2022-11-16

USMC Manpower Models Modernization

Seagren, Chad W.; Reich, Daniel; Bacolod, Marigee; Das, Arijit

Naval Postgraduate School

http://hdl.handle.net/10945/71373

NPS-DDM-22-006

NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

USMC MANPOWER MODELS MODERNIZATION

by
Chad W. Seagren, Daniel Reich, Marigee Bacolod, Arijit Das

16 November 2022

Prepared for: **Headquarters Marine Corps, Manpower & Reserve Affairs**
This research is supported by funding from the Naval Postgraduate School, Naval Research Program (PE 0605853N/2098). NRP Project ID: NPS-22-M205-A

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED START DATE	END DATE
11/16/2022	Technical Report	20211024	20221022

3. TITLE AND SUBTITLE
USMC Manpower Models Modernization

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
		0605853N/2098
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**
NPS-22-M205-A		

6. AUTHOR(S)
Chad W. Seagren, Daniel Reich, Marigee Bacolod, Arijit Das

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Naval Postgraduate School
1 University Circle Monterey, CA 93943

8. PERFORMING ORGANIZATION REPORT NUMBER
NPS-DDM-22-006

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR'S ACRONYM(S)	11. SPONSOR'S REPORT NUMBER(S)
Naval Postgraduate School, Naval Research Program, Headquarters Marine Corps, Manpower & Reserve Affairs	NRP, M&RA	NPS-DDM-22-006; NPS-22-M205-A

12. DISTRIBUTION/AVAILABILITY STATEMENT
Distribution Statement A: Approved for public release. Distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
In his planning guidance, the Commandant states “We will ... accelerate our transformation from disconnected legacy systems to an integrated data architecture that treats data as it should be –a critical resource.” This project is a bold step in support of that transformation for Manpower \& Reserve Affairs (M&RA), which possesses numerous mathematical models that support the management of the human resources development process (HRDP). Our approach demonstrates how M&RA can leverage the latest technology in data architecture and decision support models to overcome these deficiencies. We formulate modernized versions of models to fill the roles currently held by the Total Force Planning Model (TFPM) and the Enlisted End-Strength Planning Model (ESPM). In addition, we formulate a discrete event simulation model to analyze changes in structure that currently has no legacy counterpart. We successfully replicate the TFPM in Python and identify a number of improvements that should be made with respect to data formatting. An inability to obtain sufficiently accurate separation data prevents progress towards replication of the ESPM. Finally, we successfully implement a discrete event simulation of the manpower system. While initial validation attempts appear favorable, a more complete validation is prevented by the same inability to obtain appropriate data.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	C. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
U	U	U	UU	57

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
Chad W. Seagren	*(831) 656-7902*

PREVIOUS EDITION IS OBSOLETE

STANDARD FORM 298 (REV. 5/2020)
Prescribed by ANSI Std. Z39.18

THIS PAGE INTENTIONALLY LEFT BLANK

NAVAL POSTGRADUATE SCHOOL
Monterey, California 93943-5000

Ann E. Rondeau President	Scott Gartner Provost

The report entitled "USMC Manpower Models Modernization" was prepared for HQMC Manpower and Reserve Affairs (M&RA) and funded by the Naval Postgraduate School, Naval Research Program (PE 0605853N/2098).

Distribution Statement A: Further distribution of all or part of this report is authorized.

This report was prepared by:

Chad W. Seagren Senior Lecturer	Daniel Reich Assistant Professor
Marigee Bacolod Associate Professor	Arijit Das Faculty Associate – Research

Reviewed by:	**Released by:**
Raymond Jones, Chairman Dept. of Defense Management	Kevin B. Smith Vice Provost for Research

THIS PAGE INTENTIONALLY LEFT BLANK

Abstract

In his planning guidance, the Commandant states "We will ... accelerate our transformation from disconnected legacy systems to an integrated data architecture that treats data as it should be –a critical resource" (Berger 2020a, p. 15). This project is a bold step in support of that transformation for Manpower & Reserve Affairs (M&RA), which possesses numerous mathematical models that support the management of the human resources development process (HRDP). Our approach demonstrates how M&RA can leverage the latest technology in data architecture and decision support models to overcome these deficiencies. We formulate modernized versions of models to fill the roles currently held by the Total Force Planning Model (TFPM) and the Enlisted End-Strength Planning Model (ESPM). In addition, we formulate a discrete event simulation model to analyze changes in structure that currently has no legacy counterpart. We successfully replicate the TFPM in Python and identify a number of improvements that should be made with respect to data formatting. An inability to obtain sufficiently accurate separation data prevents progress towards replication of the ESPM. Finally, we successfully implement a discrete event simulation of the manpower system. While initial validation attempts appear favorable, a more complete validation is prevented by the same inability to obtain appropriate data.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 1: Marine Corps Manpower Models

In his planning guidance, the Commandant states "We will . . . accelerate our transformation from disconnected legacy systems to an integrated data architecture that treats data as it should be – a critical resource." This project is a bold step in support of that transformation for Manpower & Reserve Affairs (M&RA).

M&RA possesses numerous mathematical models that support the management of the Human Resources Development Process (HRDP). In general, these models suffer from the following:

- Models tend to be implemented in an outdated computational language.
- Models tend to require substantial human intervention to manage the interface between models, as when the product of one model is a required input for another model in the process.
- Models tend to require substantial human intervention to manage the interface between the model and sources of empirical data such as Total Force Data Warehouse (TFDW).
- Models tend to lack standardization with respect to elements common among models. For example, several models might require an estimate of the attrition rate for, say, Marines of a particular rank and occupational specialty, yet each model might contain a differently calculated rate.

In addition, there are steps in the HRDP that lack any sort of decision support tool, but could benefit from the addition of such a tool, provided the tool is sufficiently user friendly and sufficiently interoperable with other models and data sources.

Our approach demonstrates how M&RA can leverage the latest technology in data architecture and decision support models to overcome these deficiencies. We will formulate modernized versions of models to fill the roles currently held by the Total Force Planning Model (TFPM) and the Enlisted End-Strength Planning Model (ESPM). In addition, we will formulate a model to analyze changes in structure that currently has no legacy counterpart. Finally, we will thoroughly plan and describe the data architecture in which they reside, in order to demonstrate how to effectively address (a) through (d).

1.1 Study objectives

The purpose of this project is to answer the following questions:

1. To what extent may we integrate a group manpower models into a family of interconnected modern models that share a data architecture scheme in order to improve planning and the human-machine teaming with analysts?
2. To what extent may the answer to (1) provide a template to modernize the rest of the system?

Our approach executes the following lines of effort in parallel:

a. Formulate a Mathematical Program that fills the role in the HRDP that the legacy TFPM currently fills. Given the Authorized Strength Report and various budgetary and policy constraints as inputs, this model will produce an ideal inventory by Grade and MOS across the FYDP that is akin to the legacy Grade Adjusted Recapitulation (GAR). We will provide analysis as to the feasibility of implementing such a model in Python as well as the benefits such a model could bestow. In addition, we will consider ways to make the model more capable and more user-friendly than the legacy TFPM.

b. Formulate a predictive model that fills the role in the HRDP that the legacy Enlisted End-Strength Planning Model currently fills. Given historical attrition behavior and the planned accession mission, this model will provide a within-execution-year forecast to determine the likelihood that the system will meet the target enlisted end-strength at the end of the Fiscal Year. Key expected improvements relative to the legacy model include more automated interaction with TFDW, greater model transparency, and enhanced managerial relevance of output. We will provide an analysis of the feasibility of implementing such a model in R, as well as the benefits such a model could bestow.

c. Formulate a simulation model to assist in the analysis of changes in the structure of a given Occupational Field (OccFld). Given a notional GAR for that OccFld; the current inventory for that OccFld; historic attrition behavior; and the relevant Enlisted Career Force Controls; this model will produce an estimate of the resultant promotion timing for each grade, as well as an estimate of the feasibility of obtaining the desired inventory levels. We provide an analysis of the feasibility of implementing such a model in Java, as well as the benefits such a model could bestow.

d. Develop the plan for a data architecture that allows these models to seamlessly integrate with each other and take advantage of automated interaction with input data sources. We will examine the data required for each of the models and design the respective database schema, to include the various tables, integrity constraints, and loader tools. At the initial stage, we will describe a 2-tier architecture structure, though in the future we expect a 3-tier architecture with a web enabled interface.

1.2 Findings

We successfully implement a replication of the TFPM for officers in Python using a parallel processing approach. We process all 257 combinations of MOS and Officer Type with FY22 data and compare those against the official TFPM outputs and produce identical results in 244 out of the 257 cases. The other 13 instances contain discrepancies that were reviewed by a domain expert. Some discrepancies were caused by manual adjustments and others were categorized as errors. We provide detailed recommendations for improvements in four key areas. We identify significant disadvantages with the data formats we received, including inconsistent field naming conventions from one data table to another, and implicit representations that are understandable to a human analyst but insufficient for machine processing. These data formatting deficiencies increase the complexity of logic and code required for generating the TFPM.

We assess the legacy ESPM and formulate a predictive machine-learning-based ESPM. We had hoped to estimate this model to illustrate its feasibility in replacing the legacy model. Unfortunately, we encountered extreme challenges collecting the necessary data. As of the writing of this report, we are still awaiting key data elements in order to estimate this model. The model and results will be part of the FY23 study, beginning with Major Aaron Falk's March 2023 thesis.

We successfully formulate and implement a discrete event simulation model of the manpower system in Java. Given a notional target inventory for an OccFld; the current inventory for that OccFld; historic attrition behavior; and the relevant Enlisted Career Force Controls; the model provides estimates of expected promotion timing for each grade, expected accession mission, and expected retention requirements. The model easily employs data-farming techniques and lends itself to both transient and steady-state analyses. While initial validation attempts the model appears to perform well, this aspect of the project was plagued with the some of the same issues at LOE 2, and much less time was available for validation. Therefore, we recommend to continue the validation efforts into the future, both with respect to additional target years and additional communities and PMOSs. Researchers should confirm the existence of proper TTP data as well as empirical retention data. It is likely that minor refinements to the manner in which Marines in the above-zone for promotion are managed, as well as minor refinements to how reenlisted Marines are distributed across grades will prove worthwhile and improve model performance.

Ultimately, the most important finding of this project is that the records contained in TFDW might be insufficient to support building a rigorous mathematical model of attrition behavior. The Marine Corps will fail in its efforts to transform from an industrial age organization to an information age organization if the data it relies on to make the most elementary manpower management decisions, i.e. the effort to make end-strength, is flawed or non-existent.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 2:
Line of Effort One: Total Force Planning Model

2.1 Introduction

The Total Force Planning Model (TFPM) is used to generate the Grade Adjusted Recapitulation (GAR), which is the desired future personnel of the Marines by grade (rank) and Military Occupational Specialty (MOS). Manpower Plans and Policy (MP) Division produces the GAR twice annually, the same cadence as its main input, the Authorized Strength Report (ASR).

TFPM consists of a set of inputs and rules. We concentrate our testing on officer grades. An example format of the desired output is shown in Table 2.1.

	O7	O6	O5	O4	O3	O2/O1	Total
ASR							
TOTAL A-BILLETS							
TOTAL B-BILLETS							
TOTAL A+B BILLETS							
NAR BILLETS							
T2P2							
NAR							
GAR							

Table 2.1. GAR build report output format for officer ranks, by MOS and category.

The "TOTAL A+B BILLETS" line is simply the sum of the "TOTAL A-BILLETS" and "TOTAL B-BILLETS" lines. Similarly, the "NAR" line is the sum of the "NAR BILLETS" and "T2P2" lines.

To produce such outputs, numerous inputs are required, including:

- Authorized Strength Report (ASR)
- ASR modifications
- Military Occupational Specialty (MOS)
- Prisoners, Patients, Transients and Trainees (P2T2)

2.2 Querying the ASR Database

The ASR database from 2021 contains 53 fields and over 2 million rows. A subset of the fields used is shown in Table 2.2.

FY	N_GRADE	BMOS	PMOS	ASR	BRANCH	MPR_TYPE	BIL_STA	MCC
2022	O2	0952	7208	1	M	A	A	KES
2022	O3	0952	4402	1	M	A	A	K26

Table 2.2. Subset of fields in ASR database.

