Using Simulation to Help Manage the Pace of Play in Golf

MoonSoo Choi
Industrial Engineering and Operations Research, Columbia University,
New York, NY, USA {moonsoo.choi@columbia.edu}

Qi Fu
Credit Suisse, Eleven Madison Avenue,
New York, NY, USA {qf2112@columbia.edu}

Ward Whitt
Industrial Engineering and Operations Research, Columbia University
New York, NY, USA {ww2040@columbia.edu}

April 29, 2017

Abstract

To help design and manage golf courses, Whitt (2015) developed a new stochastic queueing network model of a conventional 18-hole golf course and characterized the capacity of each hole. The primitives in that model are the random times required for each group of golfers to play each stage on each hole, where a stage is an appropriate set of steps. The model is relatively high level, but the model still captures the important property that multiple groups can play at the same time on the longer holes, subject to precedence constraints. In this paper, we develop a flexible and efficient simulation algorithm of that model and apply it to study important operational issues in golf. We substantiate the characterization of hole capacities in Whitt (2015) and study: (i) the costs and benefits of a wave-up rule designed to increase the pace of play and (ii) the costs and benefits of alternative tee schedules, i.e., the intervals between successive groups starting play on the first hole. We estimate that the wave-up rule can increase the daily number of groups that can play each day by about 13%, while a two-level tee schedule rule can increase the daily number of groups that can play each day by about 3%. We show that the costs of making the tee interval too short tend to be far greater than the costs of making it too long.

Keywords: pace of play in golf, using simulation to manage golf courses, stochastic models of golf, waving up in golf, alternative tee schedules on golf courses

Contact Author: Ward Whitt, ww2040@columbia.edu

1 Introduction

In this paper we develop a flexible and efficient simulation algorithm to compute the performance measures for a stochastic model of group play on a golf course over a day and apply it to study important operational issues in golf. We use the stochastic model of group play in Whitt (2015). Our work follows previous work in this direction by Kimes and Schruben (2002), Tiger and Salzer (2004), Tiger and Ellerbrook (2016) and Riccio (2012, 2013, 2014a,b).

1.1 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 goal in golf is for each golfer to put the ball into the hole on the green using as few strokes (shots) as possible. The par value, i.e., the target number of strokes to play the hole, ranges from 3 to 5, with the values increasing with the length and difficulty of the hole. A hole is rated par 4 because good play should require 4 shots: one from the tee, one from the fairway and two more to clear the green (put it in the hole on the green). The design of golf courses varies, but there usually are 12 of the average-length par-4 (P4) holes and 3 each of the shorter par-3 (P3) holes and longer par-5 (P5) holes. Counter to initial intuition, the shortest P3 holes tend to be the bottlenecks holes, where groups of golfers experience the greatest delay.

That phenomenon can be explained by the different numbers of groups that are allowed to play on each hole at the same time. On the longer holes, more than one group can play on the hole at the same time, provided they follow appropriate rules of play, because then the groups do not interfere with each other; e.g., the group in front is not in danger of getting hit by a ball from the group behind. Typically, two groups can be playing on a P4 hole at the same time, while three groups can be playing on a P5 hole at the same time. On the other hand, a conventional P3 hole is more elementary, because only one group can play on it at the same time. That explains why the P3 holes typically are the bottleneck holes.

In an effort to increase the pace of play, some golf courses have adopted a special wave-up rule to use on P3 holes, which 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 a group has hit its tee shots and walked up to their balls near the green, they should wait before hitting their next shot 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 has already arrived and 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 par-3 holes and improve the pace of play.

However, the wave-up rule actually is somewhat complicated. Because a group waits before completing play to allow the next group to first hit its tee shots only if that next group is ready, the wave-up inevitably is not applied consistently. Some groups may experience the good fortune of being waved up by the group in front of them, but then not in turn having to wave up the group behind, because they are not yet ready. That is good for that particular group. However, the next group may have to wave up the following group, even though they themselves were not waived up. Thus, the wave-up is inevitably applied inconsistently, which introduces 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. A key to obtaining an efficient algorithm is exploiting the high-level model of group play developed in Whitt (2015).

1.2 The Simulation Study

This paper is an important complement to the earlier work in Whitt (2015), because we develop and apply an efficient simulation algorithm to compute all the standard performance measures for this model applied to n groups playing on a general 18-hole golf course, as a function of the basic model parameters and course designs.

1.2.1 A Flexible Scalable 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, we think that the level of detail in that model is appropriate for operations management issues such as the ones we study.

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, (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 i.i.d. simulation replications.

The simulation algorithm is also designed to be efficient for large scale. 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 standard course design with 12 P4 holes, 3 P3 holes and 3 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 6120 stage playing times for one day of golf.

Since the model is a stochastic model, we require multiple (independent and identically distributed) replications of our simulation. We consistently used 2000 replications. (Table 3 shows that the halfwidths 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 $6,120 \times 2000 = 12.24 \times 10^6$ 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 3 P3, 12 P4 and 3 P5 holes, but we think that the relevant number of arrangements might be only 20. Second, there are 3 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. (That becomes larger when we consider two-level tee intervals.) As is, the number of course designs becomes $20 \times 3 \times 30 = 1800$. That requires generating $(12.24 \times 10^6) \times (1.8 \times 10^3 \approx 22 \times 10^9$ stage playing times. With MATLAB, time measurements indicate that a full simulation required about one full week. We elaborate on the simulation methodology in §2 and §3 of the online supplement.

1.2.2 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 we think these model parameters are 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 (in §2.2) that make the full course roughly balanced, i.e., so that the capacities (determined in Whitt (2015)) of all holes are approximately equal. Then we apply simulation to verify that the performance is approximately independent of the order of the holes.

1.2.3 Applications of the Simulation Algorithm

Then we show how the simulation algorithm can be applied for any 18-hole course to optimally choose the tee interval (the interval between the times successive groups are scheduled to begin play) and the number n of groups that can play subject to specified constraints on the hours of daily operation and the maximum allowed expected sojourn time, i.e., the maximum expected time for any group to play on the course. We use the common goal of achieving a 4-hour limit for all groups. Simulation examples, show that there is much greater cost, in terms of the number of groups that can play each day, from choosing the tee interval too short than from choosing it too long.

