
Original Article

Bitz & Pizzas: Optimal stopping strategy for a slot machine bonus game

Noelia Oses

NOF Consulting, Zorroaga pasealekua 23, 20011 Donostia, Euskal Herria, Spain.
E-mail: Noelia@nofcon.com
Website: <http://www.nofcon.com/>

Abstract Slot machine games used to be very simple, often limited to just one spin of the reels, but have evolved to accommodate a variety of features that provide added excitement and the potential for greater wins. They also engage the players more by, quite often, requiring them to make choices. This paper presents one of these games, Bitz & Pizzas. This case study introduces a novel application of Operational Research (OR) and serves as evidence that as slot games, and their corresponding models, are becoming more sophisticated, the role of a well-trained OR professional is becoming more important and necessary in this industry. It also shows that OR practitioners are in touch with the needs of the gambling industry and understand current practices and legislation. *OR Insight* (2009) **22**, 31–44. doi:10.1057/ori.2008.4

Keywords: dynamic programming; gaming; optimal stopping of a finite-horizon Markov chain; slot machines

Introduction

Games in the infancy of the slot machine industry were very simple, often limited to just one spin of the reels (Fey, 1994). Since then, they have evolved to accommodate a variety of features. These bonus game routines provide added excitement and the potential for greater wins. Quite often they also engage the players more by requiring them to make choices or decisions. This is a good strategy on the part of the game designers, as research shows that players like feeling that a degree of skill is necessary or that they are, to some

extent, in control (Griffiths, 1993). Detailed descriptions of these features with corresponding modelling solutions are presented in Osese and Freeman (2006). The objective of this paper is to illustrate the impact of increased addition of player choices and increased sophistication of slot game features in the modelling of slots. We will discuss how the industry's needs have evolved regarding the modelling of the games, and it will be shown that Operational Research (OR) practitioners add value to gambling/gaming companies.

Mathematical models of slot games are used to calculate the probability distribution of the prizes ('win distribution' for short). The most notable result obtained from the win distribution is the 'percentage return' of the game, that is, the percentage of the collected money the game will return to the players in the long term. This is the standard measure for the classification of slots. When slot providers purchase a new slot game from a slot manufacturer, they usually also specify the percentage return they want the game to achieve in the long term. A percentage return that is too low might push players to decide against playing the game. On the other hand, if it is too high, the game provider might not obtain the desired revenue. The win distribution also allows us to obtain the hit rates and the volatility. If prizes do not occur often enough, players might get bored and stop playing the game. In addition, the distribution of the prizes is important for performing risk analysis and gambler's ruin analysis. The former is valuable for game providers' business plans and the latter is valuable for game designers to, among other things, speculate why some games are more successful than others.

The impact of increased addition of player choices

Slot games started as random games that required neither skill nor input from the player. For example, line wins and scatters, where the player wins prizes by lining up symbols according to predetermined patterns, are completely random features over which the player cannot exert any influence. Later, some bonus games that interact with the player were introduced. These can be classified into two groups: those in which the players' decisions have no impact on the long-term percentage return of the feature and those in which player decisions do have an impact on the percentage return. An example of the former are the games in which the player is presented with a number of hidden prizes and is required to choose one. Usually the hidden prizes are all chosen according to the same distribution and, therefore, the expected prize value in the long term is the same whichever particular prize is chosen each time the game is played. This paper is interested in the games of the latter group.

