

Personnel Cost Minimization Using A Stochastic Inventory Projection Model: A Prototype

Gregg Schell and Ann Parcell

with David Reese

DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

Abstract

On average, female enlisted sailors have lower continuation rates than male sailors, but the size of the difference varies by enlisted management community (EMC) and years of service. To fill requirements as the female share of accessions increases, the Navy can increase the overall number of accessions, increase retention bonuses, or both. The choices generate different costs for each EMC that depend on the required accession qualifications (e.g., recruiting effort and training intensity/time), the EMC billet structure, and the size of the gender differences in continuation rates. We present a prototype stochastic inventory projection model that helps make two main decisions for each of 5 EMCs independently: (1) number of accessions and (2) selective reenlistment bonus (SRB) levels. For different levels of the female share of accessions, the model minimizes cost while meeting manning requirements. We then employ a second cost minimization routine (i.e., a bi-level optimization) to find the cost-minimizing solution across the five EMCs. If expanded to all EMCs, the model could provide analytic support for finding cost-minimizing accession and SRB plans.

This document contains the best opinion of CNA at the time of issue.

It does not necessarily represent the opinion of the sponsor.

Distribution

DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

Public Release.

4/17/2019

This work was performed under Federal Government Contract No. N00014-16-D-5003.

Approved by:

April 2019



Molly McIntosh, Director
Navy Compensation and Policy Program
Resources and Force Readiness Division

Request additional copies of this document through inquiries@cna.org.

Abstract

On average, female enlisted sailors have lower continuation rates than male sailors, but the size of the difference varies by enlisted management community (EMC) and years of service. To fill requirements as the female share of accessions increases, the Navy can increase the overall number of accessions, increase retention bonuses, or both. The choices generate different costs for each EMC that depend on the required accession qualifications (e.g., recruiting effort), the EMC billet structure, and the size of the gender differences in continuation rates. We present a prototype stochastic inventory projection model that, for a given female share of accessions, helps make two main decisions for each of five EMCs: (1) number of accessions and (2) selective reenlistment bonus (SRB) levels that minimize total personnel costs while meeting manning requirements. We then employ a second cost minimization routine (i.e., a bi-level optimization) to find the cost-minimizing distribution of female accessions across the five EMCs (and their corresponding optimal accession-SRB solutions) for a given female share of accessions when the total number of accessions and the total SRB budget are constrained. If expanded to all EMCs, the model could provide analytic support for finding cost-minimizing accession and SRB plans.

This page intentionally left blank.

Executive Summary

In the past 30 years, the female share of enlisted accessions has increased to about 25 percent. Moreover, the Navy opened most skill specialties to women in 1994, and it opened the remaining few specialties in the last several years. On average, across all skill specialties, women have lower continuation rates than men, although the gender difference in rates varies by skill specialty (enlisted management community (EMC)) and across the years-of-service (YOS) profile. As the female share of accessions increases and is distributed across skill specialties, the Navy may have to consider a range of accession and retention strategies to fill requirements efficiently. Possible strategies include increasing accessions, increasing retention bonuses (e.g., selective reenlistment bonuses (SRBs)), or some combination of the two. To fill requirements as efficiently as possible, the Navy should choose a joint accession/retention strategy that minimizes personnel costs.

To our knowledge, the Navy does not have a model to aid in determining the cost-minimizing accession-SRB plan for a desired requirement fill rate and desired female share of accessions. Although some previous analyses have shown accession-retention personnel cost trade-offs at the aggregate level, we do not know of any model that finds joint cost-minimizing accession-SRB choices by EMC, particularly as the female share of accessions changes. One result is that there may be a tendency to increase the number of accessions as the default response to an increase in the female share of accessions. Yet personnel costs may not be minimized by increasing accessions only, and the optimal solution likely varies by EMC.

In this report, we build a prototype model to find cost-minimizing accession-SRB solutions to fill requirements when the female share of accessions is changed. Specifically, we built a stochastic inventory projection model (SIPM) that simulates the movement of every male and female sailor from accession through his or her Navy career for each of five EMCs. The model is stochastic because each sailor's career progression is dependent on probabilities for annual continuation (vice loss), and advancement.

To set the baseline inventory simulation, we use recent billets authorized (BA) by paygrade for each EMC as the benchmark for requirements, and we use recent accessions and inventory to initialize the simulation. In addition, we base our stochastic treatment of sailor career progression on recent historical rates for continuation and advancement.

Our model includes three types of personnel costs: recruiting (including enlistment bonuses (EBs)), sailor compensation (less EBs and SRBs), and SRBs. For recruiting costs, we use a simple cost structure in which the cost of accessions varies directly with the test scores required for the promise of an EMC and includes an estimate of recruiting effort, advertising costs, and EBs. For sailor compensation, we apply the annual Department of Defense (DOD) composite rates for the Navy by paygrade to each sailor in inventory, from which we subtract the EB and SRB amounts. Then we use recent SRB levels as the baseline cost of SRBs by EMC and allow the SRB levels to fluctuate in the cost minimization procedure. We assume that the retention rate response to a one-level increase in SRB can vary by EMC and follows results based on previous CNA research [1].

Our modeling and optimization approach proceeds in two stages. First, we employ the grid randomized iterative descent (GRID) search technique for each of five EMCs to find optimal accession and SRB plans. The optimal plan minimizes total personnel cost while satisfying fill requirements for a desired female share of accessions. We measure fill by the inventory-to-BA ratio (INV/BA) by paygrade for each year of the simulation. The model also calculates the average YOS by paygrade for each year of the simulation to monitor changes in average experience levels across the inventory. In the second stage, we use another cost minimization routine (i.e., a bi-level search) to find the cost-minimizing distribution of female accessions across all five EMCs when total accessions and the total SRB budget are constrained.

Our results from the first stage of modeling show that considering changes to accession and SRB levels simultaneously—in contrast to considering changes to accession and SRB levels in isolation—can yield savings for the Navy. Moreover, the cost-minimizing combinations of changing accessions and SRBs may vary by EMC. For example, for the Hospital Corpsman (HM) EMC, we increased the female share of accessions from the recent historical share of 31 percent to 36 percent and required a BA fill rate of at least 95 percent for every paygrade from E-3 through E-6. The increase-SRBs-only solution is almost 5 percentage points more costly than the increase-accessions-only solution, which, in turn, is 3.5 percentage points more costly than the optimal solution—that is, a combination of increasing accessions and changing the SRB levels. By contrast, when we increased the female share of accessions for the Aviation Ordnanceman (AO) EMC from the historical share of 35 percent to 40 percent, the model suggested that the cost-minimizing solution was to increase accessions and reduce SRBs to zero.

We then found the associated optimal (i.e., personnel cost minimizing) combination of accessions and SRB levels for five different values of the female share of accessions for each EMC. These five values are the recent historical female share of accessions, plus and minus 5 percentage points of the historical share, and plus and minus 10 percentage points of the historical share. These variations on female share of

accessions and their associated cost-minimizing accession-SRB solutions became the inputs for the second optimization routine.

Because Navy personnel planners typically face constraints on the total number of accessions and the total SRB budget, we employ a second optimization routine to find the optimal allocation of female accessions across the five EMCs. Specifically, this second stage finds the optimal allocation of female accessions across the five EMCs for a desired all-Navy female share of accessions. The optimal female accession allocation minimizes all-Navy personnel costs while meeting constraints on the total number of accessions, total SRB budget, and EMC fill requirements.

Our results from the bi-level optimization illustrate two points. First, our prototype can determine the personnel-cost-minimizing female accession allocation across the five EMCs within the set of constraints. Second, if we change the desired all-Navy female share of accessions, the constraint on the total number of accessions, and/or the constraint on total SRB budget, the model allows us to reoptimize efficiently and relatively quickly.

The results provide evidence that developing a true all-Navy SIPM/cost minimization model could be beneficial for accession and retention policy planning. In future model development, however, we must address some simplifying assumptions and modeling choices that we made to achieve a functioning inventory projection model and cost minimization routines within the time and resource constraints of this study. They include the following changes:

- **Expand the definition of meeting requirements.**
 - Ensure that the simulated inventory can meet a desired level of both sea and shore requirements.
 - Ensure that a desired level of fit is also achieved.
- **Refine the modeling of career progression.**
 - Adopt a modeling strategy to move sailors between sea and shore tours that takes into consideration the fact that notional and actual sea/shore flows differ.
 - Model more precisely the changes in time-to-train before reaching the fleet as the number of accessions increases.
 - Model more Navy career events as stochastic.
- **Include the following costs:**
 - Costs associated with encouraging sailors to go to sea and/or stay at sea longer

- Costs associated with time awaiting training and/or the infrastructure cost of reaching the limits of the current training system

Contents

Introduction	1
Motivation for the study.....	1
Study goals.....	1
Approach.....	2
A prototype inventory projection model.....	2
The cost minimization routine.....	2
Key measurement challenges.....	3
Report organization.....	4
SIPM: Our Prototype Model	5
Starting inventory.....	8
Gains to EMC.....	8
Losses from EMC.....	8
Advancements.....	9
Costs.....	9
Performance measures.....	10
Cost-Minimizing Accession and SRB Policy	11
Optimization model.....	11
Decision variables.....	11
Constraints.....	12
Objective function.....	13
Model formulation.....	13
Grid randomized iterative descent (GRID).....	13
Bi-level optimization.....	14
Results	17
Individual EMCs.....	17
HM.....	17
AO.....	19
The all-Navy bi-level optimization.....	20
Imposing all-Navy constraints.....	20

Varying the all-Navy constraints.....	21
Conclusions and Recommendations	25
Appendix A: Background for Modeling Effort	28
Appendix B: Recent Inventory Projection Models	31
Appendix C: Notional Versus Actual Sea/Shore Flows	33
Appendix D: Additional Details on GRID	35
Appendix E: Cost and Inventory Excursions.....	38
References.....	40

List of Figures

Figure 1.	SIPM high-level process	7
Figure 2.	Heat map of costs for target total female accession percentage	22
Figure 3.	All-Navy one-year and cumulative continuation rates by gender.....	28
Figure 4.	AO one-year and cumulative continuation rates by gender	29
Figure 5.	YN one-year and cumulative continuation rates by gender.....	29
Figure 6.	Percentage of sailors on sea duty by years since full duty	33
Figure 7.	Randomized iterative descents.....	36

This page intentionally left blank.

