RAND CORPORATION

KATHARINA LEY BEST, JENNIFER LAMPING LEWIS, CHRISTINA PANIS, NORAH GRIFFIN, ALICE SHIH, CHRISTIAN JOHNSON

Forecasting End Strength in the U.S. Army Reserve

An Integrated Modeling Concept

About This Report

This report documents research and analysis conducted as part of a project titled *An Integrated Modeling Concept for Forecasting End Strength in the U.S. Army Reserve*, sponsored by the Office of the Chief of Army Reserve. The purpose of this project was to develop an integrated modeling concept to support 24-month forecasts of aggregate U.S. Army Reserve end strength.

This research was conducted within RAND Arroyo Center's Personnel, Training, and Health Program. RAND Arroyo Center, part of the RAND Corporation, is a federally funded research and development center (FFRDC) sponsored by the United States Army.

RAND operates under a "Federal-Wide Assurance" (FWA00003425) and complies with the *Code of Federal Regulations for the Protection of Human Subjects Under United States Law* (45 CFR 46), also known as "the Common Rule," as well as with the implementation guidance set forth in Department of Defense (DoD) Instruction 3216.02. As applicable, this compliance includes reviews and approvals by RAND's Institutional Review Board (the Human Subjects Protection Committee) and by the U.S. Army. The views of sources utilized in this study are solely their own and do not represent the official policy or position of DoD or the U.S. Government.

Contents

Figures

Tables

Summary

The U.S. Army Reserve (USAR) is an integral part of the U.S. Army and the country's national defense. Its mission is to provide trained individuals who can serve as active duty soldiers when the mission calls for it. Well-trained service members are central to the USAR mission, and personnel and career management are critical to building a well-trained force. To support its personnel management efforts, the Office of the Chief of Army Reserve (OCAR) is seeking more accurate forecasts of aggregate end strength (the size of the force) over a 24-month period.

End strength forecasts are an important input to analyses that support recruiting and retention policy decisions, as well as resourcing and planning discussions with other Army components. However, generating these forecasts is a complex task due to the many paths into and out of USAR, as well as the diverse career fields and career trajectories of military reservists. Currently no single model is capable of providing such estimates. The RAND Corporation's Arroyo Center examined available modeling capability, identified a set of modeling tools that could estimate portions of the personnel flows into and out of USAR, and developed an Integrated Modeling Concept (IMC), a detailed plan for combining outputs from these tools, as appropriate, to enable USAR to forecast 24-month end strength in a more comprehensive way than currently exists.

Personnel Flows

Our analysis of personnel flows into and out of USAR, as well as the plan we developed to forecast end strength, centers on USAR's Selected Reserve (SELRES). From a data set of close to 1.3 million service members who appear in SELRES between fiscal year (FY) 1990 and FY 2018, we selected about 760,000 individuals with sufficient career data to analyze how Army reservists move into and out of SELRES. The career paths identified from this analysis served as the foundation for the integration plan we developed.

Over a 24-month period, three populations come together to determine total end strength, as shown in Figure S.1. The blue circle represents the SELRES population at the start of the 24-month period of interest, and the orange circle the population at the end. During this two-year period, some reservists will leave SELRES. The blue crescent on the left represents those individuals who are present when the period begins but will be absent when the period ends. Similarly, some people will join SELRES during the two-year period. The orange crescent on the right represents those individuals who are absent when the period begins but will be present when the period ends. The intersection of the circles, denoted "static population," are those reservists who were present on both the first and last days of the two-year period. Therefore, the challenge in constructing the desired end strength forecast is in estimating the number of service members who exit and enter SELRES during the 24-month period of interest—the blue and orange crescents, respectively.

Figure S.1. Forecasting SELRES Exit and Entry
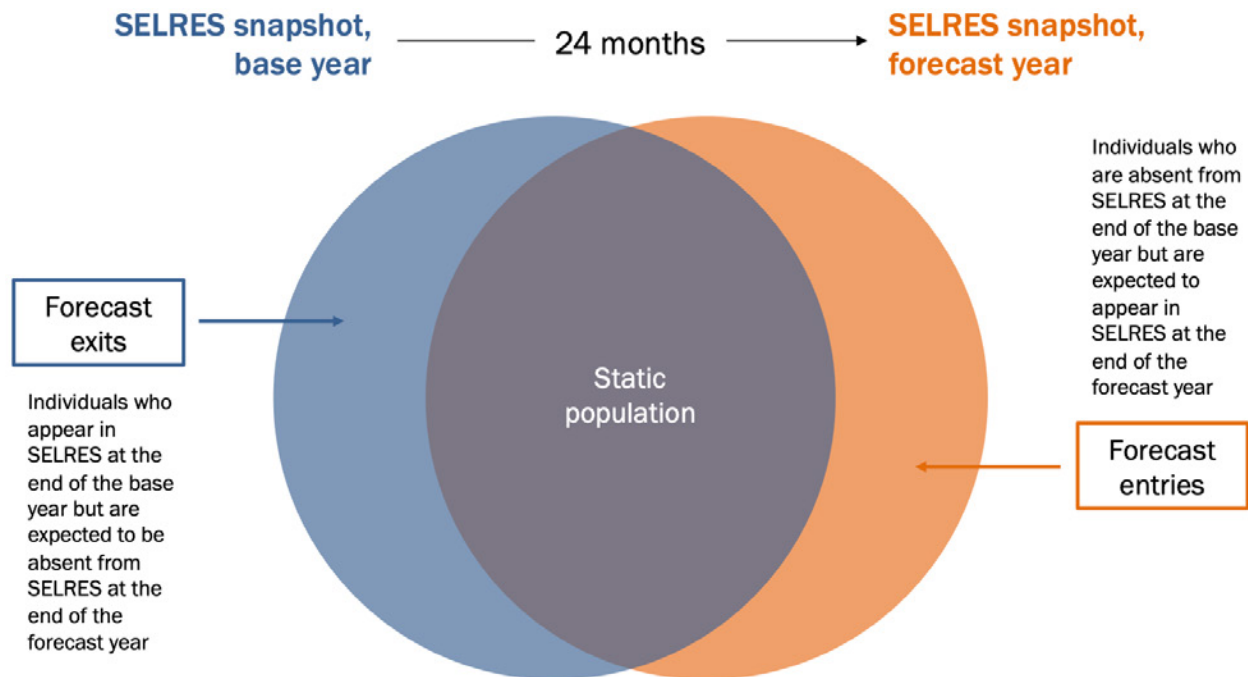
Personnel enter and exit SELRES through many different paths, which we studied in detail. Some members have no prior military service and enter directly from the civilian world; others have served in other Army components, including active duty components, the Army National Guard (ARNG), and the Individual Ready Reserve, or in the active or reserve components of the other military services. Between FY 2016 and FY 2018, for example, 55 percent of SELRES entrants had no prior service of any kind. Another 32 percent had served in the Army's active component (AC) at some point in the past. Of the remainder, the overwhelming majority had prior service in the ARNG. Among those who left SELRES during the FY 2016–FY 2018 period, nearly 70 percent spent at least six months following separation as civilians, but others continued military service in another component within the Army or elsewhere. Estimating SELRES accessions and separations along these many and often complex pathways defines the forecasting challenge.

An Integrated Modeling Concept

RAND has several modeling tools that are capable of predicting flows into and out of SELRES:

- The Recruiting Difficulty Index (RDI) is a forecasting model that provides three distinct measures of recruiting difficulty: the difference between signed enlistment contracts and the contract mission, the average number of days between contract signing and accession, and the percentage of training seats filled. It applies only to enlisted personnel under the U.S. Army Recruiting Command (USAREC) mission.

- The Reserve Recruiting Resource Model (RRRM) is an econometric model that determines the most efficient allocation of recruiters, enlistment incentives, and advertising expenditures to achieve USAREC's recruiting objectives, conditional on recruit eligibility policies and labor market conditions.
- The Reserve Recruit Selection Tool (R-RST) estimates the effects of reservists' characteristics at enlistment on first-term attrition, adverse personnel actions, reenlistment, and the costs of recruiting, training, and replacement. Like the RDI and RRRM, the R-RST applies only to enlisted personnel under the USAREC mission.
- The Dynamic Retention Model (DRM) is an econometric model that simulates the retention and cost effects of changes in compensation policy. Adapted for USAR in FY 2018, it predicts the retention decisions of Army reservists, as well as entry into USAR by those with prior Army AC service. The USAR-adapted DRM covers the vast majority of the enlisted population but considers only those officers with prior service in the Army's AC.
- A machine learning (ML) approach was developed as part of this project to address gaps left by the first four modeling tools. ML can be used to predict entry by reservists with observable histories and exit by reservists generally. The population covered by this approach is quite broad and overlaps with some of the populations covered by the other four tools.

None of the first four modeling tools can generate 24-month forecasts of USAR end strength because none was designed for that purpose. The RRRM, R-RST, and DRM are policy analysis tools, not forecasting models. They were designed to estimate the effects of policy changes on recruiting and retention outcomes, holding all other factors constant. To this end, they provide predictions that vary with the policy levers of interest (and any other assumptions the user is permitted to adjust) rather than a single forecast that incorporates the policies and conditions that are expected to materialize. The RDI, in contrast, is a forecasting model, but it delivers a set of recruiting difficulty measures, rather than a forecast of end strength per se.

Nevertheless, the modeling tools can estimate particular flows into or out of USAR and, with the appropriate care and caution, the estimates can be combined to construct a 24-month end strength forecast. At OCAR's request, RAND Arroyo Center leveraged these tools to the extent possible (given their intended purpose) to develop the IMC. As illustrated in Figure S.2, the IMC assigns modeling tools to personnel flows based on the capabilities of each tool, with ML used to fill as many gaps as possible. For example, the DRM is applied primarily to predict exit by and retained entry of enlisted USAR personnel; it makes only a small contribution to estimating officer exits and entries, where ML is relied upon more heavily. The RRRM and R-RST cover the majority of enlisted accessions: those with no prior service of any kind. The RDI is omitted from the IMC because the RRRM provides an estimate that is more directly relevant to the flow of enlisted personnel into USAR.

It is important to note that the percentages listed in Figure S.2 were computed using actual counts of separations and accessions that occurred between September 2016 and September 2018, not estimates generated by the modeling tools. We did not run the tools because the RRRM and R-RST were under development while this project was underway and operating

Figure S.2. The Application of Modeling Tools to Forecast Entries and Exits
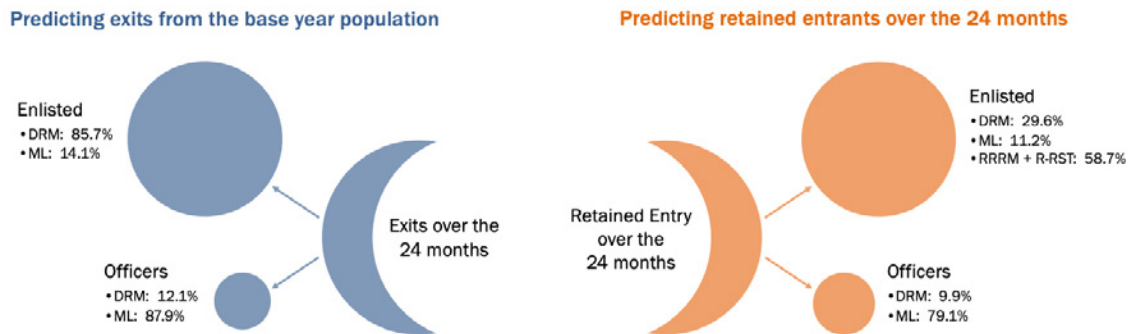
NOTES: Percentages are based on actual exit and retained entry that occurred between September 2016 and September 2018. Percentages do not add to 100 percent because of gaps in modeling capability and data errors.

the DRM was outside the project's scope. The percentages in the figure are offered for the sole purpose of illustrating the relative prevalence of each modeling tool within the IMC.

