one full week*. We elaborate on the simulation methodology in Sections 2 and 3 of the Online Companion.

4.2. The Model Parameters

There also is the question about model parameters. Our model parameters draw heavily on the previous work, especially (Riccio, 2012). The mean values here are somewhat less than the deterministic values in (Riccio, 2012), but that is compensated for by the variability that we include in our more general stochastic model. While these model parameters should be realistic, it is significant that the methods we develop still apply with other model parameters if others are deemed appropriate. Moreover, our simulation experiments show that the operational conclusions we deduce do not depend critically on the specific parameters used.

We first choose parameters (as indicated in Section 3.2) that make the full course roughly balanced, i.e., so that the capacities of all holes, as determined in (Whitt, 2015), are approximately equal. Then we apply simulation to investigate the actual performance over a full course over a day.

5. Courses With and Without the Wave-Up Rule

In order to study the wave-up rule, we conducted extensive simulation experiments comparing alternative course designs. We primarily focused on the case in which there are 12 identical P4 holes, three identical P5 holes and three identical P3 holes, considering each of the P3WU, P3 and SP3 options discussed in Section 3.2. The base case was the P3WU model, which produces a balanced course with the limiting cycle time for each hole being approximately equal to the P4 value of 6.533. In contrast, when we include the P3 holes, they are bottlenecks with limiting cycle time 8.392. The SP3 holes make the course have limiting cycle time 6.792, making the course almost balanced. For these course models, we found that the performance is approximately independent of the order of the holes. For any order of holes, we found that the performance depends on whether the course is overloaded (having a tee interval on the first hole that is less than the minimum limiting cycle time) or underloaded (having a tee interval on the first hole that is less than the minimum limiting cycle time).

We not only apply the model in (Whitt, 2015), but we test conclusions derived in that paper about the capacities of individual holes and the golf course as a hole. The capacity of each hole was defined as the maximum rate at which golfers could complete play if there were always new groups of golfers ready to play on that hole when the opportunity arises. The capacity of the entire course was then defined as the maximum of the capacities of the individual holes. Following standard queueing terminology, the course is considered overloaded, critically loaded or underloaded if the actual tee schedule makes the input rate greater than, equal to, or less the course capacity. A course was defined as “balanced” if all the hole capacities are approximately equal, and “unbalanced” otherwise. We use the simulation to verify that these notions of capacity, loading and balance are valid and useful for performance analysis on a general golf course, including a variety of hole types. We find that these notions are very important for understanding the results.

We illustrate by showing simulation estimates of the mean waiting times (before starting play) on each hole for one group (group 75) as a function of the type of P3 hole and the tee schedule. We considered many possible course designs, but we only report three here. The three course designs are (i) *the base case*, having P3 and P5 holes appearing alternately, separated by P4 holes:

454 434 454 434 454 434,

(ii) the par-5 holes first
 555 343 434 444 444 444,

and (iii) the par-3 holes first
 333 454 444 454 444 454.

The last two designs are relatively extreme deviations of the base case.

We collected results for all groups over all holes, but here we report results for only the 75th group. Additional waiting time statistics for other groups are available in Section 4.2 of the Online Companion.

5.1 Shorter Tee Intervals: Heavier Loads

First, Table 1 shows simulation estimates of the mean waiting times in minutes for group 75 before playing on each of the 18 holes for three course designs with the three kinds of par-3 holes: P3, P3WU and SP3. In all cases the tee interval is set at 7.50 minutes. From the capacity analysis reviewed above, the golf course is stable for P3WU and SP3, but unstable for P3. To highlight the differences, we show the results for the various par-3 holes in bold. For the base case, we see a significant impact of the P3WU holes on the following hole; these are highlighted in italics in Table 1. Notice that the waiting time is approximately the same at the following P4 hole as at the P3WU hole itself. This shows the impact of the complexity of the wave-up rule.

Table 2 gives the corresponding estimated percentage of the total waiting time at each hole for tee intervals of length 7.50 minutes, under which all course models are stable, with the exception of P3 models. These proportions are estimated by the estimated mean waiting time for that hole, divided by the mean of the total waiting time, and then multiplied by 100 to convert into a percentage.

Consistent with the stability analysis, the P3 holes are bottlenecks, and the waiting time at the first P3 holes grows without bound linearly in the number of holes. On the other hand, steady state is reached approximately by group 50 for the courses with P3WU and SP3 holes. The simulation results confirm that the course is approximately balanced when the P3WU and SP3 holes are used, but not when the P3 holes are used. For the unbalanced courses with P3 holes, 68% of the total expected waiting time for group 75 occurs on the first bottleneck P3 hole, which is hole 5. The other two P3 holes produce most of the remaining wait. None of the 15 non-par-3 holes have more than 1% of the total expected waiting time. In contrast, for the balanced designs with P3WU and SP3 holes, all holes have less than 12% of the total expected waiting time. Even though the performance is consistent with a balanced design, we do see that the average waiting times are slightly larger for the P3WU and SP3 holes than at the P4 and P5 holes.

Table 1 Simulation Estimates of the Mean Waiting Times for Group 75 in Minutes Before Starting Play on Each of the 18 Holes for Three Course Designs With the Three Kinds of par-3 Holes: P3, P3WU and SP3