We then investigate the performance consequence for different designs of the par-3 holes. Our base case is a model in which the parameters for the three hole types are chosen so that the course is balanced by the definitions in Whitt (2015) for the P3WU, P4 and P5 holes. We compare that model to the associated model in which the wave-up rule is not used on the par-3 holes. We show that the balanced course composed of P3WU, P4 and P5 holes significantly outperforms the associated unbalanced course in which the wave-up rule is not used, and the P3 holes are bottlenecks. We also contrast these models with the balanced course, in which the balanced course is achieved by scaling the parameters of the par-3 hole, which we refer to as SP3 holes. We show that the performance of the balanced course achieved by the wave-up rule is not quite as good as the performance of the balanced course achieved by simply scaling the parameters. We include the model with scaling, not because we think it can be achieved, but to expose some of the costs of the wave-up rule. For example, we see that the wave-up rule tends to increase the delays at the following hole.

Thus, we carefully examine the costs and benefits of the wave-up rule on the par-3 holes. Overall, we find that the wave-up rule provides a significant performance benefit, allowing about 10 more groups to play on each day (84 instead of 74). For our model this is a 13% improvement, which is consistent with the conclusion reached by Tiger and Salzer (2004), but it is far less than the 28% difference in course capacities determined from Whitt (2015); see §2.2. We regard the 13% as a reasonable rough estimate of the performance gain from the wave-up rule on a typical golf course. We further find that scaling the par-3 hole parameters provides a slight performance benefit by allowing 86 groups to play on the field, instead of the 84. Nevertheless, the cost of the wave-up rule is slight compared to its advantage.

Finally, we study alternative tee schedules (the intervals between the times successive groups are scheduled to start play on the first hole). For simplicity, we examined a number of cases with a two-level tee policy. In particular, we assign a shorter fixed tee interval for several groups in the beginning, then we assign a longer interval for all later groups. We find that such a policy can increase the throughput slightly (by about two groups).

1.3 Contributions

In summary, our main contributions in this paper are to:

(i) develop a flexible and efficient simulation algorithm for the full stochastic model of group play on a general golf course in Whitt (2015);

(ii) verify that the definition of hole capacities in Whitt (2015) is meaningful and useful; i.e., if the course is constructed to be balanced as defined there, then the performance reflects that, being approximately independent of the order of the different holes;

(iii) formulate and implement a simulation optimization algorithm to specify the optimal number of groups to play each day and the associated tee interval, in order to maximize throughput subject to constraints on the expected sojourn times of all groups;

(iv) 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;

(v) systematically study the wave-up rule on par-3 holes by examining its impact on the performance of the golf course operations.

(vi) examine whether, and to what extent, starting off with a shorter tee interval followed by a longer tee interval can allow more groups to play per day.

1.4 Organization

In §2 we review the stochastic model of group play proposed in Whitt (2015). In §3 we report results of simulation experiments comparing alternative course designs. In §4 we introduce the simulation optimization algorithm to choose an optimal tee interval subject to constraints. In §5 we investigate the possible advantage of an uneven tee schedule, using a shorter interval and then shifting to a longer one. Finally, in §6 we draw conclusions. Additional material appears in an online companion.

2 The Stochastic Model of Group Play

In this section we review the stochastic model of group play proposed in Whitt (2015).

2.1 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 the parameters to change 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 typical. Then we review the P3 and P3WU models. For the P5 model, we refer to §6 of Whitt (2015).

2.1.1 The Steps and Stages for a Par-4 Hole

We first describe the steps of group play on a par-4 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

$$T \rightarrow W_1 \rightarrow F \rightarrow W_2 \rightarrow G. \tag{2.1}$$

The first step T is the tee shot (one for each member of the group); the second step W_1 is walking up to the balls on the fairway; the third step F is the fairway shot; the fourth step W_2 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 par-4 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:

$$(T, W_1) \to F \to (W_2, G) \tag{2.2}$$

Stage 1 is (T, W_1), stage 2 is F and stage 3 is (W_2, G). We use this particular aggregation, because it turns out to be 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 $n \geq 1$, 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, for $n \geq 1$, 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 par-4 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.

2.1.2 The Stochastic Model of Group Play on a Par-4 Hole

The description of group play on a P4 hole is the basis for a mathematical model.

Let A_n be the *arrival time* of the n^{th} group 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, $1 \le j \le 3$; these are called the *stage playing times*. The *mathematical model data* for a par-4 hole consists of a sequence of 4-tuples: $\{(A_n, S_{1,n}, S_{2,n}, S_{3,n}) : n \ge 1\}$, where the four components for each n are nonnegative random variables. We think of $\{S_{j,n}\}$, $1 \le j \le 3$, as being three independent sequences of i.i.d. random variables.

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:

$$\begin{aligned} B_n &\equiv \max\{A_n, F_{n-1}\}, \qquad T_n \equiv B_n + S_{1,n}, \\ F_n &\equiv \max\{T_n, G_{n-1}\} + S_{2,n} \qquad \text{and} \qquad G_n \equiv F_n + S_{3,n}, \end{aligned} \tag{2.3}$$

where $\equiv$ denotes "equality be definition." As initial conditions, assuming that the system starts empty, we set $F_0 \equiv G_0 \equiv 0$. The two maxima capture the two precedence constraints.

The model in (2.3) extends directly to any number of such single-hole models in series. We simply let the completion times G_n from one hole be the arrival times at the next hole.

2.1.3 The Stochastic Model of 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 par-3 hole, with or without wave-up:

$$T \to W \to G. \tag{2.4}$$

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 $S_{i,n}$ for group n on the three stages as before, the total time for group n to play the hole is $X_n = S_{1,n} + S_{2,n} + S_{3,n}$.

