ELSEVIER

European Journal of Operational Research 133 (2001) 557–565

EUROPEAN JOURNAL OF OPERATIONAL RESEARCH

www.elsevier.com/locate/dsw

Theory and Methodology

Modelling curling as a Markov process

Kent J. Kostuk [a], Keith A. Willoughby [b,*], Anton P.H. Saedt [c]

[a] *Federated Co-operatives Limited, 401-22nd Street East, Saskatoon, SK, Canada, S7K 0H2*
[b] *Department of Management, Bucknell University, Lewisburg, PA 17837, USA*
[c] *Agrotechnological Research Institute, Wageningen, Netherlands*

Received 13 July 1999; accepted 24 July 2000

Abstract

Markov processes have been used to model a variety of sports such as football, jai alai and baseball. Curling has a scoring system similar to baseball; both progress on an inning by inning basis. This suggests that curling could also be modeled as a Markov process. We will develop such a model and compare some basic strategies.

Keywords: Markov processes; Recreation; Modelling

1. Introduction

Curling has been described as a combination of bowling and chess. There are a variety of strategic nuances and subtleties, but as in most sports there are two basic strategies; play cautiously waiting for an opportunity to exploit your opposition's mistakes, or play aggressively where the "best defense is a strong offense". There are no documented instances of a probabilistic or simulation based curling model. We must develop a probabilistic model to objectively analyze our basic strategies.

Curling is a winter team sport with origins in Scotland. Each team is made up of four players and the game is played on sheets of ice 14 ft 2 in. wide and 144 ft long. At each end of the sheet is a house. The house is a set of four concentric circles of various diameters (12, 8, 4, and 1 ft).

Curling is played with circular disks of polished granite weighing approximately 45 pounds. The teams take alternate turns sliding each of their eight rocks down the sheet of ice. Teams attempt to strategically position their rocks in front of, or within the house. A team may also deliver their rocks to remove their opposition's rock(s) from play. Once all 16 rocks have been played the score is tallied. A team scores a point for each stone closer to the middle of the house than their opposition's closest stone. For example, suppose your team has one rock in the house and your

* Corresponding author. Tel.: +1-570-577-1306; fax: +1-570-577-1338.

E-mail addresses: k.kostuk@fcl.ca (K.J. Kostuk), kwilloug@bucknell.edu (K.A. Willoughby).

0377-2217/01/$ - see front matter
PII: S0377-2217(00)00202-2

opposition has three. If the opposition's rocks are closer to the middle of the house than yours, your opposition would score three points (Fig. 1). Conversely, if you have a rock closer than any of your opposition's rocks then you will score a point (Fig. 2). If no rocks are in the house at the completion of an end, this is referred to as a blank end.

There is a strategic advantage to having the last rock in an end (also known as having the hammer). Teams will flip a coin to determine who will have the last rock in the first end. In subsequent ends, the team that was scored on in the previous end receives the hammer in the next end. A standard curling game consists of 10 ends or innings. If the game is tied after 10 ends, extra ends are played until one of the teams scores at least one point.

A literature search on curling would generate information that could be classified in one of four categories: rules and instruction (CCA, 1996; Lukowich et al., 1981; Lukowich, 1993), historical (Murray, 1981; Weeks, 1995), examination of the sport's physics (Johnston, 1981; Shegalski et al., 1996; Denny, 1998), and scheduling (Kostuk, 1997). Unlike several other sports, analytical models have not been developed.

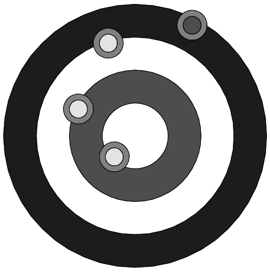

Fig. 1. An example of scoring three points.

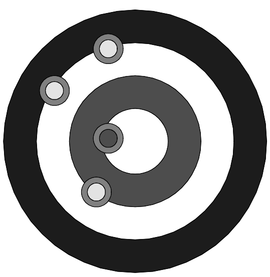

Fig. 2. Scoring a single point.