Two key considerations guided the assignment of modeling tools to personnel flows. First, the IMC aims to leverage prior work that was sponsored by USAR and the broader Army by applying existing modeling tools whenever possible. Second, the IMC preferentially includes tools that support exploration of alternative policy choices, such as those related to recruiting resource allocations (RRRM), recruit selection criteria (R-RST), and personnel compensation (DRM). These decisions were made to maximize the versatility and future applicability of the IMC to support planning into the future.

The IMC, however, is not the only alternative for estimating the personnel flows required to construct 24-month end strength forecasts using the set of modeling tools identified here. For example, all of the flows assigned to the DRM could be estimated using an ML approach instead. The DRM should be applied only if the effects of changes in compensation policy on USAR end strength are of interest. If not, substituting ML for the DRM should be considered. More generally, the "best" modeling concept depends heavily on the ways in which the concept will be used to support policy decisions.

Concluding Thoughts

Future implementation of the IMC should begin with a careful evaluation of the applications and requirements for IMC outputs throughout OCAR's planning processes, especially as these may have evolved since the writing of this report. The thorough mapping of personnel flows to modeling tools presented in this report provides some flexibility in terms of supporting alternate tool assignments to best suit future needs and identifying some of the reasons why different choices might be made. The work of developing the IMC as described here has resulted in a deep understanding and careful documentation of the applicability of various modeling tools to different personnel flows and can thus provide a foundation for future implementation that supports a variety of needs.

Acknowledgments

Abbreviations

AC	active component
AIT	Advanced Individual Training
ARNG	Army National Guard
BCT	Basic Combat Training
CMF	career management field
DRM	Dynamic Retention Model
FY	fiscal year
GRU	gated recurrent unit
IET	Initial Entry Training
IMC	Integrated Modeling Concept
IRR	Individual Ready Reserve
ML	machine learning
MOS	Military Occupational Specialty
OCAR	Office of the Chief of Army Reserve
OSUT	One Station Unit Training
PULHES	physical capacity/stamina (P), upper extremities (U), lower extremities (L), hearing and ears (H), eyes (E), and psychiatric (S)
RC	reserve component
RDI	Recruiting Difficulty Index
RRRM	Reserve Recruiting Resource Model
R-RST	Reserve Recruit Selection Tool
SELRES	Selected Reserve
USAR	U.S. Army Reserve
USAREC	U.S. Army Recruiting Command
WEX	Work Experience
YOS	years of service

1. Introduction

The U.S. Army Reserve (USAR) is an integral part of the U.S. Army and the country's national defense. Its mission is to provide trained individuals who can serve as active duty soldiers when the mission calls for it. Well-trained people are central to the USAR mission, and personnel and career management are critical to building a well-trained force. To support its personnel management efforts, the Office of the Chief of Army Reserve (OCAR) is seeking more accurate forecasts of aggregate end strength (soldiers on hand) over a 24-month horizon. Such forecasts are an important input to analyses that support recruiting and retention policy decisions, as well as resourcing and planning discussions with other Army components.

At OCAR's request, we identified a set of existing modeling tools, each able to estimate a portion of the personnel flowing into or out of USAR, and developed an Integrated Modeling Concept (IMC) that describes how outputs from the individual tools could be used together to generate 24-month forecasts of USAR aggregate end strength. The primary product of the project is an integration *plan*. Because some of the tools were under development while this project was underway, we did not operate the tools or integrate their outputs.

The Structure of the U.S. Army Reserve

We begin with a brief overview of USAR's structure and career paths, which serve as the basis for the IMC. The Army Reserve components were established during the Korean War from the Organized Reserve Corps. The two components are the Army National Guard (ARNG) and USAR. A primary distinction between the ARNG and USAR is their command structure within the U.S. government. The ARNG is part of state militias under control of governors unless called into active duty by the President to aid during federal emergencies. USAR is a federal entity that provides trained units to be available for active duty in times of threats against national security.[1] Although not a primary USAR mission, under 10 U.S.C 12304a, the Secretary of Defense has the authority to order units and individuals of USAR (as well as the Air Force Reserve, Marine Corps Reserve, and Navy Reserve) to active duty for up to 120 days "when a governor requests federal assistance in responding to a major disaster or emergency."[2]

---

[1] U.S. Code, Title 32, Section 101, Definitions, June 3, 1956; U.S. Code, Title 10, Section 10104, Army Reserve: Composition, October 5, 1994.

[2] Lawrence Kapp and Barbara Torreon, *Reserve Component Personnel Issues: Questions and Answers*, Washington, D.C.: Congressional Research Service, October 4, 2018, p. 16.

Since the Korean War, the two components have evolved to the current structure of a Ready Reserve, a Standby Reserve, and, for USAR, a Retired Reserve, as illustrated in Figure 1.1.[3] Within USAR, the Ready Reserve contains the Selected Reserve (SELRES) and the Individual Ready Reserve (IRR). SELRES is the main population of interest for this study.

Figure 1.1. Reserve Components of the U.S. Army
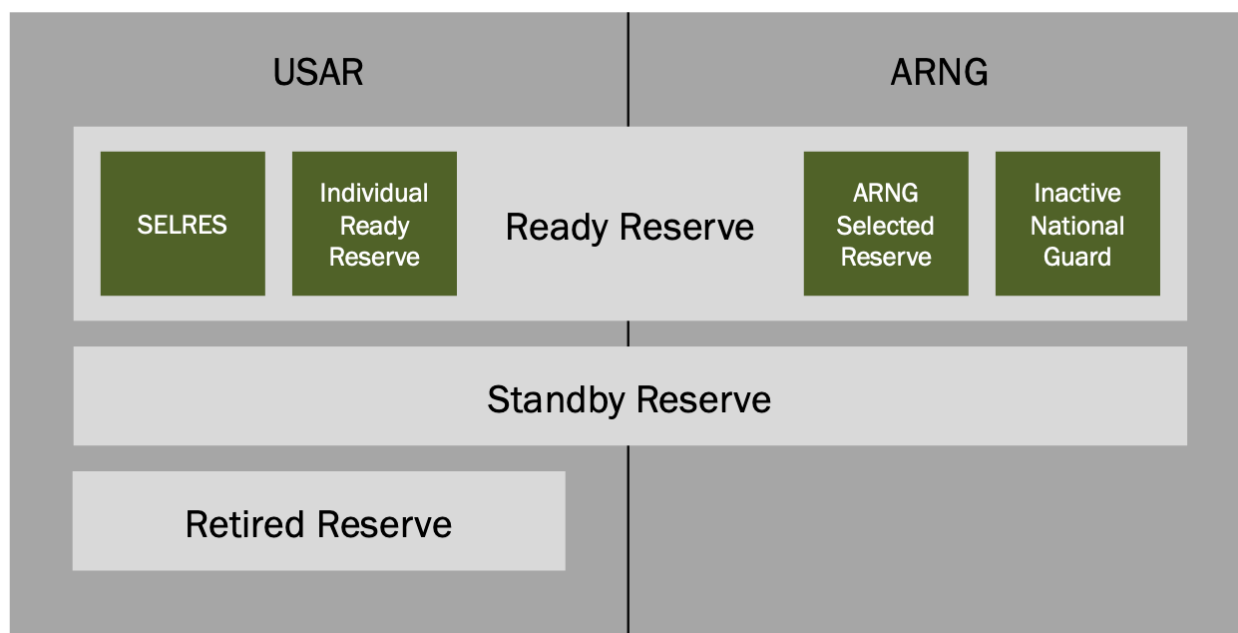
Consisting of active reserve units that conduct periodic training, SELRES makes up the majority of USAR. SELRES soldiers differ from their active component (AC) counterparts because they serve part-time. Their service obligation normally consists of one weekend of training per month (known formally as Battle Assembly, or "drills") and an annual training of two weeks. The flexibility in service obligation allows SELRES soldiers to attend school or work in the civilian sector while living and working where they choose.[4] This "dual life" allows reserve soldiers to enjoy civilian life unless called to active duty. SELRES soldiers have the option to volunteer for a deployment or active duty for operational support orders.

Also part of the USAR Ready Reserve, the IRR provides a solution for soldiers with training incompatibilities and is often a home for transitioning soldiers who need to finish their service commitments. While the IRR has no training requirements, the President has the authority to call up to 30,000 members of the IRR during times of crisis.

---

[3] Congressional Research Service, *Defense Primer: Reserve Forces*, fact sheet, last updated December 21, 2021.

[4] U.S. Army, "About the Army," webpage, undated.

The structure of the Ready Reserve in the ARNG mirrors the structure of the Ready Reserve in USAR. The ARNG Selected Reserve is equivalent to USAR's SELRES, and the Inactive National Guard is equivalent to the IRR.

The Standby Reserve also exists within both USAR and the ARNG. In USAR, the Standby Reserve provides a limited one-year placement for soldiers who have a short-term incompatibility that requires temporary suspension of training, such as a personal or medical hardship. Certain Standby Reserve soldiers participate in training for retirement points only (they cannot receive payment).[5] After the one-year period, a soldier can transition into the IRR to complete his or her military service obligation. If the soldier is not required to serve anymore, he or she can choose to discharge from USAR or transfer into the IRR or Retired Reserve.[6] The Standby Reserve in the ARNG is similar to the Standby Reserve in USAR.

Finally, the Retired Reserve, which exists only within USAR, provides a mechanism for reserve soldiers under the age of 60 with 20 years or more of active duty service to continue serving in the reserves and qualify for specific pension plans.[7]

Career Paths into and out of SELRES

To generate a plan for integrating outputs from existing analytic tools, the IMC details the career paths leading into and out of SELRES and then pairs those paths with the tools that apply. This section provides a high-level summary of the ways in which individuals generally come to be part of and leave SELRES; a more detailed analysis is provided in subsequent chapters.

There are three main paths into SELRES. First, service members may transition from the AC or another part of the reserve component (RC) into SELRES to complete their service obligations. Second, individuals may affiliate with SELRES directly from the civilian sector, having no prior military service. As a third path, service members with prior AC service may affiliate with SELRES following an extended period in the civilian sector. These soldiers, which will be treated as a separate group throughout this report, are generally referred to as "prior service-civil life."[8]

---

[5] Headquarters, Department of the Army, *Army Reserve: Assignments, Attachments, Details, and Transfers*, Washington, D.C.: Department of the Army, Army Regulation 140-10, April 25, 2018. Retirement pay in the military is based on the number of points one accrues while in service; the more points one has, the higher the retirement pay. For more on the military point system as it applies to the reserves, see Jim Absher, "Guard and Reserve Retirement," Military.com, June 22, 2021.

[6] Headquarters, Department of the Army, 2018, p. 51.