Hole	par	Base case			par-5 first			par-3 first		
		P3	P3WU	SP3	par	P3	P3WU	SP3	par	P3
1	4	0.34	0.31	0.33	5	0.26	0.28	0.27	3	66.41
2	5	0.60	0.61	0.68	5	0.71	0.76	0.71	3	12.59
3	4	0.94	0.97	1.02	5	0.88	0.92	0.93	3	9.14
4	4	0.96	0.94	0.92	3	66.64	2.03	2.45	4	0.17
5	3	65.00	1.92	2.24	4	0.16	2.02	1.10	5	0.31
6	4	0.18	1.99	1.12	3	13.19	2.15	2.46	4	0.46
7	4	0.35	1.17	1.09	4	0.18	2.10	1.17	4	0.44
8	5	0.41	1.02	0.94	3	9.52	2.12	2.70	4	0.45
9	4	0.58	1.44	1.33	4	0.18	2.07	1.17	4	0.51
10	4	0.48	1.20	1.08	4	0.31	1.23	1.12	4	0.49
11	3	14.28	2.13	2.69	4	0.40	1.22	1.18	5	0.49
12	4	0.20	2.04	1.24	4	0.42	1.21	1.19	4	0.65
13	4	0.34	1.28	1.17	4	0.46	1.15	1.25	4	0.54
14	5	0.40	1.04	0.97	4	0.51	1.16	1.14	4	0.53
15	4	0.60	1.38	1.41	4	0.44	1.23	1.18	4	0.56
16	4	0.51	1.19	1.11	4	0.49	1.13	1.16	4	0.53
17	3	10.40	2.11	2.56	4	0.52	1.18	1.13	4	0.51
18	4	0.20	2.17	1.23	4	0.56	1.09	1.16	5	0.64
sum		97.74	24.90	23.13		95.81	25.05	23.44		95.39
										25.81
										22.82

Note. In all cases the tee interval is 7.50 minutes, which makes the P3 case overloaded, but not the others. The course designs are (i) the base case, (ii) the par-5-holes first and (iii) the par-3 holes first.

Table 2 Simulation Estimates of the Proportion of Waiting Times for Group 75 (in %) Before Starting Play on Each of the 18 Holes for Three Course Designs With the Three Kinds of par-3 Holes: P3, P3WU and SP3

Hole	Base case						par-5 first						par-3 first					
	par	P3	P3WU	SP3	par	P3	P3WU	SP3	par	P3	P3WU	SP3	par	P3	P3WU	SP3		
1	4	0.3%	1.2%	1.4%	5	0.3%	1.1%	1.2%	3	69.6%	6.5%	5.4%						
2	5	0.6%	2.5%	2.9%	5	0.7%	2.9%	3.3%	3	13.2%	10.5%	8.8%						
3	4	1.0%	3.9%	4.4%	5	0.9%	3.6%	4.1%	3	9.6%	10.8%	10.4%						
4	4	1.0%	3.8%	4.0%	3	69.6%	7.9%	10.5%	4	0.2%	7.6%	5.1%						
5	3	67.5%	7.7%	9.7%	4	0.2%	8.4%	4.7%	5	0.3%	4.3%	4.1%						
6	4	0.2%	8.0%	4.8%	3	13.8%	8.7%	11.0%	4	0.5%	5.4%	5.6%						
7	4	0.4%	4.7%	4.7%	4	0.2%	7.5%	4.8%	4	0.5%	4.6%	4.8%						
8	5	0.4%	4.1%	4.1%	3	9.9%	8.6%	11.2%	4	0.5%	4.3%	4.9%						
9	4	0.6%	5.8%	5.8%	4	0.2%	8.0%	4.9%	4	0.5%	4.5%	5.1%						
10	4	0.5%	4.8%	4.7%	4	0.3%	5.3%	4.8%	4	0.5%	4.6%	4.8%						
11	3	14.6%	8.5%	11.6%	4	0.4%	4.8%	4.7%	5	0.5%	4.0%	4.3%						
12	4	0.2%	8.2%	5.4%	4	0.4%	5.0%	5.2%	4	0.7%	5.3%	5.7%						
13	4	0.3%	5.1%	5.1%	4	0.5%	4.8%	4.8%	4	0.6%	4.6%	5.0%						
14	5	0.4%	4.2%	4.2%	4	0.5%	4.8%	4.9%	4	0.6%	4.5%	5.2%						
15	4	0.6%	5.5%	6.1%	4	0.5%	4.5%	4.8%	4	0.6%	4.7%	5.2%						
16	4	0.5%	4.8%	4.8%	4	0.5%	4.7%	5.1%	4	0.6%	4.5%	5.1%						
17	3	10.6%	8.5%	11.0%	4	0.5%	4.8%	5.0%	4	0.5%	3.9%	4.7%						
18	4	0.2%	8.7%	5.3%	4	0.6%	4.6%	5.0%	5	0.7%	5.3%	5.9%						
sum		100%	100%	100%		100%	100%	100%		100%	100%	100%						

Note. In all cases the tee interval is 7.50 minutes, which makes the P3 case overloaded, but with the other two underloaded. The course designs are (i) the base case, (ii) the par-5-holes first and (iii) the par-3 holes first.

Table 3 gives the corresponding estimated standard deviations of the waiting times for group 75 for the three course designs. In all cases the tee intervals are 7.50 minutes, under which all course models are stable, with the exception of P3 models. The last two rows give simulation estimates for the standard deviation of the sum of the waiting times on all 18 holes and the halfwidth of the 95% confidence interval for the mean of the sum, labeled HW. The HW is computed as $[ID]EQ16/[ID]$ where s is the estimated standard deviation, because the number of replications is 2000.

Table 3 shows that the halfwidths of the estimates of the total waiting times are about 1% of the mean itself. The statistical precision is far less for the individual holes. Table 3 shows that P3 holes have the highest standard deviations among the three hole types, while the other two hole types have approximately the same level of standard deviations.

5.2. Longer Tee Intervals: Lighter Loads

We now show that the performance is quite different when all three par-3 holes make the course underloaded. Paralleling Table 1, Table 2 gives the estimated mean waiting times for the longer tee intervals of 8.50 minutes, under which all course models are stable. Table 4 shows that the average waiting times are much lower with the longer tee interval. Since the P3 holes are now underloaded, steady-state is achieved for them too about by group 50. However, the fact that the P3 holes are bottlenecks is again clearly evident from the tables. Just as in Table 1, for the base case, we see a significant impact of the P3WU holes on the following hole; these are highlighted in italics in Table 2. Again, the waiting time is approximately the same at the following P4 hole as at the P3WU hole itself. More detailed statistics for the longer tee intervals of 8.50 minutes, which include analogs of Tables 2 and 3, are available in Section 4.1 of the Online Companion.