For games in which player choices have an impact on the percentage return, this (and, of course, the win distribution) depends on the player's play strategy and presents the modeller with the question of how to model the player's



strategy, that is, what assumptions to make about the player's decisions. The first published analysis of such a bonus game was the model for the Hi-Lo bonus (Freeman, 1998). This author suggested analysing four different, hypothetical player strategies (not including the optimal strategy) and then estimating the player distribution to calculate the percentage return. This has the obvious disadvantage that different modellers can propose different sets of player strategies and easily arrive at different distributions of players; thus the approach is subjective and can lead to different results. It is clear that, for the percentage return to have a consistent meaning across all games, some sort of standard approach must be defined and thus slot game modelling has evolved from using simulations and analysis of hypothetical player strategies to the increased use of analytical results and objective approaches. In the author's commercial experience, the new, generalized trend is to model those games where the player strategies may alter the percentage return following the optimal strategy to calculate the win distribution and analyse the worst-case scenario for the machine operator. This is a legal requirement in Nevada (Nevada Gaming Control Board, 1959), which usually sets the trend in this industry. This approach provides a unique, objective figure for the percentage return (the maximum percentage return of the game or minimum house edge) independent of the modeller analysing the game (as well as an objective win distribution). This approach also allows game providers to know what their minimum revenue will be (in terms of percentage of money played) independent of the factors outside the machine, like player's behaviour towards risk or strategy, provided enough games are played (confidence intervals are usually taken as a guide). This paper presents the case study of the Bitz & Pizzas bonus game and aims to illustrate how bonus games of this type are analysed in the industry nowadays following this latter objective approach as opposed to Freeman's (1998) approach, therefore showing that OR practitioners understand the gambling industry's current needs and practices.

The impact of the increased sophistication of bonus games

Slot machine games have evolved from simple, spinning reels games with only line and scatter wins to games with an eclectic variety of features. In the days of the former, it was easy for game designers to have a spreadsheet model of a game that they could reuse for other games, just by modifying the prizes and reel distributions as necessary. Modelling of the latter is not so straightforward. Nowadays bonus features often need custom modelling using OR techniques, like stochastic dynamic programming, and they often require knowledge of stochastic processes, Markov chains frequently (see Oses and Freeman (2006) for examples). These models often also have to be programmed in a computer, thus requiring the modeller to be a competent programmer. It can then be

concluded that these games need to be modelled by a well-trained OR professional, and the case study presented in this paper is evidence of this. Had there not been an OR practitioner in the company, Bitz & Pizzas would not, most likely, have been developed, as it is necessary to produce a mathematical model of the game when requesting a licence from the Nevada Gaming Commission, where this game was licensed. Thus, OR professionals add value to gambling companies by allowing them to develop a wider range of games.

The Bitz & Pizzas (Figure 1) bonus game presents the player with a succession of prizes in sequential order. At each stage, the player has to decide whether to take the prize on offer and end the game or reject it and move on to the next offer. The player does not have the option of choosing to take previously rejected prizes, and the maximum number of offers that the player can have is limited and predetermined and known to the player. When the player reaches the maximum number of offers without having accepted any of them, the last offer is auto-collected, that is, he automatically receives the last offer as his prize.

This paper is structured as follows. The next section describes the bonus game in detail. In the subsequent section, the game is modelled as a Markov chain (Norris, 1997; Grimmet and Stirzaker, 2001), and dynamic programming (Bellman, 1957; White, 1969) is used to find the optimal stopping strategy for maximizing the expected prize (Ross, 1983; Bather, 2000). The theory developed is applied to the game shown in Figure 1, and its optimal stopping strategy and associated expected value are obtained in the further section. The last section provides the conclusions. The model described in this paper was developed by the author and used as part of the model submitted to the Nevada Gaming Commission to obtain the licence for Bitz and Pizzas Slots.

Bitz & Pizzas

Bitz and Pizzas Slots (Figure 1) is a traditional three-reel slot machine game released in 2005 by Barcrest USA. This game features line wins, a jackpot for maximum bet and the Bitz and Pizzas bonus (Figure 2). The latter is the focus of attention in this paper. This bonus is triggered when a 'Bonus' symbol lands on the payline with maximum credits bet.