2.1.4 The Stochastic Model of 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 a group has hit its 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 has not yet arrived at the hole, then the current group immediately starts 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 E_n, defined by

$$E_n \equiv \{A_n \leq \max\{W_{n-1}, G_{n-2}\} < T_n\}, \qquad (2.5)$$

and let E_n^c be its complement, $n \geq 1$. If group n is the last scheduled group, then let $A_{n+1} \equiv \infty$ (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:

$$\begin{aligned}
B_n &\equiv \max\{W_{n-1}, G_{n-2}\}1_{E_n} + \max\{A_n, G_{n-1}\}1_{E_n^c}, \quad n \geq 2, \quad B_1 \equiv A_1, \\
T_n &\equiv B_n + S_{1,n}, \quad W_n \equiv T_n + S_{2,n}, \quad n \geq 1, \quad \text{and} \\
G_n &\equiv \max\{W_n, G_{n-1}\}1_{E_{n+1}^c} + T_{n+1}1_{E_{n+1}} + S_{3,n} \\
&= \max\{W_n, G_{n-1}\} + S_{1,n+1}1_{E_{n+1}} + S_{3,n},
\end{aligned} \qquad (2.6)$$

where 1_E is the indicator function of the event E, i.e., $1_E = 1$ if E occurs and $1_E = 0$ otherwise. As initial conditions, again assuming that the system starts empty, we set $W_0 \equiv G_0 \equiv G_{-1} \equiv 0$. If n is the last group, then, instead, $G_n \equiv \max\{W_n, G_{n-1}\} + S_{3,n}$. (The event E_{n+1} cannot occur.)

Note that the expression for B_n involves G_{n-2}, because only two groups can be playing on the hole at the same time. Observe that the event E_n in (2.5) actually simplifies. By the first two lines of (2.6),

$$T_n = B_n + S_{1,n} \geq W_{n-1} + S_{1,n} > W_{n-1},$$

so that $E_n = \{A_n \leq W_{n-1} \vee G_{n-2}\}$. 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.

2.1.5 The Stochastic Model of Group Play on a Par-5 Hole

A par-5 hole is considerably more complicated than a par-4 hole, primarily because three groups can play at the same time instead of two. We model the par-5 hole to have seven steps (instead of five steps in par-4 holes), and we group the steps into five stages (instead of three stages in par-3 holes). In particular, we represent the Par-5 hole as:

$$(T, W_1) \rightarrow F_1 \rightarrow W_2 \rightarrow F_2 \rightarrow (W_3, G) \tag{2.7}$$

.

The following is the mathematical representation of the six-part recursion for the par-5 hole:

$$\begin{aligned}
B_n &\equiv \max\{A_n, F_{1,n-1}\}, \qquad T_n \equiv B_n + S_{1,n}, \\
F_{1,n} &\equiv \max\{T_n, F_{2,n-1}\} + S_{2,n} \qquad \text{and} \qquad W_{2,n} \equiv F_{1,n} + S_{3,n} \\
F_{2,n} &\equiv \max\{W_{2,n}, G_{n-1}\} + S_{4,n} \qquad \text{and} \qquad G_n \equiv F_{2,n} + S_{5,n},
\end{aligned} \tag{2.8}$$

Again, more detailed description on the mathematical model of Par-5 holes are available in §6 of Whitt (2015).

2.2 The Stage Playing Time Distributions

For all stages of all holes, we assume that the stage playing times $S_{i,n}$ are mutually independent random variables with a symmetric triangular ($Tri(m, a)$) distribution, as in §4.2 of Whitt (2015), but we also use the modification to allow for a lost ball in the first stage of each hole, as discussed in §4.3 of Whitt (2015). The parameters depend on the stage and the hole type.

The triangular $Tri(m, a)$ distribution has the parameter pair (m, a); it 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 $m \pm a$, 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 $a^2/6$. 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 $m^2/24$).

In particular, for all P3, P4 and P5 holes, we initially let the mean values of the triangular distributions stage playing times (for the 3, 3 and 5 stages) 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 S_2 is symmetric on $[0.5, 3.5]$, having mean 2.0, whereas for the fourth stage of a P5 hole, the density of S_4 is symmetric on $[0.0, 2.67]$, 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 $p = 0.05$ 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 $13.333 + 0.200 = 13.533$.

From Theorem 1 and §4.3 of Whitt (2015), we can calculate the limiting cycle time in (6) of Whitt (2015) is 6.5325 for a P4 hole. (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; see §2.3 of Whitt (2015). The reciprocal of the limiting cycle time is the capacity of the 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 in (6) of Whitt (2015) 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 6.5325 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 1.0177 to get the adjusted mean vector $(4.0709, 2.0355, 2.0355, 1.3569, 4.0709)$, 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, $1/6.5325$. 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, this 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 $6.5325/8.167 \approx 0.8$ 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 6.7925. Hence, this SP3 model is still slightly a bottleneck, being about 6% larger than 6.5325. 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.

3 Comparing Alternative Course Designs

We conducted extensive simulation experiments comparing alternative course designs. We primarily focused on the common case in which there are 12 identical P4 holes, 3 identical P5 holes and 3 identical P3 holes, considering each of the P3WU, P3 and SP3 options in §2.2. 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 6.5325. 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 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. The course designs are (i) the base case, having P3 and P5 holes appearing alternately, separated by P4 holes:

$$4\,5\,4 \quad 4\,3\,4 \quad 4\,5\,4 \quad 4\,3\,4 \quad 4\,5\,4 \quad 4\,3\,4,$$

(ii) the par-5-holes first:

$$5\,5\,5 \quad 3\,4\,3 \quad 4\,3\,4 \quad 4\,4\,4 \quad 4\,4\,4 \quad 4\,4\,4$$

and (iii) the par-3 holes first:

$$3\,3\,3 \quad 4\,5\,4 \quad 4\,4\,4 \quad 4\,5\,4 \quad 4\,4\,4 \quad 4\,5\,4$$

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

Of course, the 75th group is not the only group that exhibits substantial difference in waiting times between having $P3$ holes and having $P3WU$ or $SP3$ holes. Although the earlier groups may not experience as much difference as the 75th group, all groups do exhibit some degree of differences in waiting times. In fact, the waiting time differences become larger for the later groups following the 75th group. However, we only report the results of the 75th group in this paper to keep our discussion brief. Additional waiting time statistics for other groups are available in §4.2 of the Online Companion.