List of Tables

Table 1. SIPM parameters and data sources	5
Table 2. HM results for an increase in female accession shares (31% to 36%)	18
Table 3. AO results for an increase in female accession shares (35% to 40%).....	19
Table 4. Excursions for AO optimal accession and SRB policies.....	39

This page intentionally left blank.

Glossary

AO	Aviation Ordnanceman
ASN(FM&C)	Assistant Secretary of the Navy (Financial Management and Comptroller)
BA	billets authorized
DOD	Department of Defense
EB	enlistment bonus
EMC	enlisted management community
EMF	Enlisted Master File
EN	Engineman
EPA	enlisted program authorizations
ETF	Enlisted Tracking File
FC	Fire Controlman
FCA	Fire Controlman (Aegis)
FYDP	Fiscal Year Defense Plan
GRID	grid randomized iterative descent
HM	Hospital Corpsman
INV/BA	inventory-to-BA ratio
IPM	Inventory Projection Model
LIMDU	limited duty
MMSW	Machinist's Mate (Surface Warfare)
MPTE	Manpower, Personnel, Training, and Education
NEC	Navy Enlisted Classification
NeMMo	Navy Enlisted Management Model
OS	Operations Specialist
PACT	Professional Apprenticeship Career Track
PG	paygrade
SIPM	stochastic inventory projection model
SRB	selective reenlistment bonus
STF	Street to Fleet
TFMMS	Total Force Manpower Management System
TIG	time in grade
YN	Yeoman
YOS	years of service
YSFD	year since reaching full duty

This page intentionally left blank.

Introduction

This report is part of a larger study on Navy and USMC personnel costs sponsored by the Assistant Secretary of the Navy (Financial Management and Comptroller) (ASN(FM&C)). In this report, we address the study tasking on gender differences in Navy career progression and the implications for manning and personnel costs.

Motivation for the study

The female share of accessions has increased over the last two decades and could increase in the future. In many specialties, women have lower continuation rates than men. If the Navy takes no policy actions to counter lower continuation rates, the result could be lower manning rates. Thus, to achieve manning targets as the female share of accessions increases, it may be necessary to increase the total number of accessions, provide incentives to increase retention, or both. However, we must analyze the personnel costs associated with each policy to know which approach may be most cost effective.

Study goals

We focus our analysis on answering the question, how will an increase in the female share of accessions affect manning and personnel costs? Typically, inventory projection models are used to determine the manning implications of a change in the number or mix of accessions. They do not typically model changes to accession and retention incentives jointly, and, in particular, they do not typically identify cost-minimizing solutions to manning the fleet across possible combinations of accession and retention policies.

Differences in female and male sailor continuation across the years-of-service (YOS) profile may stem from differences in initial training outcomes, occupation assignment/choice, reenlistment rates, and advancement rates.¹ It is important to

¹ Appendix A provides additional background on male/female career differences across the Navy enlisted YOS profile.

model these differences across sailors' careers and across enlistment management communities (EMCs) because they can affect manning and costs differently as the female share of accessions increases. To that end, we built a simulation model that follows every female and male sailor from accession throughout his or her career. Then, we enhanced our inventory projection model with a routine that determines the personnel cost-minimizing policies for accessions and selective reenlistment bonuses (SRBs) as the female share of accessions changes.

Approach

A prototype inventory projection model

Developing a cost-minimizing inventory projection model is an expansive and complicated effort. We had hoped to capitalize on previous inventory projection modeling efforts—either by the Navy and/or CNA—but we found that, because they were not designed to be used in conjunction with a cost minimization routine, none was appropriate for this analytical effort. Thus, we built a prototype model that can be enhanced and expanded in the future. Our stochastic inventory projection model (SIPM) prototype consists of five Navy EMCs that vary by sea intensity, differences in continuation rates for men and women, and in the female share of accessions and inventory.² We initialize the model with recent EMC inventory levels by paygrade and use recent historical data (e.g., continuation rates, total number of accessions in the EMC, female share of accessions, and SRB levels) to move individual sailors through their careers. The model simulates EMC manning by paygrade and YOS annually over the next 10 years.

The cost minimization routine

We use recruiting costs (including EBs), Department of Defense (DOD) composite rate costs (less EBs and SRBs) for enlisted Navy personnel by paygrade, and SRB costs per sailor to compute annual personnel costs of the simulated EMC inventories across the simulation period. We apply a cost minimization routine that determines the least expensive combination of accession and SRB-level policies that satisfies the desired manning targets over the entire simulation period.

The user can change the female share of accessions for an EMC to a level different from the recent historical level and rerun SIPM and the cost minimization routine to determine the least expensive accession-SRB policy combination that meets

² The five EMCs are MMSW (Machinist's Mate (Surface Warfare)), OS (Operations Specialist), AO (Aviation Ordnanceman), YN (Yeoman), and HM (Hospital Corpsman).

requirements. Note that, although a main goal for this study is to determine the cost-minimizing solution as the female accession shares change, the model can also be used to find the cost-minimizing female share of accessions (to include the possibility of a minimum female share of accessions) as manning targets change or as the desired total number of accessions changes.

Key measurement challenges

Managing the enlisted inventory to meet requirements is complicated by the fact that the personnel system includes some self-correcting mechanisms, such as advancement rates (and ultimately retention rates) in lower paygrades that increase or decrease in response to increases or decreases in loss rates in higher paygrades. It also requires some interventions when self-correcting mechanisms fail to fully close requirement gaps, including offering special and incentive pays to encourage sailors with needed skills to stay in the Navy, to go to sea sooner than (or stay at sea longer than) the sea/shore rotation plan, or to take particularly challenging or unpopular assignments. Choosing the level of detail to include in a model of the system is also challenging. We describe three key measurement challenges below.³

Defining manning targets

Manning targets should align with achieving a desired level of readiness. This should include targets for fill rates at sea and shore. It should also include a desired level of fit, or the level at which those with specialty skills described by Navy enlisted classifications (NECs) fill billets that require those skills.

Modeling all critical aspects of career flow

How sailors rotate between sea and shore tours is a key part of achieving manning targets. There are notional sea/shore rotation patterns (also known as the sea/shore flow table) that help guide sailor assignments. However, personnel records show that the actual length of sea and shore tours and, in some cases, the sequence of sea and shore tours differ from the sea/shore flow table, especially around the second sea tour and after. The challenge is whether to model the notional sea/shore flow or the historical actual sea/shore flow.

Including all costs associated with achieving manning targets

The main costs of ensuring that there are enough trained personnel in inventory to achieve manning targets are recruiting and training costs, compensation common to all sailors, and retention incentives, such as SRBs. Having a certain number of sailors in inventory, however, is not always sufficient to achieve manning targets. The Navy

³ Appendix B contains a more detailed discussion of inventory projection models.

must also ensure that sailors with the right training and experience are filling billets that require such qualifications while allowing sailors to have time ashore periodically throughout their careers. Sometimes this requires assigning sailors to sea billets at times when the sea/shore flow table suggests that they should be assigned to shore. Thus, a model that helps identify personnel-cost-minimizing solutions must consider the costs of encouraging sailors to go to sea, stay at sea longer, or fill unpopular billets when needed (i.e., non-SRB special and incentive pays).

In addition, as we model the flow of sailors through the first part of their careers (i.e., from “street to fleet”), we must consider how that flow—and the costs of the flow—may change as the total number of accessions changes. For example, an increase in the total number of accessions may increase the time awaiting training and therefore reduce the time that a sailor may spend in an operational billet before completing his or her obligation. A complete model should include this detail.

To construct a prototype model with a functioning cost-minimizing routine within the study resource constraints, we had to make some simplifying assumptions concerning the key measurement challenges:

- We use the total inventory-to-billets-authorized (INV/BA) ratio by paygrade (i.e., a fill rate) as the requirements constraint to meet. The constraint in the model only ensures that there are enough total sailors (including unassignable sailors) to meet a certain fill rate. Likewise, we have not yet incorporated fit measures into the model.
- Similarly, we do not move sailors through a sea/shore rotation plan (either notional or actual).
- We increase recruiting costs as the number of accessions increases, but the only personnel costs that we include while sailors are in pre-fleet training are the composite rate costs. We did not have other per-sailor training costs that could potentially accompany an increase in accessions, such as an increase in schoolhouse costs or time awaiting training.

Report organization

This report has four main sections. The section following this introduction describes the model and the data used in our simulation and cost minimization procedures. The second main section presents our prototype inventory project model and discusses the methodology for identifying the cost-minimizing accession and SRB policies. The third section summarizes our results. The fourth section provides conclusions and recommends for next steps for expanding the modeling effort.

