THE INSTITUTE OF STATISTICS

THE CONSOLIDATED UNIVERSITY
OF NORTH CAROLINA

EASILY DETERMINING WHICH
URNS ARE "FAVORABLE"

Gordon Simons
University of N.C.
Chapel Hill, N.C. 27514

Mimeo Series #1596

February 1986

DEPARTMENT OF STATISTICS
Chapel Hill, North Carolina

Easily Determining Which Urns are "Favorable"

by Gordon Simons

University of North Carolina at Chapel Hill

ABSTRACT

The optimal sampling strategy for an urn, containing known numbers of plus and minus ones, can be simply described with the use of a empirically justified rule, based upon what appears to be a legitimate third-order asymptotic expansion of "the optimal stopping boundary" as the urn size goes to infinity. The rule performs exceedingly well. There is a known first-order asymptotic expansion due to Shepp. The reader is invited to try to justify a second-order asymptotic expansion of a type described by Chernoff and Petkau. The evidence presented in its support is very persuasive.

AMS 1980 Subject Classifications: Primary 62L15
Keywords: optimal stopping, asymptotic expansions, stopping boundary

The author's work is supported by the National Science Foundation, Grant Number DMS-8400602.

0. Motivation. The author's interest in the urn problem described below was spawned by a desire to apply the methods of approximation developed in [7] and [8] to a rather different type of optimal stopping problem. While this desire was largely unfulfilled, the author found that: (i) a stopping rule based upon an asymptotic approximation described by Shepp (1969) performs well for most urns, (ii) a simple modification, motivated by Chernoff and Petkau (1976), performs substantially better, and (iii) a second, empirically-fashioned modification performs stupendously; it hardly ever makes a mistake: Two urns of about 1.5 billion considered are misclassified.

It is certainly hoped that someone with a theoretical bent will be inspired, by the evidence, to find justifications for these modifications; the numerical evidence leaves no doubt that the first modification is the right thing to do.

The practical value of working with the urn problem needs to be stressed. The backward recursion, based on the dynamic equation, is so simple that the optimal rule can be worked out *exactly* for very large urns. And without such calculations, the empirical aspects of this paper would be impossible.

A further motivation of the author is to focus attention on a neglected issue that arises when one is working with a discrete-state stochastic process: The concept of "optimal stopping boundary" lacks the usual precision of meaning; there are *many* different "optimal boundaries." All of them partition the lattice of observable points into the same two well-define sets of "optimal continuation points" and "optimal stopping points." (Observable points *within* an optimal boundary must be in both sets.) So the issue is not one that should concern the practitioner. But the possibility of different optimal boundaries raises, for the theoretician, the awkward possibility of different and mutually inconsistent "asymptotic approximations."

An illustration is provided by the asymptotic approximation (7) below, which only applies to *some* optimal boundaries. There are other optimal boundaries with equally valid asymptotic approximations *which do not satisfy* (7).

Since *any* reliable asymptotic description of *any* optimal boundary will do, it does not seem appropriate to single out a particular optimal boundary, *a priori*, and call it "the optimal boundary."

The optimal boundaries arising in [3], [7] and [8] are *not* unique. But the issue is ignored in the first paper; the object is to pursue particular approximations. (See Hogan (1985) for some helpful insight.) And the issue does not arise *directly* in the latter papers. For the main focus of attention is on discovering optimal continuation and optimal stopping points, not optimal boundaries. Here, the non-uniqueness issue is addressed, but, as in [3], particular approximations are pursued.

<u>1. The urn problem.</u> One may draw at random without replacement from an urn containing m minus ones and p plus ones, and stop whenever one wishes. The object is to obtain a large sum. The values m and p are known at the outset, and one is free not to draw at all. The task is to determine whether the urn is "favorable", i.e., whether its expected return $R(m,p)$, under optimal stopping, is (strictly) positive. Letting $\mathscr{C}$ denote the class of (m,p) urns for which $R(m,p)$ is positive, Shepp (1969) observed that an optimal policy is to stop as soon as one reaches a depleted urn not in $\mathscr{C}$.