Baseball was the first, and is likely the most analyzed of all sports. Various modelling techniques have been applied; statistical analysis (D'Esopo and Lefkowitz, 1977), Markov processes (Howard, 1960; Bukiet et al., 1997; Trueman, 1977) and dynamic programming (Bellman, 1977). Other sports which have been modelled include: Australian rules football (Clarke and Norman, 1998), cricket (Clarke and Norman, 1999), tennis (Hannan, 1976), track and field (Stefani, 1996), American football (Casti, 1971), basketball (Carlin, 1996; Schwertman et al., 1996), jai alai (Byrne and Hesse, 1996), snooker (Percy, 1994) and hockey (Schutz and Liu, 1996).

Curling has a scoring system similar to baseball (but much more similar to bowls); points are tallied at discrete intervals during play. This suggests that one of the models used to analyze baseball could be applied in a similar fashion to curling. We have elected to model curling as a Markov process, and will use this model to compare basic strategies.

The rest of this paper will be divided into four sections. Section 2 will develop our model formulation. The structure of a Markov process will be discussed, and its components will be introduced in the context of our curling model. Section 3 of the paper will illustrate how we developed our transition probability matrix. With the model fully developed, Section 4 will compare basic strategies. This will be followed by a discussion of the results and a conclusion in Section 5.

2. Curling as a Markov process

For a problem to be modeled as a Markov process we must define a series of states, the probability of moving from state to state, and establish the time interval over which these transitions will occur (Hastings, 1973). The classic example of a Markov process is a frog sitting in a pond full of lily pads (Howard, 1960). Each pad represents a state. At the end of a fixed time interval the frog will jump from pad i to any other pad j with probability $p(i,j)$. The difficulty in defining the state space for any Markov process is that it must both be detailed enough to clearly illustrate what is being modeled but at the same time be concise enough that the problem does not become intractable.

For a Markov process timing is of the essence. The Markov process has a fixed time interval. For curling there are two natural representations; a shot by shot model of the game's progression, or an end by end representation. We have selected the end by end representation.

All strategy in curling hinges on two pieces of information; namely, the score and whether or not the team has the hammer. In actuality the absolute score is not as important as the difference in score between the two teams. Our state space will be described by the tuple $\{i,j\}$ where i is the difference in score and j indicates possession of the hammer (1 means that the team has the hammer). Note that for this notation to work, one team must be selected as the benchmark (or frame of reference). Imagine that we have teams A and B, and B is selected as our benchmark. Whenever B is leading A, i is a positive value. Conversely, when A leads B, i is a negative value. The second index, j, indicates whether team B has the hammer. If B has the hammer, $j = 1$ otherwise $j = 0$. For further clarification, if B is leading A by two points without the hammer, the system's state is $(2, 0)$ (Fig. 3).

Teams will concede defeat before ten ends if they feel their opponent's lead is insurmountable. This suggests that extreme winning or losing positions do not exist. Consequently we can bound the state space for convenience without sacrificing realism in our model. We can limit the range of possible scores with lower (LB) and upper bound (UB) values. The only additional limitation is that a team cannot have the hammer if the score is equal to the UB and a team must have the hammer if the score equals the LB. The best way to describe these two additional constraints is by example. The only way to 'reach' the LB score is to have a point taken from you (i.e. your opponent scores). If a point is scored on you, you will have last shot in the next end. Consequently, if $i = LB$, then $j = 1$. A similar argument exists for $i = UB$. This reduces the size of the state space (N) from $N = 2(UB - LB + 1)$ to $N = 2(UB - LB)$.

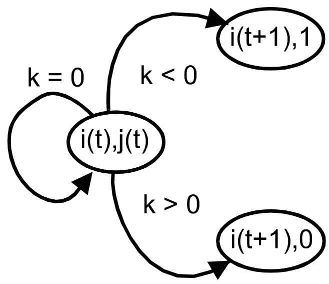

Fig. 3. Markov chain representation of scoring transitions from state $\{i,j\}$.

Next we must describe the state transitions. Let us define k as the number of points scored in an end. If k is positive you have scored a point, if it is negative your opposition has scored. The tuple describing our starting state is $\{i(t), j(t)\}$. The score at the beginning of the next end is equal to the present score plus the number of points scored during the end $[i(t+1) = i(t) + k]$. Unfortunately we can not describe the transition from $j(t)$ to $j(t+1)$ as succinctly. There are six possible events (Table 1) which can be consolidated into three rules (Table 2).