[7] Rod Powers, "Military Reserves Federal Call Up Authority: From Reserves to Active Duty," The Balance Careers, last updated July 30, 2019; Headquarters, Department of the Army, 2018, p. 51.

[8] A previous RAND report defined prior service–civil life gains as pertaining to someone who "join[s] more than six months after leaving [active duty] *and* more than eight years after initial accession." See Jennie W. Wenger, Bruce R. Orvis, David Stebbins, Eric Apaydin, and James Syme, *Strengthening Prior Service–Civil Life Gains and Continuum of Service Accessions into the Army's Reserve Components*, Santa Monica, CA: RAND Corporation, 2016.

When transitioning out of the RC, soldiers often maintain an IRR status for a few years to complete their service obligations. The number of years that they remain in IRR status varies depending on the initial contract, and they can elect to remain in the IRR even if they have completed their service obligation. AC soldiers have the option to transition directly into the IRR to complete their military service obligation, but the service obligation takes longer to complete with that option.

End Strength Management in the U.S. Army Reserve

Recruiting and retention outcomes depend on the decisions of individual reservists. Several factors affect these decisions, including a reservist's proclivity for military service, compensation and benefits offered by USAR, career opportunities available in other Army components and other services, and career opportunities in the civilian sector. Accordingly, USAR end strength management must consider not only the policy levers USAR has at its disposal but also the recruiting and retention policies of the other Army components and military services, as well as economic conditions in the civilian labor market.

Predictive modeling can help USAR prepare for the environment it is likely to face. Tools that accurately forecast aggregate end strength are one such predictive capability. Knowing how many soldiers are expected to be in USAR in 24 months factors into a variety of USAR recruiting and retention policy decisions, such as allocation of recruiting resources, eligibility policies and waivers, bonuses and other compensation adjustments, geographical considerations, and retention policies and waivers. The knowledge also helps USAR prepare for resourcing and planning discussions with the other components.

Existing analytic tools do not provide this predictive capability. Those that were designed to generate forecasts predict outcomes that relate to end strength but do not predict end strength per se. For example, the Recruiting Difficulty Index (RDI) estimates the difference between signed enlistment contracts and the contract mission, the average number of days between contract signing and accession, and the percentage of training seats filled.[9] Other tools were simply not designed to generate forecasts but rather to predict changes in outcomes that may result from changes to specific policies of interest, holding all other factors constant. For example, the Dynamic Retention Model (DRM) simulates the retention and cost effects of changes to the compensation system; the model was not designed to account for changes in retention that may result from factors other than changes in compensation.[10]

In sum, there is no comprehensive, end-to-end model that forecasts end strength while accounting for the effects of changes to USAR policy and other contextual factors that may

---

[9] Jeffrey B. Wenger, David Knapp, Parag Mahajan, Bruce R. Orvis, and Tiffany Berglund, *Developing a National Recruiting Difficulty Index*, Santa Monica, Calif.: RAND Corporation, RR-2637-A, 2019.

[10] Beth J. Asch, Michael G. Mattock, and James Hosek, *Effects of the Blended Retirement System on United States Army Reserve Participation and Cost*, Santa Monica, Calif.: RAND Corporation, RR-2591-A, 2019.

influence end strength. This deficiency limits USAR's ability to determine how best to adjust recruiting and retention policies in order to achieve end strength targets in a cost-effective fashion.

The Integrated Modeling Concept

To address this deficiency in USAR's analytical capability, OCAR asked the RAND Corporation's Arroyo Center to leverage existing tools to the extent possible, given their intended purpose, and develop an approach that can be used to forecast end strength over a 24-month horizon. The IMC is the product of that effort: a detailed plan for integrating the outputs of existing tools, as appropriate, to enable USAR to generate the desired end strength forecasts. At the highest level, the IMC consists of three components:

1. *A set of personnel flows that represent the various paths individuals take into and out of USAR.* These flows are the individual pieces of the end strength prediction problem that must be solved.
2. *A set of analytic tools that could be applied to predict the magnitude of each personnel flow.* The tools vary with respect to model type and structure, data requirements, the policy questions they can address, the subpopulations to which they can be applied (e.g., enlisted members and officers), and the parts of the career pipeline they cover (e.g., accessions and separations).
3. *A detailed mapping between personnel flows and applicable tools.* The mapping is the primary product of this study. It lays out which tools can be used to predict each personnel flow and identifies gaps in the modeling landscape.

Because some of the tools were under development while this project was underway, we did not implement the IMC; that is, we did not operate the tools or integrate their outputs as stipulated by the mapping.

Report Overview

We describe these three components in the remainder of this report and, where applicable, discuss new methods that could be used to address modeling gaps. Chapter 2 provides a detailed description of the study population and introduces the personnel flows that must be tracked to construct the IMC. Chapter 3 reviews the modeling tools whose outputs must be integrated to construct the IMC. Chapter 4 links the personnel flows and tools, culminating in a detailed plan for integrating the outputs of the tools, as appropriate, to enable USAR to generate the desired 24-month end strength forecasts. Chapter 5 discusses options and considerations for potential future implementation of the IMC.

2. Personnel Flows and Career Paths

In this chapter we describe the data sample used in our analyses and present three different views of the personnel flows into and out of USAR:

1. *Immediate sources and destinations of SELRES members.* The first view depicts the sources from which individuals most recently entered SELRES and the destinations of individuals upon leaving SELRES. The view is myopic in that it is limited to the six months that precede entry into SELRES and the six months that follow exit from it.
2. *Military service histories of SELRES exit and entry populations.* The second view depicts the military service histories of those who enter and exit SELRES over a longer time frame. The view accounts for 27–29 years of service (YOS) history to place each reservist in one of six categories.
3. *Career paths of SELRES members.* The third view depicts the results of an automated clustering exercise that identifies groups of reservists whose behaviors during their time in SELRES are similar. The most salient characteristics of each group, including career length, prior military service, occupation or career field, and demographic characteristics, are described.

These three views lay the foundation for the more complex taxonomy specified by the IMC. The mapping of existing modeling tools to personnel flows presented in Chapter 4 relies on an intersection of the subpopulations identified in this chapter, as well as a few other considerations.

Data Description

Forecasting SELRES end strength requires an understanding of the pattern of movements among Army reservists. To identify these patterns, we compiled a comprehensive data set that captures information about this population throughout their military careers. The data set includes information on service within ARNG, the Regular Army, and USAR, as well as service in branches of the military other than the Army.

Using longitudinal data obtained from the Defense Manpower Data Center, we identified close to 1.3 million (1,273,149) unique service members who had been in SELRES between fiscal year (FY) 1990 and FY 2018. The data set we compiled drew heavily from the Defense Manpower Data Center's Work Experience (WEX) file. The file provides information on AC and RC service members, logging all major changes in an individual's career: gains, losses, and promotions. RAND has access to a truncated version of the WEX file with data going back to FY 1990. All individuals present from FY 1990 onward are included in the truncated file, but their histories prior to FY 1990 are not included. Consequently, there are individuals in the data set for which we could not identify the point in time at which military service began. The problem is more acute for earlier cohorts. Since the WEX file is no longer updated, we used four other Defense Manpower Data Center files to gather information on more recent service history: the

Active Duty Master File, the Reserve Master File, and the corresponding transaction files. To obtain demographic information for the individuals in the data set, we used the Defense Enrollment Eligibility Reporting System, which provides information on marital status, location, ethnicity, educational level, birth date, and gender. The demographic data are also truncated; they are not available for individuals whose military service ended prior to FY 2000. Finally, we added activation and deployment data to the monthly analysis file, which already included information on service, grade, component, reserve category, primary occupation code, and unit identification code.

Personnel Flows

As the foundation for developing an integration plan, we examined the flows of individuals entering and exiting SELRES between FY 2016 and FY 2018. Three populations come together to determine total end strength at the end of the 24-month period, as represented in Figure 2.1. In the figure, the blue circle represents the SELRES population at the start of the 24-month period, and the orange circle represents the population at the end of the 24-month period. The blue crescent on the left represents the portion of the population that is present on September 30, 2016, but absent on September 30, 2018—that is, these individuals exited USAR at some point during the 24-month period. The orange crescent on the right represents the reverse: reservists who joined SELRES at some point during the 24-month period and were retained long enough to be present on September 30, 2018. The overlap in the middle represents the population that was present on both September 30, 2016, and September 30, 2018.

To better understand the population that left SELRES as represented by the blue crescent, and the population that entered SELRES as represented by the orange crescent, we examined the

Figure 2.1. SELRES Exit and Entry Populations, FY 2016–FY 2018

various paths reservists followed as they exited and entered SELRES during the FY 2016–FY 2018 period. Figure 2.2 depicts the more immediate sources and destinations of Army reservists, while Figure 2.3 provides a longer view of the military service histories of those who enter and exit SELRES. We provide additional details about each figure below.

The top portion of Figure 2.2 shows the sources from which individuals most recently entered SELRES. For each individual who was present in SELRES on September 30, 2018, but not on September 30, 2016 (the orange crescent in Figure 2.1), we identified the month and year of entry and examined the individual's military service history over the preceding six months to determine the immediate source of entry. Over half of the entrants had been civilians for at least six months prior to joining SELRES. The remaining entrants joined directly from the IRR, active duty service in any branch of the military, the Air National Guard or ARNG, or the reserves of a branch other than the Army.

The bottom portion of Figure 2.2 shows the destinations of individuals upon leaving SELRES. For each individual who was present in SELRES on September 30, 2016, but not on September 30, 2018 (the blue crescent in Figure 2.1), we identified the month and year of separation and examined the individual's service records over the six months that followed to determine the immediate destination. Nearly 70 percent of exiting reservists spent at least six months immediately following separation as civilians, while the majority of the remainder transitioned to the IRR. However, this distribution may be affected by censoring; for those exiting SELRES