To compute the "ASR" line in the GAR build report shown in Table 2.1, the MOS in the "BMOS" field is used. When a record's primary MOS in field "PMOS" matches the billet MOS in field "BMOS", that record is counted as an A-billet; otherwise, it is counted as a B-billet. T2P2 entries are marked with the code "000" in the "MCC" field and are not included in the "ASR", "TOTAL A-BILLETS" or "TOTAL B-BILLETS" counts in Table 2.1.

Several filters for the ASR database, illustrated in Table 2.2, are required to determine the counts in the GAR Build Report. The "FY" field can be used to focus on a particular fiscal year. The "ASR" field may contain blank entries, but only entries with a value of "1" are included. Only entries in the "BRANCH" field with a value of "M" for Marines are included. Only entries with an "MPR_TYPE" of "A" are included. The billet status "BIL_STA" can be either active "A" or reserve "R". The "N_GRADE" field specifies rank.

To produce an officer report, one could query all officer grades. However, only a subset of those applies to a particular combination of MOS and officer type. This subset can be found in the MOS input shown in the template in Table 2.3.

MOS	Off Type	Low Grade	High Grade	Primary/Non-Primary	Skill Family
0952	1	O2	O5	N	ONO08
0952	2	O3	O5	N	ONO08

Table 2.3. Template for MOS metadata.

Performing a lookup of the entry in Table 2.2 against the officer types in Table 2.3 bins the first entry against category 1 officer type, because the grade of O2 is not contained in the grade range for category 2. However, the second entry in Table 2.2 has a grade of O3, which is contained in both officer type grade ranges, so this record may be double-counted when summarizing MOS and category combinations in the GAR report in Table 2.1.

Pseudocode for querying the ASR database is as follows:

```
SELECT N_GRADE, COUNT(N_GRADE)
FROM {TABLE_NAME} WHERE
FY = {YEAR}
N_GRADE IN {GRADE_LIST}
AND PMOS IN {MOS_LIST}
AND ASR = 1
AND BRANCH = M
AND MPR_TYPE = A
AND BIL_STA IN (A,R)
AND MCC = 000 {if T2P2}
AND MCC <> 000 {if not T2P2}
AND BMOS = {MOS}
AND PMOS = {MOS} {if A_BILLETS}
AND PMOS <> {MOS} {if B_BILLETS}
GROUP BY N_GRADE
```

where curly brackets denote parameters that must be set in the query or logic conditions that must be applied.

2.3 Amendments to the ASR

The rules outlined in section 2.2 produce initial counts for the "ASR", "TOTAL A-BILLETS", "TOTAL B-BILLETS", "TOTAL A+B BILLETS" and "T2P2" lines in the GAR Build Report shown in Table 2.1. However, these counts are then modified using manual inputs provided by an assigned analyst.

2.3.1 ASR Line Item Modifications

ASR modifications are noted as a series of line items that include values for "MOS", "Officer Type", "Grade" (equivalent to the "N_GRADE" field in the ASR) and whether the modification is an addition (denoted by "+") or subtraction (denoted by "-"). The template used for ASR modifications is shown in Table 2.4.

Grade type	MOS	Mod type	Force	Location	O7/W5/E9	O6/W4/E8	O5/W3/E7	O4/W2/E6	O3/W1/E5	O2/O1/E3	E3	E2/E1
1	8006	-	Non	C				1	1			
1	430	+	Non	C				1				

Table 2.4. Template for ASR modifications.

2.3.2 TOTAL A-BILLETS Line Item Modifications

In cases where the "ASR" and "TOTAL A-BILLETS" line counts are equal, we apply the modification described in section 2.3.1 to the "TOTAL A-BILLETS" line as well, but clarification is required on the correctness of this approach.

2.3.3 TOTAL B-BILLETS Line Item Modifications

"TOTAL B-BILLETS" modifications are noted in a series of line items that include values for "BMOS", "Officer Type", "Low Grade", "High Grade", "Alloc Type" ("N" for number, "P" for percentage or "All"), "Amount", and a list of "MOS Spec" and corresponding "Share". The template used for B-billet modifications is shown in Table 2.5.

Officer Type	BMOS	Low Grade	High Grade	Alloc Type	Amount	Conus	O/S	Dpld	Force	User Group	MOS Spec 1	Share 1	MOS Spec 2	Share 2	...
1	952	O4	O4	N	13	Yes	Yes	Yes	All	No	0102	1	0202	1	...
1	953	O2	O5	N	125	Yes	Yes	Yes	All	Yes	GINS3	125			

Table 2.5. Template for B-billet modifications.

When the "Group" field value is set to "Yes", the template shown in Table 2.6 is required to lookup the group details.

Officer Type	Group Name	MOS	Primary Shares
1	GINS3	7509	12
1	GINS3	7518	24
1	GINS3	7523	16
1	GINS3	7532	34
1	GINS3	7557	14

Table 2.6. Template for MOS group shares, where the "Primary Shares" for a given "Group Name" are percentages that must sum to 100.

2.3.4 T2P2 Line Item Modifications

"T2P2" modifications are handled by four templates, one for each of the subcategories "training", "transient", "patients" and "prisoners".

Training Shares

The training modifications are noted as a series of line items, where each line item has associated values for "MOS", "Officer Type" and "Grade". Records with "Training Type" equal to "E" are added to the specified "MOS" and subtracted from one or more associated "Intended MOS" and "Intended Amount" pairs provided. The template used for training modifications is shown in Table 2.7.

Officer type	MOS	Training Type	Grade	Amount	Intended MOS 1	Intended amount 1	Intended MOS 2	Intended amount 2	...
1	7201	E	O2	102	7204	18	7208	37	...

Table 2.7. Template for training share modifications.

Transient, Patients, Prisoners
The transient, patient and prisoner modifications are noted as a series of line items, where each line item has associated values for "MOS", "Officer Type" and grade. The template used for each of these three subcategories' modifications is shown in Table 2.8.

Officer Type	MOS	O2 & O1	O3 W1	O4 W2	O5 W3	O6 W4	O7 W5
1	0402	1	1	1	0	0	0

Table 2.8. Template for transient, patient and prisoner modifications.

2.4 Computational Experiments and Results

We implement the rules described above in a Python program with parallel processing. We process all 257 combinations of MOS and Officer Type with FY22 data and compare them against the official TFPM outputs. With respect to the "ASR", "TOTAL A-BILLETS" and "T2P2" line items, we produce identical results in 244 out of the 257 cases. The other 13 instances contained discrepancies that were reviewed by a domain expert. Some were caused by manual adjustments and others were categorized as errors. Parallel processing with 8 threads on a Macbook Pro reduced the run time from 253 seconds down to 87 seconds, an efficiency gain of approximately 3x.

2.4.1 Line Items Not Computed

We are unable to compute the "NAR BILLETS" and "GAR" line items, because both rely on proprietary models to which we did not have access. Additional templates required include "exclusions", "free billets shares" and "NAR GAR Limits". Without the "NAR BILLETS" line item, we were also able to compute the "NAR" line item, which is simply the sum of the "NAR BILLETS" and "T2P2" line items.

While we were able to compute initial "TOTAL B-BILLETS" counts from the ASR database, we were unable to complete the line item modifications. We discovered at a late stage of the project that we were missing the template for redistributing B-billet shares. Upon receiving the required data, we found that the data representation was incompatible with the structure of our code base and would require a significant overhaul to process the percentage-based allocation type option. We discuss this further in Section 2.5.1.

2.5 Recommendations

Our efforts in implementing routines for producing the desired TFPM outputs brought into focus aspects about the data inputs that are not currently handled in an ideal format. We provide detail and recommendations for future development below.

2.5.1 Count vs. Percentage Based Representations

Table 2.5 contains a field "Alloc Type" that takes values of "N" for number, "P" for percentage or "All". This data representation is incomplete for the latter two options of "P" and "All", and requires referencing external tables. While these options may be convenient for a human analyst, machine processing requires a reference to the actual counts. These counts should be retrievable within the database and preprocessed to appear explicitly within Table 2.5.

2.5.2 Inconsistent Field Name Conventions

Field names are inconsistent between various inputs. From a processing standpoint, mapping the various naming conventions for equivalent fields of data adds unnecessary complexity. For example, "Off Type" is used in Table 2.3, "Grade type" is used in Table 2.4, and different casings of "Officer Type" are used in Tables 2.5 and 2.8. A single naming convention should be selected and used in all tables.

Another example of inconsistent field names is in the treatment of grade. In the ASR database shown in Table 2.2, "N_GRADE" is used to denote rank, whereas the MOS metadata in Table 2.3 contains the fields "Low Grade" and "High Grade". However, in this example, updating naming conventions would require modifying data representations, which we discuss next.

2.5.3 Implicit vs. Explicit Data Representations

Data file formats we processed contained data representations that are convenient for human-readability, but are incomplete for machine processing. For example, one can read the "Low Grade" and "High Grade" fields in the MOS metadata in Table 2.3 and infer that a range of grades is being specified, where the range includes both endpoints listed explicitly and all grades between than are not listed but are implicit. Machine processing of such data requires rules be established to interpret all grades contained between two endpoints. A more structured approach to data handling could remove the need for these rules, by explicitly representing all grades of data. For example, Table 2.3 could be rewritten in the format of Table 2.9, where all grades are represented explicitly.

MOS	Off Type	N_GRADE	Primary/Non-Primary	Skill Family
0952	1	O2	N	ONO08
0952	1	O3	N	ONO08
0952	1	O4	N	ONO08
0952	1	O5	N	ONO08
0952	2	O3	N	ONO08
0952	2	O4	N	ONO08
0952	2	O5	N	ONO08

Table 2.9. Template for reformatting the MOS metadata from Table 2.3 with explicit representation of all grades.

Another benefit to the representation in Table 2.9 is that a single field name is used for grade, allowing it to be consistent with the naming convention of the ASR database shown in Table 2.2. Specifically, both now contain the field "N_GRADE", rather than the inconsistent field names of "Low Grade" and "High Grade", which required interpretation.

A third benefit to the representation in Table 2.9 is that it allows for the inclusion of non-contiguous grade sets, for example, including O3, O4 and O6, while excluding O5. The latter example would be impossible to implement in the original data representation in Table 2.3.

The downside of the representation in Table 2.9 is that it contains redundant data, which is both wasteful in terms of storage space and inconvenient in terms of human readability. Each of these downsides can be resolved. To improve the storage efficiency, a relational database structure can be implemented, by splitting the data in Table 2.9 into the two tables shown in Table 2.10.

An SQL query can then be used to rejoin the two tables in Table 2.10 and reconstruct the data in 2.9 using a "JOIN" command on the newly introduced "MO_ID" field. Viewing the original format in 2.3 is also achievable in a number of ways, if desirable for human readability. One possible approach would be to construct a function that takes a set of "N_GRADE" values and summarizes it. This function could be used on any data table in which "N_GRADE" values are present, another advantage of ensuring data representations are consistent from one table to another.

MO_ID	MOS	Off Type	Primary/ Non-Primary	Skill Family
09521	0952	1	N	ONO08
09522	0952	1	N	ONO08

MO_ID	N_GRADE
09521	O2
09521	O3
09521	O4
09521	O5
09522	O3
09522	O4
09522	O5

Table 2.10. Template for reformatting the MOS metadata from Table 2.3 with explicit representation of all grades efficiently using two relational tables.

While we focus on the MOS metadata in Table 2.3 in this section, our recommendations can be applied directly to the similar treatment of grades for B-billets in Table 2.5.

2.5.4 Field vs. Record Data Representations

Several of the line item modification templates contain multiple fields instead of additional records with a single field value.

Grade Representations in Fields vs. Records

Table 2.4 contains fields for "O7/W5/E9", "O6/W4/E8", etc. These compound field names are understandable when read by an analyst familiar with the data, who is aware that the meaning is the field will contain either officer rank, warrant officer rank or enlisted rank data, but not all three simultaneously. Similar naming conventions appear in Table 2.8. We will proceed by focusing on an improved representation of Table 2.4, but the manner of restructuring can be directly applied to Table 2.8 as well.