SIPM: Our Prototype Model

To analyze the effects of the varying female accession percentages on an EMC’s manning and personnel costs, we developed a prototype stochastic inventory projection model. SIPM uses historical data and Navy policies to simulate the careers of individual sailors. Since our model focuses on manning and its associated costs, we require individual sailor modeling for fidelity in our projections. By using stochastic simulation, we can also capture the variability in outcomes due to uncertainty in sailor careers and their response to Navy policy (e.g., continuation behavior under different SRB levels). Simulation enables us to compare performance metrics, such as fill rates, across different manning policies by scenario (e.g., a larger versus a smaller female share of accessions). Because the model forecasts the outcomes for individual sailors, we can aggregate these into forecasts of the Navy inventory and its characteristics over time.

SIPM uses an annual periodicity over a specified time horizon for a specific EMC chosen by the user. The model initializes with the current inventory of sailors and their characteristics. Through a series of computations of gains, losses, and advancements, we predict future inventory and compute performance metrics. The data driving the SIPM come from the Enlisted Master File (EMF) and the Total Force Manpower Management System (TFMMS).

Table 1 summarizes the key model parameters and data sources, and Figure 1 illustrates the overall process flow of the simulation model. In the subsections that follow, we describe each of these computations in more detail.

Table 1. SIPM parameters and data sources

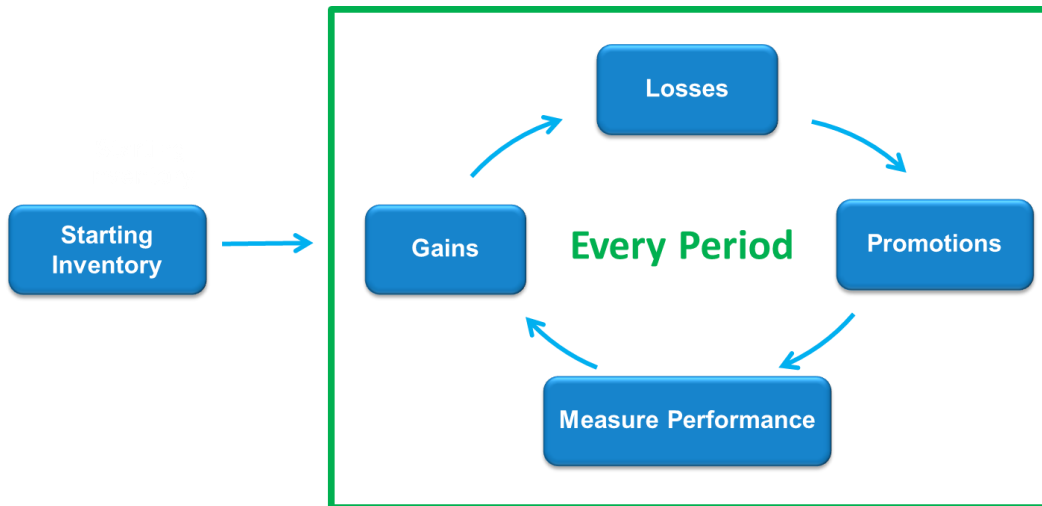
Event/data element	Data source	Calculation	Rules/how applied
Starting inventory	EMF	Counts of sailors by gender with information on paygrade (PG), YOS, and time in grade (TIG) by EMC	Initializes model runs. Although the model is initialized with sailor counts, each sailor is followed as a separate entity throughout his/her career.
Accessions	EMF	Average annual accessions by PG, gender, and EMC (Sep. 2013-June 2017)	Gains to inventory with 0 YOS at start of each year

Event/data element	Data source	Calculation	Rules/how applied
Recruiting costs^a	NIT/SERCO Replacement Cost of Sailor Model [2], [3]	No additional calculations	Costs vary only by test scores and are incurred upon accession.
Street-to-fleet time	CNA's Enlisted Street to Fleet (STF)	Average STF time based on historical relationships between enlistment program/rating promised and EMCs for those who reach the fleet	Applied to non-Professional-Apprenticeship-Career-Track (non-PACT) sailors by EMC and gender
PACT gains	EMF	Average number of gains by EMC	Gains to inventory at start of each year
Continuation rate	CNA's Enlisted Tracking File (ETF)	Average annual rate by gender, YOS, and EMC (Sep. 2013-Sep. 2015); modified based on SRB level for zones A-C for the EMC	Randomly assigned to sailors using Bernoulli distribution centered at the computed average annual rate (Pre-fleet losses are determined here.)
SRB costs	Pay Table	Monthly pay multiplied by SRB multiplier by EMC and zone	Incurred when SRB level exceeds 0 and is paid in full immediately, at reenlistment
SRB effect on continuation rate	2002 CNA study [1]	Based on EMC group	SRB level compared with historical average SRB (Sep. 2011-June 2017) to determine change in SRB that translates to change in continuation rate
Billets authorized	TFMMS	Total billets authorized by PG and EMC	Denominator in INV/BA ratios (i.e., fill rates)
Advance-ment	N/A	Promote-to-vacancy based on BA and inventory by PG	Sailor must satisfy minimum TIG for promotion eligibility; eligible sailors sorted so that those with longest TIG are promoted first
Pregnancy	EMF	Average annual rate by YOS and EMC (Mar. 2014-Mar. 2017)	Randomly assigned to sailors according to a multinomial distribution of friction statuses (i.e., pregnancy, LIMDU, and unassignable) based on computed annual rates
Limited duty (LIMDU)	EMF	Average annual rate by gender, YOS, and EMC (Mar. 2014-Mar. 2017)	Randomly assigned to sailors according to a multinomial distribution of friction statuses
Unassignable	EMF	Average annual rate by gender, YOS, and EMC (Mar. 2014-Mar. 2017)	Randomly assigned to sailors according to a multinomial distribution of friction statuses

Event/data element	Data source	Calculation	Rules/how applied
Personnel costs	DOD Composite Rate for Navy	FY18 less the EB and SRB portion	Varies by PG, but not by YOS; incurred at end of every year based on inventory distribution across PGs

^a. This category includes the cost of recruiting sailors, which varies only by the military entrance exam scores that recruits achieve and includes an estimate of recruiting effort, advertising costs, and EBs. It does not include noncomposite rate per-sailor training costs at schoolhouses or the potential costs of time awaiting training.

Figure 1. SIPM high-level process



To run a scenario using SIPM, the model requires the following data from the user:

- EMC
- Forecasting horizon (in years)
- Female percentage of annual accessions
- Annual gains to EMC
- SRB levels for zones A, B, and C

The last three required inputs are the essential parameters for scenario analysis using SIPM. SIPM can estimate the effect on manning and costs of adjusting the female accession percentage, annual accessions, and SRB levels. For example, a natural starting point for the user-required inputs are the recent historical female percentage of accessions, annual accession levels, and SRB levels for each zone for the EMC.

Starting inventory

Once the SIPM user has chosen an EMC, we initialize the model with the September 2017 starting inventory. From the EMF, we pull these data, which include the following:

- Gender
- Paygrade (PG)
- Time in current paygrade (TIG)
- Years of service (YOS)

Gains to EMC

Next, we add gains to the EMC. At the beginning of every simulated year, we add the user-specified number of sailors to the EMC. The gender distribution of the accessions is determined by the user-specified female accession percentage. Because some accessions receive credit for certain types of education and other achievements, paygrades of accessions can vary from E-1 to E-3. We base the PG distribution of the gains on historical data. The sailors enter with zero YOS and require the historically based average street-to-fleet (STF) time to complete training.

SIPM also adds Professional Apprenticeship Career Track (PACT) sailors to EMC inventories at the beginning of every year. The count, distribution, and female share of PACT sailors are allocated to EMCs based on historical data.

Note that this treatment of gains is different from a traditional inventory projection model in which losses are executed first, then advancements, and then accessions. We built our model specifically to avoid this default to accession planning. Instead, the user specifies a number of gains to initialize the scenario, but the optimization routine allows the number of gains to increase or decrease to minimize costs given the constraints of desired fill rates and female share of accessions.

Losses from EMC

We then execute losses. At the beginning of every simulated year, SIPM determines whether each sailor will become a loss to the EMC in that year. The loss rates are based on historical rates and are conditional on the sailor's EMC, gender, and YOS. The model assigns losses to sailors using random number generation and sampling from a binomial distribution.

The loss rates are adjusted for the effect of historical SRB levels. Using estimated continuation elasticities for SRBs, we can remove the effect of historical SRB levels to obtain a “natural” continuation rate (i.e., the continuation rate if there were no SRBs). Using this natural rate, we can add in the continuation rate effect of the user-specified SRB levels. Then, at the reenlistment YOS for each zone, we have dynamic continuation rates reflecting the user-desired SRB levels. Note that our estimated continuation elasticities for SRBs are the same for men and women but vary by EMC.

Advancements

After executing the gains and losses, SIPM determines the eligibility and selection of sailors for promotion. We cannot use historical promotion rates to advance sailors because these rates are conditional on historical accession and retention policies. Since the goal of SIPM is to capture the manning implications of *changes* to accession and retention policies, we require the promotion process to be flexible and adapt to the policy specified by the user.

Eligibility for promotion is based on minimum TIG requirements for the next PG. Conditional on eligibility, selection for promotion is modeled one of two ways, depending on the PG:

1. Promotion to E-2 and E-3 is automatic after 9 and 18 months of service, respectively.
2. Promotions to E-4 through E-9 are based on fill-to-vacancy.