Figure 2.2. Immediate Sources and Destinations of SELRES Members, FY 2016–FY 2018
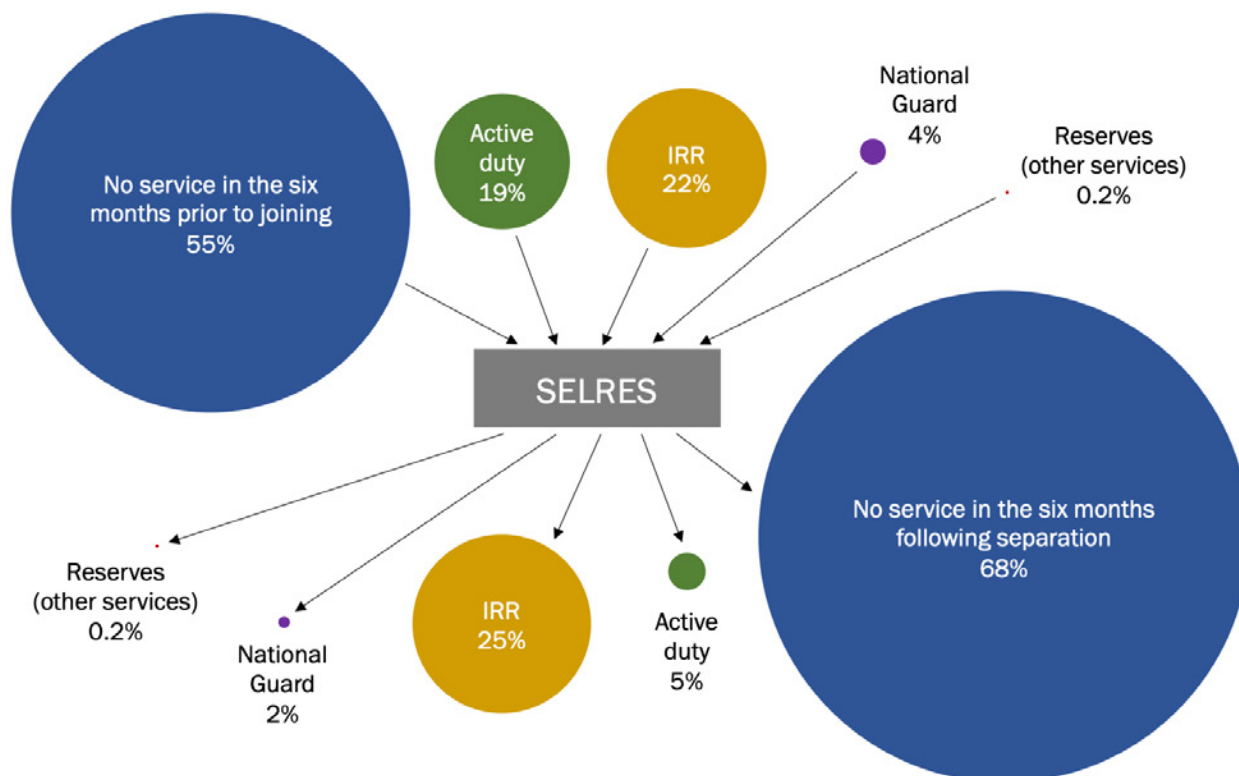
after March 31, 2018, we could not observe a full six months following separation. For a more granular breakdown of the flows depicted in Figure 2.2, see Appendix A.

Figure 2.3 examines the same entry and exit populations (the orange and blue crescents, respectively, in Figure 2.1) but accounts for the reservists' military service histories going back to FY 1990 to subdivide the populations into six mutually exclusive categories:

1. *No prior military service of any kind.* These individuals had not served in any branch of the military—AC, reserve, or national guard—prior to joining SELRES.
2. *Prior Army AC service.* These individuals had served in the Army AC at some point in time prior to joining SELRES.
3. *Prior non-Army AC service.* These individuals had AC service at some point in time prior to joining SELRES, but that service was not in the Army AC.
4. *Prior ARNG service.* These individuals had no AC service in any branch of the military prior to joining SELRES, but did have prior ARNG service.
5. *Prior non-Army Reserve or guard service.* These individuals had no AC service in any branch of the military or ARNG service prior to joining SELRES, but did have prior service in the Air National Guard or the reserves of a branch other than the Army.
6. *Prior SELRES service only.* These individuals had previously served in SELRES but had no other prior military service.

Figure 2.3. Military Service Histories of SELRES Exit and Entry Populations, FY 2016–FY 2018
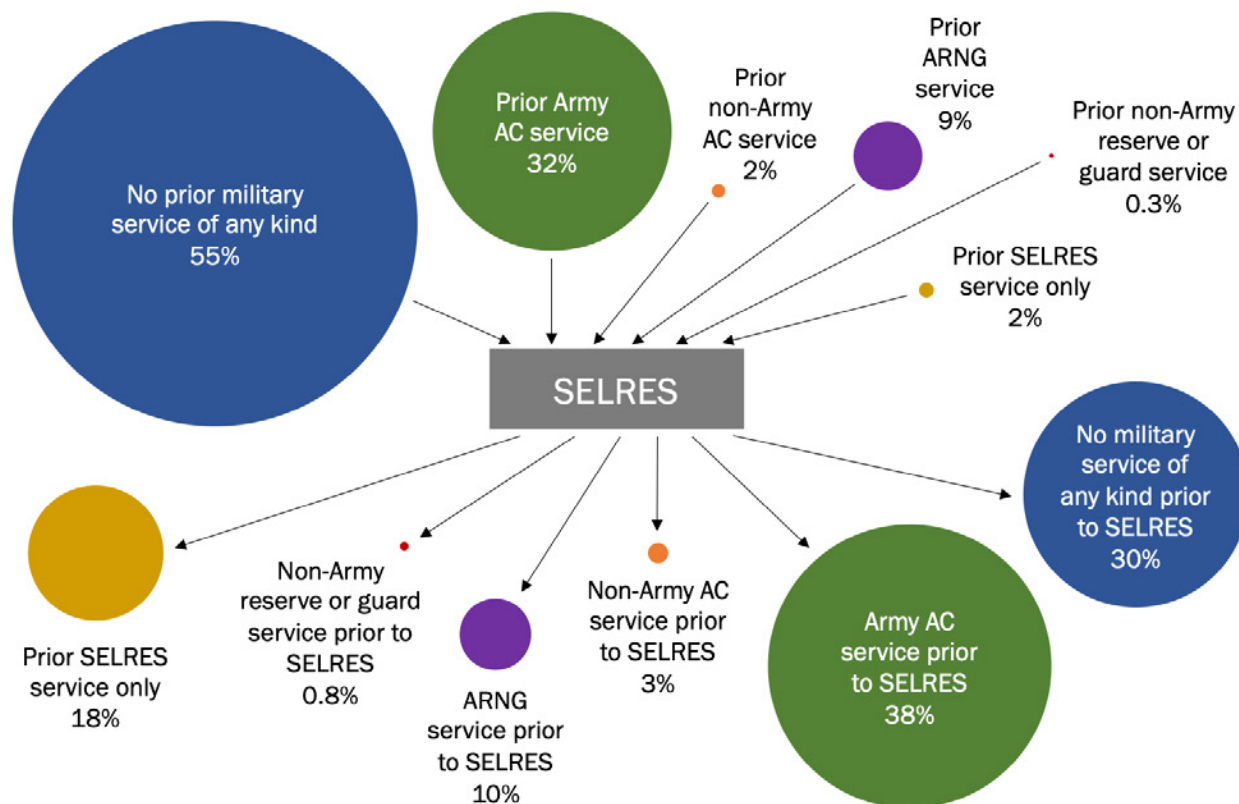
NOTE: The "No prior military service of any kind" category includes members of SELRES who experienced delayed entry. These persons appear to be in the IRR for a short period of time (six months or less) prior to joining SELRES, despite having no prior military service.

9

The top portion of Figure 2.3 shows that 55 percent of SELRES entrants had no prior military service of any kind. Of the remainder, the overwhelming majority had prior Army AC service. The bottom portion shows that the service histories of those exiting SELRES skew more toward reservists with prior Army AC service. Only 30 percent of those who separated from SELRES had no prior military service of any kind.

Career Paths

As we have noted, our population of SELRES personnel is large and diverse. The data set used for the career path analysis described in this section contains 762,549 individuals (those from the larger data set with complete information) and comprises about 3,600 unique data points on each individual.[1] While the focus of this report is on the inflows and outflows of this population to and from SELRES, these data also contain a wealth of information about individuals' behaviors *during* their time in SELRES. In order to better understand behavior patterns *during* SELRES stints, we apply automated dimensionality reduction and clustering techniques, which are described in more detail in Appendix B, along with details of the clusters and our results. The goal of this type of analysis is (1) to identify groups of individuals whose behaviors and paths *during* their time in SELRES are similar, and then (2) to describe the characteristics of these groupings. The clustering exercise performed as part of this study is a proof of concept intended to demonstrate that these kinds of techniques may be useful in detecting such groups and patterns.

In this exercise, we identified six clusters within the SELRES sample analyzed.[2] These clusters are described in Table 2.1. The driving variables that define cluster membership are related to career length and prior service status, with individual job category and demographic characteristics having smaller impacts on cluster membership. More specifically, the automated clustering exercise identifies three clusters with and three clusters without significant prior service. The three clusters in each group differ in terms of YOS, departure path, and, to a lesser extent, career management field (CMF). Demographic characteristics, such as educational attainment, age at entry into SELRES, race, and gender, do not vary as much between clusters. Age-related factors are also necessarily correlated with other variables, such as career length, YOS, and the fiscal years in which individuals joined SELRES.

---

[1] A total of 3,639 data points are used for each individual, representing one-hot encodings of static and time series variables. Specifically, each individual is described by 39 data points that contain one-hot encodings of six static variables (race, age, sex, education level, and prior service), as well as ten data points for each month that contain one-hot encodings of time series variables. *One-hot encoding* refers to the use of indicator variables: each entry contains exactly one 1, identifying the category to which the entry belongs. More information can be found in Appendix B.

[2] As explained in Appendix B, the number of groups is specified by the analyst, but exploration of alternative numbers of clusters confirms that six clusters is appropriate for this sample.

Table 2.1. Descriptions of Career Clusters for SELRES Careers

Cluster Name	Cluster Description	Percentage of Sample
One (term) and done	Primarily younger individuals with little or no prior service and shorter SELRES service (5–10 YOS).[a] SELRES departers who continue service tend to do so in the IRR.[b]	30
One term, maybe AC	Primarily individuals with little or no prior service (0–5 YOS). SELRES departers who continue service tend to do so in the Army AC in lieu of the IRR.	25
Satisfied reservists	Primarily individuals with little or no prior service and longer SELRES careers (10–22 YOS).	15
Combat arms	Primarily individuals who have significant prior service (generally coming from the IRR) and a wide distribution of YOS. Generally, these individuals serve in the infantry and in other combat arms roles and do not achieve retirement.	10
AC transfers	Primarily individuals with significant prior service. Those who continue service tend to move to/from the Army AC and generally have less than 22 YOS.	10
SELRES retirees	Primarily individuals with some prior service that tend to have long histories in SELRES and achieve retirement. Almost the entire population with >20 YOS is in this group.	10

NOTES: Every individual is assigned to exactly one cluster, meaning that the "Percentage of Sample" metric reflects the number of individuals in the cluster and that they are not weighted by career length. Cluster names are descriptive, not prescriptive: cluster membership is determined entirely by the time series and static data fields for each individual.
[a] YOS in cluster descriptions is based on YOS at the last point in time the individual was observed in our data. Therefore, YOS could refer to either YOS at departure or current YOS (for those still serving).
[b] Departures from SELRES are most commonly to the IRR or a separated status across all groups. Descriptions on departure paths in different clusters focus on deviations from this pattern.

The result from the clustering exercise that is most relevant to the IMC is a description of the paths individuals take into and out of SELRES. Figure 2.4 shows these results and links cluster membership to the personnel flows that define which models can be used to predict accession and separation for these individuals. We see that clusters with significant prior service (SELRES retirees, AC *transfers*, and those in combat arms roles) have paths into SELRES that include the IRR (green and orange), other services (blue), or other Army components (red and purple). The majority of those in the remaining three clusters do not have prior service. In terms of flows out of SELRES, paths are dominated by separation across all clusters, but significant portions of the populations of most clusters also move to the IRR. The exception here is the "One term, maybe AC" cluster, which shows individuals primarily separating (pink), continuing service (aqua), or moving to the AC (purple) in lieu of the IRR, as the cluster name implies. Appendix B contains additional results.

