

A NOTE ON OPTIMAL STOPPING FOR SUCCESS RUNS¹

BY SHELDON M. ROSS

University of California, Berkeley

1. Introduction. The following model is considered by Starr (1972): at most n tosses of a coin, having a constant probability p of coming up heads, are made. After each toss we have the option of either stopping and receiving an amount equal to the length of the terminal run of heads (that is if we were on a streak of k heads in the last k tosses then we could stop and receive k), or of paying an amount c and tossing the coin again. When n tosses have already been made we must stop.

The purpose of this note is to point out that with a simple modification the above problem fits the framework in which a one-stage look ahead policy is optimal. This yields not only an easy solution to the problem but also provides much insight. For instance, the reason for the additivity of the optimal continuation boundary, which is commented on by Starr on page 1890 (1972), now becomes clear. Also the problem may be generalized so that the terminal payoff is a more general function of the terminal run of heads, which may even also depend on the number of tosses made.

2. The optimal policy. Consider the above problem with the exception that the return when we stop after a terminal run of r heads is $f(r)$, where $f(r)$ is such that

$$f(r) - pf(r + 1) \text{ is non-decreasing in } r.$$

Define V_n to be the value to the decision maker if he is allowed to make at most n tosses before stopping and when he employs an optimal strategy, and note that V_n is non-decreasing in n . Say that the process is in state (r, j) if we are on a run of r heads and we are allowed at most j more coin tosses.

Now let us consider a modified problem which is such that when we are in any state of the form $(0, j)$, $j \geq 0$, we are forced to stop and we receive a terminal reward V_j . (That is, whenever a tail occurs we must stop but we are paid as if we acted optimally from this point on.) In this modified problem if we stop when in state (r, j) we receive $f(r)$, while if we continue for exactly one more toss and then stop, our expected return is $pf(r + 1) + (1 - p)V_{j-1} - c$. Hence the one-stage look ahead policy (see Derman and Sacks (1960), Chow and Robbins (1961) or [3], pages 137-38) is to stop at state (r, j) either if $r = 0$

Received April 1973; revised July 1974.

¹ Partially supported by the Office of Naval Research under Contract N00014-69-A-0200-1036 and the U.S. Army Research Office Durham under contract DAHC04 74G0226 with the University of California.

AMS 1970 subject classification. 62L15.

Key words and phrases. One stage look-ahead policy, optimal stopping, modified problem.