PS-OR-93-002

AD-A258 779

NAVAL POSTGRADUATE SCHOOL
Monterey, California

DTIC
ELECTE
JAN 0 6 1993
S B D

Performance Review of
The Officer Rate Generator,
Version 1.0

Robert R. Read

October 1992

Approved for public release; distribution is unlimited

Prepared for:
HQ USMC, Code MI
Washington, D.C. 20380-0001

93-00343

NAVAL POSTGRADUATE SCHOOL
MONTEREY, CALIFORNIA

Rear Admiral R. W. West, Jr.
Superintendent

Harrison Shull
Provost

This report was prepared for and funded by HQ USMC, Code MI, Washington, D.C.. Reproduction of all or part of this report is authorized.

This report was prepared by:

ROBERT R. READ
Professor
Operations Research Department

Reviewed by:

PETER PURDUE
Chairman and Professor
Operations Research Department

Released by:

PAUL J. MARTO
Dean of Research

REPORT DOCUMENTATION PAGE

1a REPORT SECURITY CLASSIFICATION UNCLASSIFIED	1b RESTRICTIVE MARKINGS
2a SECURITY CLASSIFICATION AUTHORITY	3 DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution is unlimited.
2b DECLASSIFICATION/DOWNGRADING SCHEDULE	
4 PERFORMING ORGANIZATION REPORT NUMBER(S) NPS-OR-92-002	5 MONITORING ORGANIZATION REPORT NUMBER(S)
6a NAME OF PERFORMING ORGANIZATION Naval Postgraduate School	6b OFFICE SYMBOL (If applicable) OR
7a NAME OF MONITORING ORGANIZATION HQ USMC, Code MI	
6c ADDRESS (City, State, and ZIP Code) Monterey, CA 93943-5000	7b ADDRESS (City, State, and ZIP Code) Washington, D.C. 20380-0001
8a NAME OF FUNDING/SPONSORING ORGANIZATION HQ USMC, Code MI	8b OFFICE SYMBOL (If applicable)
9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER	
8c ADDRESS (City, State, and ZIP Code) Washington, D.C. 20380-0001	10 SOURCE OF FUNDING NUMBERS PROGRAM ELEMENT NO / PROJECT NO / TASK NO / WORK UNIT ACCESSION NO

11 TITLE (Include Security Classification)

Performance Review of The Officer Rate Generator, Version 1.0

12 PERSONAL AUTHOR(S)
Robert R. Read

13a TYPE OF REPORT	13b TIME COVERED	14 DATE OF REPORT (Year Month Day)	15 PAGE COUNT
Technical Report	FROM 1/1/92 TO 9/20/92	10 October 92	34

16 SUPPLEMENTARY NOTATION

17 COSATI CODES			18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	

19 ABSTRACT (Continue on reverse if necessary and identify by block number)

The report deals with the author's review of the manpower attrition rate generator developed by Decision Science Associates, Inc. for use with the Marine Corp's manpower planning models. The major focus is on the implementation and documentation of the author's methodology for treating the rather vast number of small personnel inventory cells present. That methodology features two main steps: an algorithm for the aggregation of extremely small but moderately sized cells; and the use of modern "shrinkage" methods applied to the groups in order to provide statistically stable attrition rates. These methods are described and suggestions for expediting their use are made.

The use of these methods in other communities is considered and groundwork is laid for such extensions. The question of default values for some model parameters has been studied and a framework to treat these issues has been developed.

20 DISTRIBUTION/AVAILABILITY OF ABSTRACT ☒ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT ☐ DTIC USERS	21 ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED
22a NAME OF RESPONSIBLE INDIVIDUAL ROBERT R. READ	22b TELEPHONE (Include Area Code) (408) 646-2383 / 22c OFFICE SYMBOL OR/Re

DD FORM 1473, 84 MAR 83 APR edition may be used until exhausted
All other editions are obsolete

Abstract

The report deals with the author's review of the manpower attrition rate generator developed by Decision Science Associates, Inc. for use with the Marine Corps's manpower planning models. The major focus is on the implementation and documentation of the author's methodology for treating the rather vast number of small personnel inventory cells present. That methodology features two main steps: an algorithm for the aggregation of extremely small but similar (in terms of attrition behavior) cells into groups of moderately sized cells; and the use of modern "shrinkage" methods applied to the groups in order to provide statistically stable attrition rates. These methods are described and suggestions for expediting their use are made.

The use of these methods in other communities is considered and groundwork is laid for such extensions. The question of default values for some model parameters has been studied and a framework to treat these issues has been developed.

DTIC QUALITY INSPECTED 5

Accession For	
NTIS GRA&I	☑
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

PERFORMANCE REVIEW OF THE OFFICER RATE GENERATOR, VERSION 1.0

1 INTRODUCTION

The Officer Rate Generator (ORG) was developed by Decision Systems Associates, Inc. (DSAI) in 1991, reference [10]. The personnel at DSAI performed an execellent and thoughtful job in its conception and development. The product is a general purpose one, and feedback based upon the experiences of the various consumers is valuable for its improvement and the extension of its capabilities.

The author and his students have been working on the "small cell" problem for the Officer Rate Generator (ORG) for a number of years; see reference [4]. Implementation of the work was recommended in 1990, specifically the versions in [2,3], and such was made a part of the ORG system [10]. Most developments have "growing pains" and the present review was authorized under Document Number M00027-92-WRA8054 for code MI of HQUSMC in 1992. This paper is the report of that work. An interim report was issued in March [10]. It treated an implementation error and provided a first draft for documenting the "small cell" section of ORG. DSAI reports that the implementation error has been corrected. The final version of the "small cell" section is contained herein as Section 3. It is offered to replace Appendix C of reference [10].

Throughout the paper it is assumed that the reader has some familiarity with the ORG and is also familiar with the the statement of work (SOW-RRREAD) in the authorizing document. The bulk of the report deals with system performance. We

begin in Section 2 with some rather general input and output items which are a bit cosmetic, but nonetheless are of a time saving nature for most users. This is followed by a deeper discussion of the organizational structure of ORG and the impact that it may have upon anticipated use.

Much of the performance evaluation is concerned with the interface between the details of implementation of the small cell methodology and the setting in which it was tested and is known to perform well. The author believes that substantial improvement can be made in these areas; the case is presented. This analysis leads us to some new requirements and, because of this, some enhancements to the existing methodology are included.

A rather important item is the manner of use of a set of attrition rates. For example, it appears that users of the Steady State Promotion Model apply rates that discriminate grade and years commissioned service (YCS), but not MOS (military occupation specialty). The argument given is "we don't promote by MOS." The implication is that MOS is therefore not a useful discriminating tool when forecasting the number of attritions (vacancies) by grade and YCS. There is danger in this attitude and the issue is discussed in Section 4.