We now need a technique for mapping our tuple $\{i(t), j(t)\}$ to a point in our solution state space. Based on the state space description and the

Table 1
The six possible scoring transitions

	Event	K	$j(t)$	$j(t+1)$	$i(t+1)$
(1)	Score with the hammer	$k > 0$	1	$j(t) - 1$	$i(t) + k$
(2)	Scored on with hammer	$k < 0$	1	$j(t)$	$i(t) + k$
(3)	Blank with hammer	$k = 0$	1	$j(t)$	$i(t) + k$
(4)	Score without hammer (a steal)	$k > 0$	0	$j(t)$	$i(t) + k$
(5)	Scored on without hammer	$k < 0$	0	$j(t) + 1$	$i(t) + k$
(6)	Blank without hammer	$k = 0$	0	$j(t)$	$i(t) + k$

Table 2
Consolidating the six events into three basic rules

Event	k	$j(t)$	$j(t+1)$	$i(t+1)$
1 and 4	$k>0$	0 or 1	0	$i(t)+k$
2 and 5	$k<0$	0 or 1	1	$i(t)+k$
3 and 6	$k=0$	0 or 1	$j(t)$	$i(t)+k$

definition of the state transition Fig. 3 provides a general description for the transition between states.

Let us define the size of the state space as N and n as an element of the state space. As was discussed earlier, $N = 2(UB - LB)$. The transition probability matrix (TPM) will have dimensions of $N \times N$. The key is to define a constant such that $\{LB, 1\}$ is mapped to state 1 and $\{UB, 0\}$ is mapped to state N. We can formulate a functional relationship $n = f(i, j, z)$ where z is a constant. We can determine z via the equation $n = z + j + 2i$. The result is $z = -2LB$. Thus we have a universal relationship $n = -2LB + j + 2i$ or $n = 2(i - LB) + j$.

So how do we generate our TPM? The initial state is a function of i and j (as stated above). The transition from $\{i(t), j(t)\}$ to $\{i(t+1), j(t+1)\}$ is dictated by k. Thus, the probability of making the transition is equal to the probability of k occurring. By generating a list of all feasible triples of $\{i, j, k\}$ we can determine the TPM. Table 3 is a partial enumeration of all permutations of $\{i, j, k\}$.

For clarity we will describe the columns in Table 3. The first two columns represent the current point differential $[i(t)]$ and whether or not the team has the hammer $[j(t)]$. The third column indicates the state mapping described by our function $n = 2[i(t) - LB] + j(t)$. The fourth column (k) is the number of points the team scores. $P(k|j(t))$ is the probability of scoring k points given the team has or does not have the hammer. The remaining three columns describe the subsequent state; $i(t+1)$ is determined by the relationship $i(t+1) = i(t) + k, j(t+1)$ is determined as shown in Table 2, and, the Destination state column is determined by the same functional relationship which determined the Initial state column.

3. Transition probability matrix development

Our base data are statistical information recorded from 1985 to 1997 at the Canadian Men's Curling Championships (also known as the Brier). Table 4 is a summary of that data. Each column represents the number of points scored in an end. Probabilities $[P(k|j(t))]$ are generated from the frequency table. The rows represent the ends; the bottom row is the sum of each column. Note that a team may concede defeat at any point in the game. This fact is reflected in the last column of the table. This column reflects the number of games which played that many ends or fewer. For example there were 902 games that completed at least four ends. The number of games that progressed five ends or less was 900; less than half (410) went the full 10 ends.