For small values of m and p, it is easy and practical to determine $R(m,p)$ by induction: $R(m,0) = 0$ $(m=0,1,\cdots)$, $R(0,p) = p$ $(p=0,1,\cdots)$, and for $m,p = 1,2,\cdots$,

$$R(m,p) = \max(0, \frac{m}{m+p} R(m-1,p) + \frac{p}{m+p} R(m,p-1) - \frac{m-p}{m+p}). \tag{1}$$

Naturally, $R(m,p)>0$ when $p>m$. What is somewhat surprising is the fact that $R(m,p)$ can be positive when $m>p$. Clearly it is better not to draw from the urn when $m>p$ than it is to proceed with a policy which calls for a fixed number of draws. But, because the sampling is performed without replacement, it may be desirable to proceed with a more sophisticated policy which depends upon the outcome of the draws.

Let $n = m + p$ and $k = m - p$ (notation that will be used throughout the paper). When $n = 4,5,6,7,8,9$ and 10, it is best to draw at least once from the urn if k is no larger than $0,1,0,1,0,1$ and 2, respectively. When $n=3$ and $k=1$, it does not matter whether one refrains from drawing or one draws until the first plus one is obtained; the expected return in either case is zero. (This phenomenon occurs for no other urn with $n \leq 54{,}000$. Boyce (1973) conjectures there are no other occurrences.) Shepp (1969) has shown that there is a "boundary sequence" $b(1), b(2), \ldots$ such that $\mathcal{C} = \{(m,p): k<b(n)\}$. And he has shown that

$$b(n) = \alpha n^{1/2} + o(n^{1/2}) \quad \text{as} \quad n \to \infty , \tag{2}$$

where the coefficient $\alpha = 0.83992\ldots$ is the unique solution of the equation

$$(1-\alpha^2)\int_0^\infty \exp(\alpha x - x^2/2)\,dx = \alpha. \tag{3}$$

The sequence $b(n)$ is not unique. Since k is always an integer of the same parity as n, odd or even, each $b(n)$ must be specified within a semi-open interval $(b_\ell(n), b_u(n)]$ of length two.

Equation (2) suggests a "first-order asymptotic stopping rule":

$$\left.\begin{array}{l}\text{stop as soon as the current values}\\ \text{of } k \text{ and } n \text{ satisfy } k \geq \alpha n^{1/2}.\end{array}\right\} \tag{4}$$

We shall compare this rule with a "second-order asymptotic stopping rule",

$$\left.\begin{array}{l}\text{stop as soon as the current values of}\\ k \text{ and } n \text{ satisfy } k \geq \alpha n^{1/2} - .5\end{array}\right\}, \tag{5}$$

and with various "third-order asymptotic rules" of the form:

$$\left.\begin{array}{l}\text{stop as soon as the current values of } k\\ \text{and } n \text{ satisfy } k \geq \alpha n^{1/2} -.5 + \beta c(n)\end{array}\right\}, \tag{6}$$

where the sequence $c(n)$ goes to zero with n, and β is an empirically determined constant, depending on $c(n)$.

The " -.5 ", appearing in (5) and (6), is suggested by the work of Chernoff and Petkau (1976, page 888): When n is large and k is near the interval $(b_\ell(n), b_u(n)]$, the proportion of minus ones in the urn must be close to one-half. Consequently, as n begins to decrease, the changing of k is closely approximated by a random walk with steps sizes $\pm$ 1, the circumstance for which the correction " -.5 " is appropriate.

Note that stopping rules (5) and (6) are modest refinements of (4); the adjustments " -.5 " and " -.5 + $\beta c(n)$ ", when n is large, are both small when compared to the length of the interval $(b_\ell(n), b_u(n)]$. So if there is a boundary sequence of the form

$$b(n) = \alpha n^{1/2} - .5 + o(1) \quad \text{as} \quad n \to \infty , \tag{7}$$

there must be others which are *not* of this form. These small refinements can make a significant difference, as we will see.

Based on a careful examination of the locations of the allowable intervals $(b_\ell(n), b_u(n)]$, it seems likely that there exist optimal boundaries which are *increasing and concave*. It also seems likely that all increasing and concave optimal boundaries will agree up to a suitable third-order term. So *unique* asymptotic results might hold for "smooth optimal boundaries." The desired number of terms in the asymptotic result, i.e., the desired accuracy, would determine the amount of smoothness required.