In this bonus, there are 16 different values on display and a reel that has four possible outcomes: 'Plus One', 'Change One', 'Plus Two' and 'Change All' (Figure 2 does not show the reel, which is in the white space inside the jar held by the chef). At the start of the bonus, 1 of 16 different values lights up to show the prize on offer. The player has a choice and can take this award or spin the reel in the hope that he will increase his offer. At any time, the current offer is



Figure 1: Barcrest USA's Bitz & Pizzas slots.

the sum of all the lit values. If the player decides to spin, the reel will stop in one of these four outcomes. 'Plus One' and 'Plus Two' cause one or two of the unlit values to light up and add their value to the offer. 'Change One' causes one of the lit values to change for one of the unlit values; the former will be



Figure 2: Bitz & Pizzas bonus artwork.

unlit, the latter will light up and the offer will change accordingly. 'Change All' causes all the lit values to be replaced by unlit values. The player has a maximum of four chances of rejecting the current offer and spinning the reel. After the fourth spin, the player will automatically be awarded the resulting prize.

The 16 values of the bonus have a probability distribution of being selected associated with them. The first value at the start of the game is selected according to this distribution, and the 'Plus One', 'Change One', 'Plus Two' and 'Change All' transitions are made according to the conditional probability distributions derived from this. There is a probability distribution associated with the four possible outcomes of the reel spin.

Mathematical Model

The stochastic process that determines the prize in the Bitz & Pizzas bonus game is clearly Markovian because the next state depends only on the current state to which one or two values are added or one or all values are changed. In fact, the state space is finite and, thus, the process is a (time-homogeneous) Markov chain. At each step of the game, the player has only two options, to continue or to stop, and the objective is to maximize the prize. Therefore, this problem fits well into the optimal stopping of Markov chain problems described by Bather (2000) in which the reward received for stopping at the current state is compared with the expectation that can be achieved by allowing another transition. The model is as follows.

Let N be the number of prizes displayed on the bonus game. Let's index these prizes from 1 to N , and let A be the set of these indexes, $A = \{1, \dots, N\}$. Let Y be the random variable associated with obtaining one of these prizes, where for $a \in A$ $P(Y = a) = p_a$, $0 \leq p_a \leq 1$ and $\sum_{a \in A} p_a = 1$. Let SN be the maximum number of spins the player is allowed to have. Then the maximum number of lights that can be lit simultaneously in one offer is $R = 2 \times SN + 1$. The following condition, $2(2 \times SN - 1) < |A|$, must hold so that any of the transitions can take place at any stage during the game. For example, the maximum number of lit lights at stage SN is $2 \times SN - 1$; for a 'Change All' transition to take place, there must be another $2 \times SN - 1$ prizes that are unlit, at least, therefore we need to have at least $2(2 \times SN - 1)$ different values in the game, that is, $2(2 \times SN - 1) \leq |A|$. On the other hand, for the case in which only one spin is allowed, $SN = 1$, for a 'Plus Two' transition to take place there must be another two prizes that are unlit, at least. If $2(2 \times SN - 1) = |A|$ for $SN = 1$, that is, if there are only two values in the game and only one spin allowed, then we cannot have a 'Plus Two' transition. Therefore, generalizing, we need the number of values to be strictly greater than $2(2 \times SN - 1)$, that is, $2(2 \times SN - 1) < |A|$.

For $2 \leq r \leq R$, we define the sets

$$\hat{A}^r = \{(a_1, \dots, a_r) \in A^r : \forall i, j = 1, \dots, r \quad i < j \Rightarrow a_i < a_j\}$$

These sets are defined because the order in which the prizes have been lit is not important. For example, saying that prizes 1 and 2 or prizes 2 and 1 are lit is equivalent and will be uniquely represented by the (1, 2) element in the \hat{A}^2 set. Let the probability of the outcome of one spin of the reel being 'Plus One' be q_1 , q_2 for 'Change One', q_3 for 'Add Two' and q_4 for 'Change All', where $q_1 + q_2 + q_3 + q_4 = 1$ and $0 \leq q_i \leq 1$ for $i = 1, \dots, 4$. Let X_n be the random variable that represents the lights lit at stage n , for $0 \leq n \leq R$. Then, $\{X_n\}_{n=0}^R$ is a homogeneous Markov chain with state space $S = A \cup \hat{A}^2 \cup \dots \cup \hat{A}^R$. To make the notation easier in the following sections, let a_0 be 0.