Table 3
A partial example of enumerating all permutations of $\{i, j, k\}$

$i(t)$	$j(t)$	Initial state	k	$P(k\|j(t))$	$i(t+1)$	$j(t+1)$	Destination state
−8	1	1	0	1.000	−8	1	1
−7	0	2	−1	0.573	−8	1	1
−7	0	2	0	0.278	−7	0	2
−7	0	2	1	0.115	−6	0	4
−7	0	2	2	0.028	−5	0	6
−7	0	2	3	0.004	−4	0	8
−7	0	2	4	0.001	−3	0	10
−7	0	2	5	0.000	−2	0	12
−7	0	2	6	0.000	−1	0	14
−7	0	2	7	0.000	0	0	16
−7	0	2	8	0.000	1	0	18
−7	1	3	−1	0.149	−8	1	1

Table 4
Frequency table for the number of points scored in an end by a team with the hammer

	−5	−4	−3	−2	−1	0	1	2	3	4	5	6	
1		1	4	25	104	251	322	156	32	7			902
2		2	6	53	133	184	322	162	32	6	2		902
3		2	12	56	161	163	277	170	51	8	1	1	902
4		4	11	54	149	154	294	180	49	6	1		902
5	1	1	14	42	156	135	297	189	50	10	4	1	900
6	1	5	9	46	154	149	292	170	54	12	4		896
7	1	3	6	55	133	154	282	179	46	14	1		874
8	1	4	8	44	147	143	280	168	42	6	3		846
9			14	35	124	164	231	152	39	4	1		764
10		1	8	19	92	22	183	70	10	4	1	1	411
11				4	15	4	85	7	2	1			118
12				1	3								4
Totals	4	23	92	434	1371	1523	2865	1603	407	78	18	3	8421

Figs. 4–6 illustrate how the likelihood of scoring x points in an end changes during the course of the game. The first end distribution (Fig. 4) is essentially normal. The distribution is centered around the one point score indicating that the team with the hammer is most likely to score one point. The distributions for ends 2–9 are similar to the first end with a slight variation at the zero point mark.

In the 10th (and if necessary an 11th end) the scoring distribution is bimodal. This is because one rarely gets the opportunity to blank the last end. Blank ends occur if and only if there are no rocks in the house at the conclusion of an end. Consider a team that is trailing by i points in the 10th end. The team knows that it will need to score at least i points to prolong the match, or to emerge victorious. To accomplish this, it will surely adopt a strategy that keeps as many rocks as possible in play. Our empirical data (see Table 4) indicates blanking the 10th end occurs roughly 5% of the time.

The time varying characteristics of these distributions suggest that the standard time invariant assumption for a Markov process may not be valid. Consequently we must investigate whether we can establish a time invariant TPM which adequately represents what occurs during a game.

The time variant model is based on TPMs for each distribution in Table 4. Each TPM is applied over the appropriate time interval.

Our time invariant model is based upon the total likelihood of scoring x points in any given end. In other words, our transition probabilities

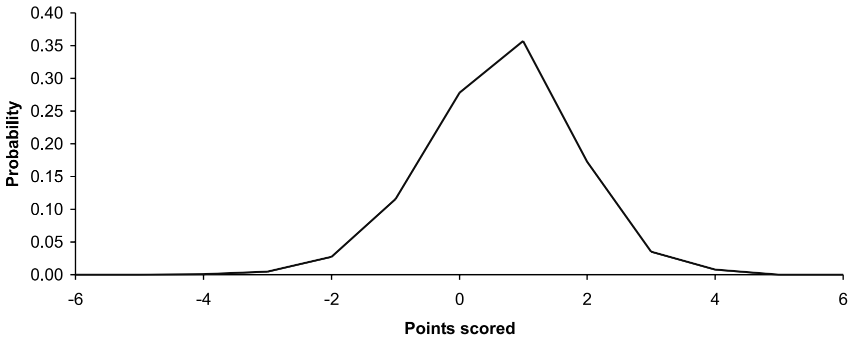

Fig. 4. Distribution for the number of points scored in the first end.