2. The favorable urns. The author has determined all of the favorable urns for urn sizes $n \leq 54{,}000$ using about 100 hours of computing time on an IBM - AT. A concise summary is given in Table I. The interpretation is as follows: If $k \leq k_o$ and $n \geq n_o$ for some pair (k_o, n_o), then $R(m,p) > 0$ and the urn is favorable for sampling. Otherwise, it is optimal not to sample. Formally, n_o is define by:

$$n_o = \min(\text{ integer } \ell \text{ of parity } k_o : R((\ell + k_o)/2, (\ell - k_o)/2) > 0). \tag{8}$$

Boyce (1973) argues, incorrectly, that these calculations can *not* be conducted with adequate precision in *floating-point* arithmetic, and he shows how they *can* be performed in fixed-point (integer) arithmetic *for small urn sizes*. Table I was produced using *double precision* floating-point arithmetic, and the results were rechecked using *single precision* floating-point arithmetic. Both gave identical results in all but one instance, a genuinely close case which simply could not be resolved with the significant digits available in single-precision arithmetic. Accumulated round-off error was not a problem. A simple variant of (1) was used, based on the function $S(m,p) = R(m,p) - (m-p)$, because, unlike (1), it can be programmed to avoid all subtractions of positive *floating-point* numbers.

Boyce's paper contains many interesting results. For instance, he shows that $(m+1,p+1)$ is an optimal continuation point whenever (m,p) is, a fact which helps make possible the concise summary shown in Table I: There are 1,458,081,000 urns of size $n \le 54{,}000$.

k_0	n_0	k_0	n_0	k_0	n_0	k_0	n_0	k_0	n_0	k_0	n_0
-1	1	32	1498	65	6081	98	13752	131	24511	164	38356
0	2	33	1591	66	6268	99	14033	132	24884	165	38823
1	5	34	1686	67	6457	100	14316	133	25261	166	39294
2	10	35	1787	68	6650	101	14603	134	25642	167	39769
3	19	36	1888	69	6847	102	14892	135	26025	168	40244
4	30	37	1993	70	7044	103	15183	136	26410	169	40723
5	43	38	2100	71	7245	104	15478	137	26799	170	41206
6	60	39	2211	72	7450	105	15777	138	27190	171	41691
7	81	40	2324	73	7657	106	16076	139	27583	172	42178
8	102	41	2441	74	7866	107	16379	140	27980	173	42669
9	129	42	2560	75	8079	108	16686	141	28381	174	43162
10	156	43	2681	76	8294	109	16995	142	28782	175	43657
11	187	44	2806	77	8513	110	17306	143	29189	176	44156
12	222	45	2935	78	8734	111	17621	144	29596	177	44659
13	259	46	3064	79	8959	112	17940	145	30007	178	45162
14	298	47	3197	80	9184	113	18259	146	30422	179	45671
15	341	48	3334	81	9415	114	18582	147	30837	180	46180
16	386	49	3473	82	9648	115	18909	148	31258	181	46693
17	435	50	3614	83	9883	116	19238	149	31679	182	47210
18	486	51	3759	84	10120	117	19569	150	32106	183	47729
19	539	52	3906	85	10361	118	19904	151	32533	184	48250
20	596	53	4057	86	10606	119	20241	152	32964	185	48775
21	655	54	4210	87	10851	120	20582	153	33397	186	49302
22	718	55	4365	88	11102	121	20925	154	33834	187	49833
23	783	56	4524	89	11353	122	21270	155	34273	188	50366
24	850	57	4687	90	11608	123	21619	156	34716	189	50901
25	921	58	4850	91	11867	124	21970	157	35161	190	51440
26	996	59	5017	92	12128	125	22325	158	35610	191	51981
27	1071	60	5188	93	12391	126	22682	159	36059	192	52526
28	1152	61	5361	94	12658	127	23041	160	36514	193	53073
29	1233	62	5536	95	12927	128	23404	161	36969	194	53622
30	1318	63	5715	96	13200	129	23771	162	37430	195	54001+
31	1407	64	5896	97	13475	130	24140	163	37891		

TABLE I

3. Empirical studies. An empirical assessment of the asymptotic rules, based on the information in Table I, will now be described. For greater precision when n is large, the more accurate value $\alpha = 0.8399236757$ will be used with all calculations; all of the digits shown are significant.