Figure 2.4. Paths into and out of SELRES, by Cluster
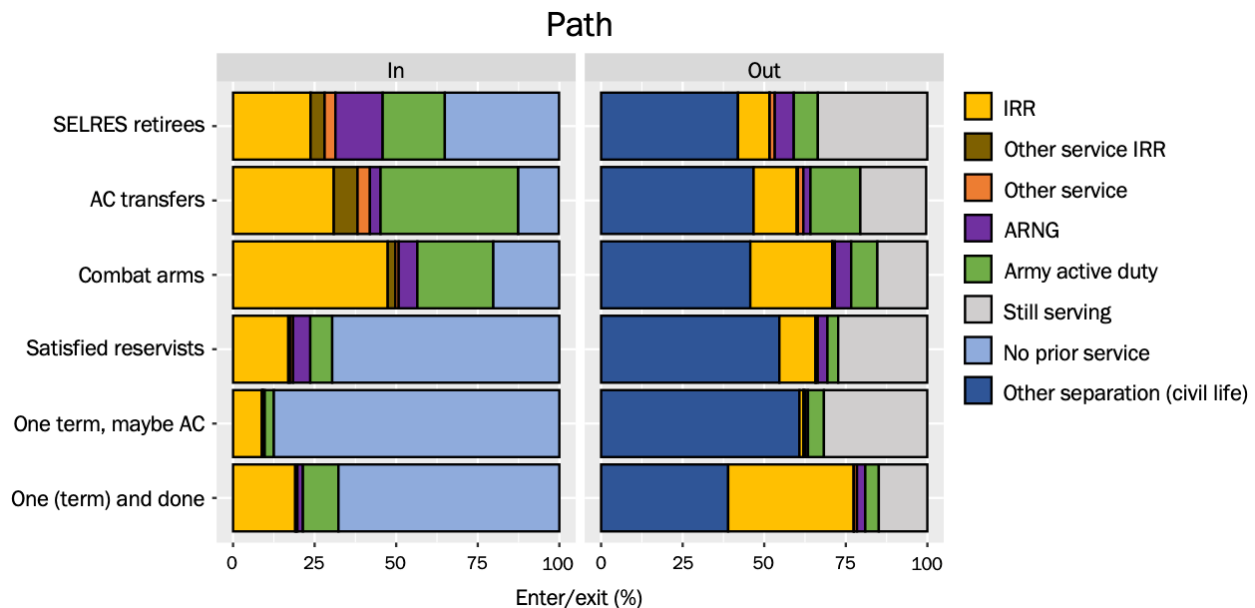
The clustering exercise is of interest primarily because it demonstrates that this approach succeeds in identifying groups that may be useful for understanding career paths.[3] Automated clustering does not work well in all applications, but this proof-of-concept exercise suggests that it may indeed work well for detecting groupings of similar SELRES careers—which may in turn inform our understanding of entry and exit into USAR and be useful in considering policies related to recruiting and retention. The details of the approach presented here could be refined in future studies to alleviate concerns, such as censoring and correlation issues.

Chapter Summary

This chapter has provided three distinct views of the personnel flows into and out of USAR: immediate sources and destinations of SELRES members, military service histories of SELRES entry and exit populations, and career paths of SELRES members. The analyses revealed the extent to which Army reservists vary in terms of both the paths that led them to SELRES and their behaviors (e.g., occupation and YOS) while in SELRES. The work described in this chapter also serves as a foundation for the integration plan presented in Chapter 4: the mapping of existing modeling tools to personnel flows relies on an intersection of the subpopulations depicted in the three views, in addition to a few other considerations. In Chapter 3 we review the set of modeling tools that could be used to forecast particular flows into and out of USAR. In Chapter 4 we describe how outputs from the tools could be integrated to forecast USAR end strength.

---

[3] By succeeding in identifying groups, we mean simply that the algorithm detects clusters that form groups with discernable differences along some of the dimensions of interest, such as YOS or prior service status.

12

3. Modeling Tools

RAND has developed several modeling tools that are capable of predicting flows into and out of SELRES:

- the RDI
- the Reserve Recruiting Resource Model (RRRM)
- the Reserve Recruit Selection Tool (R-RST)
- the DRM
- a machine learning (ML) approach.

The RDI, RRRM, and R-RST are a suite of tools that predict entry and first-term attrition of enlisted personnel under the U.S. Army Recruiting Command (USAREC) mission. The RDI is a forecasting model that provides three distinct measures of recruiting difficulty.[1] The RRRM is an econometric model that determines the most efficient allocation of recruiters, enlistment incentives, and advertising expenditures to achieve USAREC's recruiting objectives (conditional on recruit eligibility policies and labor market conditions).[2] The R-RST estimates the effects of reservists' characteristics at enlistment on first-term attrition, adverse personnel actions, reenlistment, and the costs of recruiting, training, and replacement.[3]

The DRM is an econometric model that simulates the retention and cost effects of changes in compensation policy.[4] It predicts the retention decisions of each service member over his or her career, allowing for individual differences in taste for military service and random shocks in each period. Because the DRM models individual decisionmaking, it can simulate the effects of policy changes that have no analogues in past policies.

As is noted in Chapter 1, none of the first four modeling tools can generate 24-month forecasts of USAR end strength because none was designed for that purpose. The RRRM, R-RST, and DRM are policy analysis tools, not forecasting models. They were designed to estimate the effects of policy changes on recruiting and retention outcomes, holding all other factors constant. To this end, they provide predictions that vary with the policy levers of interest (and any other assumptions the user is permitted to adjust) rather than a single forecast that incorporates the policies and

---

[1] Wenger et al., 2019.

[2] Bruce R. Orvis, Craig A. Bond, Daniel Schwam, and Irineo Cabreros, *Resources Required to Meet the U.S. Army Reserve's Enlisted Recruiting Requirements Under Alternative Recruiting Goals, Conditions, and Eligibility Policies*, Santa Monica, Calif.: RAND Corporation, RR-A1304-1, 2022.

[3] Unpublished RAND research that developed a Reserve Component version of Bruce R. Orvis, Christopher E. Maerzluft, Sung-Bou Kim, Michael G. Shanley, and Heather Krull, *Prospective Outcome Assessment for Alternative Recruit Selection Policies*, Santa Monica, Calif.: RAND Corporation, RR-2267-A, 2018.

[4] Asch, Mattock, and Hosek, 2019.

conditions that are expected to materialize. The RDI, in contrast, is a forecasting model, but it delivers a set of recruiting difficulty measures, rather than a forecast of end strength per se.

Nevertheless, the modeling tools can estimate particular flows into or out of USAR, and with the appropriate care and caution, the estimates can be combined to construct a 24-month end strength forecast. The contract production submodel of the RRRM predicts the flow into USAR of enlisted personnel with no prior military service of any kind, while the R-RST predicts the fraction of these recruits that separate from USAR during the first term. With few exceptions, the DRM predicts the flows into and out of USAR of enlisted personnel and officers with prior Army AC service. It also predicts the retention behavior of enlisted personnel with no military service outside USAR.

Because the first four modeling tools were designed for purposes other than forecasting USAR end strength, the personnel flows they cover leave some gaps. To address these gaps, we developed an ML approach to predicting entry by reservists with observable histories and exit by reservists generally. The population potentially covered by the ML approach is quite broad and overlaps some with the populations covered by the other four tools, as will be explained in more detail in later sections of this chapter.

The five modeling tools are at various stages of development. Some, like the DRM, are fully developed and have been in use for decades.[5] Others, like the RRRM and R-RST, were adapted for USAR concurrently with this research study but have versions for the Regular Army that have been in use since about 2018.[6] The ML approach is at the proof-of-concept stage; future implementation of the IMC would require full development of an ML modeling tool.

The remainder of this chapter describes each of the five modeling tools in turn. The descriptions address the purpose for which the tool was developed, the primary inputs and outputs of the tool, the sources that supply the data needed to operate the tool, and the tool's limitations, including restrictions on the populations to which the tool applies.

Army Accession Planning Tools

The RDI, RRRM, and R-RST make up a suite of modeling tools that was designed to support Army accession planning. Figure 3.1 shows the intended use for each tool and the relationships across the tools. The RDI provides a measure of recruiting difficulty over the next 24 months.

---

[5] The DRM was originally developed in the 1980s to study the retention of active duty Air Force officers. See Glenn A. Gotz and John McCall, *A Dynamic Retention Model for Air Force Officers: Theory and Estimates*, Santa Monica, Calif.: RAND Corporation, R-3028-AF, 1984. Since then the model has been applied more broadly to include both officers and enlisted personnel, the AC and the RC, and all four military services.

[6] For more information on the RRRM's Regular Army predecessor, see David Knapp, Bruce R. Orvis, Christopher E. Maerzluft, and Tiffany Berglund, *Resources Required to Meet the U.S. Army's Enlisted Recruiting Requirements Under Alternative Recruiting Goals, Conditions, and Eligibility Policies*, Santa Monica, Calif.: RAND Corporation, RR-2364-A, 2018. For more information on the R-RST's Regular Army predecessor, see Orvis et al., 2018.

Figure 3.1. The RAND Suite of Modeling Tools to Inform Accession Planning

RDI	RRRM	R-RST
Predicts future recruiting difficulty	Considers effectiveness and efficiency of alternative recruiting resource and recruit selection scenarios	Considers first-term consequences of alternative recruit selection policies

The RRRM estimates the resource costs necessary to achieve a target accession goal conditional on the recruiting environment and Army-established recruit eligibility policies. The R-RST predicts first-term attrition, adverse personnel actions, and reenlistment as a function of reservists' characteristics at enlistment. Because attrition affects the need to recruit, the R-RST also serves as an input to the RRRM. Similarly, information from the RRRM can be leveraged by the R-RST to estimate the effect of changes in recruit cohort characteristics on recruiting, training, and replacement costs.

The Recruiting Difficulty Index

The RDI, developed by the RAND Arroyo Center in FY 2017, is a forecasting model that predicts the ease with which USAREC will be able to recruit enlisted personnel over the next 24 months.[7] The RDI estimates three outcomes that reflect recruiting difficulty. The first is the difference between signed enlistment contracts and the contract mission, as a percentage of the contract mission, for high school diploma graduates with Armed Forces Qualification Test scores of 50 or higher, also known as graduate alphas. The second outcome is the average number of days between signing a contract and being accessed into the Army, and the third is the percentage of training seats filled.

To estimate these outcomes, the RDI leverages data on economic conditions, adverse events, recruiting resources, and recruit eligibility policies. Economic conditions and adverse events enter the model as exogenous variables. Recruiting resources and recruit eligibility policies are treated as endogenous policy responses to the recruiting environment that, in turn, affect the three outcomes. The criterion of mean-squared prediction error is used when estimating the model to decide which variables to include as explanatory variables. Figure 3.2 summarizes the RDI's inputs and outputs.

---

[7] Wenger et al., 2019.