We supplement the tables above by also showing histograms of the total waiting times on all 18 holes for group 75 as a function of the 3 course designs, the 3 types of par-3 holes and the 2 tee intervals 7.50 and 8.50 minutes. These are shown in Figures 1 and 2. These figures further confirm our conclusions above.

6. Simulation Optimization of the Tee Interval

We now apply the simulation to determine an optimal tee schedule, i.e., the interval between successive groups scheduled to start play on the first hole.

6.1. The Optimization Framework

To achieve that goal, we maximize the number of groups that can play on the course during each day subject to two constraints, as described in Section 2.2. We solve this problem by applying the simulation algorithm to simulate the play of 100 groups on the full 18-hole course for each value of the tee interval τ in a suitably large set. We gain further insight by first performing the optimization over n for each given τ in order to see how the number of groups that can play the course each day subject to these constraints depends on the tee interval. We also gain insight into the course design discussed in Section 5 by performing the optimization as a function of the hole order and the type of par-3 hole used.

Table 3 Simulation Estimates of the Standard Deviations of the Waiting Times Before Starting to Play on Each Hole for Group 75 for Each of the Three Kinds of par-3 Holes: P3, P3WU and SP3

hole	par	Base case				par-5 first				par-3 first			
		P3	P3WU	SP3	par	P3	P3WU	SP3	par	P3	P3WU	SP3	par
1	4	0.92	0.84	0.81	5	0.79	0.81	0.83	3	12.09	1.31	2.25	
2	5	1.24	1.10	1.15	5	1.31	1.34	1.32	3	10.49	2.63	2.97	
3	4	1.62	1.56	1.65	5	1.52	1.52	1.56	3	8.07	2.87	3.32	
4	4	1.68	1.59	1.63	3	13.12	2.00	3.12	4	0.60	2.56	1.92	
5	3	12.56	1.98	3.04	4	0.65	2.58	1.78	5	0.84	1.71	1.50	
6	4	0.63	2.61	1.82	3	10.59	2.21	3.30	4	1.07	2.10	1.97	
7	4	0.94	2.04	1.83	4	0.64	2.71	1.96	4	1.08	1.98	1.95	
8	5	0.98	1.68	1.68	3	8.40	2.09	3.54	4	1.05	1.98	1.82	
9	4	1.21	2.02	2.03	4	0.59	2.77	2.02	4	1.15	1.90	1.90	
10	4	1.15	1.96	1.90	4	0.97	2.02	1.88	4	1.10	1.99	1.90	
11	3	10.46	2.15	3.57	4	1.03	1.95	1.93	5	1.14	1.70	1.81	
12	4	0.62	2.73	1.92	4	1.15	1.95	2.01	4	1.26	1.93	2.05	
13	4	0.92	2.11	1.96	4	1.08	1.90	2.06	4	1.16	1.99	2.04	
14	5	0.91	1.83	1.58	4	1.12	1.90	1.90	4	1.24	1.88	1.96	
15	4	1.12	2.05	1.99	4	1.11	2.04	1.92	4	1.22	1.92	1.92	
16	4	1.17	1.97	2.00	4	1.11	1.84	1.90	4	1.13	1.97	1.92	
17	3	8.40	2.24	3.71	4	1.14	2.02	1.91	4	1.16	1.70	1.74	
18	4	0.61	2.64	2.05	4	1.12	1.78	1.96	5	1.35	2.06	2.14	
sum		11.91	7.25	8.05		11.59	7.47	8.11		11.45	7.77	7.80	

Note. In all cases the tee interval is 7.50 minutes, which makes the P3 case overloaded, but not the others. The course designs are (i) the base case, (ii) the par-5-holes first and (iii) the par-3 holes first. The last two rows give simulation estimates for the standard deviation of the sum of all waiting times and the half width of the 95% confidence interval for the mean.

Table 4 Simulation Estimates of the Mean Waiting Times for Group 75 in Minutes Before Starting Play on Each of the 18 Holes for Three Course Designs with the Three Kinds of par-3 Holes: P3, P3WU and SP3

Hole	par	Base case			par	par-5 first			par	par-3 first		
		P3	P3WU	SP3		P3	P3WU	SP3		P3	P3WU	SP3
1	4	0.10	0.08	0.09	5	0.09	0.09	0.08	3	5.78	0.76	0.29
2	5	0.21	0.24	0.21	5	0.31	0.28	0.28	3	7.24	1.28	0.56
3	4	0.41	0.42	0.40	5	0.43	0.42	0.38	3	6.46	1.77	0.77
4	4	0.41	0.43	0.36	3	8.39	1.16	0.91	4	0.17	1.44	0.43
5	3	8.50	1.14	0.87	4	0.17	1.39	0.47	5	0.26	0.63	0.45
6	4	0.16	1.22	0.43	3	8.16	1.30	0.98	4	0.46	0.81	0.59
7	4	0.33	0.58	0.51	4	0.15	1.34	0.49	4	0.45	0.60	0.48
8	5	0.35	0.52	0.46	3	7.36	1.36	1.03	4	0.48	0.63	0.57
9	4	0.55	0.64	0.63	4	0.17	1.33	0.59	4	0.46	0.55	0.54
10	4	0.49	0.54	0.51	4	0.29	0.66	0.58	4	0.44	0.59	0.53
11	3	9.42	1.33	1.07	4	0.38	0.64	0.54	5	0.46	0.51	0.47
12	4	0.15	1.41	0.49	4	0.40	0.57	0.51	4	0.61	0.71	0.64
13	4	0.32	0.66	0.56	4	0.43	0.59	0.50	4	0.53	0.63	0.58
14	5	0.37	0.53	0.49	4	0.44	0.58	0.51	4	0.52	0.59	0.52
15	4	0.56	0.73	0.68	4	0.47	0.57	0.53	4	0.54	0.60	0.55
16	4	0.47	0.61	0.60	4	0.47	0.59	0.62	4	0.57	0.58	0.62
17	3	8.51	1.30	1.19	4	0.46	0.57	0.60	4	0.52	0.55	0.51
18	4	0.17	1.35	0.53	4	0.49	0.56	0.60	5	0.65	0.72	0.64
sum		31.47	13.72	10.04		29.05	14.02	10.21		26.39	13.94	9.73

Note. In all cases the tee interval is 8.50 minutes, which makes all three cases under-loaded. The course designs are (i) the base case, (ii) the par-5-holes first and (iii) the par-3 holes first.