probability of the reel spin resulting in 'Plus One' is q_1):

$$\begin{aligned} P(X_{n+1} = S_{n+1} | X_n = S_n) &= q_1 \times P(Y = x | Y \neq a_1, \dots, a_r) \\ &= q_1 \times \frac{p_x}{\sum_{i \neq a_1, \dots, a_r} p_i} \end{aligned}$$

Plus two ($P_{r, r+2}$)

A state $S_n = (a_1, \dots, a_r) \in \hat{A}^r$ for some r where $1 \leq r < R-1$ can also transition to states resulting from adding another two prizes, let these be $x, y \in A$ where $x, y \neq a_1, \dots, a_r$, and assume that $x < y$ without loss of generality. Then, the ensuing state, S_{n+1} , is as follows:

- If $a_r < x < y$ then $S_{n+1} = (a_1, \dots, a_r, x, y)$, or
- If $a_r < y$ and $\exists n: 1 \leq n \leq r$ where $a_{n-1} < x < a_n$, then $S_{n+1} = (a_1, \dots, a_{n-1}, x, a_n, \dots, a_r, y)$, or
- If $\exists n, m: 1 \leq n \leq m \leq r$ where $a_{n-1} < x < a_n$ and $a_{m-1} < y < a_m$, then
- If $n = m$ then
 - $S_{n+1} = (a_1, \dots, a_{n-1}, x, y, a_n, \dots, a_r)$ if $n > 1$.
 - $S_{n+1} = (x, y, a_1, \dots, a_r)$ if $n = 1$.
- If $n < m$ then
 - $S_{n+1} = (a_1, \dots, a_{n-1}, x, a_n, \dots, a_{m-1}, y, a_m, \dots, a_r)$ if $n > 1$.
 - $S_{n+1} = (x, a_1, \dots, a_{m-1}, y, a_m, \dots, a_r)$ if $n = 1$.

$S_n \in \hat{A}^r$, therefore $S_{n+1} \in \hat{A}^{r+2}$, and, given that the probability of the reel spin resulting in 'Plus Two' is q_3 , the transition probability is

$$\begin{aligned} P(X_{n+1} = S_{n+1} | X_n = S_n) &= q_3 \times P((x, y) \in \hat{A}^2 | x, y \neq a_1, \dots, a_r) \\ &= q_3 \times \frac{p_x \times p_y}{\sum_{\substack{i, j \neq a_1, \dots, a_r \\ (i, j) \in \hat{A}^2}} p_i \times p_j} \end{aligned}$$

Change one (P_{rr} for r where $1 < r < R-1$)

For $S_n = (a_1, \dots, a_r) \in \hat{A}^r$ for some r where $1 < r < R-1$ ($S_n \in A$ will be analysed later), this state can transition to states for which one of the elements in S_n , say $a_i \in \{a_1, \dots, a_r\}$, is replaced by one of the lights not currently lit, say $x \in A: x \neq a_1, \dots, a_r$. Then, the resulting state is $S_{n+1} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_r, x)$

if $a_r < x$ or $\exists n: 1 \leq n \leq r$, where $a_{n-1} < x < a_n$ and thus

- If $n < l$ then $S_{n+1} = (a_1, \dots, a_{n-1}, x, a_n, \dots, a_{l-1}, a_{l+1}, \dots, a_r)$
- If $n = l$ or $n = l + 1$ then $S_{n+1} = (a_1, \dots, a_{l-1}, x, a_{l+1}, \dots, a_r)$
- If $n > l + 1$ then $S_{n+1} = (a_1, \dots, a_{l-1}, a_{l+1}, \dots, a_{n-1}, x, a_n, \dots, a_r)$