3.1 Shorter Tee Intervals: Heavy 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 in Whitt (2015) 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. This shows the impact of the complexity of the wave-up rule.

Table 2 gives the corresponding estimated proportion of the total waiting time at each hole for tee intervals 7.50, 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.

We do remark that, even though the performance is consistent with a balanced design, we do

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 $SP3$. 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.

	base case				par-5 first				par-3 first			
hole	par	P3	P3WU	SP3	par	P3	P3WU	SP3	par	P3	P3WU	SP3
1	4	0.34	0.31	0.33	5	0.26	0.11	0.11	**3**	**66.41**	**1.68**	**1.24**
2	5	0.60	0.61	0.68	5	0.71	0.11	0.11	**3**	**12.59**	**2.72**	**2.01**
3	4	0.94	0.97	1.02	5	0.88	0.11	0.11	**3**	**9.14**	**2.80**	**2.38**
4	4	0.96	0.94	0.92	**3**	**66.64**	**0.11**	**0.11**	4	0.17	1.97	1.15
5	**3**	**65.00**	**1.92**	**2.24**	4	0.16	0.11	0.11	5	0.31	1.12	0.94
6	4	0.18	*1.99*	1.12	**3**	**13.19**	**0.11**	**0.11**	4	0.46	1.39	1.27
7	4	0.35	1.17	1.09	4	0.18	0.18	0.11	4	0.44	1.19	1.10
8	5	0.41	1.02	0.94	**3**	**9.52**	**0.11**	**0.11**	4	0.45	1.12	1.10
9	4	0.58	1.44	1.33	4	0.18	0.11	0.11	4	0.51	0.51	1.15
10	4	0.48	1.20	1.08	4	0.31	0.11	0.11	4	0.49	1.18	1.10
11	**3**	**14.28**	**2.13**	**2.69**	4	0.40	0.11	0.11	5	0.49	1.04	0.97
12	4	0.20	*2.04*	1.24	4	0.42	0.11	0.11	4	0.65	1.38	1.30
13	4	0.34	1.28	1.17	4	0.46	0.11	0.11	4	0.54	1.18	1.15
14	5	0.40	1.04	0.97	4	0.51	0.11	0.11	4	0.53	1.16	1.18
15	4	0.60	1.38	1.41	4	0.44	0.11	0.11	4	0.56	1.21	1.18
16	4	0.51	1.19	1.11	4	0.49	0.11	0.11	4	0.53	1.16	1.16
17	**3**	**10.40**	**2.11**	**2.56**	4	0.52	0.11	0.11	4	0.51	1.01	1.08
18	4	0.20	*2.17*	1.23	4	0.56	0.11	0.11	5	0.64	1.36	1.35
sum		97.74	24.90	23.13		95.81	24.96	23.54		95.39	25.81	22.82

see that the average waiting times are larger for the $P3WU$ and $SP3$ holes than at the P4 and P5 holes. The ordering between $SP3$ and $P3WU$ depends on the case: $SP3$ produces longer waits in the base case, but not in the other two cases.

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, 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 $1.96s/\sqrt{2000}$, where s is the estimated standard deviation, because the number of i.i.d. 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.

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 $SP3$. 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.

	base case				par-5 first				par-3 first			
hole	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%

3.2 Longer Tee Intervals: Lighter Loads

We now show that the performance is quite different when all three par-3 holes make the course underloaded. Table 4, similarly to 1, gives the estimated mean waiting times for the longer tee intervals 8.50, 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 by group 50. However, the fact that the $P3$ holes are bottlenecks is again clearly evident from the tables. 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 4. Again the waiting time is approximately the same at the following $P4$ hole as at the $P3WU$ hole itself. More detailed statistics, which include analogs of Tables 2 and 3 for the longer tee 8.50, are available in the Online Companion §4.1: *Additional Wait Time Statistics for Longer Tee Interval.*

We supplement the tables above by also showing histograms of the total waiting times on all

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 $SP3$. 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

	base case				par-5 first				par-3 first			
hole	par	P3	P3WU	SP3	par	P3	P3WU	SP3	par	P3	P3WU	SP3
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
HW		0.52	0.32	0.35		0.51	0.33	0.36		0.50	0.34	0.34

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.