Section 5 contains details of certain extensions. Firstly, the aggregation system is extended to include Colonels, LDO's, and Warrant Officers. Secondly, the aggregation system is discussed more thoroughly and a table is presented showing what we have learned about the relationships between aggregation threshold parameters and various combinations of military occupational specialty (MOS) and years of commissioned service (YCS).

The producer of ORG, DSAI of Rockville, Maryland has supplied the author with the software and supporting personnel data. The author has installed it on an 80486/33 PC. All experiences relate to this installation.

2 COMMENTS ON THE INPUT/OUTPUT SYSTEM

Consider the general issue of default values for the input parameters of the rate generator. There are a number of input parameters and they all must be set with each application of the program. Such resetting of all is time consuming and annoying since in many of his sessions the analyst has but modest changes of inputs between successive applications. It would be useful to replace the present system of serial inputs with a single panel which would allow the user to maintain the inputs from the immediately preceding application and mark any changes by moving a mouse or relocating the cursor. Items such as dictionary, community, gender, zone, and loss type and likely to be constant over a rather lengthy series of studies. Much time is lost in the repeat resetting of the same values.

These issues also impact with the output reporting formats and now is a good time to make related comments about that. Let us itemize the suggestions and supporting statements.

a. The Run option is a major choice and is quite separate from the Update and Dictionary choices. Perhaps the choice of these three can be included in the form of calling the program. The DosShell and Exit options should be available at all times below a command line.

b. The issues of marking small cells and performing small cell analysis need not be

related. The small cell threshold has a default value of 20 which can be altered by the user. The effect of this is to place an asterisk next to the involved output data. I found this difficult to read on the printed output and would like to suggest a different technique, such as underlining or boldface. The small cell analysis portion has two aspects which will be discussed later in the report. At this time we will merely call the aspects "aggregation" and "Empirical Bayes"; and if selected there will be two or more parameters requiring values on the panel. They need not be tied to the small cell threshold.

c. The weighting of the data by calendar year might be given a default setting, such as equally weighting the most recent ten years.

d. The six choices of community are fine, and the default should be the unrestricted line. I wonder if this option can be inserted earlier; there is much staging of data to be done and the user might be making other choices while the computer is churning.

e. Both gender and zone should have default settings; perhaps "all" in each case.

f. The loss types provide a rich partition of the data. However, I fear that some are community dependent and it may be wise to protect the user from a false start. The distinction between voluntary and involuntary is not sharp in my mind. It likely refers to the issue of which party exercises the option. However for use in manpower models it may be wiser to separate on whether the model can anticipate the loss, or whether it must treat it as a surprise. By the way the choice of voluntary or involuntary death appears foolish and might be modified. Also it appears that manpower analysts may be interested directly in disability losses. At present they seem to be buried in other categories.

g. A separate panel for MOS selection seems quite necessary. Indeed the choices are community dependent and it would be best if the formatting reflected that. Certain broad categories may be unavailable.

h. The report formats are inefficient. There are two basic types: raw data counts and rate estimates. They should be combined within types. The counts of personnel inventory and losses can be printed as a pair, side by side, for each entry in the tables. For example, the YCS by GRADE table could be tailored to the given information. If certain grades are not represented in the community/MOS selection, then the headings should reflect this. Since there are no Colonels having YCS less than twenty, the columns should begin at twenty, etc. Like considerations apply to the rate estimates. They can be paired placing continuations and attritions side by side.

i. The previous suggestion may be difficult when the YCS by Month option is chosen. This option implies a pooling of grades, however, and this analyst regards that as not being very useful. See section 3 of the report.

j. A thought contained in f. can be carried further. Certain MOS cells and groups are not always available in the face of previous selections. The user should not be drawn into attempting to perform unproductive analyses. The female gender comes to mind as an example. A like comment holds for the YCS intervals, which are grade dependent.

3 ASPECTS OF SMALL CELL TECHNOLOGY

What is the small cell problem?

We are involved in the problem of estimating attrition probabilities for a large number of personnel cells many of which contain but a few units (inventory or stocks of personnel). Moreover there is a preponderance of low rates (i.e., small attrition probabilities). The use of flat rates can lead to some rather rigid, and in some cases foolish, interpretations. For example consider two cells, one with three officers and a second with 30 officers. Suppose neither cell experiences any attritions. The flat rate for both cells is zero. Yet the two zero values do not mean that same thing. Would you be willing to give essentially infinite odds against there ever begin an attrition in these cells? Moreover, whatever odds you might be willing to give, would they be the same for both cells? At some time in the future we may be contemplating having 25 officers in each of these same two cells. The use of zero probabilities of attrition for each is extreme and unwise.

Consider also the other extreme. It would be far less significant if all the members of the first cell were to attrite than if nearly all members of the second attrited. Also, intermediate values of flat rates are rather rigid for small cells as they are integral multiples of the cell size.

Modern multiparameter estimation techniques can be used to give relief to these problems. They exploit the communalities that exist among similar cells.

How do we deal with the small cell problem?

The basic idea involves the use of similar cells to provide information useful in developing an attrition probability for a small cell. By "similar" cells we mean cells that have a communality of aspects and whose long-term historical records exhibit

comparable attrition behavior. The common aspects issue is decided largely by the organizational structure and is driven by items such as officer type, grade, MOS group, etc. The comparable attrition behavior issue is treated by mining and analyzing the data of years past.

The details of development can be found elsewhere. Our immediate goal is to provide the reader with a sense of understanding as to how the small cell methodology system looks for similar cells. The rules are complicated, and it seems best to walk through a few examples in order to illustrate the mentality of what the rules are doing. To this end we have chosen the community of Fixed Wing Pilots; an interesting community characterized by high training costs, low attrition, and not large stocks (personnel inventory) of officers.

Table 1 shows the five year average of stocks and losses broken out by grade, YCS, and MOS. A scan of the table shows a number of points. First of all the blanks indicate zero inventory. Such cells may be empty for organizational reasons and hence called structural zeros; or they may be empty by happenstance and hence called sampling zeros. Second, the entries are stock-losses pairs. For example the first row of lieutenants, having YCS = 3 and MOS = 7507 shows 4 personnel and 0 losses. Third, there are no colonels. Fourth, the lieutenants do not appear until YCS = 3, they do not stay long (as lieutenants) and there is almost no attrition. Next, there are many small cells. There are 261 nonempty cells and, using 15 as a threshold for small cells, there are 240 small cells; over 90%. Of course other threshold values could be used but 15 seemed to be toward the lower end of usable values.

There are two aspects of Table 1 that may be mystifying the reader. One is the break out by SML code and the other is the boxing of selected YCS levels. These are aids to the selection of "similar" cells and will be explained shortly.