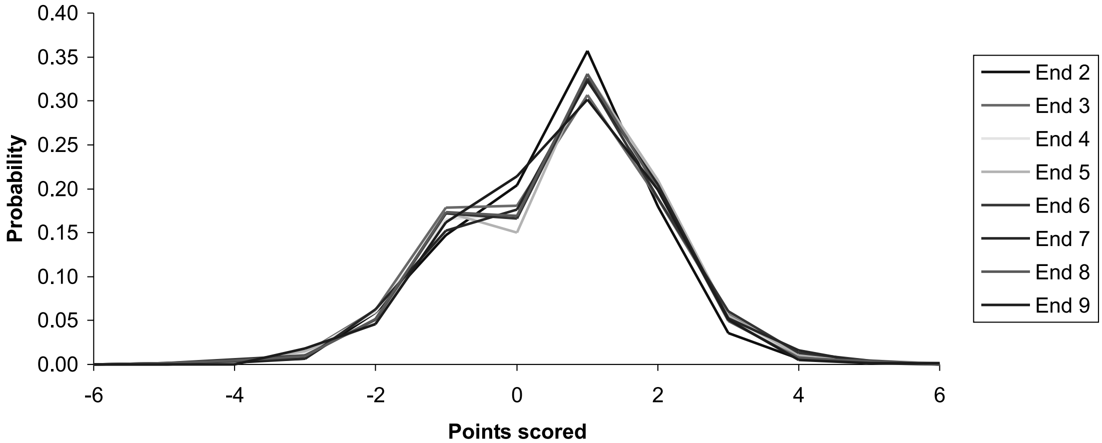

Fig. 5. Distribution for points scored during the middle ends (2–9).

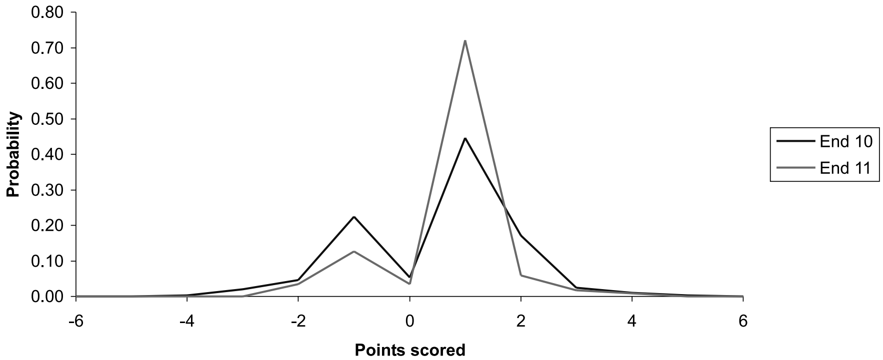

Fig. 6. Distribution for the number of points scored in the last and extra ends.

are based upon the bottom row of Table 4. Comparing the results of the time invariant and variant models we can determine the effectiveness of a time invariant model.

4. Analysis

All games start in a scoreless tie. The only difference from game to game is whether the baseline team wins the toss (and consequently has the hammer). Consequently our initial state will be either $\{0, 0\}$ or $\{0, 1\}$. The likelihood of reaching any one of the states at the end of the game is determined by the TPM(s) applied in the analysis.

The first question we wished to address was whether a time invariant model would adequately represent the outcome of a match. Figs. 7 and 8 illustrate the likelihood of being in each state (with and without the hammer) for both the time invariant and variant models. By observation, one can see that there is little difference between the two models. Quantitatively the expected outcome for each model is that winning the toss is worth approximately one point over the course of a game. As shown in Table 5, our time invariant model suggests that the expected point differential (at the conclusion of the game) of winning the toss is 1.028 points. The expected point differential with the time variant approach is 1.155 points. From

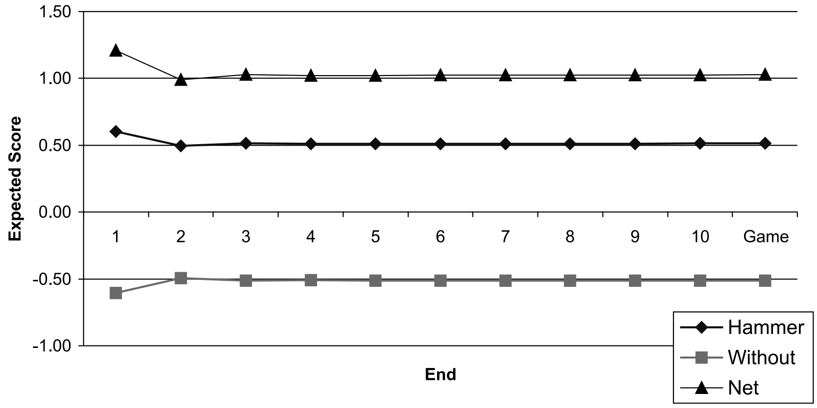

Fig. 7. Expected score with the invariant model.