Then, $S_{n+1} \in \hat{A}^r$ and the transition probability, knowing that the probability of the reel spin resulting in 'Change One' is q_2 , is

$$\begin{aligned} P(X_{n+1} = S_{n+1} | X_n = S_n) &= q_2 \times P(Y = a_l | a_l \in \{a_1, \dots, a_r\}) \\ &\quad \times P(Y = x | x \notin \{a_1, \dots, a_r\}) \\ &= q_2 \times \frac{p_x}{\sum_{i \neq a_1, \dots, a_r} p_i} \times \frac{p_{a_l}}{\sum_{j \in \{a_1, \dots, a_r\}} p_j} \end{aligned}$$

Change all (P_{rr} for r where $1 < r < R-1$)

For $S_n = (a_1, \dots, a_r) \in \hat{A}^r$ for some r where $1 < r < R-1$ ($S_n \in A$ will be analysed later) and $S_{n+1} = (b_1, \dots, b_r) \in \hat{A}^r$ where $b_i \in A - \{a_1, \dots, a_r\} \forall i = 1, \dots, r$ (the $2(2 \times SN - 1) < |A|$ assumption guarantees that this is possible), the transition probability is as follows.

$$\begin{aligned} P(X_{n+1} = S_{n+1} | X_n = S_n) &= q_4 \times P((b_1, \dots, b_r) \in \hat{A}^r | b_1, \dots, b_r \notin \{a_1, \dots, a_r\}) \\ &= q_4 \times \frac{p_{b_1} \times \dots \times p_{b_r}}{\sum_{\substack{(i_1, \dots, i_r) \in \hat{A}^r \\ i_j \neq a_1, \dots, a_r \forall j=1, \dots, r}} p_{i_1} \times \dots \times p_{i_r}} \end{aligned}$$

Change one and change all for states in A (P_{11})

For $S_n = x \in A$ and $S_{n+1} = y \in A$ where $y \neq x$, the transition probability is as follows:

$$\begin{aligned} P(X_{n+1} = S_{n+1} | X_n = S_n) &= (q_2 + q_4) \times P(Y = y | y \neq x) \\ &= (q_2 + q_4) \times \frac{p_y}{\sum_{i \neq x} p_i} \end{aligned}$$

Optimality equation, maximum expected value and the stopping rule

Let PR_i be the prize corresponding to the i th index, $i \in A$. Then for any state s in S , say $s = (a_1, \dots, a_r) \in \hat{A}^r$ for some r where $1 < r \leq R$, the prize associated with this state is $PR_s = \sum_{i=a_1, \dots, a_r} PR_i$. There are no costs associated with

the transitions. To determine the optimal stopping of this process a dynamic programming approach is employed, given that the maximum number of transitions before stopping is finite (Bather, 2000). For each state s and allowing at most SN transitions before stopping, the maximum expected reward given the initial state s is $u_s(SN)$, and by the principle of optimality

$$u_s(SN) = \max \left\{ PR_s, \sum_r p_{sr} u_r(SN - 1) \right\}$$

where p_{sr} is the transition probability, that is, $p_{sr} = P(X_{n+1} = r | X_n = s)$.

The maximum expected reward when all the allowed transitions have been made is the prize associated with the state where the chain has stopped, that is, if the chain has stopped in a state s after SN transitions, then $u_s(0) = PR_s$. Thus, the maximum expected rewards, $u_s(SN)$, can be calculated iteratively and the overall maximum expected reward corresponding to the optimized process is the weighted average of these values, that is, $EV = \sum_{i=1}^N P(Y=i) \times u_i(SN)$.