TABLE 1
FIXED WING PILOTS: EQUALLY WEIGHTED AVERAGES, 1987–1991

Lieutenants

SLM Code = 7,3,2

YCS/MOS	7500	7501	7507	7511	7522	7543	7545	7576
3			4 0					
4	1 1		17 0	2 0	1 0	1 0	1 0	1 0
5	2 0	5 0	3 0	9 0	4 0	3 0	3 0	4 0
6		1 0		1 0				1 0

SLM Code = 8,3,2

YCS/MOS	7521	7523	7524
3	5 0		1 0
4	19 0	10 0	
5	14 0	20 0	
6		1 0	

SLM Code = 9,4,2

YCS/MOS	7556	7557	7558
3	6 0		9 0
4	10 0	5 0	6 0
5	12 0	5 0	2 0
6	1 0		
7			
8		1 1	

SLM Code = 10,4,2

YCS/MOs	7509
4	10 0
5	19 0
6	1 0

SLM Code = 13,5,3

YCS/MOS	7510	7550	7575
2		1 0	
3	1 0	7 0	
4	6 0	4 0	3 0
5	1 0	2 0	2 0

Captains

SLM Code = 7,3,2

YCS/MOS	7501	7507	7511	7522	7543	7545	7576
5	1 0	1 1	3 0	1 0	1 0	1 0	1 0
6	13 0		16 0	10 1	3 0	5 0	8 1
7	16 5		17 4	14 3	2 0	5 1	10 4
8	11 5		10 5	15 4	1 0	2 1	6 4
9	7 3		6 2	13 3	1 0	1 1	3 1
10	5 1		5 0	10 2	2 0	2 0	3 0
11	4 1		6 1	7 1	3 0	2 0	3 0
12	3 1		5 1	5 1	2 1	2 0	1 1
13			2 1	1 1			

Table 1 (Continued)

SLM Code = 8,3,2

YCS/MOS	7521	7523	7527
5		6 0	
6	1 0	23 0	
7	1 0	22 3	
8	1 0	19 2	
9	1 0	19 3	
10		20 2	1 1
11		21 2	
12		17 2	
13		3 1	

SLM Code = 9,4,2

YCS/MOS	7556	7557	7558
5	2 0	3 0	
6	13 1	11 1	1 0
7	7 4	12 5	1 0
8	4 1	7 4	1 0
9	2 1	5 2	
10	1 0	4 0	
11	2 0	5 1	
12	1 1	5 1	
13		1 1	

SLM Code = 10,4,2

YCS/MOS	7508	7509
5		
6	1 0	
7	1 0	10 1
8	1 0	10 0
9	1 1	1 0
10	1 0	7 0
11	1 0	22 0
12		22 2
13		16 4

SLM Code = 13,5,2

YCS/MOS	7510	7550
6	1 0	1 0

Majors

SLM Code = 7,3,2

YCS/MOS	7501	7511	7522	7543	7545	7576
12	1 0	2 0	1 0			1 0
13	3 0	5 0	4 0	2 0	1 0	4 0
14	4 0	6 0	5 0	3 0	1 0	4 0
15	4 0	7 0	5 0	3 0	1 0	4 1
16	5 0	7 0	4 0	2 0	1 0	4 0
17	6 0	5 0	4 0	2 0	1 0	4 0
18	5 0	3 0	3 0	1 0	1 0	2 0
19	4 0	3 1	4 0	1 0	1 0	2 0
20	1 3	3 3	5 5		1 1	2 2

Table 1 (Continued)

SLM Code = 8,3,2 SLM Code = 9,4,2 SLM Code = 10, 4,2

YCS/MOS	7523	YCS/MOS	7557	YCS/MOS	7508	7509
9	1 0			9		9 2
12	2 0	12	2 1	10		1 0
13	14 1	13	6 0	12		1 0
14	16 1	14	8 0	13		7 1
15	15 0	15	10 1	14	1 0	6 0
16	11 0	16	9 0	15	1 0	6 0
17	8 0	17	8 1	16		5 0
18	2 0	18	5 0	17		5 0
		19	3 0	18	1 0	2 0
		[20]	2 2	19		2 0
				[20]	1 1	1 1

Lieutenant Colonels

SLM Code = 7,3,2

YCS/MOS	7501	7511	7522	7543	7545	7576
17	1 0	1 0	1 0			
18	4 0	3 0	3 0	1 0		1 0
19	5 0	5 0	5 0	1 0	1 0	2 0
20	7 1	7 1	6 1	1 0	1 0	3 0
21	7 3	6 1	7 1	1 0		3 1
22	6 1	5 1	5 1	1 0		2 1
23	2 1	2 1	2 2	1 0		1 1
24	1 0	1 0	1 1			
25	1 0		1 0			
[26]			1 1			

SLM Code = 8,3,2 SLM Code = 9,4,2 SLM Code = 10,4,2

YCS/MOS	7423	YCS/MOS	7557	YCS/MOS	7508	7509
17	1 0	17	1 0	14		1 0
18	8 0	18	3 0	18	1 0	4 0
19	9 0	19	5 0	19	2 0	5 0
20	10 1	20	5 1	20	2 1	6 1
21	9 2	21	4 1	21	2 0	6 1
22	6 1	22	3 1	22	1 0	5 1
23	1 1	23	1 0	23	1 1	1 0
		24	1 1	24		1 0

The SLM code system is a hiearchial partitioning of the MOS indices. The symbols SLM stand for small, large, and major. Each MOS value belongs to one and only one of 14 (this is the value for unrestricted officers, non-colonels) small MOS groups. Within each small group there is a communality of organizational aspects and historical attrition behavior. If a user cannot find adequate stocks in his selected MOS index, he can aggregate further stocks from within the same small MOS group. (Actually ORG performs the aggregation.) Each small MOS group belongs to one and only one of six large MOS groups, and each large MOS group belongs to one and only one of three major MOS groups. This system provides a path for seeking additional similar cells for aggregation, should the situation require it. The MOS indices are sorted by SLM codes in Table 1.

Another way to seek additional stocks is to consider neighboring YCS levels. The rules for doing this follow an alternating "next smaller, next larger" pattern. For example, if a user were interested in YCS = 12 but found a need to aggregate more stocks, he would successively turn to YCS = 11, YCS = 13, YCS = 10, etc. However, there are exceptions or boundaries to the application of this rule, and they are keyed to the small MOS group. In each case the YCS scale is partitioned into four sets as indicated in Table 2.