Automatic promotion to E-2 and E-3 ensures that, when a sailor has achieved the minimum TIG required for eligibility, the sailor will immediately be promoted to E-2 or E-3. This reflects the common practice of immediate and nearly guaranteed promotion during early apprentice years in the Navy.

Promotions to E-4 through E-9 are handled by fill-to-vacancy, which uses the losses of sailors in a PG to determine how many eligible sailors to promote. Using a fill-to-vacancy approach guarantees high manning percentages for these paygrades (unless there is a shortage of TIG eligible sailors for promotion) but does not allow for excess (i.e., manning above 100 percent).

Costs

To maintain the inventory, the Navy incurs the following personnel costs: (1) recruiting costs, (2) SRBs, and (3) other sailor compensation that we capture using the composite rate (adjusted to exclude the cost of EBs and SRBs). Upon accession, SIPM generates a

recruiting cost per sailor. This cost is based on test scores and, therefore, is specific to the EMC.⁴ We use three tiers of costs for recruits found in the Cost of a Sailor Model, where each EMC is associated with a low, middle, or high recruiting cost.

At each zone's *typical* reenlistment YOS, SIPM counts the number of sailors who continue to the next year.⁵ This count is then multiplied by the SRB cost (given the SRB level) for the zone. Finally, at the end of each year, the inventory count by PG is multiplied by the appropriate composite rate (less the EB and SRB portion) to determine the total personnel costs.⁶

Performance measures

At the end of the run for each EMC, SIPM reports the following performance measures:

1. Inventory by PG and FY
2. Average TIG and YOS by PG and FY
3. Gains, losses, and promotions by PG and FY
4. Fill rates (percentage of BA) by PG and FY
5. Total costs (including recruiting, SRB, and composite rate) by FY

Using these metrics, the user can compare runs of the model under different parameter values (e.g., higher versus lower female accession percentage) to estimate the manning and cost implications of different policies for the EMC. These implications include not only the immediate metrics of fill and total cost but also second-order effects, such as experience (as measured by average YOS by PG) and promotion rates.

⁴ Recall that we use the recruiting costs described in [2] and [3], which include estimates of recruiting effort, advertising, and EBs. Because these three factors may be higher for higher-scoring recruits, our recruiting costs are correlated with the length of required training.

⁵ We assume that sailors make reenlistment decisions at 4, 8, and 12 YOS. This is a simplifying assumption that can be addressed in future versions of the model.

⁶ We received cost inputs to the FY 2018 composite rates from the Resource Management Division in the Office of the Deputy Chief of Naval Operations (Manpower, Personnel, Training, and Education) (MPTE) (OPNAV N10). The level of detail in the data provided allowed us to exclude EBs and SRBs from the composite rate calculations.

Cost-Minimizing Accession and SRB Policy

Using the performance metrics computed by SIPM (i.e. costs and manning), we can use optimization techniques to determine the cost-minimizing accession and SRB levels that achieve a target fill rate (e.g., percentage of BA), given a female accession percentage.

Optimization model

The goal of the optimization model is to determine the cost-minimizing combination of accession and SRB policies that achieves a target fill percentage for an EMC. The optimization model has three parts: (1) decision variables, (2) constraints, and (3) the objective function. Given the constraints, the model determines the values of the decision variables, which minimize the objective function.

Decision variables

We consider three sets of decisions that can be varied over the planning horizon:

- Total annual accessions
- Female accession percentage
- SRB levels

Let the accession vector \mathbf{a} be the total number of accessions for the EMC in each FY of the planning horizon T . While the Future Years Defense Program (FYDP) goes through FY 2023, we use a planning horizon of FY 2030 to prevent boundary effects from influencing the optimal policy (e.g., accessing too few sailors in FY 2023 to satisfy future target fill rates since it would be the lowest cost strategy). Combined with the female accession percentage f , we then know the number of male and female accessions in each FY. Let \mathbf{B} be a matrix of SRB levels for every combination of FY and reenlistment YOS for zones A, B, and C.

Constraints

Our model incorporates four sets of constraints that place lower and upper limits on (1) annual accessions, (2) female accession percentage, (3) SRB levels, and (4) manning levels. In our notation, underbars and overbars denote lower and upper limits for the corresponding variable, respectively. The following descriptions provide more detail of the four constraints in the SIPM.

1. The lower and upper limits for annual accessions, $[\underline{a}, \bar{a}]$, are used to prevent undesirable variation in the accession policy over the planning horizon. The limits are based on the historical annual accessions and allow up to a maximum of a 10 percent change in accessions per year.
2. The lower and upper limits for the female accession percentage, $[\underline{f}, \bar{f}]$, control the allowable change in the female share over the planning horizon. While f is a decision variable, for our optimization runs, we fully constrain f such that the lower and upper limits are the same and equal to the user-specified female accession percentage. Mathematically, this means that $f = \underline{f} = \bar{f}$. By fully constraining this variable, we are fixing the female accession percentage, preventing the optimization model from changing it. This reflects our desire to understand how changes in accessions and SRBs affect manning and cost, given a female accession percentage.
3. The lower and upper limits for the SRB levels by zone, $[\underline{B}, \bar{B}]$, are set to encompass actual historical SRB levels. However, unlike the limits for annual accessions, we do not artificially limit the minimum and maximum SRB level to prevent a large change relative to historical SRBs. If we confine the SRB levels to a small range, we cannot properly use SIPM to determine trade-offs between accessions and retention for achieving the target manning level.
4. For manning levels, the model allows the user to specify manning targets by paygrade. To compute manning levels, we first must compute sailor inventory I_t over time, as a function of our decision variables. We let $g(I_0, \mathbf{a}, \mathbf{f}, \mathbf{B})$ denote the internal computations of SIPM, which are based on the initial inventory I_0 , the accession policy \mathbf{a} , the female accession percentage \mathbf{f} , and the SRB levels \mathbf{B} . Using SIPM as our inventory calculator, we can compute manning levels (e.g., fill rates), which we denote as $\mathbf{m}(I)$ with corresponding lower and upper limits, $\underline{\mathbf{m}}$ and $\bar{\mathbf{m}}$, respectively. In a cost-minimizing optimization, the lower limit for manning $\underline{\mathbf{m}}$ is the same as the target manning (e.g., 90 percent fill) since manning levels above the lower limit generate higher personnel costs (i.e., worsen the objective function). Upper limits prevent excess manning but may lead to infeasibility problems if the upper limit is set too low. For our optimization runs, we do not set a maximum manning level (i.e., $\bar{\mathbf{m}} = \infty$).

Objective function

The goal of the optimization model is to find the set of decision variable values (accessions and SRB levels) that minimize the total personnel cost over the planning horizon. Total personnel cost includes recruiting costs, composite rate costs (less EBs and SRBs), and SRBs. We denote the total cost over the planning horizon as $C(\mathbf{I}, \mathbf{B})$ because it is a function of the inventory \mathbf{I} and the SRB levels \mathbf{B} .

Since the costs are accrued over time, we compute the present value of the total costs by discounting future costs:

$$C(\mathbf{I}, \mathbf{B}) = \sum_{t \in T} \frac{F_t(\mathbf{I}, \mathbf{B})}{(1+i)^t}$$

where $F_t(\mathbf{I}, \mathbf{B})$ is the total cost in each year of the planning horizon $t \in T$ and i is the discount factor (2.9 percent). This procedure matches the guidance set forth in the Office of Management and Budget's Circular A-94 memorandum for program evaluation. Naturally, we can set $i = 0$ to optimize the unadjusted total cost.

Model formulation

We formally write the optimization model as:

$$\begin{aligned} & \text{Min}_{\mathbf{a}, \mathbf{f}, \mathbf{B}} C(\mathbf{I}, \mathbf{B}) \\ & \text{s.t. } \mathbf{I}_t = g(\mathbf{I}_0, \mathbf{a}, \mathbf{f}, \mathbf{B}) \quad \forall t \in T \\ & \quad \mathbf{m}(\mathbf{I}) \in [\underline{\mathbf{m}}, \overline{\mathbf{m}}] \\ & \quad \mathbf{a} \in [\underline{\mathbf{a}}, \overline{\mathbf{a}}] \\ & \quad \mathbf{f} \in [\underline{\mathbf{f}}, \overline{\mathbf{f}}] \\ & \quad \mathbf{B} \in [\underline{\mathbf{B}}, \overline{\mathbf{B}}] \end{aligned}$$

By virtue of SIPM's inventory calculation g and the interaction between the accession size and the female percentage, the objective function and constraints are nonlinear in the decision variables. As such, the model is classified as a nonlinear optimization problem, which is challenging to solve to global optimality. To address the problem, we employ the grid randomized iterative descent technique.

Grid randomized iterative descent (GRID)

Common solutions to nonlinear optimization problems require exploiting knowledge of the gradients of the model objective function and constraints to incrementally change the decision variables to improve the objective function value (in our case,

lowering total cost). For our formulated model, these gradients inform the rate of change for personnel cost and manning levels given a small change in accession and retention decisions. By iteratively inducing small changes in the accession and retention decision variables, we continue toward improved personnel costs until no improvement can be found. Unfortunately, given the nature of SIPM (a simulation tool with complex logic and equations), we cannot readily compute a priori gradients that would be needed for leveraging the common solution methods to nonlinear problems (e.g., Newton's descent). For example, without a closed-form equation, we do not know how the personnel costs will change if zone A SRB levels are changed from 6 to 7 because of the retention interactions with advancements. However, by running SIPM with a zone A SRB level of 6, and again with an SRB level of 7, we can compute the resulting change in personnel costs.