A sequence $\gamma(1), \gamma(2), \cdots$ will be said to be "optimal at n" if $\gamma(n) \in (b_\ell(n), b_u(n)]$, i.e., if the favorable urns of size n are those for which $k < \gamma(n)$.

The "first-order asymptotic rule" (4) performs well; the sequence $\alpha n^{1/2}$ is optimal at n about 75% of the time. When it errors, it misclassifies a single urn, of a given size, as "favorable." The "second-order asymptotic rule" (5) performs much better; the sequence $\alpha n^{1/2}-.5$ is optimal at n about 99.7% of the time when $n \leq 54{,}000$. One can expect this percentage to converge to 100% as $n \to \infty$. For the number of times $\alpha n^{1/2}-.5$ *fails* to be optimal at n appears to grow with n like $n^{1/2}\log^2 n$. When it errors, it misclassifies a single urn, of a given size, as "not favorable."

There is a simple conjecture which, if true, explains the 75% performance rate for (4), and the reason that (5) performs so much better. Consider the "empirical distribution functions"

$$F_N(x) = N^{-1}\sum_{n=1}^{N} 1(\alpha n^{1/2} - .5 \leq b_\ell(n) + x) \qquad ,x\in\mathbb{R},\ N\geq 1, \qquad (9)$$

where $1(\cdot)$ denotes the indicator function. Then $F_N(1.5) - F_N(-.5)$ is the proportion of times $n \leq N$ that the sequence $\alpha n^{1/2}$ is optimal at n. And $F_N(2) - F_N(0)$ is the same proportion for the sequence $\alpha n^{1/2} - .5$. The conjecture is that F_N has a limiting uniform distribution on the interval $[0,2]$, so that

$$F_N(1.5) - F_N(-.5) \to .75 \quad \text{and} \quad F_N(2) - F_N(0) \to 1 \quad \text{as} \quad N \to \infty. \tag{10}$$

The evidence for the conjecture is strong. See Table II, which describes the situation when $N = 54{,}000$. The entry

$$f_i = N[F_N(i/50) - F_N((i-1)/50)]$$

is the number of times $n \le 54{,}000$ that $\alpha n^{1/2} - .5 - b_\ell(n)$ is in the interval $((i-1)/50, i/50]$. Under a uniform distribution on $[0,2]$, the expected cell frequencies would be 540 for $i = 1,2,\cdots,100$, and would be 0 for $i = 0$. In fact, if the cells for $i = 0$ and 100 are combined, then the chi-square goodness of fit statistic becomes

$$\sum_{i=1}^{99} \frac{(f_i-540)^2}{540} + \frac{(f_0+f_{100}-540)^2}{540} \doteq 4.66, \tag{11}$$

which describes an excellent fit. Even though $F_N(0) > 0$ and $F_N(x) = 1$ for some $x < 2$, one still has $F_N(-.0078) = 0$ and $F_N(1.995) < 1$. The entry $f_0 = 181$ in Table II is the number of times $n \le 54{,}000$ that the sequence $\alpha n^{1/2} - .5$ *fails* to be optimal at n. As noted earlier, this count appears to grow with N like $N^{1/2}\log^2 N$. If so, then $F_N(0) = O(N^{-1/2}\log^2 N)$ as $N \to \infty$.

i	f_i	i	f_i	i	f_i	i	f_i	i	f_i	i	f_i	i	f_i	i	f_i
0	181	13	547	26	536	39	540	52	538	65	534	78	546	91	531
1	540	14	534	27	550	40	540	53	540	66	544	79	537	92	535
2	535	15	542	28	534	41	540	54	547	67	544	80	548	93	542
3	536	16	541	29	547	42	538	55	536	68	538	81	536	94	545
4	534	17	541	30	536	43	534	56	535	69	547	82	537	95	534
5	538	18	544	31	549	44	541	57	535	70	538	83	554	96	537
6	547	19	541	32	544	45	540	58	547	71	543	84	536	97	535
7	540	20	539	33	539	46	536	59	541	72	539	85	545	98	538
8	544	21	541	34	540	47	542	60	538	73	544	86	531	99	543
9	528	22	546	35	537	48	539	61	542	74	539	87	537	100	364
10	549	23	529	36	542	49	532	62	535	75	535	88	534		
11	539	24	552	37	539	50	538	63	545	76	544	89	541		
12	538	25	541	38	534	51	541	64	541	77	548	90	542		