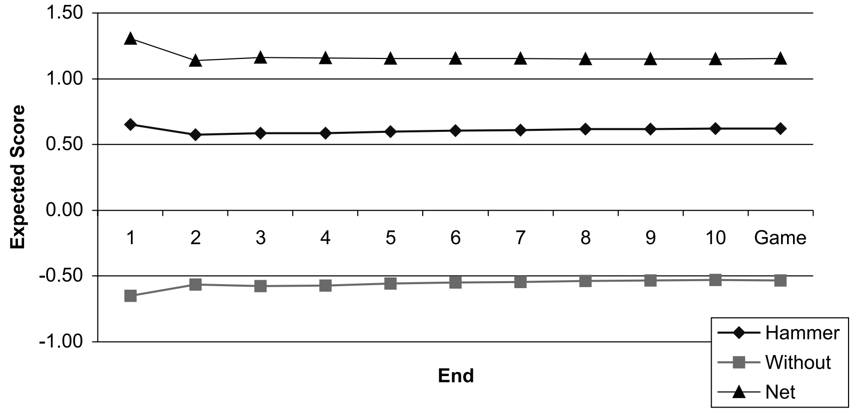

Fig. 8. Expected score with the variant model.

Table 5
Expected point differential at game end

Toss	Invariant	Variant
Win	0.514	0.623
Lose	−0.514	−0.532
Net	1.028	1.155

our perspective this signifies that a time invariant model is an adequate representation of the sport of curling.

In order to provide statistical support for the value of winning the toss, we computed a 95% confidence interval for the time invariant model. Using the number of times each particular tally was obtained (the bottom row of Table 4), we used a bootstrap method (Efron and Tibshirani, 1993) to generate 300 new sets of point estimations. Our 95% confidence interval for the value of winning the toss was [1.02598, 1.03002], with a mean of 1.028. The fairly narrow spread in this interval provides substantial confidence that the value of winning the toss is just over one point.

The distribution that we have observed for tenth end scoring is intriguing. The significant ‘spike’ at the one point and minus one point level shown in Fig. 6 appears to indicate a myopic strategy in the last end; a sufficient number of points scored in the short term will guarantee a win. From this observation arose the question, 'what would happen if a team played this way the whole game?' We wanted to compare the results of this extreme strategy with the more traditional approach to the game. By applying the 10th end TPM as our sole distribution we found that the expected points differential with the hammer was 0.404 and without was −0.404 for a difference of 0.808.

This was a counter-intuitive result. The myopic strategy was considered a gambling approach, and one would have expected a gambling or more offensive approach to generate a larger scoring differential. Upon further reflection it was realized that this distribution did not provide the potential gains of a gambling strategy; most of the distribution was focused around getting a single point, but yet still possessed the significant downside potential of a gambling strategy.

The three models we discussed so far have implicitly stated that the two teams being modeled are of equal strength; the likelihood of scoring a point and giving up a point is influenced by who has the last shot. Few if any teams consider this to be the case. We selected a pair of distributions representing the likelihood of our baseline team scoring x points when they have the hammer and scoring y points when they do not. The expected scoring differential under this scenario was significantly different from our previous models. Under these circumstances the expected differential with the hammer was 5.64 and 4.60 without. So although the expected outcome for this team was a win regardless of whether they won or lost the toss (unlike the previous models), the value of winning the toss was still approximately one point.

5. Conclusions

This paper has shown that curling, in a fashion similar to other 'discrete-event' sports, can be modeled as a Markov process. We have illustrated the outcomes of respective strategies. It would appear that teams generally enjoy an advantage of just over 1 point when beginning the game with the hammer. This result has been developed through various modeling scenarios such as myopic full-game strategies and teams of unequal abilities.