TABLE 2

SETS OF NEIGHBORING YCS LEVELS

Small MOS Group	First Set	Second Set	Third Set	Fourth Set
7,8,14	(1,6)(8,19)	(7)	(20,25)	(26,30)
9,10	(1,5)(8,19)	(6,7)	(20,25)	(26,30)
All others	(1,3)(6,19)	(4,5)	(20,25)	(26,30)

All sets in the partition are intervals, save the first. It appears that the second set is capturing the early career departures while the third and fourth sets represent types of retirements. The reader may have sharper interpretations, but our only concern is to draw attention to the roles played by these sets in the search for "similar" cells. The changing YCS sets are marked with boxes in Table 1.

Caveat. The preceding details apply to unrestricted regular officers. Some details might be different for other communities. Details belong in the Dictionary portion of ORG. ORG provides for user selected changes in the Dictionary details.

The next issue is the order of application of our options when seeking to aggregate similar cells. The initial point is the choice of a community and a grade. (In our first example this will be fixed wing pilots who are lieutenants.) The user may want to treat the entire set or he may be interested solely in a specified MOS and YCS pair. Let us assume the latter, for now, in order to continue. There is a user supplied threshold, t_0, that is used to decide whether the cell is small or not. If the stocks equal or exceed t_0, the cell is not small and no aggregation of similar cells is required. Otherwise we search the small MOS group, holding YCS fixed, and adjoin cells in order to achieve or exceed the inventory threshold, t_0. Often there is more than one way to do this. For technical reasons it is preferable that the aggregated inventory remain close to t_0 rather than to adjoin a relatively large cell which would produce an excessive aggregated stock figure.

It is possible (and not uncommon) for one to exhaust the entire small MOS group before gathering sufficient stocks. The next source of similar cells is to be found in the YCS set. The alternating manner of choosing nearby YCS levels has been described earlier. The original YCS level and small MOS group determine the YCS set to be used (see Table 2). But one must not go outside of the set. When adjoining cells from

a neighboring YCS level one uses all of the MOS values in the small MOS group.

Suppose that all of this is done and the user is still short of his inventory threshold. The rule is to return to the original YCS level and expand to the large MOS group. If more stocks are still necessary then one looks for them in the neighboring YCS values as before.

The pattern is repeated if sufficient stocks cannot be found in the large MOS group and the entire YCS set. The rule is to return to the original YCS value and search the major MOS group for officers of that YCS value. Again one can adjoin neighboring YCS levels as needed. At this point the system of seeking similar cells must stop. The boundaries of major MOS groups and YCS sets will not be crossed. Nor will the grade be augmented. The user will be informed. He may wish to try a different t_0.

The above rules can be illustrated with some examples using Table 1. The grade is lieutenant and $t_0 = 15$.

Example 1. MOS = 7550 and YCS = 3

Using N for the accumulated inventory we see the $N = 7$ at the beginning. This is inadequate so we adjoin MOS = 7510 and N is increased to 8. Next we expand in the YCS scale (YCS = 2), yielding one more; $N = 9$. To continue on the YCS scale we must search YCS = 6 because it is the next priority value in the first YCS set. But there are none, so we must reset YCS = 3 and search the entire small MOS group 13, etc.

Example 2. MOS = 7511 and YCS = 5

The initial value of N is 9 from Table 1. Six more are required and they can be found using either of the MOS pairs (7500,7522) or (7500,7576). Either one of these achieve the value of 15 and is preferable to choices that produce excessive totals, e.g.,

(7501,7500), etc. This example draws attention to the fact that the rules do not produce unique choices. It is believed that the variabilities introduced because of this are rather small.

Thus far we have described a somewhat elaborate way to get an attrition rate for a single cell. (Indeed if that were our goal and flat rates are to be used, then threshold numbers in the range of 30 or more should be used.) Generally we want rates for entire collections of cells and the previously described superstructure is quite useful for this purpose.

Modern multiparameter estimation methods exploit the communality of the cells and produce estimators with smaller MSEs (mean squared errors) than traditional (flat rate) estimators. For a fixed community and grade consider the problem of estimating attrition rates for all of the nonempty cells over a specified set of MOSs and YCSs. We have seen that the inventory numbers for problems of this type can be rather small so let us agree to aggregate cells to the t_0 level in order to assure some modicum of stability. It is important that there is no overlap in these aggregated cells. That is, each original cell must be a part of one and only one aggregated cell. Then we can safely assume statistical independence among the aggregated cells.

Below is a worksheet illustration showing how this extended aggregation algorithm works. It refers to captains in SLM code 7,3,2. The entries are inventory numbers by YCS set 1 and the last two digits of the MOS index. Again $t_0 = 15$. The first pass will treat each row individually and form as many row aggregates as possible. The used cells are boxed on the worksheet and the (YCS, MOS) pairs are listed below, left, together with their total stocks, ΣN. Notice the first row has none; this is acceptable.

WORKSHEET 1

YCS/MOS	01	07	11	22	43	45	76
5	1	1	3	1	1	1	1
6	[13]		[16]	[10]	[3]	[5]	8
8	[11]		10	[15]	1	2	[6]
9	[7		6]	[13	1	1]	[3]
10	5		[5	10]	2	2	3
11	4		[6	7]	3	[2]	3
12	[3		5	5	2]	2	1
13			2	1			

First Pass Pairs	ΣN	Finalization	Total Stocks/Losses
(6,11)	16	(5,01)(5,07)(5,22)	19/1
(6,22)(6,45)	15	(5,11)	18/1
(6,01)(6,43)	16	(5,43)(5,45)(5,76)	19/0
(8,22)	15	(8,45)	17/5
(8,01)(8,76)	17	(8,43)	18/9
(9,22)(9,43)(9,45)	15		15/4
(9,01)(9,11)(9,76)	16		16/6
(10,11)(10,22)	15		15/2
(11,11)(11,22)(11,45)	15	(11,01)	19/3
(12,01)(12,11)(12,22)(12,43)	15	(12,45)(12,76)(13,11)(13,22)	21/7

At this point we have 10 aggregated cells accounting for a total of 155 officers. The straggler group contains 58. Up to three fully aggregated sets can be formed from them. The remaining will be distributed over existing aggregates. The method suggested has some arbitrary aspects and is certainly not unique. The principals employed are

(i) It is not necessary to look outside of small MOS group 7.

(ii) The aggregates to be formed should be selected from a YCS interval of minimal width.

(iii) Once all aggregates are formed, the final allocations should be in the same YCS value, or if necessary in a neighboring one. This should be done so as to minimize the unevenness of the new cell sizes.

Notice that **(iii)** gives priority to "same or near" YCS value over a "general evenness" of cell sizes.

Let us make a second pass to form up to three more new cells. The first step is to find the largest straggler, note its YCS value (say y_0) and search rows $y_0 - 1$ and $y_0 + 1$ for uncommitted cells. Refer to Worksheet 2. It contains the unaggregated stocks left from the first pass, and we repeat the first pass procedure but with YCS expansion. Call it the second pass.