The stopping rule not only depends on the current state the chain is at but also, because the time horizon is finite, on the stage. At any stage and any state, stop the process if the maximum expected reward of this state at this stage is equal to the state's prize, that is, the set of stopping states for each stage n ($0 \leq n < SN$) is the following:

$$Q(n) = \{s \in S : u_s(SN - n) = PR_s\}$$

Application

This section applies the above theory to the Bitz & Pizzas bonus game of Figure 2 to obtain numerical results. There are 16 prizes on display (N). For indexes 1–16 ($A = \{1, \dots, 16\}$), the associated prizes $PR_1 - PR_{16}$ are 1000, 200, 100, 50, 30, 25, 22, 20, 15, 12, 10, 8, 6, 5, 4 and 2. The probability distribution associated with these prizes is

$$\begin{aligned} p_1 &= \frac{10}{400}, p_2 = \frac{10}{400}, p_3 = \frac{15}{400}, p_4 = \frac{20}{400}, \\ p_5 &= \frac{25}{400}, p_6 = \frac{30}{400}, p_7 = \frac{30}{400}, p_8 = \frac{30}{400}, \\ p_9 &= \frac{40}{400}, p_{10} = \frac{40}{400}, p_{11} = \frac{40}{400}, p_{12} = \frac{35}{400}, \\ p_{13} &= \frac{20}{400}, p_{14} = \frac{20}{400}, p_{15} = \frac{20}{400}, p_{16} = \frac{15}{400} \end{aligned}$$

And the probability distribution associated with the reel outcomes is

$$q_1 = \frac{35}{100}, q_2 = \frac{35}{100}, q_3 = \frac{15}{100} \text{ and } q_4 = \frac{15}{100}$$

These two probability distributions have been changed to protect data not available in the public domain. The maximum number of spins the player is allowed to have (SN) is 4. Under these conditions, the overall maximum expected prize of the optimized process is 258.48. In stage 0, after the initial offer, the process should only be stopped when the top prize (the prize with index 1, that is, 1000) is on offer. Thus, the stopping set for stage 0 is $Q(0) = \{1\}$. In stage 1, after one spin has been used, the process should be stopped only when the current offer contains the top prize; thus,

$$Q(1) = \{1\} \cup \{(i, j) \in \hat{A}^2 : i = 1\} \cup \{(i, j, k) \in \hat{A}^3 : i = 1\}$$

In stage 2, the process should be stopped when the current offer contains the top prize or it is one of the following offers 2, (2, 3) or (2, 3, 4), that is,

$$\begin{aligned} Q(2) = & \{1\} \cup \{(i, j) \in \hat{A}^2 : i = 1\} \\ & \cup \{(i, j, k) \in \hat{A}^3 : i = 1\} \cup \{(i, j, k, l) \in \hat{A}^4 : i = 1\} \\ & \cup \{(i, j, k, l, m) \in \hat{A}^5 : i = 1\} \cup \{2, (2, 3), (2, 3, 4)\} \end{aligned}$$

Finally, in stage 3, the process should be stopped for those offers specified in this stage's stopping set:

$$\begin{aligned} Q(3) = & \{1\} \cup \{(i, j) \in \hat{A}^2 : i = 1\} \cup \{(i, j, k) \in \hat{A}^3 : i = 1\} \\ & \cup \{(i, j, k, l) \in \hat{A}^4 : i = 1\} \cup \{(i, j, k, l, m) \in \hat{A}^5 : i = 1\} \\ & \cup \{(i, j, k, l, m, n) \in \hat{A}^6 : i = 1\} \\ & \cup \{(i, j, k, l, m, n, o) \in \hat{A}^7 : i = 1\} \cup \{(2, (2, 3), (2, 4), (2, 5)\} \\ & \cup \{(2, 3, k) \in \hat{A}^3 : 4 \leq k \leq 9\} \cup \{(2, 3, k) \in \hat{A}^3 : 13 \leq k \leq 16\} \\ & \cup \{(2, 3, 4, l) \in \hat{A}^4 : 5 \leq l \leq 8\} \cup \{(2, 3, 4, l) \in \hat{A}^4 : 13 \leq l \leq 16\} \end{aligned}$$

The computational calculation of the strategy, or the set of optimal actions for each state at each stage, requires repeating the same calculations again and again. Therefore, as much of the necessary data as possible must be ready before the strategy calculation algorithm starts. The necessary data include the state space and the conditional probabilities. It is possible that not all these

data can be fit in the memory, trying to do so could give rise to a stack overflow runtime error. In such case, as much of the data as possible should be prepared in advance.