Additional research into this area appears to be warranted. Specifically, our model has analyzed end-by-end results as a curling game progresses. Is it possible to develop Markov models that examine the behavior of a curling game during a given end? Can one use these models to suggest the appeal of different strategies? Should curling teams use different tactics at separate occurrences during a single end of play? Two generally accepted strategies are 'conservative' and 'aggressive'. The former strategy involves teams trying to keep as few rocks in play as possible (so-called 'keeping the house clean'). Teams try to remove their opposition's rocks from play. An aggressive strategy suggests that teams keep as many of their rocks in play as possible (possibly by positioning some of their rocks in front of the house as 'guards'). The complexity of this modeling effort would arise due to the seemingly infinite number of possible states (rock position and shots remaining) during an end.

Acknowledgements

We acknowledge Brian Cassidy, University of New Brunswick, for his assistance in gathering the data.

References

Bellman, R. 1977. Dynamic programming and markovian decision processes, with application to baseball. In: Optimal Strategies in Sports. North-Holland, Amsterdam, pp. 75–85.

Bukiet, B., Harold, E.R., Palacios, J.L., 1997. A Markov chain approach to baseball. Operations Research 45 (1), 14–23.

Byrne, P.J., Hesse, R., 1996. A Markov chain analysis of the game of jai alai. Mathematics Magazine 69 (4), 279–282.

CCA. 1996. The Rules of Curling. Canadian Curling Association, Gloucester, Canada.

Carlin, B.P., 1996. Improved NCAA basketball tournament modeling via point spread and team strength information. The American Statistician 50 (1), 39–43.

Casti, J., 1971. Optimal football play selections and dynamic programming: A framework for speculation. Technical Note, Project PAR284-001, Systems Control, Inc., Palo Alto, CA.

Clarke, S.R., Norman, J.M., 1998. When to rush a behind in Australian rules football: A dynamic programming approach. Journal of the Operational Research Society 49, 530–536.

Clarke, S.R., Norman, J.M., 1999. To run or not: Some dynamic programming models in cricket. Journal of the Operational Research Society 50, 536–545.

Denny, M., 1998. Curling rock dynamics. Canadian Journal of Physics 76 (4), 295–304.

D'Esopo, D.A., Lefkowitz, B., 1977. The distribution of runs in the game of baseball. In: Optimal Strategies in Sports. North-Holland, Amsterdam, pp. 55–62.

Efron, B., Tibshirani, R.J., 1993. An Introduction to the Bootstrap. Chapman & Hall, New York.

Hannan, E.L., 1976. An analysis of different serving strategies in tennis. In: Management Science in Sports. North-Holland, Amsterdam, pp. 125–135.

Hastings, N.A.J., 1973. Dynamic Programming with Management Applications. Butterworths, London, UK.

Howard, R.A., 1960. Dynamic Programming and Markov Processes. MIT Press, Cambridge, MA.

Johnston, G.W., 1981. The dynamics of a curling stone. Canadian Aeronautics and Space Journal 27 (2), 144–161.

Kostuk, K.J., 1997. A decision support system for a large multi-event tournament. Infor 35 (3), 183–196.

Lukowich, E., Folk, R., Gowsell, P., 1981. The Curling Book. Modem Press, Saskatoon, Canada.

Lukowich, E., 1993. Power Curling. McGraw-Hill Ryerson, Toronto, Canada.

Murray, W.H., 1981. The Curling Companion. Collins, Toronto, Canada.

Percy, D.F., 1994. Stochastic snooker. The Statistician 43 (4), 585–594.

Schutz, R.W, Liu, Y. 1996. The effectiveness of overtime in the national hockey league. In: 14th Trienniel Conference of the International Federation of Operational Research Societies. Unpublished.

Schwertman, N.C., Schenk, K.L., Holbrook, B.C., 1996. More probability models for the NCAA regional basketball tournaments. The American Statistician 50 (1), 34–38.

Shegalski, M.R.A., Niebergall, R., Walton, M.A., 1996. The motion of a curling rock. Canadian Journal of Physics 74, 663–670.

Stefani, R.T., 1996. Evaluating achievement differences between male and female athletes using Olympic winning. In: 14th Trienniel Conference of the International Federation of Operational Research Societies. Unpublished.

Trueman, R.E., 1977. Analysis of baseball as a Markov process. Optimal Strategies in Sports, North-Holland, Amsterdam, pp. 68–76.

Weeks, B., 1995. The Brier: The History of Canada's Most Celebrated Curling Championship. Macmillan, Toronto, Canada.