Figure 3.2. Inputs and Outputs of the Recruiting Difficulty Index
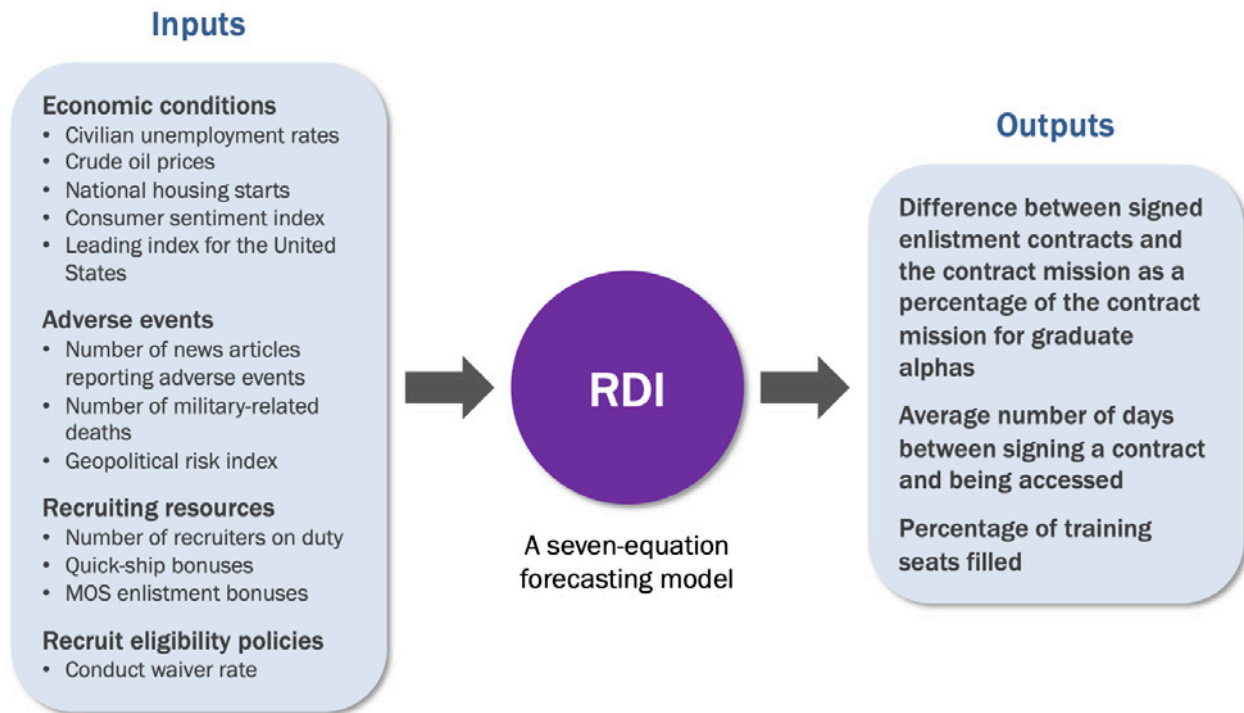
SOURCE: Wenger et al., 2019.
NOTE: Graduate alphas are enlistees with a high school diploma and an Armed Forces Qualification Test score of 50 or higher.

The input data are collected from a number of sources and broadly divided into three categories: military data, national economic and demographic data, and national military-related measures. Data on recruiters and contract missioning are drawn from databases maintained by USAREC. Data on contracts written and recruit accessions are collected from databases maintained by the U.S. Army Human Resources Command. Data on training seats available and filled are drawn from a database maintained by the U.S. Army Training and Doctrine Command, and data on the total training seats planned for future months are drawn from the accession mission letter issued by the Office of the Deputy Chief of Staff, G-1, U.S. Army. Data on Basic Military Pay are collected from the Defense Finance Accounting Service.

National economic and demographic data are also obtained from multiple sources. Unemployment rates are measured using the Current Population Survey administered by the U.S. Census Bureau. Crude oil prices reflect the price of West Texas Intermediate crude oil. National housing starts are based on projections from the Survey of Construction and the Building Permits Survey administered by the U.S. Census Bureau. The Consumer Sentiment Index is produced by the University of Michigan, and the Leading Index for the United States is produced by the Federal Reserve Bank of Philadelphia.

National military-related measures include reports of adverse events in the news, counts of military-related deaths, and a geopolitical risk measure. The number of Associated Press news

stories mentioning deployments, injuries and deaths, medical support and well-being, military crime and improprieties, and Middle Eastern conflicts was obtained using LexisNexis. Counts of military-related deaths are obtained from the iCasualties database, and the geopolitical risk measure is provided by Dario Caldara and Matteo Iacoviello.[8]

Although the RDI provides an estimate that bears on the flow of enlisted personnel into USAR—namely, the percentage difference between signed enlistment contracts for graduate alphas and the contract mission for them—the RRRM provides an estimate that is more directly relevant. As we explain in the next section, the contract production submodel of the RRRM predicts the total number of USAR contracts for enlisted personnel with no prior service of any kind.

The Reserve Recruiting Resource Model

The RRRM is an econometric model that determines the most efficient allocation of recruiters, enlistment incentives, and advertising expenditures to achieve USAREC's recruiting objectives, conditional on recruit eligibility policies and labor market conditions.[9] It was developed by the RAND Arroyo Center in FY 2020, adapted from a similar tool that had been developed for the Regular Army in FY 2017.[10]

The RRRM estimates two distinct but related submodels: a contract production function and a cost allocation function. The contract production function predicts the number of enlistment contracts as a function of recruiting resources, recruit eligibility policies, and labor market conditions. The cost allocation function computes the resourcing costs paid to achieve the fiscal year's enlistment contracts and accessions. An optimization algorithm is layered over the two submodels to determine the efficient allocation of recruiters, enlistment incentives, and advertising expenditures to achieve USAREC's recruiting objectives. Figure 3.3 summarizes the RRRM's inputs and outputs.

The input data are collected from a number of sources and broadly divided into three categories: military data, advertising data, and economic and demographic data. Military data, including the number of recruiters on duty, enlistment incentives, waiver rates, contracts written, and recruit accessions, are drawn from databases maintained by Human Resources Command and USAREC and from the accession mission letter issued by the Office of the Deputy Chief of Staff, G-1, U.S. Army. National television advertising data—both impressions and costs—are provided by the Army Marketing and Research Group and the Army's advertising agency. Data on labor market conditions are collected from multiple sources, including the U.S. Bureau of

---

[8] Dario Caldara and Matteo Iacoviello, *Measuring Geopolitical Risk*, Washington, D.C.: Board of Governors of the Federal Reserve System, International Finance Discussion Paper 1222, February 2018.

[9] Orvis et al., 2022.

[10] Knapp et al., 2018.

Figure 3.3. Inputs and Outputs of the Reserve Recruiting Resource Model

Inputs

Recruiting resources
- Number of recruiters on duty
- Enlistment incentives
- Advertising impressions and costs

Recruit eligibility policies
- Number of prior-service contracts
- Percentage of high-quality contracts
- Conduct waiver rate
- Medical waiver rate

Labor market conditions
- Woods and Poole projections of qualified military
- Civilian unemployment rates
- Consumer sentiment index
- Civilian-to-military wage ratio
- Minimum wage by state

RRRM

An econometric model for estimating contract production and costs with an optimization algorithm overlaid

Outputs

Number of enlistment contracts

Number of accessions

Resourcing costs

Efficient allocation of recruiting resources

SOURCE: Orvis et al., 2022.

Labor Statistics, the U.S. Department of Defense's *Compensation Greenbooks*, David Neumark at the University of California–Irvine, the University of Michigan, and the U.S. Census Bureau (Neumark, undated).

As we discuss in Chapter 4, the contract production submodel of the RRRM is most relevant for our purposes. It can be used to predict execution year accessions of enlisted personnel with no prior service, *given a specified resourcing plan*. The RRRM does not cover enlisted recruits with prior service of any kind or officer recruits from any source.

The Reserve Recruit Selection Tool

The R-RST estimates the effects of reservists' characteristics at enlistment on first-term attrition, adverse personnel actions, reenlistment, and the costs of recruiting, training, and replacement.[11] The set of characteristics analyzed by the tool is quite broad; it includes demographic factors, education level and aptitude, physical- and health-related factors, and medical and conduct waivers. Because recruit eligibility policies alter the characteristics of the recruit cohort, the R-RST enables the user to examine strategically the trade-offs among the behavioral and cost outcomes that may arise when changing these policies.

---

[11] Unpublished RAND research that developed a Reserve Component version of Orvis et al., 2018.

Like the RRRM, the R-RST was developed in FY 2020, having been adapted for the reserves from an analogous tool for the Regular Army that had been developed three years earlier.[12] As an initial step, logistic regression was used to estimate the relationships between recruit characteristics and each of the outcomes. Only those characteristics that were found to be statistically significant after controlling for other factors were incorporated in the tool itself. The user of the tool selects the desired distribution of characteristics for a particular recruit cohort. The tool, in turn, computes the weights that must be applied to each soldier in the cohort to produce the distribution selected. It then averages over the weighted soldiers to estimate the new cohort outcomes and costs. Note that the R-RST does not estimate causal effects; instead, it leverages correlations between recruit characteristics at enlistment and the behavioral and cost outcomes of interest. Figure 3.4 summarizes the R-RST's inputs and outputs.

Figure 3.4. Inputs and Outputs of the Reserve Recruit Selection Tool
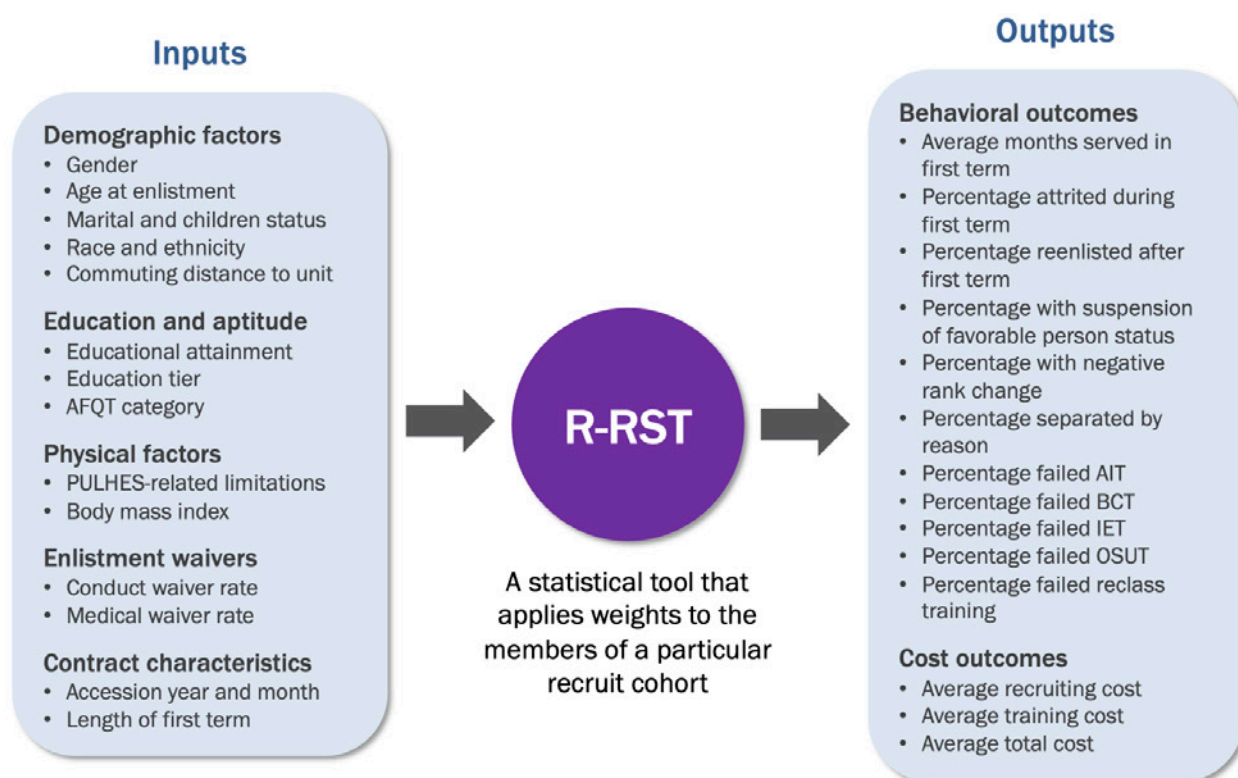
SOURCE: Unpublished RAND research that developed a Reserve Component version of Orvis et al., 2018.