Figure 1 — A 3 × 3 collage of histograms of total waiting times for group 75 over 2000 simulation replications. Each row represents three different course designs: (i) the base case, (ii) the par-5-holes first and (iii) the par-3 holes first. Each column represents three different types of par-3 holes in the golf course: *P3*, *SP3*, and *P2WU*. In all of the nine histograms, the tee interval is 7.50 minutes.

Table 5 shows the maximum number of groups that can play each day as a function of (i) the tee interval on the first hole, (ii) the hole order, either the base case or the “par-3 first” and (iii) the type of par-3 hole used. The tee intervals are allowed to range from 5.00 minutes (very overloaded) up to 9.50 minutes (under-loaded). The optimal tee intervals for each case are shown in Table 5 in bold type. Additional maximum throughput statistics for other hole orders are available in Section 5.1 of the Online Companion.

6.2. Important Insights from Table 5

We can draw several important conclusions from Table 5. First, for the unbalanced course with three P3 holes, the optimal tee interval is approximately equal to the




Figure 2 — A 3×3 collage of histograms of total waiting times for group 75 over 2000 simulation replications. Each row represents three different course designs: (i) the base case, (ii) the par-5-holes first and (iii) the par-3 holes first. Each column represents three different types of par-3 holes in the golf course: P3, SP3, and P3WU. In all of the nine histograms, the tee interval is 8.50 minutes.

limiting cycle time, which makes the course critically loaded. However, for the balanced courses with P3WU holes, the optimal tee interval is slightly greater than the limiting cycle time, so that the course is slightly underloaded, with traffic intensity (about 0.90). Thus, we conclude here that, with P3WU holes, it does not suffice to assume that the course is critically loaded. Instead, it should be underloaded.

The situation is less clear for the SP3 holes. Expanding the level of detail, we found that the maximum number of groups with SP3 holes was 70 with $\tau = 6.85$, 83 with $\tau = 6.90$ and 87 with $\tau = 6.95$. Thus, we conclude that the optimal tee interval is covered by $= 6.79/6.95$ (about 0.977). Consequently, we conclude that the optimal course design with SP3 holes also can be considered to be approximately critically loaded. Therefore, we conclude that the wave-up rule is primarily responsible for the deviation from critical loading. This conclusion has important

Table 5 The Maximum Number of Groups That Can Play Each Day as a Function of (i) the Tee Interval on the First Hole, τ (ii) the Hole Order, and (iii) the Type of par-3 Hole Used

Tee interval	Throughput for base case				Throughput for "3 3 3" case			
	P3	P3WU	SP3	P3	P3WU	SP3	P3	SP3
5.00	11	10	16	11	10	16	11	16
5.50	13	12	19	13	12	20	13	20
6.00	15	15	26	15	15	26	15	26
6.50	18	21	41	19	21	40	19	40
7.00	23	42	87	23	41	87	23	87
7.10	24	63	87	25	62	87	25	87
7.20	25	84	86	27	84	87	27	87
7.30	27	84	86	29	84	86	29	86
7.40	30	83	85	31	83	86	31	86
7.50	33	82	85	33	82	85	33	85
7.60	34	82	84	36	82	84	36	84
7.70	37	81	83	40	81	83	40	83
7.80	41	80	82	45	80	82	45	82
7.90	46	79	81	51	79	81	51	81
8.00	53	78	81	59	78	81	59	81
8.10	62	78	80	71	78	80	71	80
8.20	74	77	79	74	77	79	74	79
8.30	74	76	78	74	76	78	74	78
8.40	74	75	77	74	76	77	74	77
8.50	74	75	76	74	75	76	74	76
9.00	71	71	72	71	71	72	71	72
9.50	68	68	69	68	68	69	68	69
τ^*	8.39	6.53	6.79	8.39	6.53	6.79	8.39	6.79

Note. Two hole orders are used: the "base case", and the "par-3 first" case. The optimal tee intervals are in bold, while the critical tee interval τ^* from Whitt (2015) for that kind of par-3 hole is shown at the bottom.

implications for an analytical approximate performance analysis of a balanced golf course developed by (Fu & Whitt, 2015), because that approximation assumed that the network (i) is balanced and (ii) can operate at or near critical loading. (Fu & Whitt, 2015) tested that with a simulation in which all 18 holes were P4 holes, and found it was remarkably successful. That approximation should be relevant for more general courses provided that the courses are balanced and can operate near critical loading. Table 5 supports that approximation for balanced courses without P3WU holes, but the complexity of the wave-up rule does not permit critical loading. Unfortunately, courses with P3 holes tend to be unbalanced, whereas courses with P3WU holes cannot operate at critical loading. Hence, the simple analysis may not be so relevant.

We see that P3WU allows 10 more groups to play than P3, so that we have a good quantitative measure of the increased efficiency of the wave-up rule. This increase is 13%, which is consistent with (Tiger & Salzer, 2004), but significantly less than the 28% capacity difference determined from (Whitt, 2015). That difference can at least partly be attributed to the complexity of P3WU holes.

Third, we see that the SP3 allows 3 more groups to play than P3WU, so that the complexity of the wave-up rule has some cost, even when the theoretical capacity of P3WU is greater than SP3, as can be seen from at the bottom. From the close agreement from the left and right sides of Table 5, we also see that the order of the holes is relatively unimportant.