A more explicit representation on Table 2.4 could be achieved by restructuring with a single field name "N_GRADE", as shown in the template provided in Table 2.11

Grade type	MOS	Force	Location	N_GRADE	N_GRADE_Quantity
1	8006	Non	C	O4	-1
1	8006	Non	C	O3	-1
1	430	Non	C	O4	1

Table 2.11. Template for reformatting the ASR modifications from Table 2.4 with a standardized field name for grade.

Another advantage of the representation in Table 2.11 is that it removes the need for the "Mod type" column in Table 2.4, by incorporating this information in the newly introduced "N_GRADE_Quantity" column.

Notice, the representation in Table 2.11 again allows for using the "N_GRADE" field in a manner consistent with Tables 2.2, 2.9 and 2.10.

MOS Representations in Fields vs. Records

Table 2.7 contains fields for "Amount", "Intended MOS 1", "Intended amount 1", "Intended MOS 2", "Intended amount 2", etc. Processing data in this format requires logic that first identifies all field names differening only by a number as equivalent. For example, "Intended MOS 1" and "Intended MOS 2" represent the same type of information. The only difference between them is an arbitrary ordering. The next logical rule that is required for processing is establishing pairings between fields, for example, "Intended MOS 1" and "Intended amount 1" are connected information. This may be easily understood by an individual analyst viewing the data, but is not explicit for machine processing. An alternative representation is provided in Table 2.12.

Officer type	Source MOS	Source Amount	Training Type	Grade	Target MOS	Intended Amount
1	7201	102	E	O2	7204	18
1	7201	102	E	O2	7208	37
1	7201	102	E	O2	...	...

Table 2.12. Template for reformatting training share modifications from Table 2.7 with a single field for the target MOS and a corresponding field for the intended reallocation amount.

The representation is Table 2.12 contains redundant information, which is in inefficient from a storage perspective. It also contains inconsistent naming conventions for "MOS" fields, in order to differentiate between the source and target for the line item modifications. Both these issues can be addressed by restructuring the information in two relational tables, as shown in Table 2.13.

The left table in Table 2.13 represents only information regarding the source MOS and the right table represents only information regarding the target MOS, so standardized field names of "MOS" and "Amount" are sufficient for labeling the data without any modifiers or numbering. An SQL query can be used to rejoin the two tables and reconstruct the data in 2.12, using a "JOIN" command on the newly introduced "MO_ID" field. Viewing the original format in 2.7 is also achievable, if desirable for human readability.

MO_ID	Officer type	MOS	Amount	Training Type	Grade
72011	1	7201	102	E	O2

MO_ID	MOS	Amount
72011	7204	18
72011	7208	37
72011	...	...

Table 2.13. Template for efficiently reformatting the training share modifications from Table 2.7 with standardized fields for MOS, using two relational tables.

While we focused on the training data in Table 2.7 in this section, our recommendations can be applied directly to the "MOS Spec" and corresponding "Share" fields for B-billets in Table 2.5.

2.6 Discussion

We partially summarize the process for generating the TFPM in this chapter. However, identifying a complete and exact procedure is less straightforward, due to both an inability to access the official TFPM code and a lack of documentation. In lieu of these, we rely on a reverse-engineering approach to implement our Python program. Though a series of trial-and-error runs, we compare our outputs against an official set of outputs that we received from the project sponsor, identifying mismatches in results and updating our processing routines until such discrepancies were either resolved, determined to have been caused by manual adjustments, were a product of undisclosed modeling assumptions, or in the case of B-billet adjustments, introduced unanticipated complexity that was not resolvable at a late stage of the project. Our description of the TFPM model is based on our findings, rather than on any direct knowledge of the official model itself.

We identify significant disadvantages with the data formats we received, including inconsistent field naming conventions from one data table to another, and implicit representations that are understandable to a human analyst but insufficient for machine processing. These data formatting deficiencies increase the complexity of logic and code required for generating the TFPM. We provide detailed recommendations for reformatting the data sources in a manner that is explicit, consistent and efficient.

CHAPTER 3: Line of Effort Two: Enlisted End-Strength Planning Model

3.1 Introduction

Our second line of effort formulates a predictive model that fills the role in the Human Resource Development Process (HRDP) that the legacy Enlisted End-Strength Planning Model (ESPM) currently fills.

Our best understanding is that the end-strength planner currently uses an aggregate 3-year trailing average to predict monthly non-end-of-active-service (NEAS) attrition of enlisted Marines by Fiscal Year (FY). NEAS attrition includes losses due to recruit failures at bootcamp, retirements, and a variety of other reasons. NEAS prediction models are currently being implemented using basic spreadsheet modeling in Microsoft Excel.

Our line of effort improves upon this trailing average forecasting model by developing a machine-learning model that predicts first-term NEAS attrition using individual Marine data contained in existing USMC personnel databases.

Data provided to us indicate the spreadsheet model using aggregate trailing averages to forecast within-execution-year attrition has a calculated error rate of approximately 4 to 5% for NEAS overall attrition and even higher prediction error for NEAS-Base. NEAS-Base includes retirements, bootcamp attriters, and a few other reasons. For example, using FY18-FY20 average to predict the same month in FY21 results in a 16% over-estimation of NEAS-Base.

Estimating a more precise forecast of NEAS attrition is important for downstream planners and commands who need this information for fiscal planning, recruiting goals, and other tasks. Our model will provide a within-execution-year forecast of NEAS attrition, which together with the planned accession mission, will enable planners to determine the likelihood that the system will meet the target enlisted end-strength at the end of the FY.

Unfortunately, we still have not received all of the data required to estimate a modernized ESPM model. The rest of this Chapter outlines the model we have developed and describes its estimation and validation we intend to conduct in FY23.

3.2 Proposed ESPM Model

We will develop and estimate logistic regression models of the following form for each month t:

$$P(y_{it} = 1) = \frac{e^{X_{it}\beta_t}}{1 + e^{X_{it}\beta_t}} \tag{3.1}$$

where the outcome variable $y_{it} = 1$ if enlisted Marine i NEAS attrites in month t and is 0 otherwise, and the vector X_{it} captures all relevant personnel information of i as of month t.

To optimally select which variables in the vector X to include as predictors of NEAS attrition y, we employ the machine learning technique of Least Absolute Shrinkage and Selection Operator (LASSO). We will also explore alternative algorithms such as random forest and alternative specifications of the outcome y to predict NEAS attrition for administrative reasons, misconduct, physical conditions, and others.

3.3 Model Validation

As with machine learning algorithms, we will validate how well our models predict the outcome by splitting the data into training and validation samples. To consider how well our predictive model performs relative to the legacy ESPM, we build a model in the same spirit of the trailing average ESPM. More precisely, we will include in the training sample data from the same month in the last three FY's to determine which variables have the most predictive power and use these estimates to build a predictive model of attrition. We will then validate the predictive model on the validation sample, using data in the same month in the current FY, to determine how well the model performs at correctly predicting NEAS attrition in the current FY-month.

As an example, to predict NEAS attrition in July 2018, we estimate equation 3.1 using data from July 2015, July 2016, and July 2017, and use LASSO logit to optimally select predictors. We then take estimates of the parameters (i.e. the vector of $\hat{\beta}_t$'s in equation 3.1 for selected predictors), apply them to data of Marines i in July 2018, and calculate their predicted probabilities of attrition. We then compare this predicted attrition rate against the actual attrition of Marines in July 2018 to assess how well the model performs at correctly predicting NEAS attrition.

We plan to calculate the sensitivity (fraction of true attriters correctly predicted) and specificity (fraction of true nonattriters correctly predicted) of our predictive model. Suppose we have estimated equation 3.1 using LASSO logit and past 3 FY data in some month t. We then calculate the predicted probabilities of attrition in the same month in the current FY using data of Marines i who are in the force in that month:

$$p_i = \frac{e^{X_{it}\hat{\beta}_t}}{1 + e^{X_{it}\hat{\beta}_t}} \tag{3.2}$$

To compare this predicted attrition probability p (which range from 0 to 1) with the observed binary outcome y_i, the probability p_i must be converted to a value of either 0 or 1. Suppose average attrition in month t across all FYs is $\bar{y}$. When $p_i < \bar{y}$, the individual is predicted to not attrite, so we set $\hat{p}_i = 0$. In contrast, when $p_i \geq \bar{y}$, the individual is predicted to attrite, and we set $\hat{p}_i = 1$. Comparing the predicted ($\hat{p}_i$) and actual (y_i) values allows us to classify every Marine i in that month into one of 4 mutually exclusive categories based on their values of $\hat{p}_i$ and y_i.

When the actual and predicted values of $\hat{p}_i$ and y_i both equal to 1, we classify this individual as a true positive. When the individual is a true attrite ($y_i = 1$) but our model did not predict it ($\hat{p}_i = 0$), the case is a false negative. We then calculate sensitivity, the fraction of true attriters correctly predicted, as the number of true positive cases divided by the sum of true positive and false negative cases.

Similarly, when the actual and predicted values of $\hat{p}_i$ and y_i are both equal to 0, we classify this individual as a true negative. When the individual did not attrite ($y_i = 0$) but our model predicts it ($\hat{p}_i = 1$), the case is a false positive. We calculate specificity as the number of true negative cases divided by the sum of true negative and false positive cases.

In comparing models, one of the goals of prediction is to maximize true positive and true negative cases, and minimize false negatives and false positives. If the model perfectly identifies attriters and non-attriters, both sensitivity and specificity equal 1.

In addition to sensitivity and specificity, we will also examine criteria such as the area under the curve (ROC) across models.

3.4 Looking Forward to FY23 Study

In FY22 we assessed the legacy ESPM and formulated a predictive machine-learning-based ESPM. We had hoped to estimate this model to illustrate its feasibility in replacing the legacy model. Unfortunately, we encountered several challenges in retrieving data. As of the end of FY22 and the writing of this report, we are still awaiting key data elements in order to estimate this model. The model and results will be part of the FY23 study, beginning with Major Aaron Falk's March 2023 thesis.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 4: Line of Effort Three: Simulation Model

4.1 The manpower process

In this section, we develop a discrete event simulation of the Marine Corps manpower process and implement it in java. The model is capable of considering as many as all enlisted Marines or as few as the Marines in a single PMOS. The model is capable of both transient and steady-state analysis and provides estimates of measures such as mean time to promote for each grade, as well as estimates of accession and retention levels.

At any point in time, we know exactly the number of Marines on active duty. We have a plethora of information on each of these Marines, but right now it is sufficient to know each Marine's name, rank, and PMOS. First, we concentrate on the list of Marines in a particular PMOS or community. Next, we can imagine how that list might compare to a similarly constructed list of active duty Marines one year in the future. First, each Marine still on active duty would be a year older, that is, they would each have an additional year of service. Next, some individuals from the first list would have left active duty, either through EAS attrition or otherwise unplanned attrition.[1] Next some individuals will have changed grades. While most of these grade changes are in the positive direction, some would be in the lower direction. Some Marines will have reenlisted, while other Marines might have accessed into this community.

There are a multitude of ways in which the inventory of Marines evolve over a given period of time. Furthermore, most of these processes do not occur uniformly over the year. Accessions tend to be higher in the summer months and lower in the beginning of the calendar year. Reenlistments in a particular MOS cease as soon as the boat-spaces are filled. While selection boards for promotions occur at a fixed time each year, actual promotions are spread across the fiscal year but depend on vacancies caused by attrition and promotions in higher ranks. Thus, any model of the inventory management process must be a distillation of what is a complicated reality.

4.2 Simplifying assumptions

In an effort to develop a model that effectively navigates the tradeoff between tractiblility and realism, we must make some simplifying assumptions. They are as follows:

[1]Note, any Marine who executes a lateral move, e.g. at the FTAP, would count as an attrite from this community.

- Attrition
  - The probability a Marine agent attrites is the same for all agents in a given Grade and YOS.

The model provides greater fidelity with respect to attrition rates than models that have the same attrition rate for anyone in the same grade. Partitioning attrition rates by Grade and YOS is standard practice, though, there could be additional variable that might cause Marines in the same Grade and YOS category to experience different attrition rates.

- Reenlistments
  - The current career force is the sum of all Marine agents with four or more Years of Service. (Or five years of service is the TOE is five years.)
  - The target career force is the sum of all E4s to E9s in the GAR.
  - FTAP Mission = Target - Current
  - Among Marines in YOS 4, select reenlistees and attrite all others.
  - Fair-share reenlistees across grades E3, E4, and E5.
  - The model does not account for or allow for enlistment extensions.