WORKSHEET 2 (second pass)

YCS/MOS	01	07	11	22	43	45	76
5	1	1	3	1	1	1	1
6							[8]
8			[10]		1	2	
9							
10	[5]				[2	2	3]
11	4				[3]		3
12					2	1	
13			2	1			

Second Pass Pairs	ΣN	Finalization	Total Stocks/Losses
(8,11)(6,76)	18		18/6
(10,01)(10,43)(10,45)(10,76)(11,43)	15	(11,76)	18/1

The largest value is 10 and $y_0 = 8$. The 8 in row$(y_0 - 1)$[1] can be adjoined. Eliminate these and repeat the process. The next largest is 5 and $y_0 = 10$. The other stragglers in this row will bring the total to 12. We search rows $y_0 - 1$ and $y_0 + 1$ and select one of the threes in YCS = 11. Eliminate these and discover that a third new cell cannot be formed under the rules.

It remains to assign the rest. We can use the same row, the collection of rows created by the new aggregation in the second pass, or adjacent rows if necessary. Thus each member of row 5 must be assigned to new cells in row 6 or the second pass new cell (8,11)(6,76). It will be done to even up the sizes. The assignments appear to the right of the first and second pass tableaus.

There are 12 cells in the aggregated system and the losses in the new cells are to the right of the slash. In all there are 213 officers and 45 losses; the global loss rate is 21.1%. Details are gathered in Table 3.

[1] $y_0 - 1$ is 6 in this case because the value 7 is in a different YCS set.

TABLE 3
AGGREGATED CELLS
FIXED WING PILOTS WHO ARE
CAPTAINS AND IN SLM CODE 7,3,2

Cell	Composition	Stocks/ Losses
1	(6,11)(5,01)(5,07)(5,22)	19/1
2	(6,22)(6,45)(5,11)	18/1
3	(6,01)(6,43)(5,43)(5,45)(5,76)	19/0
4	(8,22)(8,45)	17/5
5	(8,01)(8,76)(8,43)	18/9
6	(9,22)(9,43)(9,45)	15/4
7	(9,01)(9,11)(9,26)	16/6
8	(10,11)(10,22)	15/2
9	(11,11)(11,22)(11,45)(11,01)	19/3
10	(12,01)(12,11)(12,22)(12,43)(12,45)(12,76)(13,11)(13,22)	21/7
11	(8,11)(8,76)	18/6
12	(10,01)(10,43)(10,45)(10,76)(11,42)(11,76)	18/1

Summary: Forty-five original cells; 12 aggregated cells; total stocks and losses = 213/45.

The second phase of the small cell technology applies an Empirical Bayes (or shrinkage) method to the aggregated cells. Several variations of this method were tested by Misiewicz [3]. Generally they perform nearly the same. The preference for the variation described below is slight.

Generally shrinking methods provide us with a compromise between the stability provided by the grand mean of all cells and the detail (but instability) available from the individual cells. The situation can be illustrated graphically

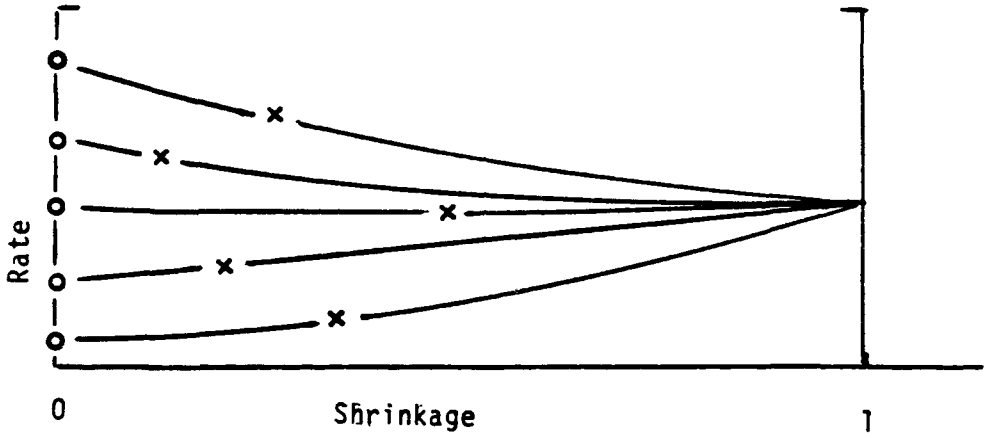

The individual cell rates are marked on the left axis at "zero" shrinkage. The grand mean ordinate is the single rate marked at full shrinkage. The paths of shrinkage can be curved on straight lines. In our system they are curved because the Freeman-Tukey transform is applied prior to a linear shrinkage and inverted to the original scale afterwards. The dots on the shrinkage paths provide the final values. The ordinate values are the rates and the abscissas mark the amount of shrinkage. The amount of shrinkage varys because of the uneven variance of rates from cell to cell. Generally the amount of shrinkage is low when the cell rate variance is small, and high when the cell rate variance is large.

The computational details involve an iteration technique to solve for the amount

of shrinkage. Let there be K cells with pairs (N_i, Y_i) representing the personnel inventory and the number of losses, respectively. Prior to shrinkage we apply the Freeman-Tukey transformation for binomial data,

$$X_j = \frac{1}{2}\sqrt{N_j + .5}\left\{\sin^{-1}\left[\frac{2Y_j}{N_j+1}\right] + \sin^{-1}\left[\frac{2(Y_j+1)}{N_j+1} - 1\right]\right\}$$

and its variance is approximated by (see [3])

$$V_j = a(X_j + C_j)^b(X_j + C_j - 1)^c$$

where

$$a = 1.6835 \qquad b = -0.8934 \qquad c = 0.9881$$

and

$$C_j = \frac{\pi}{2}\sqrt{N_j + 5}.$$

Then V_j is replaced by $\min(1, V_j)$. This approximation to the variance of the transformed value is valid for $Y_j \leq N_j/2$. The counter case is managed using symmetry. The adjustment is as follows: Use X_j^* in the displayed formula for V_j where $X_j^* = X_j$ (i.e., the original X_j) if $Y_j \leq N_j/2$. But if $Y_j > N_j/2$ then replace Y_j with $N_j - Y_j$ when computing X_j^* from the formula for X_j. The caveat applies only to the computation of V_j. The replacement applies nowhere else.

Some scaling is needed before moving on. Let

$$XT_j = X_j/\sqrt{N_j + .5}$$

In those instances involving more than one year of input data, we will have values for N_j, Y_j, X_j, XT_j for each of the years. Then the XT_j must be averaged over the years, producing XTB_j for $j = 1, \cdots, K$.