Fourth, we see a sharp decrease in the number that can play as we decrease the tee interval from its optimum, whereas we see only a slow decrease as we increase the tee interval from its optimum value. Thus, we see that there is a much greater penalty from choosing the tee interval too small than from choosing it too large.

Finally, we include a set of color-coded visualization figures, as illustrated by Figure 3, to aid the understanding of dynamic behavior of the sojourn time constraint (from the top figure in Figure 3) and the departure time constraint (from the bottom figure in Figure 3). More details are available in Section 3 of the Online Companion.

6.3. An Alternative View

We now present an alternative view of Table 5 in Figure 4 below. Figure 4 shows plots of the optimal number of groups that can play each day as a function of the tee interval for five different balanced course designs. In each case the par-3 holes are all P3WU holes, so that the limiting cycle time is 6.53 minutes. As in Table 5, the maximum number that can play is achieved for tee intervals that lie between 7.10 and 7.50 minutes, so that the course should be slightly underloaded. Second, the rate of decrease in throughput levels is noticeably steep when golf courses are increasingly over-loaded, but the decrease rate starts to level out when golf courses become increasingly under-loaded. Finally, the graph shows that golf courses, regardless of their hole sequence, will show similar trends of change in throughput levels across different tee intervals as long as all the golf courses are balanced.




Figure 3 — Sojourn times (above) and departure times (below) as a function of the group number (y-axis) and the tee interval (x-axis) for a golf course with P3 holes (no wave-up) in the base course (454-434-454-434-454-434). For both the top and bottom tables, the darker color indicates larger numerical value. For the top figure, we see 7 different colors, which respectively denote the sojourn time ranges of $[0,230]$, $(230,240]$, $(240,250]$, $(250,260]$, $(260,270]$, $(270,280]$ and $(280, \infty)$ minutes. The sojourn times below the thick black lines denotes violation of the sojourn time constraint. For the bottom figure, we see 2 different colors, which respectively denote departure times under 840 minutes and departure times over 840 minutes. The number of groups is bounded by the red region according to the departure-time constraint.




Figure 4 — The maximum number of groups that can play each day as a function of the tee interval under various balanced golf course designs, all with P3WU holes. Here is a list of hole sequences of each of the five course designs above: Design 1 (“4 5 4” course): 454-343-454-343-454-343 (this is our base case). Design 2 (“5 5 5” course): 555-343-434-444-444-444. Design 3 (“3 3 3” course): 333-454-444-454-444-454 (concentrated in the beginning). Design 4 (“3 3 3”): 444-454-444-454-444-454 (concentrated in the middle). Design 5 (“3 3 3”): 454-444-454-444-333 (concentrated in the end).

7. A Two-Level Tee Schedule

We have so far assumed that the tee intervals between playing groups are constant. We now report results of extensive simulation experiments showing that more groups can play each day, subject to the same constraints, if the tee intervals start small and increase over the day. However, for operational simplicity we now consider only two fixed tee intervals, a shorter tee interval τ_1 to be used for the first v groups, and then a longer interval τ_2 thereafter.

We report our results in a table and in a graph for the two-level tee schedule study. Table 6 shows throughput optimization results, where we optimize over τ_2 for three given pairs of (τ_1, v) : (6.00,20), (6.50,20), and (7.00,20); i.e., the first

Table 6 The Maximum Number of Groups That Can Play Each Day With a Two-Level Tee Schedule, as a Function of (i) τ_1 the Tee Interval for the First $v = 20$ Groups, (ii) τ_2 , the Tee Interval for all Later Groups, and (iii) the Type of P3 Hole Used

$v = 20$		$\tau_1 = 6.00$			$\tau_1 = 6.50$			$\tau_1 = 7.00$			$\tau_1 = 8.00$			$\tau_1 = 8.20$		
τ_2	P3	P3WU	SP3	P3	P3WU	SP3	P3	P3WU	SP3	P3	P3WU	SP3	P3	P3WU	SP3	P3
7.00	15	16	42	18	24	87	23	46	87	—	—	—	—	—	—	—
7.10	15	16	58	18	25	88	23	84	87	—	—	—	—	—	—	—
7.20	15	16	88	18	46	88	23	85	87	—	—	—	—	—	—	—
7.30	15	20	88	18	86	87	23	85	87	—	—	—	—	—	—	—
7.40	15	38	88	18	85	87	23	84	86	—	—	—	—	—	—	—
7.50	15	50	88	18	85	87	24	84	86	—	—	—	—	—	—	—
7.60	15	56	87	18	84	86	24	83	85	—	—	—	—	—	—	—
7.70	15	61	87	18	84	86	25	83	85	—	—	—	—	—	—	—
7.80	15	65	86	18	83	85	25	82	84	—	—	—	—	—	—	—
7.90	15	66	86	18	83	85	26	82	83	—	—	—	—	—	—	—
8.00	15	67	85	18	82	84	27	81	83	53	—	—	—	—	—	—
8.10	15	69	85	18	82	83	28	80	82	60	—	—	—	—	—	—
8.20	15	69	84	18	81	83	30	80	82	71	—	—	—	—	—	—
8.30	15	71	83	18	80	82	32	79	81	74	—	—	—	—	—	—
8.40	15	70	83	18	80	81	37	79	80	74	—	—	—	—	—	—
8.50	15	70	82	18	79	81	55	78	80	74	—	—	—	—	—	—
8.60	15	70	81	18	79	80	74	77	79	74	—	—	—	—	—	—
8.70	15	70	81	18	78	80	74	77	79	74	—	—	—	—	—	—
8.80	15	70	80	42	77	79	74	76	78	74	—	—	—	—	—	—
8.90	19	70	79	54	77	78	74	76	77	73	—	—	—	—	—	—
9.00	31	69	79	60	76	78	74	75	77	73	—	—	—	—	—	—
9.50	15	66	76	42	74	75	73	73	74	71	—	—	—	—	—	—
10.00	15	63	73	40	71	72	71	70	71	69	—	—	—	—	—	—
τ^*	8.39	6.53	6.79	8.39	6.53	6.79	8.39	6.53	6.79	8.39	6.53	6.79	8.39	6.53	6.79	8.39

Note. For the hole order, the base case is used. The optimal tee intervals are in italicized bold font, while the critical tee interval τ^* for that kind of par-3 hole is shown at the bottom.