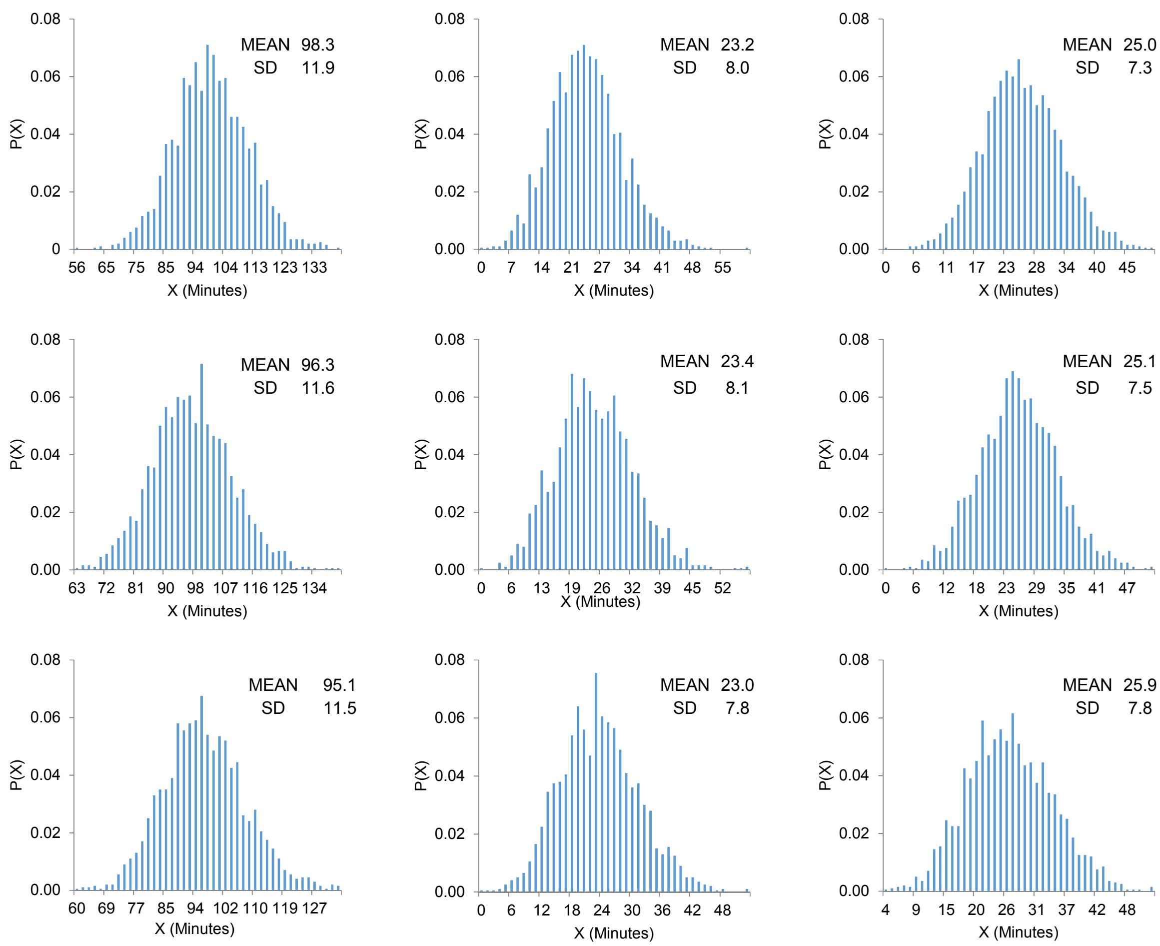

Figure 1: A 3 x 3 collage of histograms of total waiting times for group 75 over 2000 simulation replications. Each row represents three different course designs: (i) the base case, (ii) the par-5-holes first and (iii) the par-3 holes first. Each column represents three different types of par-3 holes in the golf course: $P3$, $SP3$, and $P3WU$. In all of the nine histograms, the tee interval is 7.50 minutes.

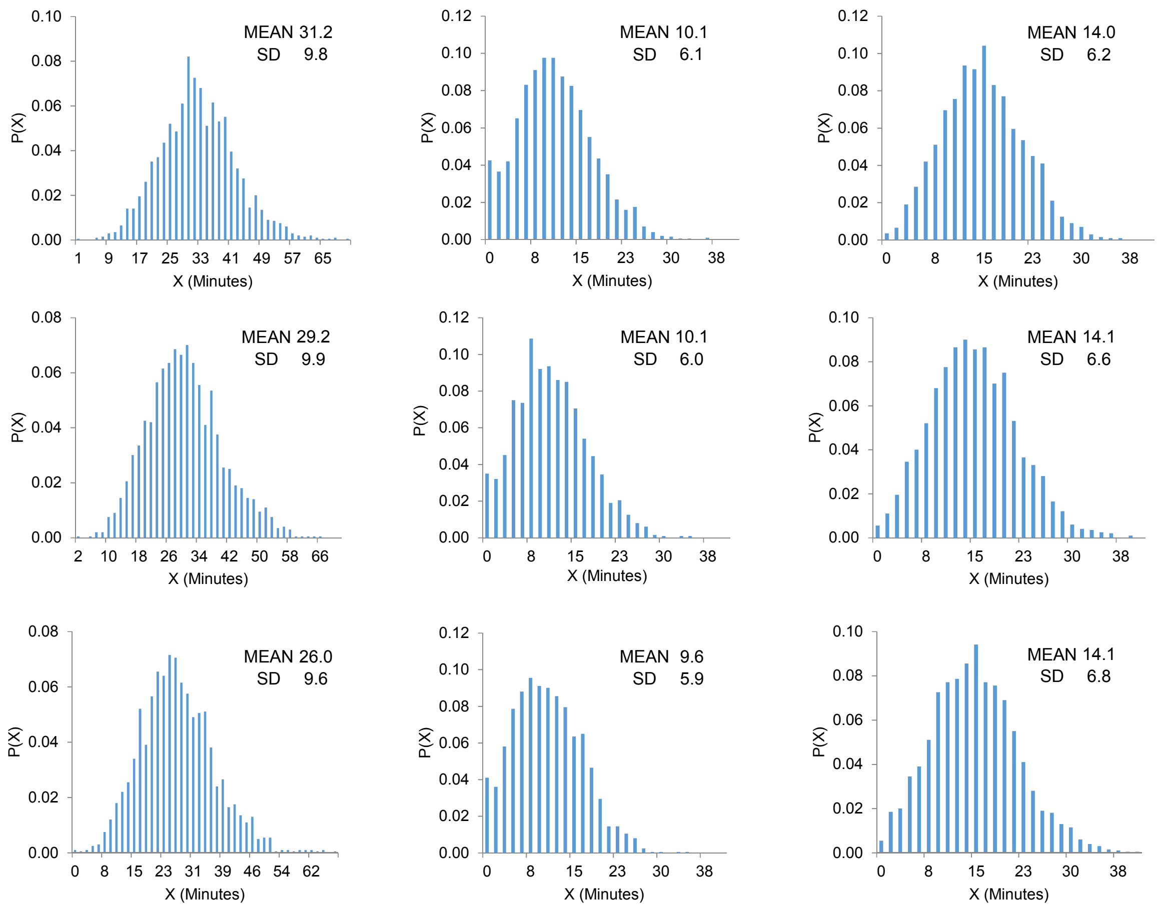

Figure 2: A 3 x 3 collage of histograms of total waiting times for group 75 over 2000 simulation replications. Each row represents three different course designs: (i) the base case, (ii) the par-5-holes first and (iii) the par-3 holes first. Each column represents three different types of par-3 holes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

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 $SP3$. 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.

	base case				par-5 first				par-3 first			
hole	par	P3	P3WU	SP3	par	P3	P3WU	SP3	par	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

4 Simulation Optimization of the Tee Schedule

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.