NOTES: "Education tier" refers to whether the recruit has a traditional high school diploma (Tier 1) or a General Educational Development certificate (Tier 2). AIT = Advanced Individual Training; BCT = Basic Combat Training; IET = Initial Entry Training; OSUT = One Station Unit Training; PULHES = physical capacity/stamina (P), upper extremities (U), lower extremities (L), hearing and ears (H), eyes (E), and psychiatric (S).

---

[12] Orvis et al., 2018.

The data required to operate the R-RST are collected exclusively from U.S. Army sources. Personnel records are obtained from the Total Army Personnel Database. Data on contracts written and recruit accessions are collected from databases maintained by Human Resources Command, and data on enrollment in and completion of training classes are drawn from a database maintained by the U.S. Training and Doctrine Command. Annual estimates of recruiting and training costs per recruit are obtained from the Office of the Deputy Chief of Staff, G-1, U.S. Army. Additional cost data are obtained from the RRRM, which accounts for changes in recruiting resource requirements due to changes in recruit cohort characteristics, the recruiting environment, and the size of the recruiting requirement.

As we discuss in Chapter 4, the R-RST's first-term attrition estimates are most relevant for our purposes. While the RRRM can be used to predict the flow into USAR of enlisted personnel with no prior service of any kind, it cannot provide an estimate of the fraction of these recruits that will be retained during the 24-month forecast window. The R-RST provides such a retention estimate.

The Dynamic Retention Model

The DRM simulates the retention and cost effects of changes to the compensation system. It is based on a mathematical model of individual decisionmaking over the life cycle of the service member in a world with uncertainty and in which members have heterogeneous preferences for active or reserve service. Once estimated, the model parameters are used to simulate active and reserve retention in the steady state and in the transition to the steady state under alternative compensation systems, assuming no change in any other factor.

The RAND Arroyo Center updated and extended the DRM in FY 2018 to predict the effects of the Blended Retirement System on retention within USAR and ARNG and on the flow of service members from the Regular Army to USAR.[13] As a consequence, the DRM can now estimate the effects of compensation changes in one component, such as the Regular Army, on participation behavior and costs in other components, such as USAR. Figure 3.5 summarizes the DRM's inputs and outputs.

The RAND Arroyo Center estimated the parameters of the DRM using individual-level data on ARNG, Regular Army, and USAR members who joined the Army as non-prior-service entrants in 1990 or 1991. The Defense Manpower Data Center provided a longitudinal record of each member's active and reserve service through 2015, which permitted the researchers to track the participation behavior of the 1990 and 1991 cohorts over a period of up to 26 years. Average annual compensation by YOS was computed using data published by the Office of the Under Secretary of Defense for Personnel and Readiness, Directorate of Compensation and by the

---

[13] Asch, Mattock, and Hosek, 2019.

Figure 3.5. Inputs and Outputs of the Dynamic Retention Model

Inputs

Individual-level longitudinal service records
- Age
- Gender
- Service or component
- Pay grade
- Reserve category code
- Primary service occupation code
- Secondary service occupation code
- Duty service occupation code
- Unit identification code
- Entries to the armed services
- Exits from the armed services

Compensation
- Regular military compensation by YOS
- Reservist compensation by YOS
- Civilian wages by education level and years of work experience

DRM

A structural econometric model that simulates the retention effects of changes to the compensation system

Outputs

Participation behavior in each period
- Probability of continuing in the Regular Army
- Probability of transitioning from the Regular Army to the ARNG or USAR
- Probability of continuing in the ARNG
- Probability of continuing in USAR

Cost outcomes
- Continuation pay costs for USAR by fiscal year

SOURCE: Asch, Mattock, and Hosek, 2019.

Office of the Assistant Secretary of Defense, Reserve Affairs.[14] Average annual earnings for full-time male workers were computed by education level and years of work experience using data collected from the U.S. Census Bureau.

The DRM outputs most relevant for our purposes are the estimates of USAR retention and the estimates of USAR accessions for those with prior service in the Regular Army. Note that these estimates are *conditional on the compensation regime specified*. Despite its breadth, the DRM does not cover several segments of the reservist population. For example, reservist officers who have no prior service in the AC or who are in the medical, legal, or chaplaincy career fields fall outside the scope of the DRM.

Because the DRM and three recruiting tools were designed for purposes other than forecasting USAR end strength, they do not cover the entire SELRES population. In the next section, we describe an ML approach developed as part of this study to address the gaps left by existing modeling tools.

---

[14] Office of the Under Secretary of Defense for Personnel and Readiness, Directorate of Compensation, *Selected Military Compensation Tables*, Washington, D.C., 2007; Office of the Assistant Secretary of Defense, *Official Guard and Reserve Manpower Strengths and Statistics: FY 2007 Summary*, Washington, D.C., January 2007.

The Machine Learning Model

The modeling tools described in the preceding sections have many strengths. First, the tools are rooted in a long history of econometric manpower modeling at RAND and across the Army and Department of Defense communities. They have been applied in a variety of settings, informing important decisions, such as the structure of the Blended Retirement System, the allocation of recruiting resources, and the calibration of recruit eligibility policies. The tools provide much more than estimates of accessions, separations, and end strength; they also predict how policy changes may affect recruiting and retention outcomes, holding all other factors constant. Some, like the DRM, incorporate behavioral models that allow the analyst to make inferences about *causal relationships* and estimate outcomes under counterfactual policy scenarios.

While understanding the underlying mechanisms is valuable, the capability often requires significant start-up costs in terms of data provision, parameter estimation, and calibration against the ever-changing real world. The primary goal of this project requires a much simpler model, one that focuses only on recruiting and continuation to predict 24-month end strength. Modeling the underlying mechanisms might not be necessary. In this section we introduce a proof-of-concept, ML-based, data-driven model that can be used to estimate many of the personnel flows into and out of USAR in a more economical fashion.

As described in previous sections, and in further detail in Chapter 4, the existing set of modeling tools does not cover the entire SELRES population. With sufficient historical data an entirely data-driven model could theoretically predict the entry and exit decisions of any population while remaining agnostic about the factors driving those decisions, at least to the extent that historical SELRES careers are good predictors of future careers.[15] Accordingly, an ML approach could address some of the gaps left by existing modeling tools, as well as provide an alternative approach to estimating some of the personnel flows that are already covered.

The proof-of-concept ML model developed for this project is similar to analysis conducted by the Institute for Defense Analyses in 2019 but makes adjustments that render the approach suitable for our specific problem: predicting 12- and 24-month retention of service members in SELRES.[16] ML encompasses a broad array of techniques, but our proof-of-concept model uses a neural network to predict individual-level retention behavior because we felt that this type of model would best be able to digest complex time series data. Neural network models take in

---

[15] Note that we consider only SELRES careers themselves in this discussion, not any exogenous economic factors that might also influence the decisions of individual reservists. Additional factors could be incorporated into future analyses.

[16] Julie Pechacek, Alan Gelder, Cullen Roberts, Joe King, James Bishop, Michael Guggisberg, and Yev Kirpichevsky, "A New Military Retention Prediction Model: Machine Learning for High Fidelity Forecasting," paper presented at the 94th Annual Conference of the Western Economic Association International, San Francisco, June 2019.

a set of characteristics (input data) and automatically fit complex nested equations to predict outcomes of interest. In this case, the neural network takes in static and time-varying characteristics of individual SELRES members and delivers estimated relative retention (dropout) probabilities. We train the neural network using the demographic characteristics and career histories of a sample of the SELRES members appearing in our data set. The neural network then optimizes its parameters to estimate the probability that a member will have exited USAR by several future deadlines. Using this set of individual probability estimates, we can then compute the expected number of retained individuals (losses). Figure 3.6 summarizes the inputs and outputs of the model. A more detailed description of the neural network is provided in Appendix C.

Figure 3.6. Inputs and Outputs of the Machine Learning Model
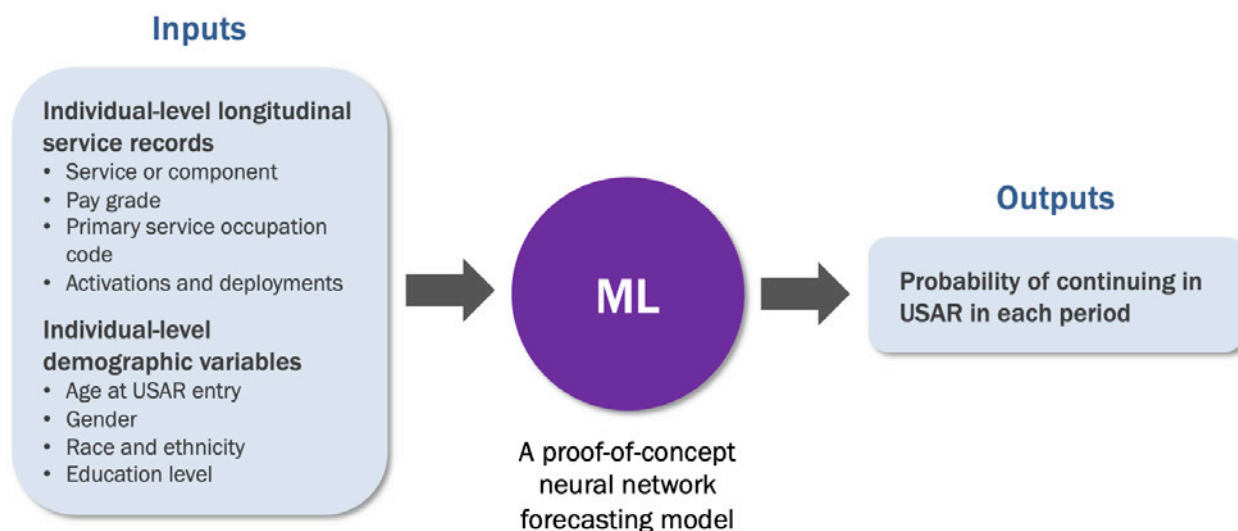
It is important to note that the ML model is not well suited to performing counterfactual analysis. Some level of counterfactual analysis may be possible by biasing the inputs to the model and performing sensitivity analyses, but this approach is unlikely to produce robust results. For example, if the user is interested in understanding the impact of a future economic recession on retention behavior, the model could be retrained by oversampling SELRES members who joined in 2009 and observing how the predictions change. However, given all the other differences in the recruiting environment between 2009 and the present, such results might not be particularly accurate.