20 tee intervals are 6.00, 6.50, and 7.00 minutes respectively, while the remaining ones are τ_2 minutes, which is being optimized.

Figure 5 compares the throughput optimization results obtained when $v = 10$ instead of $v = 20$.

7.1. Evaluating the Effectiveness of Uneven Tee Schedules

Overall, Table 6 and Figure 5 suggest that having a two-level tee schedule makes it possible to have slightly higher throughput levels.




Figure 5 — The maximum number of groups that can play each day with two-level tee schedule, as a function of (i) T_1 , the tee interval for the first v groups, (ii) T_2 , the tee interval for the later groups, and (iii) the type of P3 hole used (i.e. P3, P3WU, or SP3 holes). The cut-off level $v = 10$ is shown by the solid line, while the cut-off level $v = 20$ is shown by a the dashed line. For the hole order, the base case is used. The critical tee interval T^* for that kind of par-3 hole is shown at the top.

In Section 6 we found that the highest throughput levels that can be attained for a single tee interval using P3, P3WU, and SP3 holes were 74, 84, and 87. The optimal tee values (or range of tee values) for these throughput levels were [8.20, 8.50], [7.20, 7.30], and [7.00, 7.10]. In contrast, with a two-level tee schedule, the corresponding highest throughput levels attained were 74, 86, and 88. From the results reported in Table 6 and Figure 5, we see that there is a gain of two groups for P3WU holes and a gain of one group for SP3 holes. The maximum throughput levels stayed the same at 74 for P3 holes.

The optimal (v , τ_1 , τ_2) values were (20, 7.0, [8.60, 9.00]) for P3 holes and (20, 6.50, 7.30) for P3WU holes. For SP3 holes, both (20, 6.0, [7.20, 7.50]) and (20, 6.50, [7.10, 7.20]) yielded maximum throughput. Now, we note that the throughput results are non-degenerate for each scenario; in other words, in each scenario, maximum throughput is attained by multiple ranges of tee times, rather than a single tee time that is a single number.

Figure 5 complements Table 6 by showing that having a two-level tee schedule with $v = 10$ instead of $v = 20$ allows the maximum throughput to be attained by a wider range of tee intervals. For $v = 10$, the maximum levels are just as for $v = 20$ before, the highest throughput levels (i.e. 74 for P3, 86 for P3WU, and 88 for SP3, as described in above paragraph) are attained in all nine graphs in the matrix. In other words, the results suggest that setting $v = 10$ provides a more robust and flexible two-level tee schedule solution than setting $v = 20$.

8. The Advantages of First Hole as a Bottleneck Hole

Our results above confirm the capacity analysis in (Whitt, 2015). As predicted, with the stage-playing-time distributions in Section 3.2, the P3 holes make the course unbalanced, with the P3 holes bottlenecks, having capacity 8.39, whereas the P3WU and SP3 holes make the course roughly balanced, where there should be no serious bottlenecks, with course capacity 6.53 minutes. Hence, to study the possible advantage of putting a bottleneck hole first, we primarily want to focus on the unbalanced model with P3 holes. We also consider the P3WU and SP3 hole models for contrast.

For the unbalanced course designs with P3 holes, the course may be overloaded or underloaded. Our simulations confirm that the course with P3 holes is overloaded when the tee interval is 7.50 minutes, but underloaded when the tee interval is 8.50 minutes, as predicted because it exceeds the theoretical capacity associated with a tee interval of 8.39 minutes. In particular, Table 2 shows that group 75 experiences about 65% of its total waiting time at the first P3 hole, wherever it appears, when the tee interval is 7.50 minutes, which makes the course overloaded with the P3 holes as bottlenecks. The proportion of the delay is only slightly less when the P3 hole is first.

In contrast, Table 2 shows that the percentage drops to about 8% and 11% for all the P3WU and SP3 holes, for which the course is theoretically balanced and slightly underloaded. The observed percentages exceed the 5.6% that would hold if the delay were divided evenly over all 18 holes, but it does not differ too radically. We conclude that the characterization of capacity operates approximately.

Table 7 shows the comparable results when the tee interval is 8.50 minutes, for which all models are theoretically underloaded. Table 7 shows that, for the

Table 7 Simulation Estimates of the Proportion of Waiting Times for Group 75 (in %) Before Starting Play on Each of the 18 Holes for Three Course Designs With the Three Kinds of par-3 Holes: P3, P3WU and SP3

Hole	par	Base case			par-5 first			par-3 first		
		P3	P3WU	SP3	par	P3	P3WU	SP3	par	P3
1	4	0.30%	0.60%	0.90%	5	0.30%	0.60%	0.80%	3	21.90%
2	5	0.70%	1.80%	2.10%	5	1.10%	2.00%	2.80%	3	27.40%
3	4	1.30%	3.00%	3.90%	5	1.50%	3.00%	3.80%	3	24.50%
4	4	1.30%	3.10%	3.60%	3	28.90%	8.30%	8.90%	4	0.60%
5	3	27.00%	8.30%	8.60%	4	0.60%	9.90%	4.60%	5	1.00%
6	4	0.50%	8.90%	4.00%	3	28.10%	9.30%	9.60%	4	1.70%
7	4	1.10%	4.20%	5.10%	4	0.50%	9.60%	4.80%	4	1.50%
8	5	1.10%	3.80%	4.50%	3	25.30%	9.70%	10.10%	4	1.70%
9	4	1.70%	4.70%	6.30%	4	0.60%	9.50%	5.80%	4	1.70%
10	4	1.60%	3.90%	5.10%	4	1.00%	4.70%	5.60%	4	1.90%
11	3	29.90%	9.70%	10.70%	4	1.30%	4.60%	5.20%	5	1.70%
12	4	0.50%	10.30%	4.80%	4	1.40%	4.10%	5.00%	4	2.20%
13	4	1.00%	4.80%	5.50%	4	1.50%	4.20%	4.90%	4	2.00%
14	5	1.20%	3.90%	4.90%	4	1.50%	4.10%	5.00%	4	1.80%
15	4	1.80%	5.30%	6.70%	4	1.60%	4.10%	5.20%	4	1.90%
16	4	1.50%	4.50%	6.00%	4	1.60%	4.20%	6.10%	4	2.10%
17	3	27.00%	9.50%	11.80%	4	1.60%	4.10%	5.90%	4	2.00%
18	4	0.60%	9.80%	5.30%	4	1.70%	4.00%	5.90%	5	2.50%
sum		100%	100%	100%		100%	100%	100%		100%