4.1 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. 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 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 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

$$\text{maximize } n \text{ such that}$$

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

$$G(\tau, n) \leq \delta = 840$$

$$\text{for} \quad n \in \mathbb{Z} \quad \text{and} \quad \tau \geq 0.$$

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 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 §3 by performing the optimization as a function of the hole order and the type of par-3 hole used.

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 5.0 minutes (very overloaded) up to 9.5 minutes (underloaded). 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 Appendix.

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 (ii) the type of par-3 hole used. 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.

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

4.2 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 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 ($\rho \approx 6.53/7.30 \approx 0.90$). 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 $\rho \approx 6.79/6.95 \approx 0.977$. 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.

Second, 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 and 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 τ^* at the bottom. Fourth, from the close agreement from the left and right sides of Table 5, we 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 in 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). More details are available in §3 of the Online Companion.

4.3 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 6.5325 minutes. As in Table 5, the maximum number that can play is achieved for tee intervals that lie between 7.10 and 7.50, 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

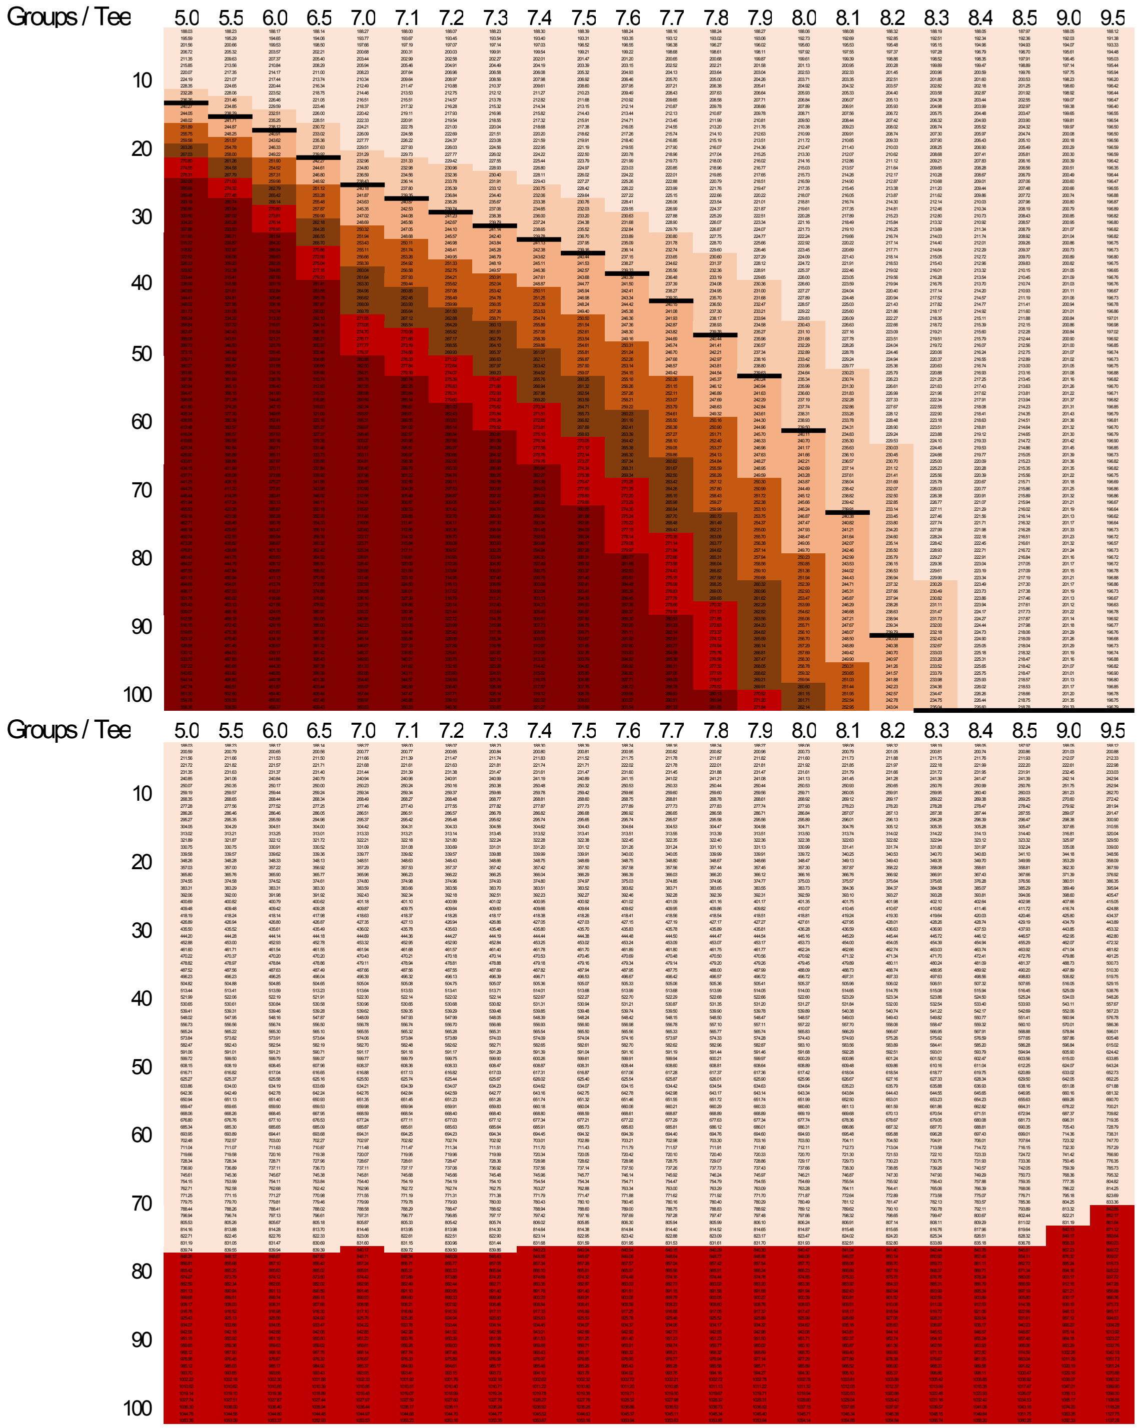

Figure 3: Sojourn times (above) and departure times (below) as a function of the group number (y-axis) and the tee interval (x-axis) for a golf course with P3 holes (no wave-up) in the base course (454-434-454-434-454-434). For both the top and bottom tables, the darker color indicates larger numerical 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 denote 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.