To solve the model formulation for determining cost-minimizing accession and SRB policies that achieve a target manning level, we developed a GRID search. GRID is an approximation technique that leverages the principles of nonlinear optimization algorithms without requiring a priori calculations of gradients. Instead, we exploit the monotonicity of the objective function and inventory levels (with respect to accession and retention) to make improved steps toward a lower personnel cost. Fundamentally, personnel costs increase as accessions and SRBs increase. Hence, we know that decreasing accessions and/or decreasing SRBs will decrease total personnel costs (i.e., improve the objective function). Similarly, we know that increasing accessions and SRBs will also increase total manning. Therefore, decreasing accessions or decreasing SRBs will reduce manning and lead to potentially infeasible solutions (i.e., manning levels below the specified manning target). We do not know, however, the a priori rate of change in costs and manning when changing accession levels and SRB levels. This requires running SIPM iteratively to numerically determine the objective function and manning values. Appendix D contains additional detail on our use of GRID.

Note that the algorithm is not guaranteed to find globally optimal solutions. Our discretization and reduction of the decision space, combined with the stochastic output of SIPM and the random selection of descents, leads to the identification of locally optimal solutions. While we cannot guarantee global optimality, we can always compare the identified solution against the current accession and SRB policies to determine if the GRID solution is a significant improvement.

Bi-level optimization

Minimizing total costs for each EMC when constraints are placed on the female share of accessions and manning levels can help the Navy move toward more efficient EMC-specific accession and SRB policies. Navy personnel planners, however, also typically face all-Navy constraints on the maximum number of accessions and/or the total SRB budget as well as a desired female share of total accessions. To incorporate these

constraints into a broader cost minimization problem, we employ a second, or bi-level, optimization routine.

The bi-level optimization goal is to distribute the female accession percentage across the EMCs such that the Navy meets its manning goals for every EMC, achieves the desired target female share of accessions, and does not exceed constraints on the maximum number of accessions and the total SRB budget while minimizing total personnel costs. For example, suppose that the Navy has a target overall manning, target female percentage accession goal (represented by \tilde{f}), a maximum SRB budget (represented by \tilde{s}), and a maximum accession limit \tilde{a} . The problem becomes how to distribute female accessions across the EMCs to meet all constraints at minimum cost. In general, the optimal solution will be to distribute female accessions to the EMC that can most readily increase female accessions without requiring large additional costs to satisfy manning targets.

Mathematically, the problem can be solved as a bi-level, nonlinear knapsack optimization problem with the following form:

$$\begin{aligned} & \min_f \sum_{i \in I} c_i^*(f_i) \\ & \frac{\sum_{i \in I} f_i a_i^*(f_i)}{\sum_{i \in I} a_i^*(f_i)} \geq \tilde{f} \\ & \sum_{i \in I} s_i^*(f_i) \leq \tilde{s} \\ & \sum_{i \in I} a_i^*(f_i) \leq \tilde{a} \\ & f_i \in [\underline{f}_i, \bar{f}_i] \quad \forall i \in I \end{aligned}$$

This formulation borrows the notation of the single EMC cost minimization, but with additional notation to specify the EMC. Each EMC is denoted by i and has a female percentage accessions f_i , optimal accession plan a_i^* , optimal SRB cost s_i^* , and optimal total cost c_i^* . The optimal accession, SRB cost, and total cost depend on the female accession percentages and are derived by solving the single EMC cost minimization model described earlier. The female accession percentages are the decision variables for this Navy optimization model and are bounded by minimum and maximum per EMC, $f_i \in [\underline{f}_i, \bar{f}_i]$.

The bi-level optimization provides three policy levers to study when seeking to minimize the Navy's total cost: (1) change the target Navy percentage of female accessions \tilde{f} , (2) change the maximum SRB budget \tilde{s} , and (3) change the maximum annual accessions \tilde{a} . Each of these policy levers corresponds to changes in the

constraints of the bi-level optimization model. By increasing the maximum SRB budget, by increasing the maximum annual accessions, and/or by decreasing the target aggregate percentage of female accessions, the bi-level problem becomes less constrained and total costs to the Navy may decrease.

Results

Using the simulation and optimization approaches described earlier, we considered case studies to assess the effect of changes to female accession policy on individual Navy EMCs. For each EMC, we parameterized and ran SIPM, using the GRID technique to determine the cost-minimizing solution. Next, we used the bi-level optimization approach to model a mini-Navy (i.e., a five-EMC Navy) to understand the policy levers that leadership can pull and their likely effect on manning and personnel costs.

The methodology presented earlier is a new modeling venture for CNA. The results presented in this section are preliminary and are intended primarily as illustrations of the model's capabilities, the types of questions it can answer, and the decision support it could provide.

Individual EMCs

To understand the model capabilities, we developed SIPM and used GRID for five EMCs: Hospital Corpsman (HM), Operations Specialist (OS), Machinist's Mate (Surface Warfare) (MMSW), Aviation Ordnanceman (AO), and Yeoman (YN). We selected these EMCs because they vary by recent average female accession and inventory shares and sea intensity. Because it is so new, SIPM needs some nuanced development/adjustment each time we include new EMCs. In this section, we present the results for HM and AO only, but we use results for all five EMCs as input to the bi-level optimization. In addition, for the AO EMC, we performed deeper dives to understand how model assumptions affect our optimal solution. See Appendix E for those results.

HM

For HM, we use the SIPM and the GRID search to find cost-minimizing solutions when the female share of accessions is increased by 5 percentage points (from 31 to 36 percent) and the desired fill rate at every paygrade from E-3 through E-6 is 95 percent. The starting conditions include the following:

- FY 2014–2016 female share of accessions for HM: 31 percent
- FY 2014–2016 average annual accessions: 2,672

- Current average SRB levels across NECs for zone A, B, C, respectively: 3, 3, 2
- 2016 average (E-3 through E-6) BA fill rate: 92 percent
- 2016 E-3 and E-4 average YOS: 1.9 and 3.4 years, respectively
- Female and male survival to YOS 6: 45.9 and 45.3 percent, respectively

We began by simulating future inventories and calculating personnel costs when there is no change in the starting conditions for accessions or SRBs. Then, we considered three scenarios: (1) allow change in SRBs only, (2) allow change in accessions only, and (3) allow changes in both SRBs and accessions. For each scenario, we required a 95 percent fill rate for each paygrade, E-3 to E-6, in FY 2026. Unlike the three scenarios, our first simulation (i.e., no change in policy) is not required to achieve the 95 percent fill rate. For this reason, the no-change policy is best viewed as a baseline, not necessarily as the choice of a feasible solution for a decision-maker.

We used the GRID technique to search for the best solution in each scenario given the variables allowed to change and the fill rate constraint. The first two scenarios are designed to reflect the current practice of considering policy levers in isolation—that is, finding the cost-minimizing solution to a forecasted fill rate shortfall by changing either accessions or SRBs but not both. The last scenario shows the potential reductions in cost when both policy levers are used simultaneously. Table 2 presents a summary of the results.

Table 2. HM results for an increase in female accession shares (31% to 36%)

Baseline and scenarios	Annual accession plan		SRB level by zone			FY26 BA fill rates by paygrade (E-3/E-4/E-5/E-6)	FY26 total BA fill rate (E-3+)	Total cost (2017–2026)	% change in cost (measured from baseline)
	FY17	FY18+	A	B	C				
Baseline ^a	2,672	2,672	3	3	2	61/100/100/100%	85%	\$18.33B	NA
(1) SRBs	2,672	2,672	8	5	0	95/100/100/100%	98%	\$20.58B	12%
(2) Accessions	2,672	3,206	3	3	2	96/100/100/100%	98%	\$19.69B	7%
(3) SRBs & accessions	2,538	3,553	1	0	0	95/100/100/100%	98%	\$18.96B	3%

Source: CNA calculations from SIPM and GRID search.

^a The baseline (i.e., no-change), a projection using recent historical data as the starting point, is not considered feasible because the E-3 fill rate is only 61 percent. We present the option only as a reference point for the scenarios, all of which are feasible.

The third scenario—in which accessions and SRBs are allowed to change—produces the lowest cost solution. Note, however, that the optimal solution increases accessions

in FY 2018 and future years by 33 percent compared to FY 2014–2016 average accessions (i.e., 2,672 to 3,553). By contrast, the first scenario represents an extreme constraint on accession growth (i.e., no growth allowed). In this scenario, SRBs are increased substantially for zones A and B, leading to an additional \$1.6 billion in total costs over 10 years compared to the optimal solution (\$20.58 billion compared with \$18.96 billion). Finally, the second scenario, in which SRB levels are held constant at recent historical levels and accessions are allowed to change, yields a 10-year total cost of \$19.69 billion, a \$0.73 billion increase from the optimal solution.

AO

We repeat these scenarios for AOs by again modeling a 5 percentage point increase in the female share of accessions, in this case from 35 to 40 percent. The starting AO conditions include the following:

- FY 2014-2016 average female share of accessions: 35 percent
- FY 2014-2016 average annual accessions: 924
- Current average SRB levels: 1, 1, and 0 for zone A, B, C, respectively
- 2016 average (E-3 through E-6) BA fill rate: 95 percent
- 2016 E-3 and E-4 average YOS: 2.4 and 4.2 years, respectively
- Female and male survival to YOS 6: 40.3 and 43.2 percent, respectively

We summarize the model results in Table 3.