TABLE II

It should be mentioned that the sequence $\alpha n^{1/2} - .5 - b_\ell(n)$ is "uniformly distributed *modulo 1*". See Hlawka (1984), pages 17,18, and 23. In fact, it can be shown that it is uniformly distributed *modulo 2*. So F_N has a limiting uniform distribution on $[0,2]$ if, and only if, $F_N(2) - F_N(0) \to 1$ as $N \to \infty$.

All of the 181 mistakes of (5), when $n \leq 54{,}000$, are caused because $\alpha n^{1/2} - .5$ is slightly too small; what is needed in (6) is a positive correction "$\beta c(n)$" with $c(n)$ converging to zero with n. There is no theory to guide the choice of $c(n)$. But an effective empirical approach is to make a reasonable choice, such as $c(n) = n^{-1/2}$, and then try to fit β so that the resulting sequence $\alpha n^{1/2} - .5 + \beta c(n)$ is optimal for $1 \leq n \leq N$, with N chosen as large as possible. Some results are shown in Table III, where "$\log_2$" refers to an iterated logarithm: "log log". The apparent winner is $c(n) = n^{-1/2}\log^2 n$. The corresponding β satisfies the inequality

$$0.008890 \leq \beta \leq 0.008976. \tag{12}$$

By using any β in this interval together with $c(n) = n^{-1/2}\log^2 n$, the third-order rule described in (6) will perform flawlessly *for every urn of size* $n \leq 30{,}836$, for a total of 475,475,702 urns!

Choice of $n^{1/2}c(n)$	1	log n	$\log^2 n$	$\log^2 n \cdot \log_2(n+2)$	$\log^2 n / \log_2(n+2)$
Maximum possible N	100	849	30,836	1495	17,305

TABLE III

Even if $c(n) = n^{-1/2}\log^2 n$ goes to zero at the right rate, one must be cautious with inequality (12). For there is likely to be a fourth-order term, leading to a "fourth-order asymptotic stopping rule", which slightly perturbs the calculations that give rise to (12). There is evidence for this: For β near the lower end of (12), the sequence $\alpha n^{1/2} -.5 + \beta n^{-1/2}\log^2 n$ *fails* to be optimal at an $n \leq 54{,}000$ a total of five times. Near the upper end of (12), the total drops to three. And for $\beta = 0.009165$, *outside the range shown in* (12), the total is *only two*.

There is an analog to (9) which takes into account the additional term "$\beta c(n)$". As with (9), the empirical distributions apparently have a limiting uniform distribution on [0,2]. The corresponding analog of Table II suggests a much better fit near the endpoints of [0,2], and a comparable (excellent) fit elsewhere.

REFERENCES

[1] Boyce, W. (1973). On a simple optimal stopping problem. Discrete Mathematics. **5** 297-312.

[2] Chernoff, H. (1965). Sequential tests for the mean of a normal distribution iv (discrete case). Ann. Math. Statist. **36** 55-68.

[3] Chernoff, H. and Petkau, A.J. (1976). An optimal stopping problem for sums of dichotomous random variables. Ann. Prob. **4** 875-888.

[4] Hlawka, E. (1984). *The Theory of Uniform Distribution* (A B Academic Publishers, Berkhamsted).

[5] Hogan, M. (1985). Comments on a problem of Chernoff and Petkau. Unpublished.

[6] Shepp, L.A. (1969). Explicit solutions to some problems of optimal stopping. Ann. Math. Statist. **40** 993-1010.

[7] Simons, G. (1986). The Bayes rule for a clinical-trials model for dichotomous data. Ann. Statist. To appear.

[8] Simons, G. and Wu, X. (1986) On Bayes tests for $p \leq 1/2$ versus $p > 1/2$: Analytic approximations. To appear in a volume honoring Professor Herbert Robbins.