The next step involves the simultaneous computation of A, the variance of the prior distribution of the mean, and XBB, the grand mean on the transformed scale. What follows can accomplish this.

$$A \longleftarrow A - \left\{ K - 1 - \sum_{j=1}^{K} \alpha_j (XTB_j - XBB)^2 \right\} \Big/ \sum_{1}^{K} \alpha_j^2 (XTB_j - XBB)^2$$

with

$$\begin{aligned} \alpha_j &= 1/[A + Var(XTB_j)] \\ \gamma_j &= \alpha_j \Big/ \sum_{1}^{K} \alpha_j \\ XBB &= \sum_{1}^{K} \gamma_j XTB; \\ Var(XTB_j) &= V_j/(.5 + N_j). \end{aligned}$$

This is an iterative algorithm that can be initialized with $A = 0$. Stop if $A < 0$ (and set $A = 0$) or when the value of A stabilizes to within epsilon, say $\epsilon = 10^{-3}$. Experience with this algorithm suggests that at least five iterations should be allowed prior to stopping. There appears to be instances in which the growth of a positive value for A is rather slow, and premature stopping with A near zero can produce poor results [DSAI, private communication].

The shrinkage step (Empirical Bayes) is a convex combination of the transformed cell means and the grand mean. Thus

$$XEB_j = \frac{A}{A + Var(XTB_j)} + \frac{Var(XTB_j)}{A + Var(XTB_j)} XBB$$

and the inversion to the original scale (probability of attrition) is

$$\hat{q}_j = \frac{1}{2}\{1 + \sin(XEB_j)\}$$

To assure that the reader understands the steps, let us walk through a simple example. There are eight cells, $K = 8$. There are both extremely small cells and extremely large ones, in fact one cell has zero inventory. There is only one year, so that $XTB_j = XT_j$. The data (and ultimate solutions) are

Cell	1	2	3	4	5	6	7	8
Stocks	0	4	60	360	540	550	700	10,000
Losses	0	4	50	300	500	540	670	790
Flat Rates	0.0*	1.00	0.8333	0.8333	0.9259	0.9818	0.9571	0.0790
Emp. Bayes	0.6229	0.9453	0.8273	0.8333	0.9248	0.9807	0.9562	0.0791

The flat rate for the first cell is zero by convention; 0/0 is indeterminate. The Empirical Bayes rates for the small cells (first two) are developed from the other cells. They reflect the common population information of all the cells. The two methods provide essentially the same rates for the six large cells.

The transform and variance of transform values are

Cell	1	2	3	4	5	6	7	8
XTB	.3483	0.111	0.0165	0.0028	0.0019	0.0018	0.0014	0.0001
$Var(XTB)$	0.0	1.107	0.7154	0.7273	1.017	1.294	1.151	−1.001

These are fed into the iterative algorithm then, with the result that

$$A = 0.5713 \quad \text{and} \quad XBB = 0.6554$$

Thus the shrinkage step produces

Cell	1	2	3	4	5	6	7	8
XEB	0.2[illegible]3	1.099	0.7137	0.7269	1.015	1.292	1.149	−1.001

and these values are converted to the Empirical Bayes rates above using the inversion formula.

The reader may also be interested in the final value for α and γ.

Cell	1	2	3	4	5	6	7	8
α	1.098	1.717	1.701	1.742	1.745	1.745	1.740	1.75
γ	0.0823	0.1298	0.1286	0.1316	0.1318	0.1319	0.1319	0.1323

Now let us finish the example of captains who are fixed wing aviators. The results of aggregation were presented in Table 3. Table 4 contains the attrition rates. There are three entries per cell in order to compare the results of differing methods. The first entry is the raw flat rate. Notice how stark they are. The second entry is the flat rate computed after aggregation. They exhibit a smoothness that supposes the use of an aggregation technique. The third entry is the Empirical Bayes shrinkage attrition rate. They exhibit even more smoothness. Note that none of them are zero.

TABLE 4

FIXED WING PILOTS; CAPTIONS; SMALL MOS GROUP 7; FIRST YCS SET; FLAT RATES/POST AGGREGATION FLAT RATES/EMPIRICAL BAYES

YCS/ MOS	7501	7507	7511	7522	7543	7545	7576
5	0.0 0.053 0.103	1.0 0.053 0.103	0.0 0.056 0.107	0.0 0.053 0.103	0.0 0.0 0.014	0.0 0.0 0.014	0.0 0.0 0.014
6	0.0 0.0 0.014		0.0 0.053 0.103	0.1 0.056 0.107	0.0 0.0 0.014	0.0 0.056 0.107	0.125 0.333 0.295
8	0.455 0.5 0.403		0.5 0.333 0.295	0.267 0.294 0.269	0.0 0.5 0.403	0.5 0.295 0.269	0.667 0.5 0.403
9	0.429 0.375 0.318		0.333 0.375 0.318	0.231 0.267 0.251	0.0 0.267 0.251	1.0 0.267 0.251	0.333 0.375 0.318
10	0.2 0.056 0.107		0.0 0.133 0.168	0.2 0.133 0.168	0.0 0.056 0.107	0.0 0.056 0.107	0.0 0.056 0.107
11	0.25 0.158 0.18		0.167 0.158 0.18	0.143 0.158 0.18	0.0 0.056 0.107	0.0 0.158 0.18	0.0 0.056 0.107
12	0.333 0.333 0.299		0.2 0.333 0.299	0.2 0.333 0.299	0.5 0.333 0.299	0.0 0.333 0.299	1.0 0.333 0.299
13			0.5 0.333 0.299	1.0 0.333 0.299			

4 DISCUSSION OF ISSUES RELATING TO DEFAULT SETTINGS AND USER MODEL INTERFACE

The question of default settings for the small cell analysis is intertwined with the issue of speeding the performance. It would be foolish and ineffective to set the defaults independently of the problem, or of the analysis being made. It is important for the user to realize that the techniques provide simultaneous rate estimates for a collection of cells. It was never intended to apply small cell analysis to individual cells one by one. In what follows I will attempt to provide some general guidelines which can be applied once the overall problem (i.e., the overall collection of cells) is defined.

The small cell methodology is described in Section 3. We emphasize that it has two phases: aggregation and shrinkage. The first phase, aggregation, is the most critical. In general, there are two issues in aggregating a collection of cells. The first is to have and adequate minimal number, t_0, of personnel in each cell; the second is to have an adequate number of cells, k_0, remaining so that the shrinkage technique can be effective. The latter number can be as low as five, but at least ten is preferred and there is no objection to much larger values, such as fifty or so.

The inventory threshold, t_0, is treated as inviolate; all aggregated cells will have at least this many personnel. The total number of cells is treated more loosely. That is, the value k_0 is viewed as an approximate or desirable guide, but we are flexible in accepting a lesser number in an operational analysis.