The manpower planners at M&RA have traditionally had a challenge in translating the GAR, which is given by grade and PMOS, into YOS. Thus, the target levels for the career force are difficult to derive from the GAR in a precise manner, since these levels depend on YOS. Our assumption seems like an appropriate one given this limitation.

- Promotions
  - Promotees to Grade_j equals Target for Grade_j - Current number in Grade_j.
  - Form zone according to Selection Opportunity for that grade, in descending order of seniority.
  - Twice passed Marines are not considered for promotion to E6 and above.

This assumption ensures the system unflinchingly promotes to all vacancies. While neglecting twice-passed Marines for promotion to E6 and above, additional empirical research into how to handle promotions to E4 and E5 would likely prove fruitful.

- Accessions
  - Accessions = Target End-Strength - Current End-Strength
  - All Marine agents enter as E3s with 0 years of service.

These assumptions should be rather innocuous. It is possible for a Marine to make E4 in their first year and the frequency with which that occurs in a given PMOS is an empirical question. The model could be easily updated to conform to such a possibility if necessary.

4.3 The Model

We implement our model in Java and make use of the REPASTJ libraries. We choose REPAST not necessarily because our target system particularly lends itself to an agent-based model, rather the visualizations available in that library makes debugging easier and improves situational awareness of how the model behaves.

We intend for the model to be able to consider a specific PMOS or Occupational Field, though it could handle the entire enlisted population, if asked. The model consists of Marine agents that flow between various states during each time-step. Taking care to model each Marine as an agent allows us to precisely measure how they flow through the system. Namely, doing so allows us to measure promotion timing and effectively enforce Enlisted Career Force Controls. In addition, using agents enables us to more faithfully model the promotion process. The Marine objects' instance variables include an accession time-stamp, an array of promotion time-stamps, and a count of the number of times the Marine has been passed over for promotion at the current grade.

Figure 4.1. Screen capture of model display

The core of our model is a network of states that consists of every feasible Grade and Year of Service combination for the community under consideration. Most PMOSs have a path from E1 to E9. For such a PMOS, our model implements grades E3 to E9. In our model, the grade E3 consists of all Marines E3 and below. This allows us to abstract away from modeling junior Marines due to the amount of churn, both up and down, in those ranks.

While the number of grades in the model depends on the structure of the MOS under consideration, the number of Years of Service states available for each grade depends on the Enlisted Career Force Controls. More specifically, since high-year tenure for, say, E6s is 20 years, there exists 20 states at the grade of E6 (YOS 0 through YOS 19). While it would generally be impossible for a Marine in reality to promote to the rank of E6 prior to, say, five or six years of service, we keep such "junior" nodes in the model because it costs very little to do so and we do not want to preclude the model from discovering that a proposed grade structure might require absurdly fast promotions at a particular grade.

Each Grade-YOS object in our model consists of an ArrayList that contains all the Marines in that state. Other instance variables include an attrition rate, as well as a uniform random variate generator.[2] Finally, our model also contains Grade objects and YOS objects. These objects are not necessary for the functioning of the model, but they do help to simplify some of the interactions in the model. For example, when it is necessary to, say, sum the number of Marine agents in the grade E4, rather than querying each of the Grade-YOS states, the method might query the Grade object for that rank.

4.3.1 Model processes map to manpower processes

Our model implements the following processes that map to the manpwer processes we mention above.

1. **Age**

At the beginning of each time-step, every Marine object ages up to the next Grade-YOS state for that Grade. For example, the ArrayList of Marine objects in, say, state E5 and 5 YOS shifts to state E5 and 6 YOS. In the display, the Grade-YOS objects are arranged from left to right in ascending order of YOS, so another way to think of this process is that all Marines simply shift to the state to the right of its current state. For those Grade-YOS states at the maximum allowable YOS for that grade, the agents are essentially attrited in that they are dropped from the model since they have met high-year tenure.

2. **Attrite**

Next, each Grade-YOS state conducts attrition. The Grade-YOS object iterates through each Marine object in its ArrayList and conducts a Bernoulli Trial to determine if that Marine should be attrited. The attrition rate for each Grade-YOS is a user input to the model and is constant for each Marine object in that state. We considered other methods of identifying Marines to attrite, such as shuffling the ArrayList and simply removing the requisite number of objects from the top of the newly shuffled list, however, surprisingly it turns out that the Bernoulli Trial method tends to be faster.

[2]This random number generator is from cern.jet.random. For more information see: https://dst.lbl.gov/ACSSoftware/colt/api/cern/jet/random/package-summary.html

3. **Reenlist**

The next step in the model is to initiate the reenlistment process. We must conduct reenlistments prior to promotions to ensure we are able to fill vacancies created by those Marines objects that do not reenlist. In order to determine the FTAP mission, the model takes as the target number of career force Marines to be the sum of all E4s through E9s in the GAR. The current career force is the sum of all Marine objects with four or more YOS (or five YOS for those PMOSs with five year Terms of Enlistment). The FTAP Mission is the difference between the sum of the target career force and the sum of the current career force.

If the FTAP mission is negative, this means that no Marine crosses over at the FTAP so all Marines in YOS 4 (or YOS 5, as required) are dropped (attrited). A note is made of this problematic and unsustainable incident. If the FTAP mission is positive, the model must determine whom to select for reenlistment. If the mission exceeds the number of Marine agents currently residing in YOS 4, then all agents are declared to have successfully reenlisted and a note is made to denote the unsustainability of that outcome. Finally, if the model must select some Marines to reenlist and some to attrite, it fair shares those reenlistments across the grades, E3, E4, and E5. More specifically, the model draws the same proportion of Marine agents from each grade that is sufficient to satisfy the FTAP requirement.

4. **Promote**

Once the three preceding processes create vacancies throughout the state-space, the model executes the promotion process. The model starts with promotions to E9, or the highest grade in the model, and works down in succession to the lowest grade. In each case, we calculate the number of required selectees to a particular grade by subtracting the number of agents currently in that grade, regardless of YOS, from the GAR target for that grade. The model calculates the required zone size by dividing the number of required selectees by the selection opportunity for that grade.

To form the in-zone population, the model starts by forming a list of every agent in the below grade by seniority, with those Marines with the highest YOS at the top of the list. Any Marine who has been twice passed is dropped from the list. Then, starting from the top of this list, the requisite number of Marines are taken to form the in-zone population. To be clear, this "in-zone" population is actually comprised of Marine agents on their first and possibly second selection boards. If the required zone size exceeds the number of eligible Marine agents, the in-zone population will be every agent in the below grade (less those with 2ps).

For the actual selection process, the model shuffles the ArrayList that contains the in-zone population and simply pulls the required number of selectees from the top of the newly shuffled list. The selected agents are pulled from their current Grade-YOS state and are added to the next higher Grade with the same YOS. The model adds promotees to the bottom of their new Grade-YOS state, in order to maintain relative measures of seniority. Those agents who were not selected are "marked" with a pass-over flag and remain in their current state. If the number of required promotees happens to exceed the zone size, then every Marine agent is selected for promotion. While policy precludes such an event from happening in reality, the fact that an unsustainable incident occurs in the model will be evident in the inventory gap output, as well as the promotion timing assessment.

5. **Access**

Finally, the model must determine the number of accessions and induct the requisite number of Marine objects. The accession mission is calculated by taking the difference between the target end-strength and the "current" end-strength that holds at the beginning of the accession process. If the mission is less than or equal to zero, the accession process is halted and the model moves to the next step. If the mission is positive, then the requisite number of Marine objects are created and loaded into the state E3 with 0 YOS. If it is possible for a significant number of Marines in this MOS to become an E4 in the first 12 months of their enlistment then it is possible for the model to add some to the grade of E4, as well.

4.3.2 Inputs

The model requires a number of inputs. In its current instantiation, the analyst places the inputs in an Excel workbook that the model reads in. However, we discuss how to improve this interface and automate both the collection and preprocessing necessary in the following section.

1. **GAR Targets by Grade**

One of the most important inputs is a GAR target for each grade in the MOS or community under consideration. This informs the model of the structure of the community under consideration, i.e. which grades to model, as well as the target inventory at each grade.

2. **Attrition behavior by Grade and YOS**

Attrition behavior by Grade and YOS is the only input that requires substantial preprocessing in that these attrition rates are not readily available and must therefore be constructed. The general steps to calculate these rates are as follows:

1. From TFDW, pull the end of FY inventory for the PMOS or community in question for the last two to six Fiscal Years. The required variables are EDIPI, Current Paygrade, PMOS, and Armed Fores Active Duty Base Date.

2. For each FY snapshot, categorize each Marine into a state according to Current Paygrade and completed YOS.
3. Compare each FY snapshot with the snapshot from the following year. If a Marine is in the first snapshot and the second snapshot, count them as a "continue" from that Grade and YOS. If a Marine is in the first snapshot but not the second snapshot, count them as an attrite from that Grade and YOS.
4. The attrtion rate for a particular Grade and YOS during a given FY is the number of attrites from that Grade and YOS during that FY and divided by the total number of Marines who began that FY in that Grade and YOS.
5. Care must be taken to appropriately aggregate the estimated attrition rates across FYs. In addition, care must also be taken to determine the appropriate window of time to look back on in order to develop these rates. However, two to three years is typically satisfactory.

The only exception to this process is for attrition for those *from* YOS 3 (or YOS 4 for a PMOS with a 5 year TOE). Under normal circumstances, a Marine who starts a year in YOS 3 and continues to YOS 4 has reenlisted. But the model handles reenlistments as an endogenous manpower process. Therefore, for YOS 3, the input empirical attrition rate should be for non-EAS attrition *only*, rather than for all categories of attrition.

3. **Term of Enlistment**

The model must know the term of enlistment in order to appropriately handle the reenlistment process. While most PMOSs tend to have four-year TOEs, there are some PMOSs with five year TOEs. The model is not currently capable of modeling a PMOS comprised of Marines with different TOEs, though that could be built in at a later date. In the event we wish to use the model to simulate such a PMOS, one technique might be to first run the model assuming all Marines are on a four year TOE, then run the model assuming all Marines are on a five year TOE, and then compare the results.

4. **Enlisted Career Force Controls**

The primary role of the Enlisted Career Force Controls is to communicate to the model how many YOS states to make for each grade. Note that the convention is that the Marine has x completed Years of Service. Suppose that high-year tenure for Staff Sergeants is 20 years of service. That means, for E6, the model will create YOSs 0 through 19. The other part of the ECFCs are selection opportunities for each grade, which the model also uses in the development of the in-zone populations for promotion purposes.

5. **Initial inventory (optional)**

Finally, the user may provide an initial inventory by Grade and YOS. Notice that the process by which one calculates the attrtion behavior in (2) above also enables the construction of an initial inventory.

An initial inventory is necessary if the objective of the analysis is to examine the transient behavior of the community in the near future. For example, one might conduct transient analysis to determine the effects of, say, forcing a community to meet a particular GAR target within a short number of years in the future. This sort of analysis will seek to uncover barriers to success in attempting a significant change in structure.

Alternatively, one might construct a steady-state analysis to try to uncover the long-run feasibility of changes in structure. In this case, one might run the model for 1000 years before assessing the outputs. This gives the analyst a sense of what the system will converge to if left to its own devices. While both of these sorts of analyses can be helpful, one technically only needs an accurate initial inventory for the former.

4.3.3 Managerially relevant outputs

In this section, we outline the outputs of the model. After each timestep, the model reports on the below Measures of Performance. In its current instantiation, the model creates a comma-separated values file that contains the output for each time-step of each run of the model. We employ R in post-processing in order to create the tables and graphs shown below.

1. **Mean Promotion Timing for each Grade**

The primary output is the mean promotion timing to each grade. The model measures this and reports it for each time-step. This allows the analyst to compare expected promotion timing with the timing goals contained in the ECFCs. Supportable and sustainable changes to structure should not require significant deviations from desired promotion timing. Perhaps most important, if a particular community already exhibits problematic promotion timing for a given grade, whether the proposed change tends to improve or exacerbate the promotion timing for that grade could be highly relevant from a managerial perspective.

2. **FTAP Mission**

Another output is the FTAP mission, as executed, for each time-step. Inordinately high or low FTAP missions, especially relative to the expected FTAP population, could be evidence against the supportability of a proposed grade structure for a PMOS.