Table 3. AO results for an increase in female accession shares (35% to 40%)

Baseline and scenarios	Annual accession plan		SRB level by zone			FY26 BA fill rates by paygrade (E-3/E-4/E-5/E-6)	FY26 total BA fill rate (E-3+)	Total cost (2017–2026)	% change in cost (measured from baseline)
	FY17	FY18+	A	B	C				
Baseline ^a	924	924	1	1	0	86/100/100/100%	95%	\$6.52B	NA
(1) SRBs	924	924	7	0	1	96/100/100/100%	98%	\$7.00B	7%
(2) Accessions	961	999	1	1	0	95/100/100/100%	98%	\$6.68B	2%
(3) SRBs & accessions	906	1,051	0	0	0	97/100/100/100%	99%	\$6.62B	1%

Source: CNA calculations from SIPM and GRID search.

^a The baseline (i.e., no-change), a projection using recent historical data as the starting point, is not considered feasible because the E-3 fill rate is only 86 percent. We present the baseline only as a reference point for the scenarios, all of which are feasible.

When the female share of accessions for AOs is increased by 5 percentage points, the optimal solution is to increase accessions and decrease SRBs relative to the baseline, as displayed in the third scenario. The estimated personnel costs over FY 2017 through FY 2026 in the third scenario is \$6.62 billion, which is about 1 percent less expensive than the second scenario (allow changes in accessions only) and nearly 6 percent less expensive than second scenario (allow changes in SRBs only).⁷

The all-Navy bi-level optimization

Imposing all-Navy constraints

Finding cost-minimizing accession-SRB plans for a given share of female accessions for each EMC is useful, but policy-makers also need to know how to optimize female accessions across all five EMCs when the total number of accessions and the total SRB budget are constrained. We use bi-level optimization to derive solutions to this all-Navy (i.e., the five EMCs) constrained problem. Given bi-level optimization's primary all-Navy constraints—(1) target female accession percentage, (2) minimum desired manning levels, (3) maximum annual total accessions, and (4) maximum SRB budget—the optimization seeks to distribute the female accessions across the EMCs to minimize the Navy's total personnel cost. Each EMC then reacts to the assigned female accession and acts optimally, per the SIPM and GRID solutions. The solution to the bi-level optimization is multifaceted, composed of the optimal distribution of the female accessions and the optimal accession and SRB policies of each EMC.

Practically, we took the following steps.

- We used GRID to determine the cost-minimizing accession and SRB policies for each of our five EMCs for different female accessions.⁸ That is, we found the cost-minimizing accession and SRB policies for each EMC for five different assumptions for the female share of accessions: the average historical female share of accessions for the EMC and plus and minus 5 and 10 percent of the average historical female share. We then interpolated between the five results to expand our first-stage optimal solution set to 40 for each of the five EMCs.

⁷ To explore whether the GRID search process produces intuitively consistent results, we also find optimal results for AO when recruiting costs and starting inventory vary. See Appendix E for details.

⁸ As a reminder, the five EMCs are MMSW, OS, AO, YN, and HM. Although we present results for HM and AO only in the previous subsection, we compute results for all five EMCs and use them in this second stage of optimization.

- We combined the cost-minimizing accession and SRB policies for each EMC under the various female share of accession scenarios and used bi-level optimization to determine the best distribution of female accessions across these five EMCs.

Note that the second stage optimization for our five-EMC Navy is not particularly complex. That is, in the second stage of our bi-level routine, we select from the optimal results generated by varying the female share of accessions for each of the five EMCs such that the overall cost-minimizing female accession allocation across the five EMCs meets all of the problem constraints.

It is clear, however, that a relatively simple exercise for a five-EMC Navy with a small range of variation in the female share of accessions becomes more challenging to solve when dozens of EMCs and many possible variations in the share of female accessions are involved. Moreover, to be especially useful as a decision support tool, the model must also be able to find optimal solutions quickly and efficiently when we change any of the problem constraints (the desired fill rate, the total number of accessions, the total SRB budget, etc.). Our prototype model provides such an approach.


Varying the all-Navy constraints

The Navy typically faces total accession and total SRB budget constraints, but, over the course of a personnel planning cycle, these constraints might change. To visualize how changes in the accession and SRB budget constraints affect minimum personnel costs, we developed a heat map showing the expected minimized total cost for each combination of maximum accessions and maximum SRB budget. Figure 2 shows the costs when the target female accession percentage is 33 percent.

Each cost corresponds to the optimal distribution of the female accessions across the five EMCs, given the constraint on annual accessions and the SRB budget. The cost is the sum of the costs for the individual EMC's optimal accession and retention policies given the distributed female accession share. Dark green areas in the figure indicate lowest total cost, while dark red areas indicate highest total cost across all five EMCs.

For example, consider the point on the graph where the maximum SRB budget is \$2.10 billion and the maximum annual accessions is 5,650. The point is in the orange area of the graph, where the minimized total personnel costs are approximately \$41.93 billion. If the Navy were allowed to increase accessions, the total personnel costs across all five EMCs decrease (moving from the orange area to the green area). Specifically, as the accession constraint increases from 5,650 to 5,750, the total annual personnel costs decrease to approximately \$41.80 billion.

Figure 2. Heat map of costs for target total female accession percentage



White areas in the figure denote infeasible combinations of constraints on accessions and SRB budgets. An infeasible combination occurs when the female accessions cannot be distributed across the five EMCs in such a way that all of the EMCs achieve their target fill rates without requiring higher accessions or SRBs than allowed. For example, consider the point on the graph where the maximum SRB budget constraint is \$1.90 billion and the maximum annual accessions constraint is 5,650. This point falls well into the white area of the graph, indicating that it is not possible to achieve 95 percent fill rates for all paygrades (E-3 through E-6) across all five EMCs when the target all-Navy female share of accessions is 33 percent, maximum annual accessions are constrained to 5,650 and the maximum SRB budget is \$1.90 billion. More generally, this example helps demonstrate that leadership can use the heat map to identify those

combinations of total accessions and SRB budgets that do not allow for healthy manning across all EMCs.⁹

The nature of constraints in optimization models leads to the phenomenon in our model that, when the Navy is allowed to access more sailors or pay higher SRBs, the Navy can potentially find lower cost solutions. By relaxing the constraints, the bi-level optimization can prescribe better solutions. Thus, the heat map can be used to identify cost-saving solutions relative to current policy. For instance, let's return to our first example in which the current annual accessions and SRB budget are 5,650 and \$2.10 billion, respectively, which is well within the feasible region. From this point in the figure, we can either allow higher annual accessions or increase the maximum SRB budget and observe the change in total cost. Since there are feasible solutions under either of these increases, we know that the optimal distribution of female accessions with the new constraints will allow the EMCs to achieve their target manning levels.

For our mini-Navy, it appears that the most significant changes in total cost come from increasing the maximum annual accessions (i.e., moving up the vertical axis). Changes in maximum SRB (i.e., moving to the right on the horizontal axis) have minimal effect on total costs for a given maximum annual accession plan. Note, however, that the current version of the model does not capture all costs and constraints, including recruiter and schoolhouse capacity and costs associated with ensuring desired manning levels at sea. Therefore, we would not advise informing policy using the current version of the heat map.

The axes of the heat map reflect the maximum accessions and SRB budget, not the actual accessions or SRB costs incurred to obtain the optimal solution. On one hand, we see that, for high maximum annual accessions, increases in the maximum SRB budget (left to right on the horizontal axis) do not change the optimized total cost (i.e., it is always dark green). This is because the optimal solution does not require the additional SRB budget (i.e., optimal SRBs for each EMC do not need to be increased). On the other hand, for a fixed maximum SRB budget, we see significant changes in the total cost (gradient of color from red to green) as we increase the maximum annual accessions. This implies that, as the constraint on annual accessions is relaxed, the optimal solution is to recruit more sailors.

Figure 2 is one of three heat maps we can generate from the bi-level optimization solutions. The bi-level optimization has three constraints; however, in a two-dimensional figure, we can show changes in total costs only because two constraints

⁹ The bi-level optimization results are based on stochastic outcomes from SIPM. Therefore, choosing maximum annual accessions and SRB budgets along the feasibility boundary (i.e., adjacent to the infeasible region) puts the Navy at risk of not achieving its manning targets. Further modeling and analysis is required to perform robust stochastic optimization that incorporates the variance of the SIPM outcomes into the bi-level model's objective function. Until such upgrades are completed, we recommend constraint selection well inside the feasible region.

are allowed to change while the third is fixed. Therefore, we can create a heat map where the maximum annual accession is fixed and show how cost changes as the maximum SRB budget and target female accession percentage change. Similarly, we could fix the maximum SRB budget and show how total costs change as the other two constraints change. The type of heat map to use for policy analysis will depend on which constraints are considered by policy-makers to be changeable and which are assumed to be fixed.

Conclusions and Recommendations

Inventory projection models are typically used to determine if manning targets can be met in the future. To date, however, these models have not explicitly been developed for finding the cost minimizing accession and SRB policies that achieve the desired manning targets. The result is that Navy leadership may not be receiving all the decision support necessary to make efficient policy decisions.

To help close this decision support gap, we developed a prototype SIPM to simulate the careers of female and male sailors in five EMCs. We applied a cost-minimizing routine (i.e., a GRID search) to find the least expensive combination of accession and SRB levels that achieve a desired fill rate for a given female share of accessions for each of the EMCs. We then applied a bi-level optimization routine to find the cost-minimizing accession and SRB levels that achieve desired fill rates across all five EMCs when the total number of accessions and the total SRB budget are constrained.