A number of cross validation computations have been made in order to understand the relationship between t_0 and k_0. The results are quite varied. For example, when dealing with LDO and Warrant officers the attrition behavior appears to be quite stable; both t_0 and k_0 can be small and the rates validate well. On the other hand

there are cell collections that are not as stable. Some of the aviation sets in the unrestricted community serve as examples. The captains in the fixed wing pilot group that was chosen to illustrate the small cell methodology is one of the more variable sets as measured by the cross validation technique.

The relationship between the aggregation parameters has been studied to the extent possible during the current period. The results are summarized in Section 5.

The author is concerned about the practice of ignoring MOS when performing attrition analysis for the Steady State Promotion Model. When this is done, one is in effect aggregating all the MOS cells for each YCS level. (Whether or not this is followed by a flat rate or by the use of a shrinkage technique is not material for the moment.) Thus the practice is not necessarily bad. The difficulty is that there may be over aggregation. This issue was one of the very first treated in our researches. See Tucker [5; performance of estimators marked MAXLIKE, p 66ff] and Robinson [6; performance of the AGG class of estimators, p 39ff]. A single rate is being multiplied by a total current inventory value in order to anticipate the number of leavers. The results can be quite poor when applied on a large scale.

To illustrate the issue suppose that a given grade and YCS can be partitioned into a number of MOS cells and the current personnel inventory in those cells is the set of integers $\{N_i\}$. Suppose further that the ORG has produced a corresponding set of attrition rates $\{q_i\}$. A detailed forecast of the total number of leavers is $\sum N_i q_i$. This value generally will be much sharper than $q \sum N_i$ where q is a global attrition rate obtained by over aggregation.

If the overall group is not too large then this particular form of aggregation might work reasonably well. For example, the captains in SLM group 7,3,2 form a small set and it appears to work well for them. But generally the practice should be

discouraged. At best it should be subject to study. It may be useful to develop special aggregation rules for use with the Steady State Promotion Model.

There are some general rules that should be adhered to when aggregating. The first is to hold the grade fixed. The various grades should be aggregated separately. Since aggregation is the single most time consuming operation performed by the rate generator, it would be wise to recognize this in the organization of the end use calculations. Thus overall system performance would be more efficient if the consumer models were organized to function serially by grade.

The user should be aware that the aggregation methods utilize the yearly average personnel inventory. The input selection panel allows the user to influence this process by choice of years and weights. This provides him with an opportunity to utilize any special information that he may have concerning his immediate analysis.

It is generally true that sharper rates are developed if the aggregation scheme can be restricted to the small MOS group containing the targeted set of cells. A gating technique for user consideration of this point can be found in [4], and is reviewed here. Its use also serves to speed up the performance. The first step is to scan the small MOS/ YCS combination for total inventory, NT, and total number of cells, KT; then total number of cells, k_1, whose inventory exceeds t_0 and the total inventory, N_1, in those cells. The aggregation method will accept the k_1 cells as they are and attempt to produce combinations such that cells containing the $NT - N_1$ remaining personnel are aggregated smoothly into larger cells which satisfy the t_0 threshold. An ad hoc estimate of the total number of cells can be computed prior to the actual aggregation. We use

$$k_1 \text{ plus } 90\% \text{ of the ratio } (NT - N_1)/t_0$$

If this value is judged to be sufficiently close to the goal k_0 then the user can

continue with the aggregation. (Of course the final number of cells may be different from the estimate.) Otherwise cells from the large MOS group must be included in the aggregation pool. When this latter choice is made the entire aggregation process begins afresh. The above gating technique can be applied again, etc.

It is in the nature of Empirical Bayes estimation that the collection of rates function well for global use concerning the entire pool of cells. There is a trade off, however, in that generally there will be cells whose individual rates do not serve as well for issues that concern those cells in isolation from the others. It is because of this that we tend to use personnel thresholds and numbers of aggregared cells that are a bit smaller than those suggested by the use of global validation measures. In other words, we tend to err in favor of giving relief to the latter problem. In our studies we have used values of t_0 and k_0 each starting with ten and all combinations of each advancing upward from ten in increments of five until the value 40 is achieved. There does not appear to be any fixed pair of numbers that can be recommended uniformly. On the other hand, the general performance does not change drastically with changes in these values. We are inclined to suggest a strategy: Start with $t_0 = 20$; perform the gating computation that keeps the aggregation in the small SLM group; accept the result if the number of cells is estimated to be at least ten; if this number is estimated to be between five and nine, then lower t_o to 15 and repeat the gating computation. If the test still fails then the user must choose between the act of lowering t_0 to ten with the implied repeat of the gating, and the act of expanding the pool to include the large MOS group. This choice is best made based upon knowledge of the sizes of the pools.

A number of applications of ORG require monthly or fractional year lead time forecasts. The simple use of models based upon the assumptions of independence or

stationarity have not tested successfully. Nor have our attempts to extend the work in [1] which developed some seasonal adjustment techniques of the type that go with the family of exponential smoothing models.

Subsequent discussions with MI and DSAI personnel have revealed a likely source of the difficulty. The time window for the annual meetings of the promotion boards is rather large and the date of their meeting varies from year to year. This impacts upon the attrition dates, especially for junior officers who are twice passed over. Having identified an important cause for the instability, it has been proposed that this information be gathered and included in the development of some Categorical Data Analysis models that include months as a variable. Such models can be superimposed upon our annual rate models. A validation scheme would be developed to measure their efficacy.

5 RESULTS OF PARAMETER STUDIES

Table 5 below contains the results of our study on the effect of aggregation input parameters upon the efficacy of the small cell rate estimation. It was not possible to perform a complete study, but it does extend the one contained in [3]. The first member of the pair is the inventory threshold; all cells must have at least that many personnel. The second is the approximate number of aggregated cells that will result. The use of these numbers assure that the small cell option will function and return a useable answer. The entries are restricted to selected grade and YCS pairs, but the values themselves are rather stable, suggesting that more general use would not be unduly risky.

One may infer some behavioral points. A group containing many personnel but few MOS cells can tolerate a large aggregation threshold. Groups with somewhat fewer

personnel but more MOS cells should be given a lower threshold, but more aggregated cells can be expected. The smaller groups must be given a lower threshold (note that 10 is common in the table) and it may occur that not many aggregated cells result. In a number of these latter cases it is necessary to use the large and perhaps the major expansion groups in order to get a useable result. These are marked with an asterisk. All others did not require expansion beyond the small group. In any case one might apply rules of this type in the selection of threshold values.

Also marked are some cases that exhibited some weakness in the cross validation comparisons. It is especially annoying when these occur in conjunction with full expansion. Small groups 11, 12, and 13 are trainees of sundry types and contain few personnel beyond the grade of lieutenant.

Table 6 contains the grouping information that allows the small cell technology to be extended to the LDO, Warrant Officer, and Colonel communities. Generally these communities behave very well with regard to attrition, and the assignments were made largely in an ad hoc fashion by making skill comparison with the regular officer groupings. Each community has its own YCS interval by grade.