3. **Accession Mission**

The model also outputs the accession mission for each time-step. Like the FTAP mission, inordinately high or low accession missions could be evidence that the underlying structure of a PMOS is not sustainable.

4. **Inventory gaps**

Finally, the output also includes any gaps in the inventory by grade. That is, the model calculates and reports the difference between the GAR Target and current inventory at each grade at the end of each time-step. If the model identifies a substantial gap at a particular Grade, such as an inability to promote enough junior Marines to fill the vacancies, this could indicate a major problem for the PMOS structure under consideration.

4.3.4 Transient and Steady-State Analysis

Each of the measures we describe above exhibit transient and steady-state features for a given set of input parameters. Figure 4.2 shows expected Time-to-Promote for an analysis of the 35xx (Motor Transport) Occupational Field. Using input values for FY2014, we run the model for 100 time-steps and replicate this design point 100 times. Each point on a line is the mean of 100 replications of that measure at the given time-step.

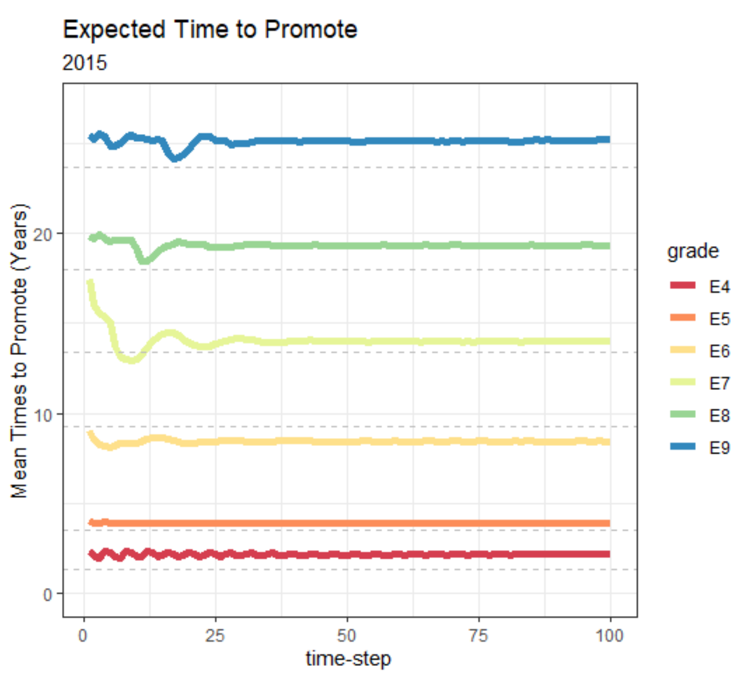

Figure 4.2. Time-plot of Expected TTP for 35xx

The primary take-away from Figure 4.2 is that there exists both a transient segment and a steady-state segment to each of these measures. Note that accession mission and FTAP mission behave similarly. In the cases examined thus far, TTP tends to achieve steady-state no later than 30 time-steps, while the other measures tend to achieve steady-state no later than 80 time-steps. In most cases, the transient values for TTP are significantly higher than the corresponding steady-state values, which suggest that even though the model starts with a valid initial inventory, it fails to handle those Marines in the above-zone properly. In other words, since the number and "location" of twice-passed or above-zone Marines is not part of the initial inputs, the model is happy to promote relatively senior individuals which biases the transient TTP estimates high. As the model endogeneously creates a representative above-zone population, the TTP values approach a steady-state.

4.4 A promising validation demonstration

The purpose of this section is to demonstrate the process by which a model of this sort could be validated. The evidence we present here suggests that the model shows promise but we are limited in that better data likely exists to which we might compare our model outputs. For example, we are unable to compare our model predictions against actual FTAP retention data. However, the primary take-away is that the model appears to behave in a favorable manner and we are optimistic that any deviations from reality can be tolerably removed by more effective calibration of the model and with the use of refined inputs.

Our general validation process is as follows. For a selected community, collect and develop the model inputs for a target year. For target year, say, 2015, we use the initial inventory at the beginning of 2015 (i.e. October 1, 2014) and the attrition behavior (weighted average) for one to three years prior (i.e. 2014, 2014 and 2013, or 2014 - 2012). We select FY2015 to FY2019 as target years so as to avoid any complicating effects of the pandemic in years FY2020 and FY2021. Because we seek to "recreate" the behavior of the system during the target year, we use the actual end-strength totals from that year for the targets. If we were using the model to predict future behavior, we would likely use the GAR projections for the year in question. We run the model for 100 time-steps and generate 100 replications of each design point. We then compare the model's prediction of target year behavior (i.e. inventory levels, TTP, accessions), with our best estimates of those values for that year.

The communities we consider in the validation process are: the entire enlisted population; the 03xx OccFld (Infantry); the 35xx OccFld (Motor Transport); and the 66xx OccFld (Aviation Supply). The enlisted population is managed at the PMOS-level rather than at the service-level, so it is not clear how well an aggregate model of the enlisted population would behave. However, if the model could be validated for use at that level, it could offer some important insights on future force design policies. We select the 03 OccFld because it is the largest and the 35 OccFld because it is approximately half the size of the 03 OccFld. Finally, we select the 66 OccFld as a representative of the smaller OccFlds.

4.4.1 Inventory gaps appear sufficiently small

We first examine the issue of inventory gaps. While gaps between desired GAR levels and end-strength inevitably occur in reality, the reliability with which the model promotes, reenlists, and accesses to vacancy leaves little excuse for any gaps between the predicted inventory levels at each grade and the actual inventory levels at each grade at the end of the target year. Table 4.1 confirms that for each of the OccFlds we consider and each of the target years, the inventory targets at each grade are met with equality.

Table 4.1. Deviations from Grade-level Targets

Community	Target	dev(E3)	dev(E4)	dev(E5)	dev(E6)	dev(E7)	dev(E8)	dev(E9)
03xx	2015	0.0	0.0	0.0	0.0	0.0	0.0	0.0
03xx	2016	0.0	0.0	0.0	0.0	0.0	0.0	0.0
03xx	2017	0.0	0.0	0.0	0.0	0.0	0.0	0.0
03xx	2018	0.0	0.0	0.0	0.0	0.0	0.0	0.0
03xx	2019	0.0	0.0	0.0	0.0	0.0	0.0	0.0
35xx	2015	0.0	0.0	0.0	0.0	0.0	0.0	0.0
35xx	2016	0.0	0.0	0.0	0.0	0.0	0.0	0.0
35xx	2017	0.0	0.0	0.0	0.0	0.0	0.0	0.0
35xx	2018	0.0	0.0	0.0	0.0	0.0	0.0	0.0
35xx	2019	0.0	0.0	0.0	0.0	0.0	0.0	0.0
66xx	2015	0.0	0.0	0.0	0.0	0.0	0.0	0.0
66xx	2016	0.0	0.0	0.0	0.0	0.0	0.0	0.0
66xx	2017	0.0	0.0	0.0	0.0	0.0	0.0	0.0
66xx	2018	0.0	0.0	0.0	0.0	0.0	0.0	0.0
66xx	2019	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Again, while this particular measure might not be as managerially relevant as others we consider below, the model convincingly passes this relatively low bar. It is important to note that these estimates also have zero variance. So, at no time in the 100 replications of each of the target years for each of the OccFlds did a grade miss their target requirement.

4.4.2 Predicted accessions appear favorable

As part of the process of deriving the attrition rates from the data, we are able to calculate reliable accession numbers for each of the target years. We compare the estimates of the accession levels for the target years in steady-state with the actual levels for those years and display them in Figure 4.3. The light-blue bars are the model estimates while the actual values are in dark blue.

One of the first things to notice in Figure 4.3 is that the expected accessions for the entire enlisted population (panel a) matches the empirical data remarkably well. Most of the predictions for the various OccFlds are in the neighborhood of the empirical levels. The standard errors for the estimates are all relatively small, so error bars in these graphs are not included since they are imperceptibly small and therefore unhelpful.

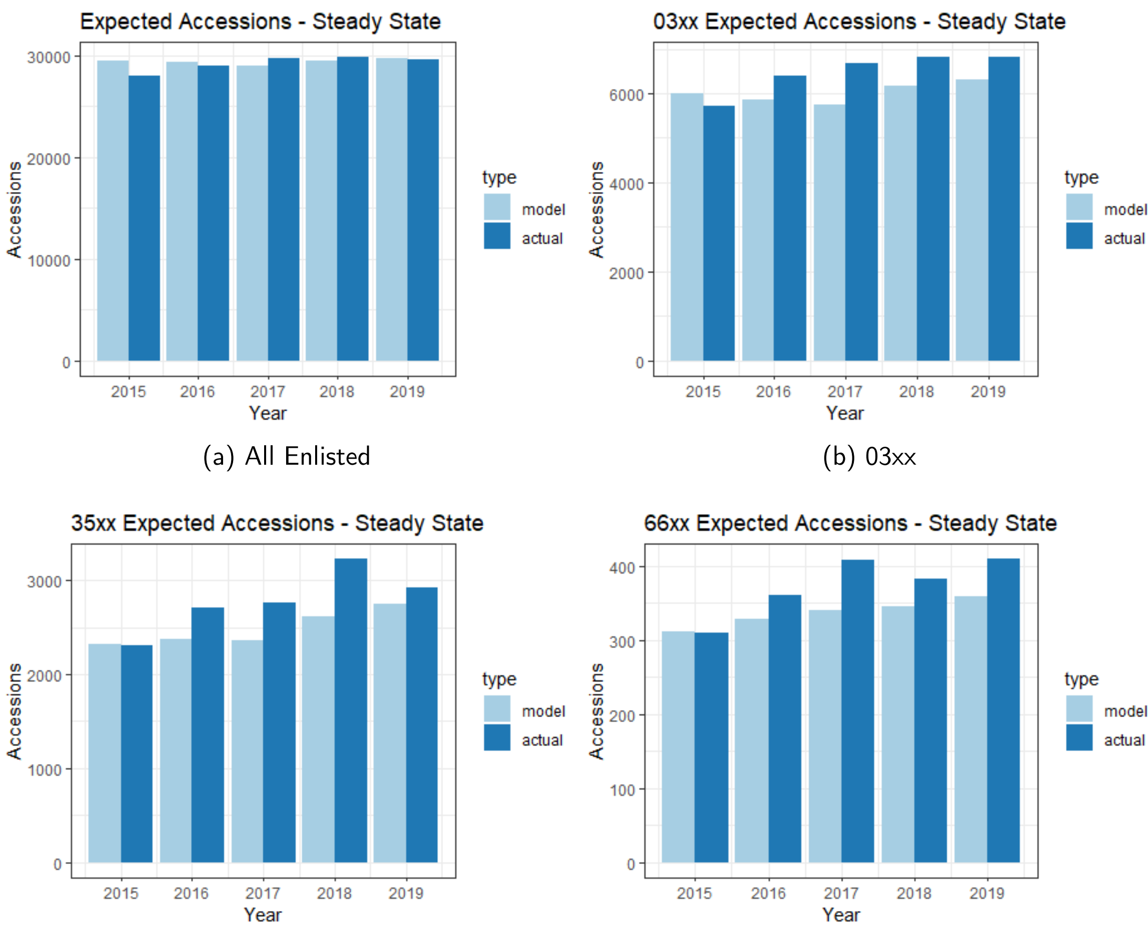

(a) All Enlisted

(b) 03xx

(c) 35xx

(d) 66xx

Figure 4.3. Expected Accessions in steady-state

Keep in mind, since the model simply observes the vacancies and selects the accession level that meets desired end-strength, this measure is more of an assessment of how the agents in the model interact with the matrix of attrition rates. Thus, as we refiene the model to better replicate the placement of Marines with respect to Grade and YOS (e.g. improve management of those above-zone for promotion) then we expect the model to perform even better with respect to this measure.

4.4.3 Mixed results for Time to Promote

A model that effectively and accurately predicts how changes in the manpower system might effect the health of a PMOS or OccFld, especially with respect to Time to Promote, would be an eminently useful model. Such a model would be especially useful for assisting managers of the HRDP during a period of great change. Thus, we compare the model's TTP prediction with empirical levels.