Our prototype SIPM and cost-minimizing routine demonstrate that, compared with considering changes to only one policy lever at a time, considering changes to both accession and SRB levels can be optimal (i.e., cost minimizing) to achieve desired fill rates. The model also allows us to produce cost-minimizing solutions for each of the five EMCs quickly and efficiently when we vary the female share of accessions.

We also showed that, when there are also constraints on the total number of accessions and total SRB budget, we can apply a bi-level optimization routine to allocate the desired female share of accessions optimally across all five EMCs. Moreover, our modeling strategy can also help determine if relaxing either the maximum accession constraint or the maximum SRB budget constraint can minimize costs within the feasible set of solutions. In addition, the bi-level optimization approach should be able to find solutions to the true all-Navy problem (i.e., feasible, cost-minimizing accession and SRB plans for *all* EMCs when the total number of accessions and the total SRB budgets are constrained).

Although our prototype model shows promise, it should not be used in policy development yet. More work needs to be done to improve our costing, inventory distribution (e.g., sea or shore), and fill rate methodologies. Rather, we documented these results to demonstrate that, to ensure the full set of possible personnel cost-minimizing solutions, we must allow for simultaneous adjustments to accession and SRB levels. In addition, it is possible to optimally allocate female accession shares

across EMCs when constraints on total accessions and the total SRB budget are imposed.

This demonstration of our prototype model lends support for development of an all-Navy enlisted SIPM to which we can employ these cost minimization techniques. To fully employ the model to support decision-making, however, we must address the following underlying assumptions and modeling choices in the prototype:


- Expand the definition of meeting requirements
 - **Ensure that the inventory can meet a desired level of both sea and shore requirements.** The current model limits the set of feasible cost-minimizing solutions to those where total inventory achieves a certain percentage of total requirements for each paygrade (e.g., the fill rate is 95 percent for each paygrade from E-3 through E-6). The next iteration of the model should limit solutions to those where a desired fill rate is achieved for both sea and shore requirements.
 - **Ensure that fit is achieved.** Similar to the preceding subbullet, we should include a desired manning target for specialty skills described by Navy enlisted classifications.
- Refine the modeling of career progression
 - **Adopt a modeling strategy to move sailors between sea and shore tours.** A critical improvement to modeling the career flow would be to move sailors between sea and shore tours. The challenge is that personnel records show that the actual length of sea and shore tours and, in some cases, the sequence of sea and shore tours, differ from the notional sea/shore flow tables, especially during and after the notional second sea tour. The adopted strategy for sea/shore rotations must consider this notional/actual disconnect.
 - **Model more precisely the changes in time-to-train before reaching the fleet as the number of accessions increases.** As a proxy for this measure in the current model, we increased recruiting costs as the number of accessions increased. However, this ignores potential increases in time to complete pre-fleet training, which can change the level of fleet manning.
 - **Model more Navy career events as stochastic.** These events should include becoming unassignable to a full duty billet for periods of time (e.g., limited duty) or to billets based on family needs or other factors that prescribe a geographic location or other distributional consideration.

- **Include all costs associated with filling the expanded definition of requirements and the improved career flow modeling.** Specifically, we suggest adding the following costs.
 - Costs associated with encouraging sailors to go to sea and/or stay at sea longer
 - Costs associated with time awaiting training and/or the cost of expanding the limits of the current training system

Appendix A: Background for Modeling Effort

In general, women do not continue in the Navy at rates as high as those of men, but the size of the difference depends on the EMC and YOS. For example, Figure 3 displays the entire enlisted Navy continuation rates for men and women averaged across FY 1991 to FY 2015. Note that one-year continuation rates at YOS 5, YOSs 7-8, YOS 14, and YOSs 20-21 are substantially lower for women than for men. These YOSs represent common reenlistment points (YOSs 14 and earlier) and becoming eligible for retirement (YOS 20).


Figure 3. All-Navy one-year and cumulative continuation rates by gender



Source: CNA calculations from the Enlisted Master File (EMF), 1991 to 2015. Excludes EMCs in Navy special warfare, submarines, diver, and diver specialties.


Essentially, each curve in Figure 3 represents a weighted average of continuation rates across all EMCs for that gender. To illustrate how these curves may be different for different EMCs, we show the results for the AO and YN EMCs in Figure 4 and Figure 5, respectively.

Figure 4. AO one-year and cumulative continuation rates by gender



Source: CNA calculations from the EMF, 1991 to 2015.

Figure 5. YN one-year and cumulative continuation rates by gender



Source: CNA calculations from the EMF, 1991 to 2015.

In Figure 4, note the large difference in the YOS curves for men and women beginning around YOS 5 and extending to about YOS 12. By contrast, as Figure 5 illustrates for the YN EMC, there is only a modest difference in the YOS curves for men and women throughout 20 YOS, indicating that, throughout their Navy careers, continuation rates for YN men and women are close.

Appendix B: Recent Inventory Projection Models

Recent modeling efforts for Navy enlisted inventory/manning include the Inventory Projection Model (IPM) developed by the Total Force Manpower, Training, and Education Requirements Division (OPNAV N12). The IPM was developed in part to help with accession and manpower planning as the female share of recruits increased. The model is deterministic, and the enlisted inventory flows through the model in groups (or cells) defined by EMC, YOS, and gender rather than as individual sailors. As we noted in [4], the IPM inputs include recent historical female and male shares of accessions for each EMC, and the model simulates EMC-specific career progression by gender. In each period, sailors are assigned to sea or shore or they are unassignable. One key underlying assumption in the model is that sailors follow the sea/shore rotation patterns described in the sea/shore flow table rather than the actual rotation patterns seen in recent historical personnel data records.

The IPM also includes enlisted requirements by EMC grouped into three paybands: apprentice (E-1 through E-4), journeyman (E-5 and E-6), and supervisor (E-7 through E-9). As each accession cohort progresses through the simulation, the IPM produces sea and shore inventory-to-requirements ratios by EMC and payband. The model can project changes in inventory-to-requirements ratios at sea and shore as the total number of accessions, the share of accessions by EMC, and/or the female share of accessions in each EMC is changed.

The Navy Enlisted Management Model (NeMMo) is another recent effort to simulate the enlisted inventory/manning and provides even more detail than the IPM. As the authors describe in [5], NeMMo is a deterministic model that advances sailors through their careers in groups defined by EMC, YOS, paygrade, gender, time remaining on contract, and sea/shore assignment. To describe requirements, NeMMo uses enlisted program authorizations (EPA).

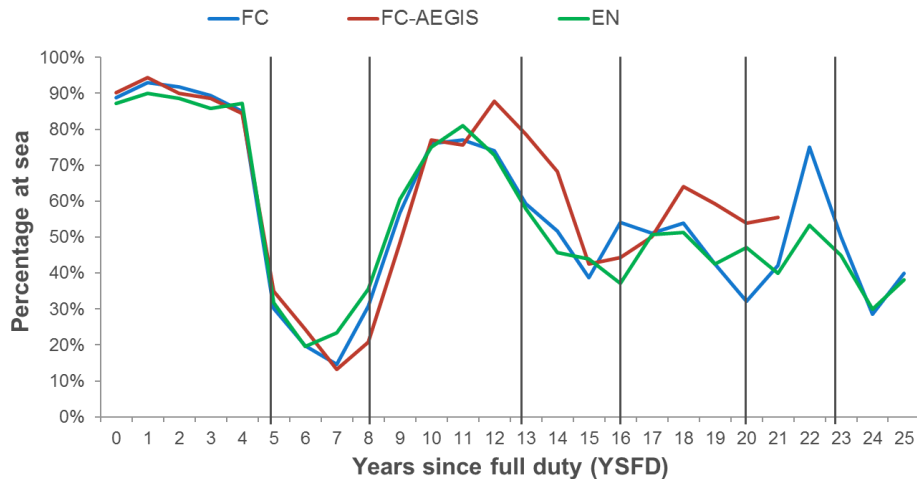
According to [5], NeMMo uses actual historical data to distribute sailors to sea and shore billets rather than relying on the notional sea/shore flow tables. The authors note, however, that they are exploring possibilities to use the best combination of actual and notional sea/shore flows in future versions of the model.

NeMMo forecasts inventory losses and gains as well as advancements. This produces key model output, including inventory counts by sea and shore by year, YOS, paygrade, gender, and time remaining on contract, as well as sea and shore inventory-to-EPA ratios by year and paygrade.

Appendix C: Notional Versus Actual Sea/Shore Flows

A recent CNA study presented results that illustrate how actual and notional sea/shore flows can differ [6]. We display those results in Figure 6. The authors use data from the end of September 2016 EMF to show the percentage of sailors in the Fire Controlman (FC), Fire Controlman (Aegis) (FCA), and Engineman (EN) EMCs that were on sea duty at each year since reaching full duty (YSFD). They compare the actual percentages to the notional sea/shore flow described in [7]. Sailors in the FC, FCA, and EN EMCs have a 60-month first sea tour followed by a 36-month shore tour, a 60-month sea tour, a 36-month shore tour, and a 48-month sea tour, which they indicate by the vertical lines in the graph.

Figure 6. Percentage of sailors on sea duty by years since full duty



Source: Adapted from Figure 6 in [6].