TABLE 5

AGGREGATION IN PARAMETERS (t_o, k_o) FOR SELECTION (grade, YCS) PAIRS IN THE REGULAR OFFICER COMMUNITY

Small MOS Group	(Lt, 2)	(Lt, 4)	(Cpt, 5)	(Cpt, 9)	(Maj, 15)	(LCol, 22)
1.	(20,5)	(20,5)	(20,5)	(15,10)*ψ	(20,5)	(20,5)
2.	(15,10)	(15,10)	(15,5)	(15,5)	(15,5)	10,5)
3.	(10,10)*	(10,5)*	(10,5)*	(10,5)	(10,5)	(10,5)
4.	(15,10)*	(10,10)*	(10,10)*	(10,10)	(10,10)	(10,10)*ψ
5.	(15,10)	(15,5)	(15,5)	(15,5)	(15,5)	(10,5)*
6.	(10,5)	(10,5)*	(10,5)	(10,5)	(10,5)	(10,5*ψ
7.	(10,10)*	(10,10)*	(10,10)*	(10,10)ψ	(10,10)	(10,10)*
8.	(10,10)*	(10,10)*	(10,10)*ψ	(10,10)*ψ	(10,10)*	(10,10)*
9.	(10,10)*	(10,10)*	(10,10)	(10,10)	(10,10)	(10,10)*
10.	(10,10)*	(10,10)*	(10,10)	(10,10)	(10,10)	(10,10)*
11.	(15,10)	(15,10)	(15,5)*	(15,5)*	—	—
12.	(10,10)	(10,10)	(10,10)*	(10,10)*	—	—
13.	—	—	—	—	—	—
14.	—	—	(10,5)*ψ	(10,5)*ψ	(10,5)*ψ	—

* Aggregation requires inclusion of the large MOS group and possibly the Major group as well.

ψ Validation results are weak.

TABLE 6

SLM GROUPINGS FOR THE LDO COMMUNITY

SLM	MOS Cells
2,1,1	1310
3,2,1	0160, 0170, 0205, 0210, 2602, 2802, 4430, 5803, 5804, 5805
4,2,1	3402, 3406, 3410, 4006, 4010, 4130, 4302, 4602, 5502, 5905
5,2,1	0430, 2105, 2110, 2120, 2301, 2305, 2340, 3010, 3050, 3070, 3102, 3302, 3510, 6004, 6302, 6502
6,2,1	6802, 7204, 7208, 7210, 7320

SLM GROUPINGS FOR THE WARRANT OFFICERS

SLM	MOS Cells
1,1,1	0306
2,1,1	0803, 1310, 1360, 1390
3,2,1	0160, 0170, 0205, 0210, 0260, 2503, 2602, 2805, 2810, 2830, 5803, 5804, 5805, 9925, 9960
4,2,1	1120, 1502, 3402, 3406, 3410, 4006, 4010, 4016, 4130, 4302, 4602, 5502, 5505, 5702, 5907, 5910, 5950, 5970
5,2,1	0430, 2120, 2125, 2305, 2340, 3010, 3020, 3102, 3302, 3510, 6004, 6302, 6502, 9002, 9803, 9805
6,2,1	6802, 7002, 7204, 7208, 7210, 7320, 7380

SLM GROUPINGS FOR COLONELS

SLM	MOS Cells
1,1,1	9904, 9906
6,2,1	9907
14,6,4	9914

TABLE 7

ADDITIONAL YCS INTERVALS BY COMMUNITY AND GRADE

Community	Grade	Interval
LDO	Lieutenant	$\leq 2, 3, \cdots, 9, \geq 10$
	Captain	$\leq 3, 4, \cdots, 12, \geq 13$
	Major	$\leq 6, 7, \cdots, 17, \geq 18$
Warrant Officers	W1	$1, 2, 3, \geq 4$
	W2	$\leq 2, 3, \cdots, 7, \geq 8$
	W3	$\leq 6, 7, \cdots, 11, \geq 12$
	W4	$\leq 10, 11, \cdots, 29, \geq 30$
Colonel	Colonel	$\leq 21, 22, \cdots, 29, \geq 30$

6 REFERENCES

[1] C. J. Mehalic, "Multiparameter Forecasting Techniques for the Marine Corps Officer Rate Generator," NPS MSc Thesis, September 1990.

[2] R. W. Larsen, "The Aggregation of Population Groups to Improve the Predictability of Marine Corps Officer Attrition Estimation," NPS MSc Thesis, December 1987.

[3] J. M. Misiewicz, "Extension of Aggregation and Shrinkage Techniques Used in the Estimation of Marine Corps Officer Attrition Rates," NPS MSc Thesis, September 1989.

[4] R. R. Read, "The Use of Shrinkage Techniques in the Estimation of Attrition Rates for Large Scale Manpower Models," NPS Technical Report, NPS55-88-006, July 1988.

[5] D. D. Tucker, "Loss Rate Estimation in Marine Corps Officer Manpower Models," NPS MSc Thesis, September 1985.

[6] J. R. Robinson, "Limited Translation Shrinkage Estimation of Loss Rates in Marine Corps Manpower Models," NPS MSc Thesis, March 1986.

[7] S. M. Stigler, "Neyman Lecture: A Galtonian Perspective on Shrinkage Estimates," *Statistical Science*, **5**, February 1990.

[8] A. C. Brandwein and W. E. Strawderman, "Stein Estimation: The Spherically Symmetrical Case," *Statistical Science*, **5**, August 1990.

[9] R. R. Read, Interim Report for MIS HQUSMC on project described by Document Number M00027-92-WRA8054, Department of Operations Research, NPS, 5 March 1992.

[10] Decision Systems Associates, Inc., Officer Rate Generator User's Manual, Version 1.0, prepared for code MI, HQUSMC, 18 November 1991.

DISTRIBUTION LIST

1.	Dudley Knox Library Naval Postgraduate School Monterey, CA 93943	2
2.	Research Office Naval Postgraduate School Monterey, CA 93943	1
3.	Defense Technical Information Center Cameron Station Alexandria, VA 22304	2
4.	Department of Operations Research Naval Postgraduate School Monterey, CA 93943-5000	1
5.	Major T. Stein Code MI HQ USMC Washington, D.C. 20380-0001	3
6.	Mr. Rod Morton Decision Systems Associate, Inc. 350 Fortune Terrace Rockville, MD 20854	3
7.	Management Applications, Inc. Dr. R. Butterworth 1530 N. Key Blvd., Suite 531 Arlington, VA 22209	1
8.	Prof. R. R. Read Code OR/Re Naval Postgraduate School Monterey, CA 93943-5000	5