Conclusions

As the slot machine industry evolves – and there is no question it will continue evolving as online casinos are becoming increasingly popular and slot games are starting to feature on mobile games – more sophisticated features are introduced in the games, which, consequently, forces the development of more sophisticated models. Thus, the role of a well-trained OR professional is becoming more important and necessary in this industry. The case study presented in this paper is evidence of this. It is the case study of the Bitz & Pizzas bonus game, which presents the player with a succession of prizes in sequential order. At each stage, the player has to decide whether to take the prize on offer and end the game or reject it and move on to the next offer. The industry assumes that players always play following the optimal strategy because this ultimately determines whether the game is profitable or not. A Markov chain model of the game has been presented and a dynamic programming approach has been used to find the optimal stopping strategy for maximizing the expected prize.

This case study is a double success story for OR. First, it shows that OR practitioners are in touch with the needs of the gambling industry and understand current practice and legislation by modelling the game following the optimal strategy. This has not always been the case in the past. Freeman (1998), for example, analysed different hypothetical and subjective player strategies and was unable to provide a definite, objective figure for the percentage return of the Hi-Lo feature he modelled. Second, it proves that OR practitioners add value to a gambling company, as had there not been an OR professional in the company when Bitz & Pizzas was being developed, the company would not, most likely, have been able to model this game and thus the game would not have been released. Therefore, OR practitioners add value to a gambling/gaming company by enabling it to develop a wider variety of games.

About the Author

Noelia Oses is a freelance consultant for the slot machine industry. Her work focuses on developing probability models of slot games to calculate the probability distribution of the prizes and adjust the maximum percentage



return in the long term or, in other words, to control the minimum house edge. Her research interests include applying OR, probability and stochastic processes techniques in this industry. Noelia has been a member of *The OR Society* (Birmingham, UK) since 1999. E-mail: email@orsoc.org.uk.

References

- Bather, J. (2000) *Decision Theory: An Introduction to Dynamic Programming and Sequential Decisions*. Chichester, UK: John Wiley & Sons.
- Bellman, R. (1957) *Dynamic Programming*. Princeton: Princeton University Press.
- Fey, M. (1994) *Slot Machines: A Pictorial History of the First 100 Years of the World's Most Popular Coin-Operated Gaming Device*. Reno, NV: Liberty Belle Books.
- Freeman, J.M. (1998) Gambling on HI-LO: An evaluation of alternative playing strategies. *Journal of the Operational Research Society* 49: 1278–1287.
- Griffiths, M. (1993) Fruit machine gambling: The importance of structural characteristics. *Journal of Gambling Studies* 9(2): 101–120.
- Grimmett, G. and Stirzaker, D. (2001) *Probability and Random Processes*. Oxford: Oxford University Press.
- Nevada Gaming Control Board. (1959) Regulation 14.040 'Manufacturers, Distributors, Operators of Inter-Casino Linked Systems, Gaming Devices, New Games Inter-Casino Linked Systems and Associated Equipment'. Adopted 1 July 1959, and Current as of March 2006.
- Norris, J.R. (1997) *Markov Chains*. Cambridge: Cambridge University Press.
- Osos, N. and Freeman, J. (2006) Hitting the jackpot with OR. *OR Insight* 19(3): 21–30.
- Ross, S.M. (1983) *Introduction to Stochastic Dynamic Programming*. London: Academic Press.
- White, D.J. (1969) *Dynamic Programming*. San Francisco: Holden-Day.