As Figure 6 shows, the actual percentage of sailors on sea duty during the notional first sea tour (i.e., YSFD 0-5) is relatively high—about 90 percent for much of the period and in each of the EMCs. Likewise, the actual percentage of sailors on their first shore tour compared with the notional first shore tour (YSFD 5-8) aligns relatively well.

During the second sea tour, however, the actual percentage of sailors on sea duty begins to diverge from the notional period (YSFD 8-13). The actual percentage of sailors on sea duty and the notional sea/shore flow continue to diverge through 20 or more YSFD.


Appendix D: Additional Details on GRID

To begin the iterative process, we initialize our algorithm with the lowest cost solution and check its feasibility using SIPM. If it is feasible, we have found the cost-minimizing solution. Otherwise, we start with the highest cost solution. If that solution is not feasible, the problem is not feasible under the given constraints and parameters, and no solution can be found. If the highest cost solution is feasible, we iteratively descend along a direction that we know will decrease costs but also decrease the chance of being a feasible solution. Since we do not know which decision variable to change to achieve lowest cost, we randomly choose a decision variable on which to make our descent (i.e., changes). Descents along the chosen decision variable can result in one of two outcomes: (1) a feasible solution at the constraint boundary for the variable or (2) an infeasible solution when the decision variable is changed. In either case, we randomly choose another decision variable for descent. Once we discover a feasible solution that is surrounded by infeasible solutions in the descent directions or at the constraint boundary for every decision, we can stop and claim a locally optimal solution (i.e., no small changes in any decision variable will decrease cost while remaining feasible). We repeat this randomized, iterative descent (starting from the highest cost solution) multiple times, each time finding a locally optimal solution. Among the set of locally optimal solutions, we then identify the lowest cost solution and select it as our optimal accession and retention policy.

To illustrate, consider the three-dimensional case where the only decision variables are the zone SRB levels. Figure 7 shows an instance of the randomized iterative descent. The large box outlined in black represents the full decision space, bounded by the minimum and maximum SRB levels for the EMC. Assume that we evaluated the lowest cost solution (the origin where all SRBs are 0) and found it was not feasible. The green circle denotes the initial solution where all SRBs equal their maximum. Since this solution is feasible, we must choose a variable for descent to find the feasible region's boundary. Imagine that we randomly chose the zone C SRB and found that, at the minimum (i.e., the zone C SRB is set to zero), SIPM returns a feasible solution. The algorithm then selects a new decision variable for descent, such as the zone B SRB. Suppose that with minimum zone B and C SRBs (i.e., each is set to zero), SIPM continues to show that the solution is feasible. In that case, the algorithm would change to the zone A SRB. Through descent, the algorithm identifies a local solution on the feasible region's boundary. This completes one run of the randomized iterative descent. As

shown by the blue circles, multiple local solutions could be identified if the random selection of descent variables were different. By running the algorithm multiple times, we attempt to identify all of these points that we can compare for lowest cost.

Figure 7. Randomized iterative descents



^a. Infeasible region is not known a priori.

The iterative descent can be very slow if we allow only small changes in the decision variables. Because we are searching for the feasible boundary, we use bisection to expedite the search along a decision variable for the feasible solution of lowest cost. Bisection evaluates two points, determines their feasibility, and uses the midpoint to iteratively find a feasible solution neighboring an infeasible solution. To use bisection, our first set of two solutions is the lowest and highest cost solution (i.e., the minimum and maximum values for the decision variables). If the lowest cost solution is also feasible, we randomly change the decision variable and begin the descent process anew from the lowest cost solution. Otherwise, bisection will rapidly identify the solution at the feasible boundary.

If we consider annual accessions to be near continuous (large range of integer values) and possibly different in every year of the planning horizon, we are left running potentially many instances of SIPM to evaluate cost and manning under the many possible accession policies. Combining these many accession policies with the many SRB combinations for all three zones leads to an unmanageable number of possible

solutions. Therefore, we shrink the decision space to five decisions—annual accessions in FY 2017 and FY 2018, plus SRB levels for zones A, B, and C. Each decision is then further discretized by (1) converting accessions into integer percentage changes in accessions relative to average historical accessions and (2) allowing SRB levels to be integers only with a maximum determined using the marginal continuation rate effect.

To help reduce computational time, we parallelize the search for cost-minimizing solutions by solving for the best SRB levels for each accession policy. Then we look over all the accession policies (given their best SRB levels) to determine our near-globally-minimized cost.

We remind readers that, for our optimization, the algorithm is not guaranteed to find globally optimal solutions. We can only identify locally optimal solutions. Although we cannot guarantee that our results are globally optimal, we can determine if the GRID solution improves on the current accession and SRB policies.

Appendix E: Cost and Inventory Excursions

In addition to identifying optimal accession and SRB policies when the female percentage of accessions increases, we examined the sensitivity of the optimal policy for Aviation Ordnancemen (AOs) to key parameters: (1) recruiting costs and (2) starting inventory fill rate. We use the AO EMC to demonstrate.

We considered two recruiting cost excursions—one in which we assume that the recruit cost is 10 percent higher than our current estimate, and the other in which we assume it is 20 percent higher than our current estimate. Higher recruit costs could reflect the possibility that recruiting becomes more difficult because of a stronger civilian economy. These scenarios are designed to determine if GRID reduces the annual accessions in favor of increased SRBs to achieve the same target manning.

We also examined the two starting inventory excursions—one in which the starting inventory is 10 percent lower than actual, and the other in which it is 20 percent lower than actual. These scenarios are designed to determine if GRID prefers higher accessions or higher SRBs to overcome a large initial deficit in manning.¹⁰

We ran these scenarios for AO under a 40 percent female accession share, and we summarize our findings in Table 4. Unlike the other scenarios, these cannot be fairly compared based on total cost. Each scenario is different in its explicit costs (higher recruiting costs) or in the starting inventory. This leads to different initial payroll costs and the costs required to achieve the target 95 percent fill rate when starting from a manning deficit. The more important comparison is the optimal accession and SRB policies across the scenarios. For example, we increase baseline recruiting costs first by 10 percent and then by 20 percent. Compared to the original optimal policy, accessions in FY 2018 and beyond are decreased (979 versus 1,051 per year) in favor of higher zone A SRBs (3 versus 0). In addition, when AOs start with lower inventory, accessions in FY 2018 and later as well as SRB levels are higher than the original optimal policy. The higher accession and retention levels reflect the increased effort required to achieve the target manning level when starting with a deficit. The optimal policies for the scenarios of lower starting inventory may change once the community has achieved a steady-state inventory level (as we have in the original scenario).

¹⁰ In both inventory excursions, we use the current estimates of recruit costs.

Table 4. Excursions for AO optimal accession and SRB policies

Excursion ^a	Annual accession plan		SRB level by zone			FY26 BA fill rates by paygrade (E-3/E-4/E-5/E-6)	FY26 total BA fill rate (E-3+)	Total cost (2017–2026)
	FY17	FY18+	A	B	C			
Original optimal scenario	906	1,051	0	0	0	97/100/100/100%	99%	\$6.62B
10% higher recruit costs	924	979	3	0	0	96/100/100/100%	98%	\$6.77B
20% higher recruit costs	924	979	3	0	0	96/100/100/100%	98%	\$6.80B
10% lower starting inventory	924	1,201	0	1	0	97/100/100/100%	99%	\$6.47B
20% lower starting inventory	1,201	1,201	10	5	2	96/100/100/100%	98%	\$7.07B

Source: CNA calculations from SIPM and GRID search.

^a. We include the original optimal scenario from Table 3 as a baseline reference. The original optimal scenario and each excursion assume that the female share of accessions increases from 35 to 40 percent.

References

- [1] Hansen, Michael L., and Jennie W. Wenger. 2002. *Why Do Pay Elasticity Estimates Differ?* CNA. D0005644.A2 Final.
- [2] SERCO. 2016. *Officer and Enlisted Replacement Costs: Requirements and Business Rules*.
- [3] SERCO. Sep. 2016. "Officer and Enlisted Personnel Replacement Costs." Briefing presented at Naval Support Facility-Arlington.
- [4] Parcell, Ann D., Hoda Parvin, and Yevgeniya K. Pinelis. 2016. *Comprehensive Female Enlisted Strategy*. CNA. DRM-2016-U-012536-Final.
- [5] Mackin, Pat, and Steve Cylke. Jun. 15, 2017. "Navy Enlisted Management Model (NeMMo): A New Enlisted Community Management Tool." Briefing presented at Naval Support Facility-Arlington.
- [6] Belcher, Steven W., Peggy Golfin, David L. Reese, and Kletus S. Lawler. 2017. *Manpower, Personnel, and Training Implications of Future Surface Ship Force Structure Plans and Ship Modernization Schedules*. CNA. DRM-2017-U-015774-Final.
- [7] NAVADMIN 190/16 Aug. 25, 2016. *Sea Shore Flow Enlisted Career Paths Update*.



This report was written by CNA's Resources and Force Readiness Division (RFR).

RFR provides analyses grounded in data to inform resource, process and policy decisions that affect military and force readiness. Our vision is to provide insight on every important resource allocation decision that our sponsors make.



3003 Washington Boulevard, Arlington, VA 22201

www.cna.org • 703-824-2000

NOBODY GETS CLOSER

TO THE PEOPLE. TO THE DATA. TO THE PROBLEM.

CNA is a not-for-profit research organization that serves the public interest by providing in-depth analysis and result-oriented solutions to help government leaders choose the best course of action in setting policy and managing operations.