Note. In all cases the tee interval is 8.50 minutes, under which all three cases are underloaded. The course designs are (i) the base case, (ii) the par-5-holes first and (iii) the par-3 holes first.

unbalanced course with P3 holes, 20–30% of the delay occurs at the first P3 hole and more than 50% occurs at the three P3 holes. The percentages for the P3WU and SP3 holes are about 8% and 11%, as before. Thus, our analysis supports the observations about putting the bottleneck hole first made on p. 97 of (Riccio, 2012). Our simulations show that the ordering of the holes matters very little if the course is approximately balanced, i.e., if the capacities of the individual holes, as defined in (Whitt, 2015), are approximately equal. On the other hand, if the course is unbalanced, with P3 holes as the bottlenecks, and if the course is slightly overloaded, then indeed, most of the delay occurs before starting play when the first hole is a P3.

If a bottleneck hole is first, then indeed the largest delays will be there. Golfers could then wait more conveniently in the clubhouse, while course managers would see the need for longer tee intervals or other remedies. If both the front nine and the back nine start at the clubhouse, so that the first hole on the back nine is hole 10, then it might be desirable to make hole 10 a bottleneck hole as well.

However, we also observe that the order of the holes has relatively little impact on the total congestion. For both overloaded and underloaded courses with P3 holes, the largest delay will occur at the first P3 hole wherever it occurs, but the total delay is not greatly affected by the order. For overloaded courses, the waiting times increase over successive groups, whereas for underloaded courses, the delays tend to reach steady state by about group 50. These observations are supported by Table 8 in this paper, as well as by Table 9 in the Online Companion.

Thus, the advantages of putting a bottleneck hole first on the course are primarily associated with the preferences of golfers and management. For given total delay, where would they like that delay to appear?

9. Conclusions

We have developed a flexible and efficient simulation algorithm for the stochastic model of group play on a general golf course in (Whitt, 2015) and applied it to study operational issues in golf. The group-stage structure greatly simplifies the analysis. Nevertheless, the simulation experiment was challenging to conduct. As explained in Section 4.1, a full design experiment required generating 22×109 stage playing times. The success in carrying out these experiments demonstrate the practical value of both the model and the simulation algorithm. We conclude that the modeling approach and the simulation tool can contribute to better design and management of golf courses.

The simulation experiments substantiate the theoretical characterization of the hole capacities in general models with the usual general hole types. For balanced courses, the simulation experiments show that the order of the hole types matters relatively little.

In Section 5 we reported the results of simulation experiments comparing alternative hole orderings and alternative versions of par-3 holes. Tables 1-4, which show the waiting times of the 75th group, provides some important insights. First, we see that the average waiting times in golf courses are substantially larger with P3 holes than with P3WU or SP3 holes (P3 holes with scaled parameters, so that the capacity is the same as for P3WU), especially when the course with P3 holes is overloaded whereas it is not with the versions.

Table 8 Simulation Estimates of the Mean Waiting Times for a Range of Different Playing Groups (e.g., 5, 10, 15, ...) Before Starting Play on Each of the Par-3 Holes With P3 Holes

Hole #	Scenario 1: base case - P3					Scenario 2: par-3 first case				
	P3 - tee: 7.50					P3 - tee: 7.50				
	h#1	h#5	par-3	h#2-18	All holes	h#1	par-3	h#2-18	All holes	
/group #	par-4			sum	sum			sum	sum	sum
5	0.3	3.7	12.2	12.4	11.8	3.9	8.0	11.8		
10	0.3	8.1	21.4	21.7	20.4	8.3	12.1	20.4		
15	0.3	12.6	28.6	28.9	27.6	12.8	14.8	27.6		
20	0.3	17.1	35.7	36.0	33.8	17.3	16.4	33.8		
25	0.3	21.5	41.6	41.9	39.6	21.8	17.8	39.6		
30	0.4	25.9	47.5	47.9	45.4	26.2	19.2	45.4		
35	0.3	30.3	53.6	53.9	51.4	30.5	20.9	51.4		
40	0.3	34.8	59.4	59.7	57.2	34.8	22.4	57.2		
45	0.3	39.2	64.9	65.3	63.0	39.4	23.5	63.0		
50	0.3	43.6	70.3	70.6	68.3	43.9	24.4	68.3		
55	0.3	48.0	75.9	76.2	74.0	48.3	25.7	74.0		
60	0.3	52.6	81.6	81.9	79.5	52.8	26.7	79.5		
65	0.3	57.0	87.0	87.3	84.7	57.3	27.4	84.7		
70	0.3	61.4	91.9	92.3	90.1	61.8	28.2	90.1		
75	0.3	66.0	97.4	97.7	95.4	66.4	29.0	95.4		
80	0.3	70.7	103.0	103.3	100.6	70.9	29.8	100.6		
85	0.3	75.1	108.0	108.3	105.7	75.3	30.4	105.7		
90	0.3	79.7	113.5	113.8	110.8	79.8	31.0	110.8		
95	0.3	84.1	118.4	118.7	116.0	84.3	31.7	116.0		
100	0.3	88.5	123.5	123.8	120.9	88.7	32.1	120.9		

Note. The courses are over-loaded in the below two scenarios (tee interval: 7.50). Scenario 1 features the base case (i.e., “4 5 4 ...”), and Scenario 2 features the par-3 first/bottleneck-first case (i.e., “3 3 ...”). The sums of the mean wait times across all 18 holes in each scenario are highlighted in bold.