First, we find that the estimates the model generates for transient analysis are generally problematic. Again, we believe this is due to the fact that the initial inventory does not identify which Marines have been passed over for promotion and we expect that refining this aspect of the model will improve performance along this margin. Therefore, in Figures 4.4-4.7 we present the steady-state estimates of TTP for each grade. The light-blue bars show the model estimates, while the actual values for those years are in dark blue. The broken grey line in each graph is the promotion timing target for that grade.

Next we plainly observe that there is substantial deviation between the promotion timing estimates and actual levels at the grades of E4 and E5 for each the experiments. Our calculations for actual mean time to promote to E4 is approximately 1.5 years, while the official promotion timing target for that grade is 2.5 years and there is similar deviation for E5s (see panels (a) and (b) of Figures 4.4-4.7). This is evidence that undermines the credibility of our empirical data. Thus, we encourage the reader to consider this section as proof of concept, rather than particularly compelling evidence in favor of validation of this aspect of the model.

Finally, we note that for estimates of promotion timing for grade E6 and higher, the model is generally in the neighborhood of the actual. Certain grade and year combinations for certain OccFlds are better than others, but in general the model appears to perform fairly well. In addition, the size of the target population does not seem to effect the general accuracy of the predictions. Finally, it is worth noting that, as above, the standard errors for these estimates are sufficiently small so as to render error bars unreadable.