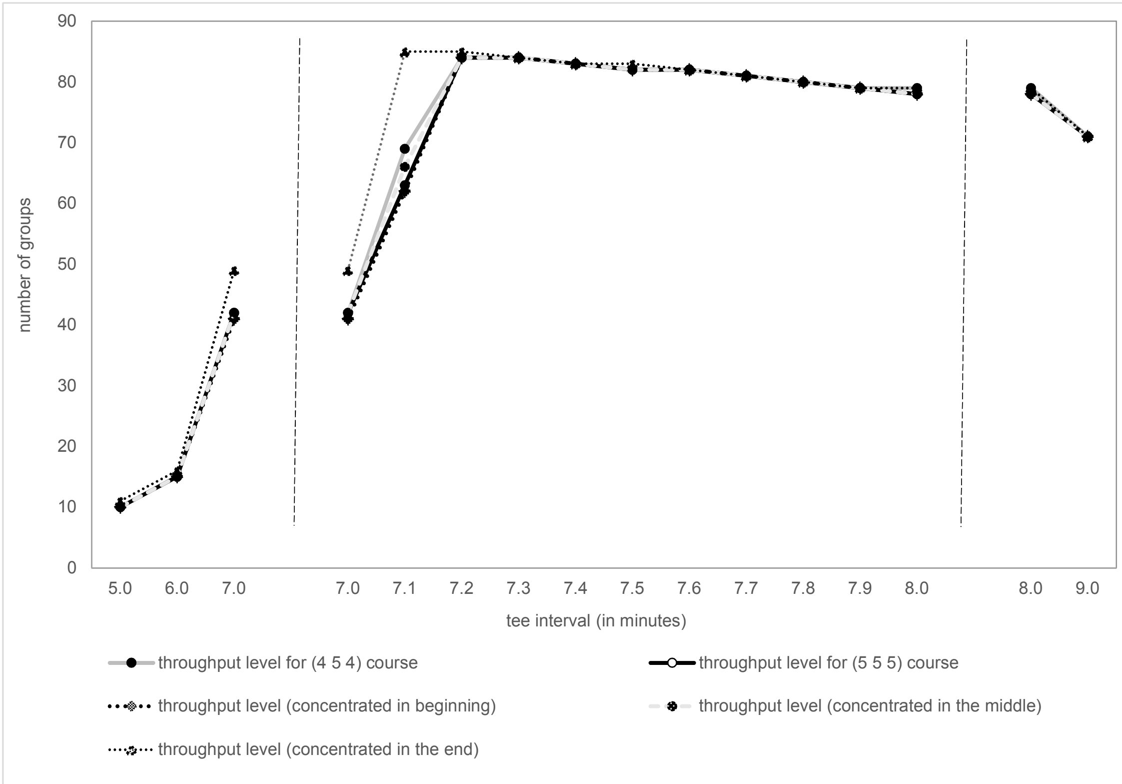

Figure 4: The maximum number of groups that can play each day as a function of the tee interval under various balanced golf course designs, all with P3WU holes. Here is a list of hole sequences of each of the five course designs above:
Design 1 (“4 5 4” course): 454 – 434 – 454 – 434 – 454 – 434 (this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

similar trends of change in throughput levels across different tee intervals as long as all the golf courses are balanced.

5 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 ν 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, ν): $(6.00, 20)$, $(6.50, 20)$, and $(7.00, 20)$; i.e., the first 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 $\nu = 10$ instead of $\nu = 20$.

5.1 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.

In Section 4 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 (ν, τ_1, τ_2) values were $(20,\ 7.00,\ [8.60, 9.00])$ for P3 holes and $(20,\ 6.50,\ 7.30)$ for P3WU holes. For SP3 holes, both $(20,\ 6.00,\ [7.20, 7.50])$ 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 $\nu = 10$ instead of $\nu = 20$ allows the maximum throughput to be attained by a wider range of tee intervals. For $\nu = 10$, the maximum levels are just as for $\nu = 20$ 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 $\nu = 10$ provides a more robust and flexible two-level tee schedule solution than setting $\nu = 20$.

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 $\nu = 20$ groups, (ii) τ_2, the tee interval for all later groups, and (iii) the type of P3 hole used. For the hole order, the base case is used. The optimal tee intervals are in bold, while the critical tee interval τ^* for that kind of par-3 hole is shown at the bottom.

$\nu = 20$	$\tau_1 = 6.00$			$\tau_1 = 6.50$			$\tau_1 = 7.00$			τ_1=8.00	τ_1=8.20
τ_2 values	P3	P3WU	SP3	P3	P3WU	SP3	P3	P3WU	SP3	P3	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	70
8.30	15	**71**	83	18	80	82	32	79	81	**74**	**74**
8.40	15	70	83	18	80	81	37	79	80	**74**	73
8.50	15	70	82	18	79	81	55	78	80	**74**	**74**
8.60	15	70	81	18	79	80	**74**	77	79	**74**	73
8.70	15	70	81	18	78	80	**74**	77	79	**74**	73
8.80	15	70	80	42	77	79	**74**	76	78	**74**	73
8.90	19	70	79	54	77	78	**74**	76	77	73	73
9.00	**31**	69	79	**60**	76	78	**74**	75	77	73	73
9.50	15	66	76	42	74	75	73	73	74	71	71
10.00	15	63	73	40	71	72	71	70	71	69	68
τ^*	8.39	6.53	6.79	8.39	6.53	6.79	8.39	6.53	6.79	8.39	8.39