When the tee interval is 7.50, which makes the P3 hole overloaded, the expected waiting time for the 75th group ranges between 95 and 97 minutes with P3 holes. Moreover, P3 holes contribute 90% of the 95–97 minutes of waiting times; 68% of the total waiting time is concentrated on the first P3 hole alone. These results are consistent with the known performance in a standard queueing network with one or more bottleneck queues.

On the other hand, the expected waiting time for the 75th group ranges only between 23 and 26 minutes with P3WU or SP3 holes, where the holes are now slightly underloaded. Making the course underloaded significantly reduces the waiting times. Nevertheless, the wave-up rule is still advantageous. For instance, when the tee interval increases to 8.50, the 75th group needs to wait between 26 and 31 minutes with P3 holes, but no more than 14 minutes with P3WU or SP3 holes.

In Section 6 we used simulation to find the optimal tee schedule according to the optimization framework proposed in Section 1.1.2. Table 5 and Figure 4 provide important insights, which are summarized in Section 6.2. We perform sensitivity analysis with the simulation optimization to show that performance degrades much more rapidly if the tee interval is too short than if it is too long. Overall, we find that having a P3WU hole allows about 10 more groups to play (84 instead of 74) than having a regular P3 hole.

We also observed that there is a slight cost of the wave-up rule in terms of throughput (measured in terms of the total number of group that can play on the course each day, subject to the constraints in Section 3.1.2), because SP3 holes can allow 2 more groups to play than a P3WU hole, even though the capacity of the P3WU hole is slightly larger. Given that the cost is slight and is primarily due to extra variability caused by the inconsistent waving up (because the following group may or may not be ready to play), the advantage of the wave-up rule over P3 holes still remains significant. Again, the throughput levels attained with each type of par-3 hole remain consistent across all course designs. Not only is the maximum throughput achieved when a golf course is slightly underloaded, but that throughput is more severely penalized when the traffic intensity is overloaded. That conclusion is corroborated by Figure 4, which shows (i) the steep slope with the tee between 5.00 and 7.00, and (ii) noticeably flat slope with the tee between 7.00 and 9.00 minutes.

In Section 7 we investigated the effects of having an uneven tee schedule. In particular, we explored the impact of having tighter tee schedule for the first few groups and looser tee schedule for the later groups. For the operational simplicity, we ran experiments with two fixed intervals. In order to draw meaningful conclusions, we performed the simulation over a large number of alternatives. We observed that having a two-level tee schedule increases the throughput level for P3WU and SP3 holes by 2 groups, while leaving no change to the throughput level for P3 hole.

Finally, in Section 8 we studied the impact of placing a bottleneck P3 hole first on an unbalanced course. As suggested by (Riccio, 2012), our results show that most of the delay will occur at the first hole, before the groups start to play. That can make management more aware of the need for adjustments in the tee interval and provide more convenient waiting for golfers in the clubhouse. On the other hand, our results also show that the hole sequence has little impact on the overall congestion level of golf courses.

Overall, we have shown that the stochastic model proposed in (Whitt, 2015) can be effectively simulated and used to investigate ways to manage the pace of play on a general golf course. There are many directions for future research. Especially promising is fitting the model to data from golf courses.

It remains to collect data, but we need more refined data than just the starting and finishing times for each group on each hole, as in (Riccio, 2014b). With better data, we would then directly estimate the stage playing times and then directly test whether they properly characterize the overall pace of play.

Acknowledgments

MoonSoo Choi and Qi Fu performed this work while they were graduate students at Columbia University. Ward Whitt thanks NSF for research support (NSF grants CMMI 1265070 and 1634133).

References

- Fu, Q., & Whitt, W. (2015). Analyzing the pace of play in golf. *Journal of Sports Analytics*, 1(1), 43–64. doi:10.3233/JSA-150004
- Hopp, J., & Spearman, M.L. (1996). *Factory physics*. Chicago, IL: Irwin.
- Kimes, S., & Schruben, L. (2002). Golf revenue management: a study of tee time intervals. *Journal of Revenue and Pricing Management*, 1(2), 111–120. doi:10.1057/palgrave.rpm.5170014
- Riccio, L. (2012). Analyzing the pace of play in golf: The golf course as a factory. *International Journal of Golf Science*, 1, 90–112. doi:10.1123/ijgs.1.2.90
- Riccio, L. (2014a). *Golf's pace of play bible*. New York, NY: Three/45 Golf Association.
- Riccio, L. (2014b.). The status of pace of play in American Golf. New York, NY: Three/45 Golf Association.
- Serfozo, R. (1999). *Introduction to stochastic networks*. New York, NY: Springer.
- Tiger, A.A., & Ellerbrook, J.E. (2016). Improving golf pace of play using time study analysis: Influencing factors on the green and tee box. *International Journal of Golf Science*, 5(2), 135–151. doi:10.1123/ijgs.2016-0006
- Tiger, A.A., & Salzer, S. (2004). Daily play at a golf course: using spreadsheet simulation to identify system constraints. *INFORMS Transactions on Education*, 4(2), 28–35. doi:10.1287/ited.4.2.28
- Tiger, A.A., Trent, C., & Haney, J. (2015). Improving golf course flow using operations research. *Journal of Management Information and Decision Sciences*, 18(2), 81–93.
- Walrand, J. (1988). *Introduction to queueing networks*. Englewood Cliffs, NJ: Prentice Hall.
- Whitt, W. (2015). The maximum throughput on a golf course. *Production and Operations Management*, 24(5), 685–703. doi:10.1111/poms.12289