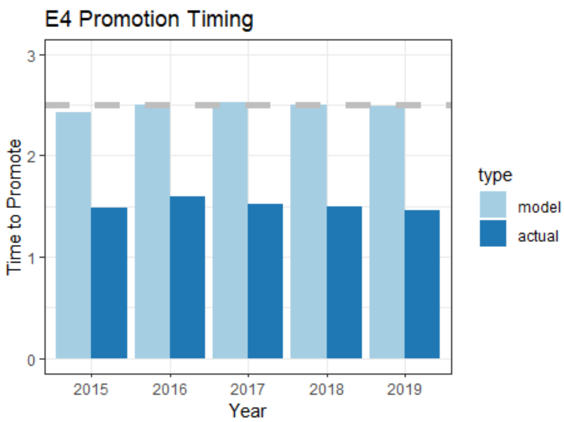

(a) E4

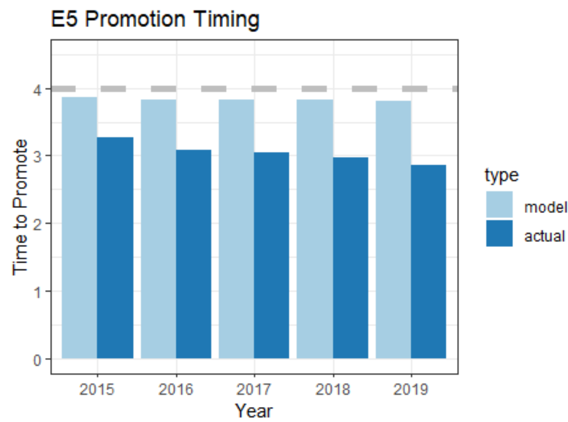

(b) E5

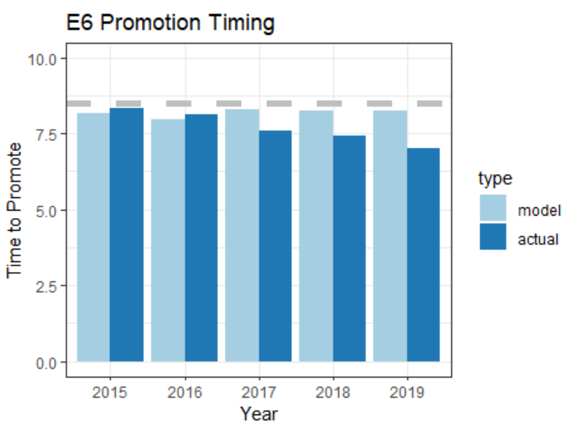

(c) E6

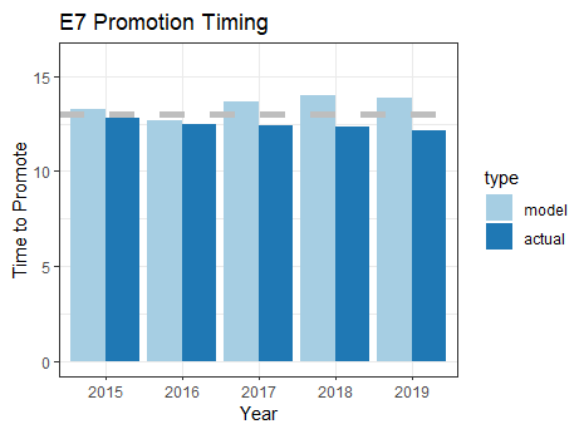

(d) E7

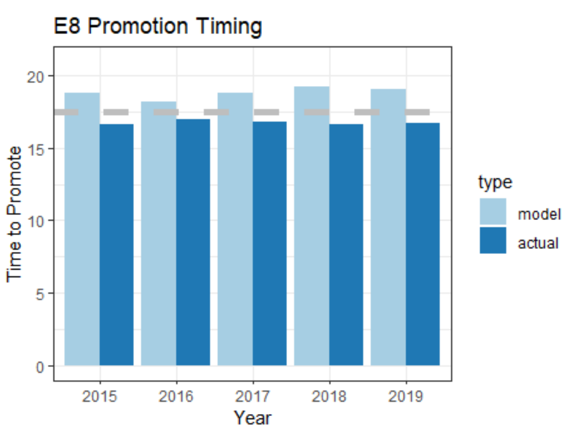

(e) E8

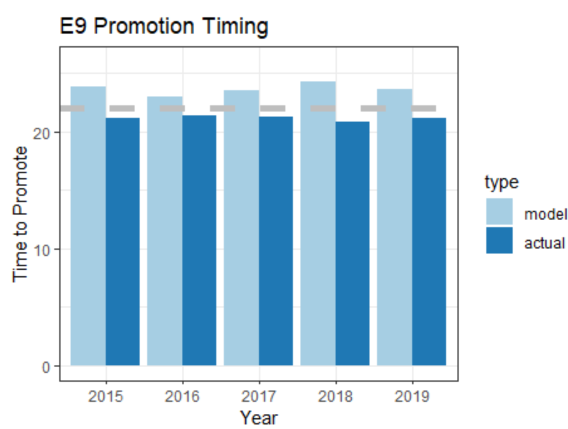

(f) E9

Figure 4.4. Expected Time-to-Promote by Grade for all PMOSs

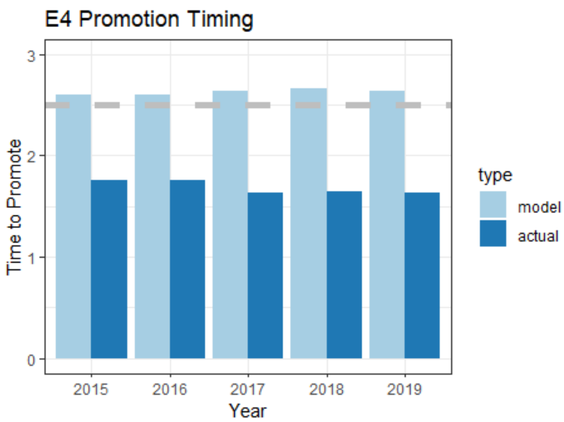

(a) E4

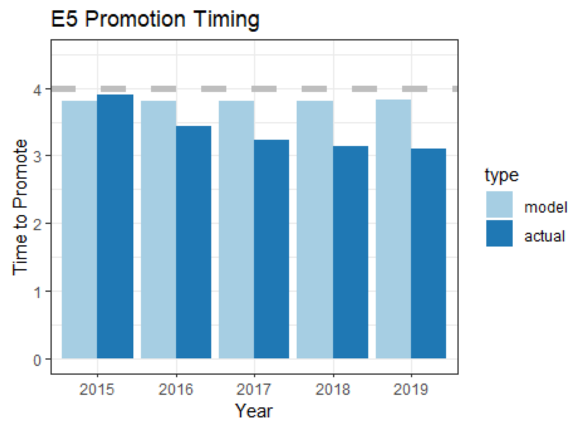

(b) E5

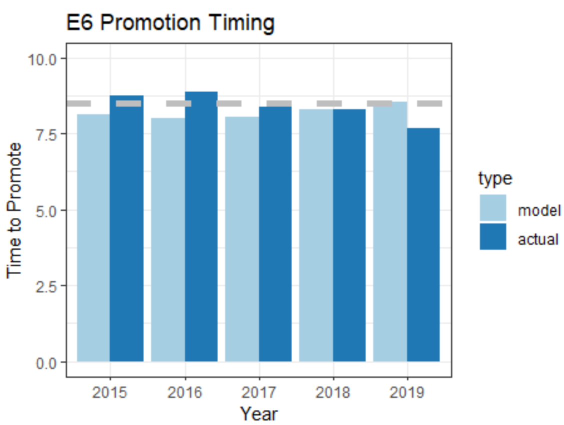

(c) E6

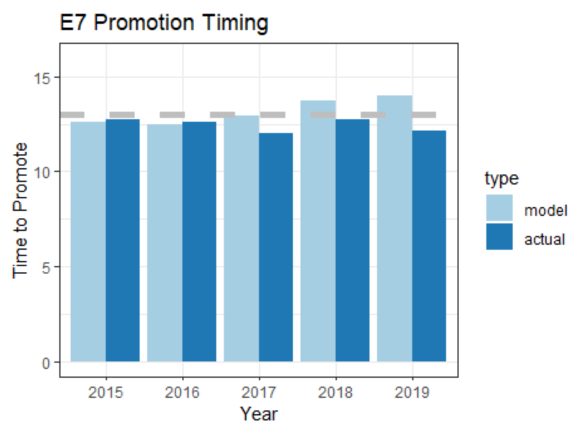

(d) E7

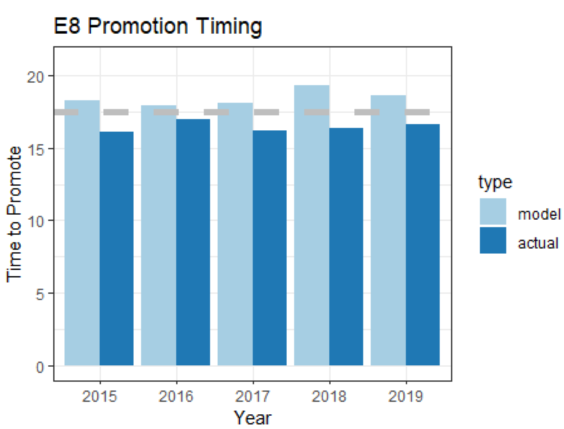

(e) E8

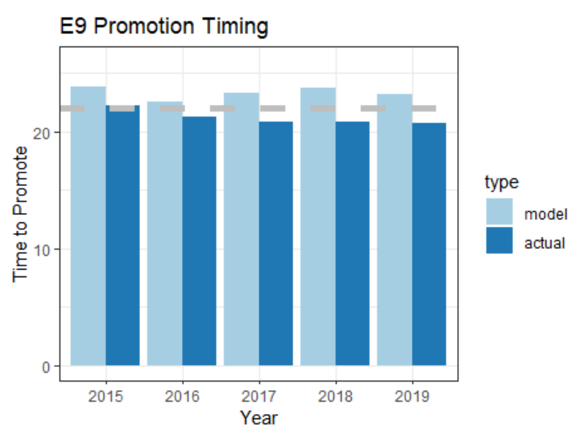

(f) E9

Figure 4.5. Expected Time-to-Promote by Grade for 03xx

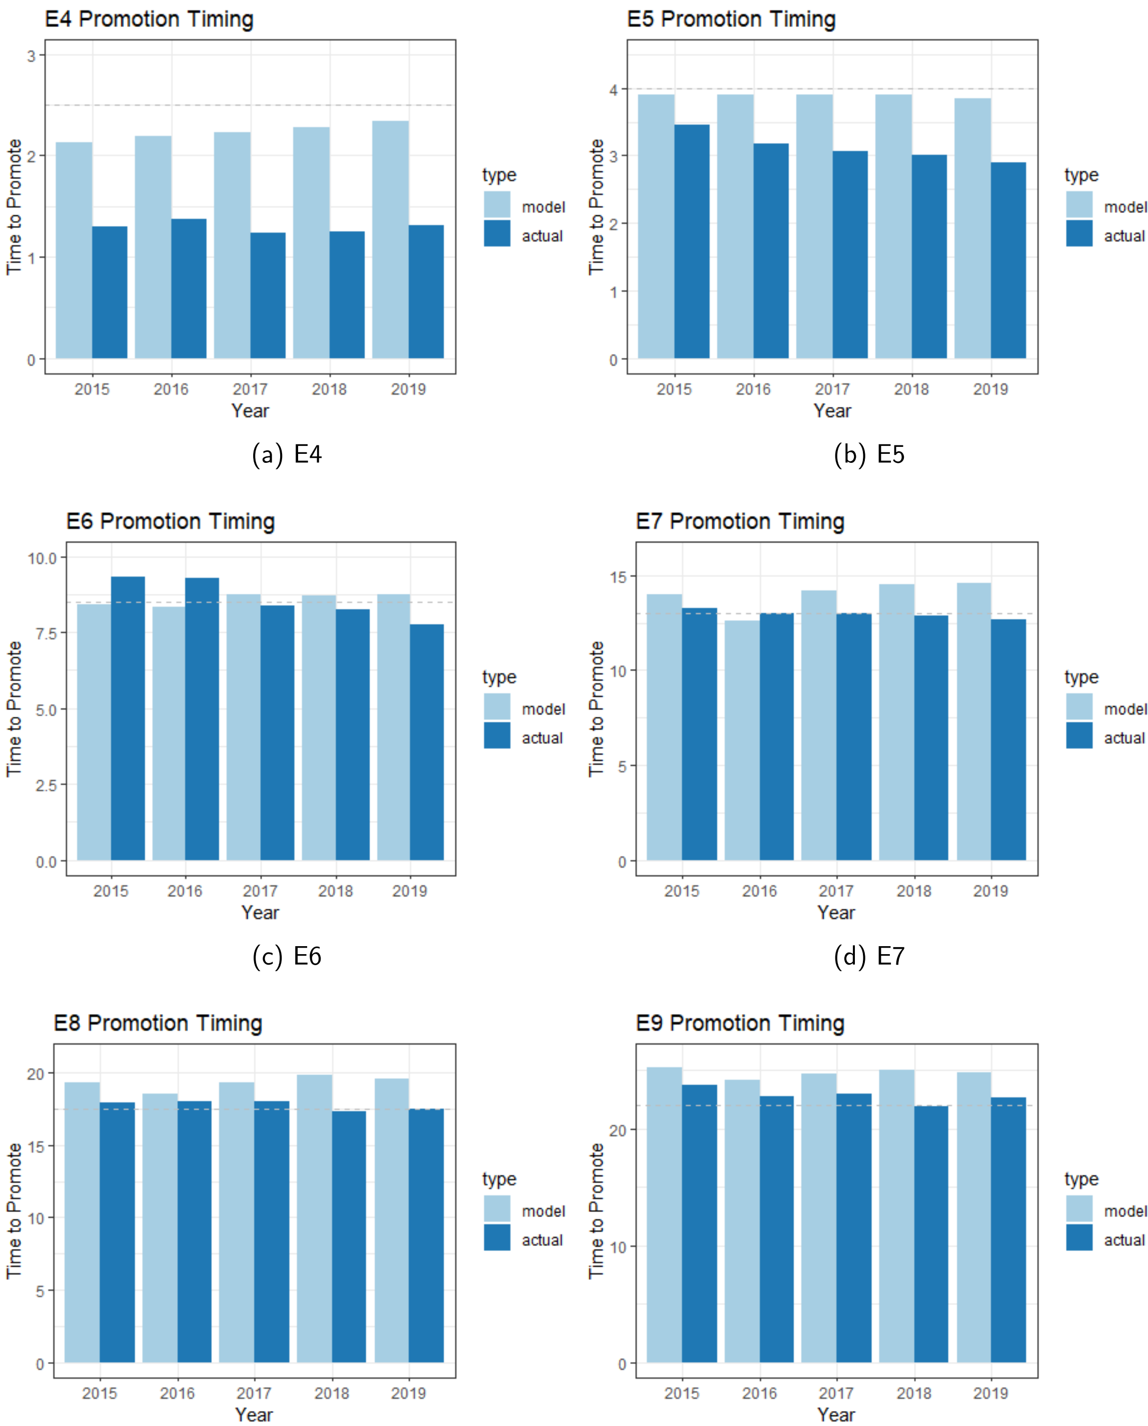

(a) E4
(b) E5
(c) E6
(d) E7
(e) E8
(f) E9

Figure 4.6. Expected Time-to-Promote by Grade for 35xx

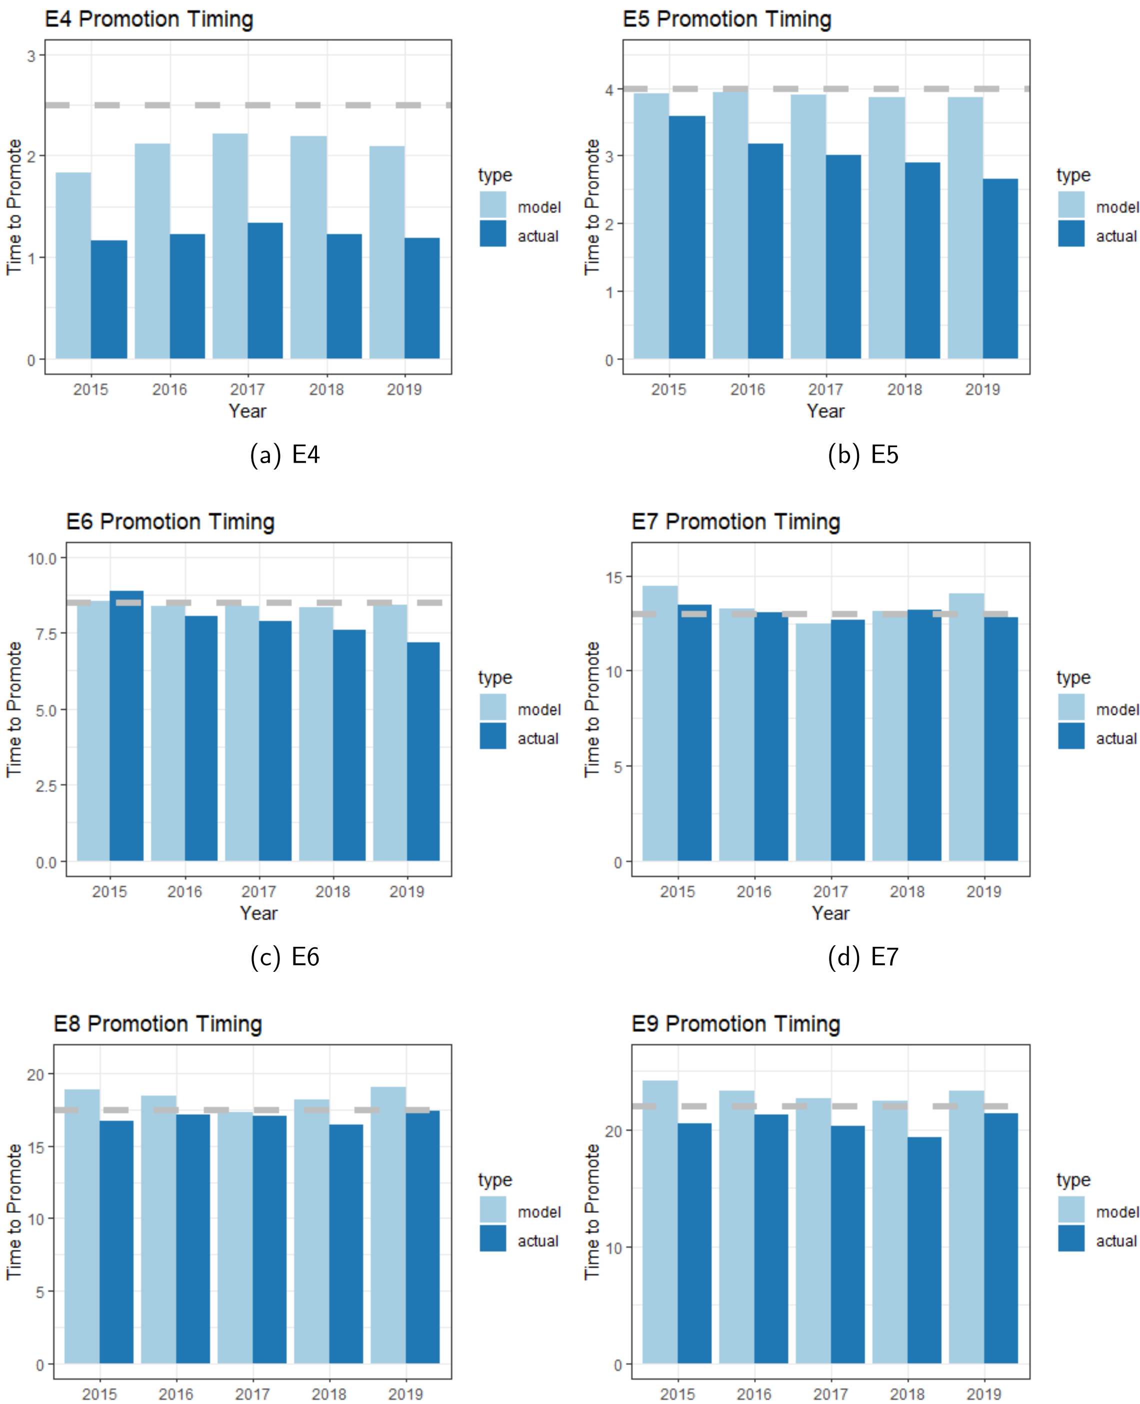

(a) E4 (b) E5

(c) E6 (d) E7

(e) E8 (f) E9

Figure 4.7. Expected Time-to-Promote by Grade for 66xx

4.5 An excursion in support of Force Design 2030

In *Force Design 2030* and elsewhere, the Commandant of the Marine Corps alludes to the fact that, in order to achieve the objectives he lays out, the Marine Corps will need an older, more mature enlisted force (Berger 2020b, 2021). That is, the Marine Corps must shift to having a larger share of enlisted end-strength as part of the career force, rather than as first-term enlistees. Tweedy (2022) is one attempt to assess the practicality of this initiative from a manpower perspective. In this section we demonstrate how to employ our model to make such an assessment.

We implement our model using the latest empirical data (i.e FY2021) from the 03 OccFld. We set as target the GAR requirements for FY2027 that we obtain from Tweedy (2022). We run the model for 100 time-steps and replicate that process 100 times. Figure 4.8 is a time-plot of the TTP for each grade. Note, the dotted grey lines are the promotion timing targets for each grade.

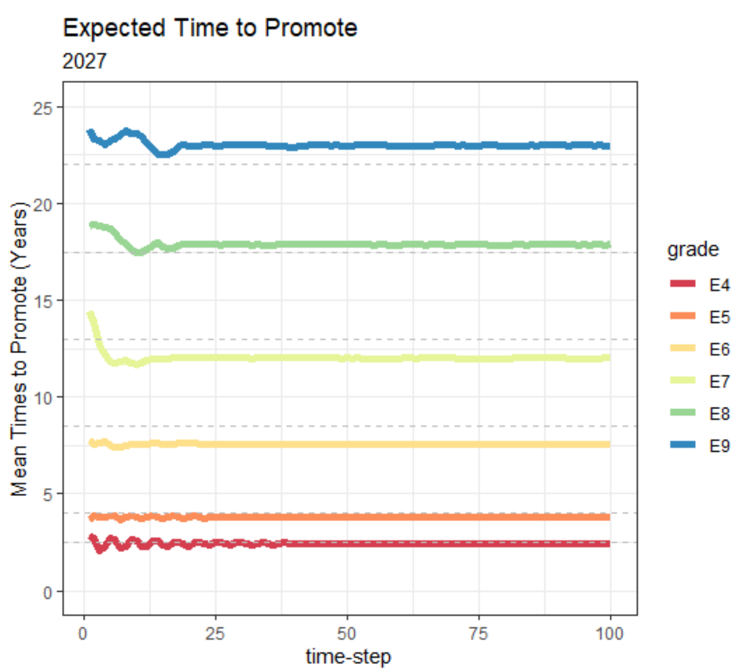

Figure 4.8. Time-plot of Time-to-Promote estimates

First, we notice that the model does converge to steady-state with respect to these measures. Next we notice that the TTP measures converge to levels relatively close to the promotion timing targets. Figure 4.9 presents a clearer picture of the deviation of the predictions from the targets.