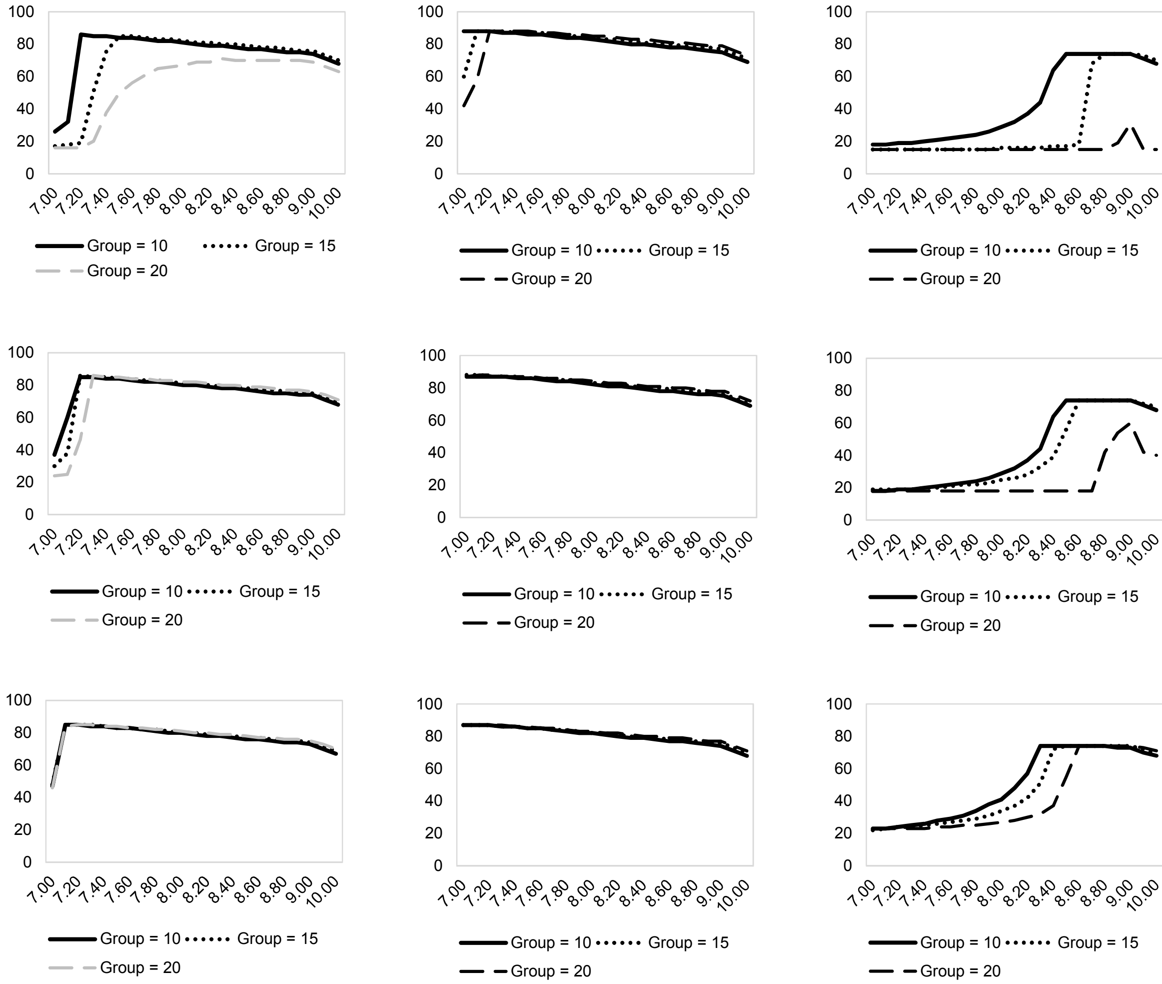

Figure 5: The maximum number of groups that can play each day with two-level tee schedule, as a function of (i) τ_1, the tee interval for the first ν groups, (ii) τ_2, the tee interval for the later groups, and (iii) the type of P3 hole used (i.e. P3, P3WU, or SP3 holes). The cut-off level $\nu = 10$ is shown by the solid line, while the cut-off level $\nu = 20$ is shown by a the dashed line. For the hole order, the base case is used. The critical tee interval τ^* for that kind of par-3 hole is shown at the top.

6 Conclusions

We have developed a simulation algorithm to simulate the stochastic model of group play on a golf course over a day proposed by Whitt (2015) and we have conducted extensive simulation experiments to study important operational issues for golf courses. 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 all conclusions, the simulation experiments show that the order of the hole types matters relatively little.

In §3 we reported 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, especially when the course is overloaded. When the tee interval is 7.50, the 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. On the other hand, the waiting time for the 75th group ranges only between 23 and 26 minutes with $P3WU$ or $SP3$ holes. We also note that 68% of the total waiting times 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.

Making the course under-loaded 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 to wait between 26 and 31 minutes with $P3$ holes, but no more than 14 minutes with $P3WU$ or $SP3$ holes.

In §4 we used simulation to find the optimal tee schedule according to the optimization framework proposed in §4.1. Table 5 and Figure 4 provide important insights, which are summarized in §4.2. 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. Again, the throughput levels attained with each type of Par-3 holes 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 §5 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. 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. 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. A similar approach may also be useful for other service systems.

Acknowledgement

We thank Lucius J. Riccio for suggesting this line of research and for helpful comments. We thank NSF for research support (NSF grants CMMI 1265070 and 1634133).

References

d

Kimes S, Schruben L (2002) Golf revenue management: a study of tee time intervals. *Journal of Revenue and Pricing Management* 1(2):111–119.

Riccio L (2012) Analyzing the pace of play in golf: the golf course as a factory. *International Journal of Golf Science* 1:90–112.

Riccio L (2013) *Golf's Pace of Play Bible* (New York: The Three/45 Golf Association).

Riccio L (2014a) Three/45 Golf Association web page: http://www.three45golf.org/.

Riccio L (2014b) The status of pace of play in american golf, 31 page research paper on the Three/45 Golf Association web page: http://www.three45golf.org/.

Tiger AA, Ellerbrook JE (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.

Tiger AA, Salzer S (2004) Daily play at a golf course: using spreadsheet simulation to identify system constraints. *INFORMS Transactions on Education* 4(2):28–35.

Whitt W (2015) The maximum throughput on a golf course. *Production and Operations Management* 24(5):685–703.