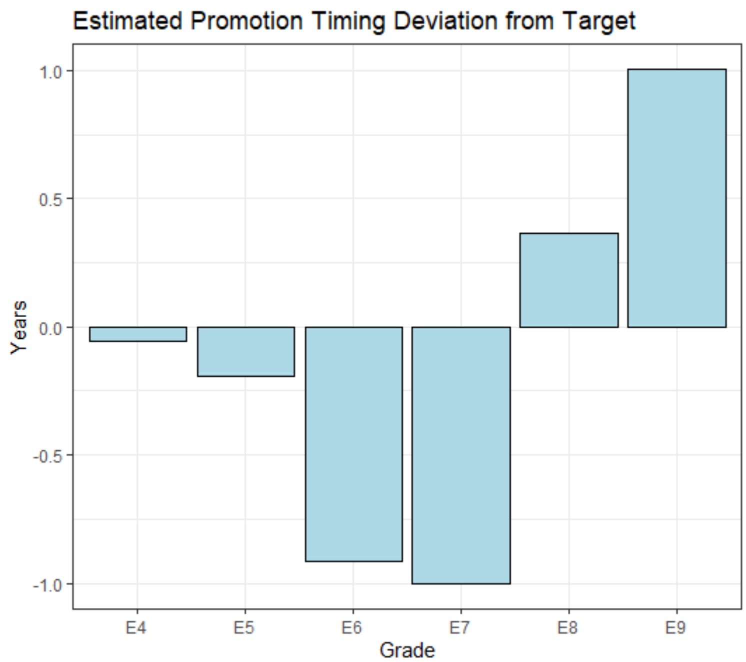

Figure 4.9. Estimates of Deviations from Promotion Timing Targets

The model predicts that the new inventory targets will result in mean promotion timing to E6 and E7 nearly one year early, while mean promotion to E9 approximately one year late. The promotion timing deviations for the other grades appear tolerably small. At this point, it would be up to the analysts and the manpower managers to determine if these deviations are acceptable. A PMOS that deviates too far from the targets is assigned a different in-zone selection opportunity for promotions to that grade. Typically, a PMOS must experience deviations in excess of plus or minus one year in order to be assigned a different selection opportunity. Thus, the evidence from the model suggests that the proposed changes to "age the force" in the 03 OccFld are likely sustainable from the perspective of promotion timing.

Finally, we assess the policy from the perspective of accession and retention levels. We need to determine if the new structure might induce a accession or retention requirements that are unsustainable. Figure 4.10 is a time-plot of the predicted accession, FTAP Mission, and FTAP population.

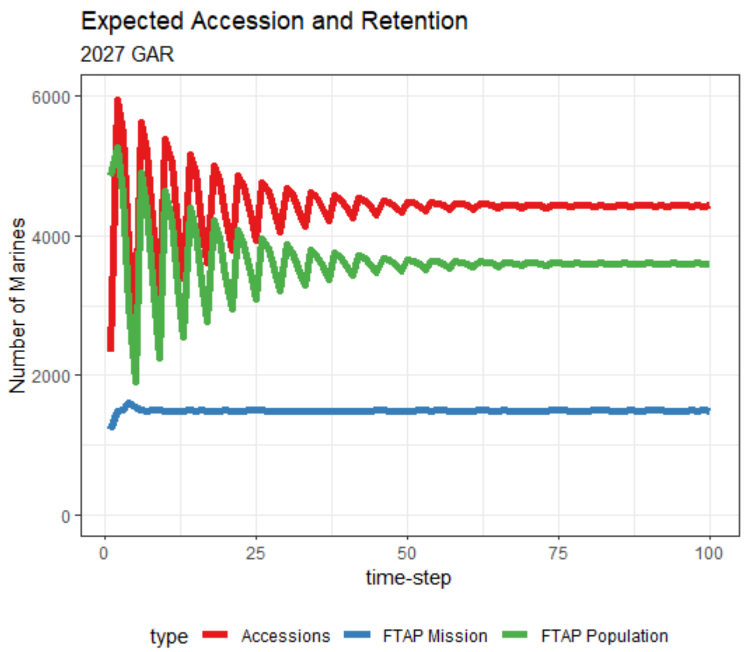

Figure 4.10. Time-plot of estimated accession and retention requirements.

First, we notice that the level of the FTAP mission level is quite stable. Next, when we compare the FTAP mission to the steady-state of the FTAP population prediction, we see it is on the order of 40%. While historically, FTAP proportions are typically in the neighborhood of 20-30%, this value is within expected levels and consistent with the Commandant's expectations of retaining a higher proportion of first-term Marines at the FTAP. Finally, we observe that the steady-state prediction for accessions is approximately 4,500. The historical levels for this OccFld is typically greater than 6,000. This prediction is also consistent with expectations concerning the effects of the maturing the force initiative.

Ultimately, this section is merely a demonstration of how analysts may employ the model to gain insight and quantitative evidence in support of a relevant policy issue. We must qualify this analysis and warn that until the model is more fully validated, the reader should refrain from drawing strong conclusions with respect to this policy. However, this section is indicative of the sorts of insight that could be gained.

4.6 Discussion

In this section, we develop a model of the manpower process which could be used to support a variety of managerial decisions. While upon initial examination it appears to perform well, this aspect of the project was plagued with the some of the same issues at LOE 2, and much less time was available for validation. Therefore, we recommend to continue the validation efforts into the future, both with respect to additional target years and additional communities and PMOSs. Researchers should confirm the existence of proper TTP data as well as empirical retention data. It is likely that minor refinements to the manner in which Marines in the above-zone for promotion are managed, as well as minor refinements to how reenlisted Marines are distributed across grades will prove worthwhile and improve model performance.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 5:
Data Architecture and Discussion

5.1 A word about data architecture

In this paper, we identify several analytic models as candidates for modernization. It is worthwhile to give a brief overview of how one might implement these models in a modernized system architecture.

The first objective would be create an analytics data store to serve as the basis from which input parameters would be drawn. For example, an analytics data store for the TFPM would consist of the relevant version of the ASR. This requires close coordination with the relevant Database Administrator staff to identify the appropriate source tables and, among other concerns, to develop the scripts to extract the data and to schedule the scripts that update the data store to execute at the right time in order to ensure proper synchronization. The extracted data then has to be loaded into databases for analysis. Moving data over the network can slow down network access speeds so dedicated connections needs to be established for fast and reliable operations.

The automated scripts must be tested and installed on each server which hosts the needed database tables. If the end users are internal to the enterprise, then a 2-Tier system needs to be designed for them to interact with the system. As users work from remote locations then a 3-Tier system to enable data access outside the enterprise has to be researched. A 3-Tier system allows access over the web using a web-server to communicate with browsers and apps, and process the information before passing it on to a database hosted in the intranet.

The loaded data should be checked for data integrity. This requires code that automates the steps of data cleaning. After data are stored in a database schema it has to be reviewed by subject matter experts. For users to interact with data in servers custom applications need to be designed that run on desktops, laptops, and handheld devices.

Once the analytics data store is created, the newly developed models can be tested in the new system. This entails building a sandbox server with browser access. Once the sandbox testing shows stability (in Analysis) a production version can be built.

Working with any legacy system means that the right community (staffing) has to be put together to understand existing systems and identifying what is relevant for analysis. Often there is existing codebase that can be leveraged for planning the next steps. This will involve code review and testing with sample data before using for future work.

Analytics data can grow over time and exceed the capacity of a database, when one has to find cost effective cluster architectures. One such option is the Hadoop Distributed File System (HDFS) that uses a divide and conquer approach to process large amounts of data. An industry best practice is to use HDFS as the first destination for the data and then the community can create subsets for analysis on databases. The assumption is that databases have a rich tools kit of software for analysis. This will involve HDFS and the database being able to communicate and the users accessing both systems as needed.

5.2 Findings and Discussion

Ultimately, the most important finding of this project is that the records contained in TFDW might be insufficient to support building a rigorous mathematical model of attrition behavior. The Marine Corps will fail in its efforts to transform from an industrial age organization to an information age organization if the data it relies on to make the most elementary manpower management decisions, i.e. the effort to make end-strength, is flawed or non-existent.

LOE 1: We successfully implement a replication of the TFPM for officers in Python using a parallel processing approach. We process all 257 combinations of MOS and Officer Type with FY22 data and compare those against the official TFPM outputs and produce identical results in 244 out of the 257 cases. The other 13 instances contain discrepancies that were reviewed by a domain expert. Some discrepancies were caused by manual adjustments and others were categorized as errors. We provide detailed recommendations for improvements in four key areas. We identify significant disadvantages with the data formats we received, including inconsistent field naming conventions from one data table to another, and implicit representations that are understandable to a human analyst but insufficient for machine processing. These data formatting deficiencies increase the complexity of logic and code required for generating the TFPM.

LOE 2: We assess the legacy ESPM and formulate a predictive machine-learning-based ESPM. We had hoped to estimate this model to illustrate its feasibility in replacing the legacy model. Unfortunately, we encountered extreme challenges collecting the necessary data. As of the writing of this report, we are still awaiting key data elements in order to estimate this model. The model and results will be part of the FY23 study, beginning with Maj Aaron Falk's March 2023 thesis.

LOE 3: We successfully formulate and implement a discrete event simulation model of the manpower system in Java. Given a notional target inventory for an OccFld; the current inventory for that OccFld; historic attrition behavior; and the relevant Enlisted Career Force Controls; the model provides estimates of expected promotion timing for each grade, expected accession mission, and expected retention requirements. The model easily employs data-farming techniques and lends itself to both transient and steady-state analyses. While initial validation attempts the model appears to perform well, this aspect of the project was plagued with the some of the same issues at LOE 2, and much less time was available for validation. Therefore, we recommend to continue the validation efforts into the future, both with respect to additional target years and additional communities and PMOSs. Researchers should confirm the existence of proper TTP data as well as empirical retention data. It is likely that minor refinements to the manner in which Marines in the above-zone for promotion are managed, as well as minor refinements to how reenlisted Marines are distributed across grades will prove worthwhile and improve model performance.

5.3 Recommendations for Future Research

First, the Marine Corps should make steps to ensure that the reason(s) that a Marine separates from active duty is properly archived and is easily retrievable for purposes of analysis. The difficulty in collecting this and other relevant data prevented any real progress towards replication of the Enlisted End-strength Planning Model (ESPM) (line of effort [LOE] 2) and hampered validation efforts of the model we analyze in LOE 3. Given this lack of progress, we recommend continuing work on the ESPM.

Upon conclusion of Aaron Falk's March 2023 thesis, the model he develops and analyzes in that thesis will be delivered to Manpower Reserve Affairs (MRA). Due to statutory restrictions on the appropriate uses of Navy Research Program funding, the other models we analyze in this project cannot be delivered. We recommend MRA seeks other channels to replicate this research and develop models that can be lawfully delivered and implemented.

THIS PAGE INTENTIONALLY LEFT BLANK

Acknowledgements

Thank you to Mark Ramirez for assistance in explaining the TFPM process, gathering the required data, reviewing our work, providing feedback on discrepancies identified between our results and the official TFPM, and for fielding many questions throughout the project.

THIS PAGE INTENTIONALLY LEFT BLANK

List of References

Berger DH (2020a) *Commandant's Planning Guidance: 38th Commandant of the Marine Corps* (Marine Corp HQ Washington DC United States).

Berger DH (2020b) Force design report 2030. Technical report, Marine Corp HQ Washington DC United States.

Berger DH (2021) Force design report 2030: Annual update. Technical report, Marine Corp HQ Washington DC United States.

Tweedy MD (2022) *The Infantry Maturity Quotient: An Analysis of the Manpower Effects of Maturing the Marine Infantry*. Master's thesis, Naval Postgraduate School, Monterey, CA.

THIS PAGE INTENTIONALLY LEFT BLANK

Initial Distribution List

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California