Figures 3.7 and 3.8 present results from an illustrative application of the proof-of-concept ML model, which was used to predict 12- and 24-month retention of enlisted reservists who were in SELRES in September 2015, had entered SELRES between FY 2002 and FY 2013, and

had no prior service outside SELRES.[17] There were 78,082 such individuals with sufficiently populated data. Figure 3.7 compares estimated and actual retention of these reservists over the 12 months that elapsed between September 2015 and September 2016. Figure 3.8 presents the same over the 24 months between September 2015 and September 2017. In both figures, the orange line represents the actual number of retained reservists, while the blue bar represents the forecast generated using the ML approach.

Figure 3.7. Estimated Versus Actual Retention of Enlisted Reservists with No Prior Service, September 2015–September 2016
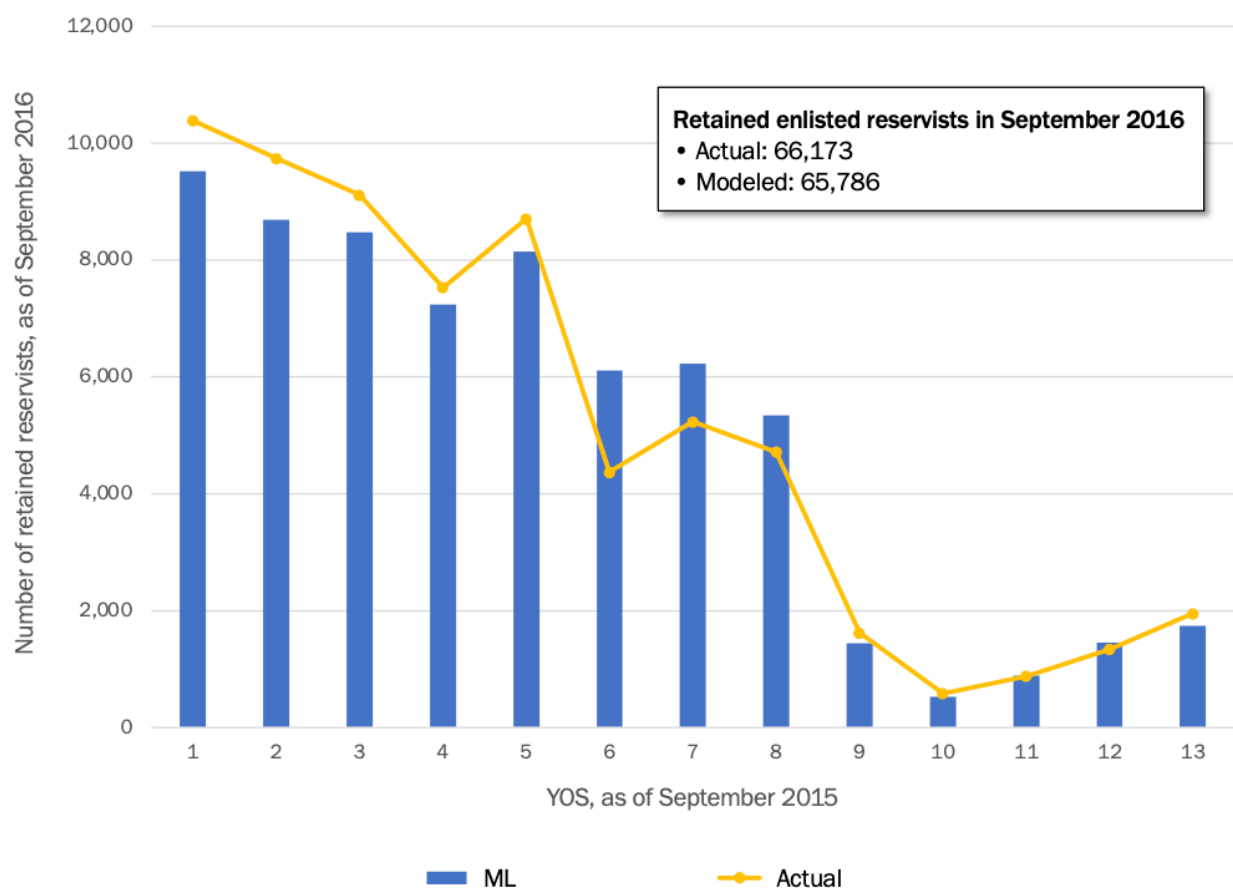
NOTES: YOS is a measure of how long an individual has been in military service and is based on the entry point of service reported in the WEX file. Reserve service years are counted in the same way as active service years. Thus, the variable is indicative of presence in military service rather than of qualifying retirement years. Adjustment to reserve service years based on documented qualifying events is possible, though documentation is inconsistent. This issue is outside the scope of this illustrative example.

---

[17] As was explained in Chapter 2, we cannot observe histories that precede FY 1990 because the WEX file is truncated.
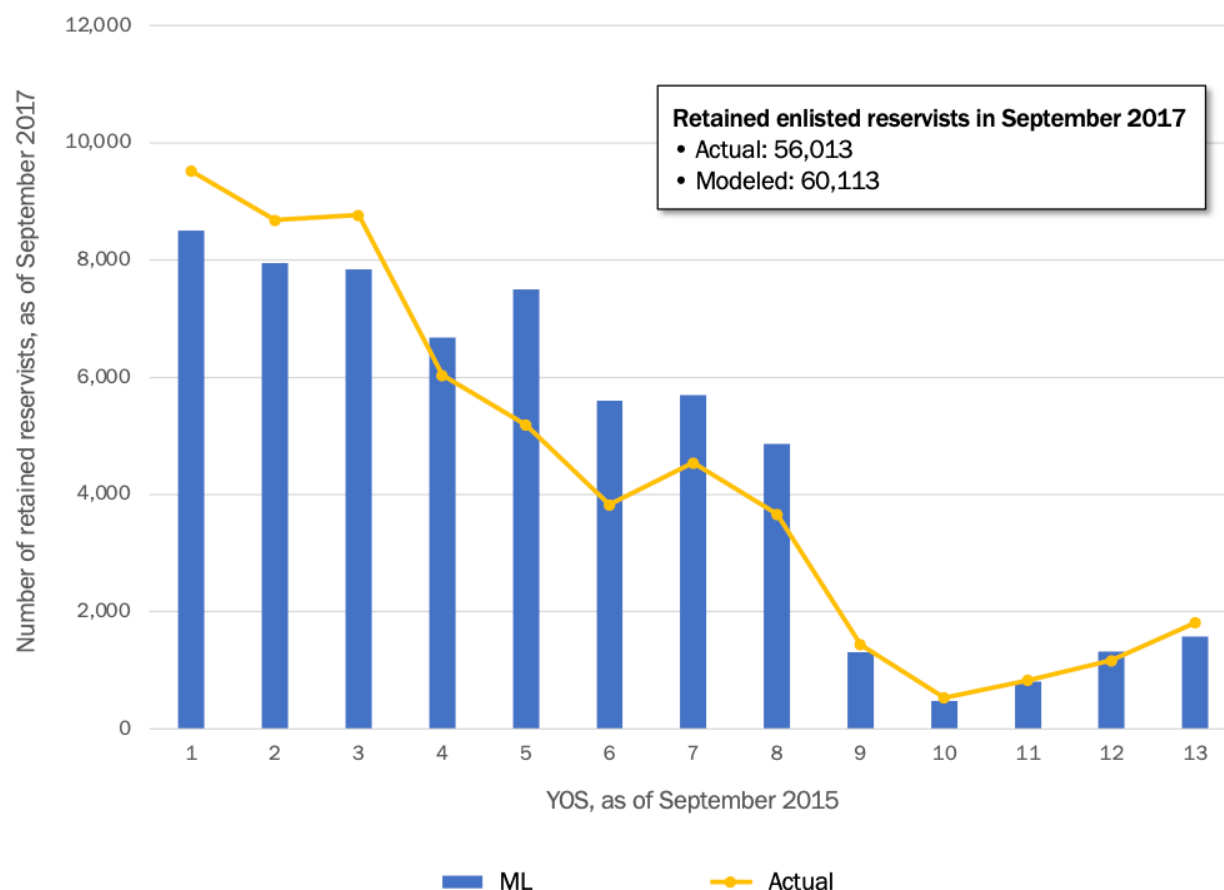
Figure 3.8. Estimated Versus Actual Retention of Enlisted Reservists with No Prior Service, September 2015–September 2017

NOTES: YOS is a measure of how long an individual has been in military service and is based on the entry point of service reported in the WEX file. Reserve service years are counted in the same way as active service years. Thus, the variable is indicative of presence in military service rather than of qualifying retirement years. Adjustment to reserve service years based on documented qualifying events is possible, although documentation is inconsistent. This issue is outside the scope of this illustrative example.

The overall retention predictions are reasonably accurate. For the 12-month period, the neural network predicted retention of 65,786 members when the actual count was 66,173 members, a difference of only 0.6 percent. We caution, however, that this example is probably overly optimistic. For the 24-month period, the model predicted retention of 60,113 members when the actual count was 56,013 members, an overestimate of 7.3 percent.[18]

_____

[18] A complete goodness-of-fit analysis and evaluation of model errors is outside the scope of this illustrative exercise. Future implementation of an ML model as part of the IMC would require retraining of the model and adjustment of parameters, as well as consideration of additional variables that could improve model fit. Formal evaluation of the model's predictive capability and errors should be conducted on this fully developed model. Such an analysis should also evaluate the extent to which errors are correlated across the various models used to predict accessions and retention for different populations covered by the IMC. Full implementation of the IMC should include estimation of the uncertainties of individual models, as well as an estimate of overall uncertainty, expected to be based on a sum of the uncertainty in quadrature of the individual pieces.

Although the framework described here covers only separations from USAR, it can be extended to cover certain categories of accessions—namely, the inflows from populations that have prior military experience, and thus available career history data. For example, the neural network could be modified to output a matrix that predicts both the probability that a reservist separates from USAR and the immediate destination of the separating reservist (e.g., civilian life, Army AC, ARNG). We could then train clones of our neural network on Army AC and ARNG populations and use the trained networks to predict the number of service members who will transfer from the Army AC and ARNG to SELRES over the next 24 months.

While similar ML approaches can be applied to other portions of the SELRES end strength prediction problem, it is important to note that the test case considered here—enlisted reservists whose data fields were complete—may represent a best-case scenario for an ML model. Because the population of enlisted SELRES members with sufficient observable career histories is large, the data sample used to train the neural network could also be large. This may not be the case for other applications.

Chapter Summary

In this chapter we have described five models with a wide range of capabilities and varying abilities to forecast USAR end strength. The primary differences among these models are their purposes and hence the populations they cover, the degree to which they are able to estimate the effects of alternative personnel policies on recruiting and retention, and the appropriateness of using them for the purpose of generating forecasts of the SELRES population. While each model has applicability, none covers the full range of personnel flows into and out of USAR, and some may be inappropriate for the purpose of developing the forecasts requested by OCAR. An ML approach could provide the desired forecasts but would not permit assessments of the effects of changing the allocation of recruiting resources, recruit selection policies, or the compensation system on the SELRES population.

4. The Integrated Modeling Concept

Having reviewed the flows of personnel into and out of SELRES and the collection of available models, we turn our attention to constructing the IMC, a plan for integrating outputs from the various models to generate 24-month end strength forecasts. Figure 4.1 depicts the overarching approach to developing these forecasts. It mirrors Figure 2.1 but shifts the time frame forward so that the blue circle represents the current SELRES population and the orange circle represents the population 24 months into the future.

Figure 4.1. Forecasting SELRES Exit and Entry
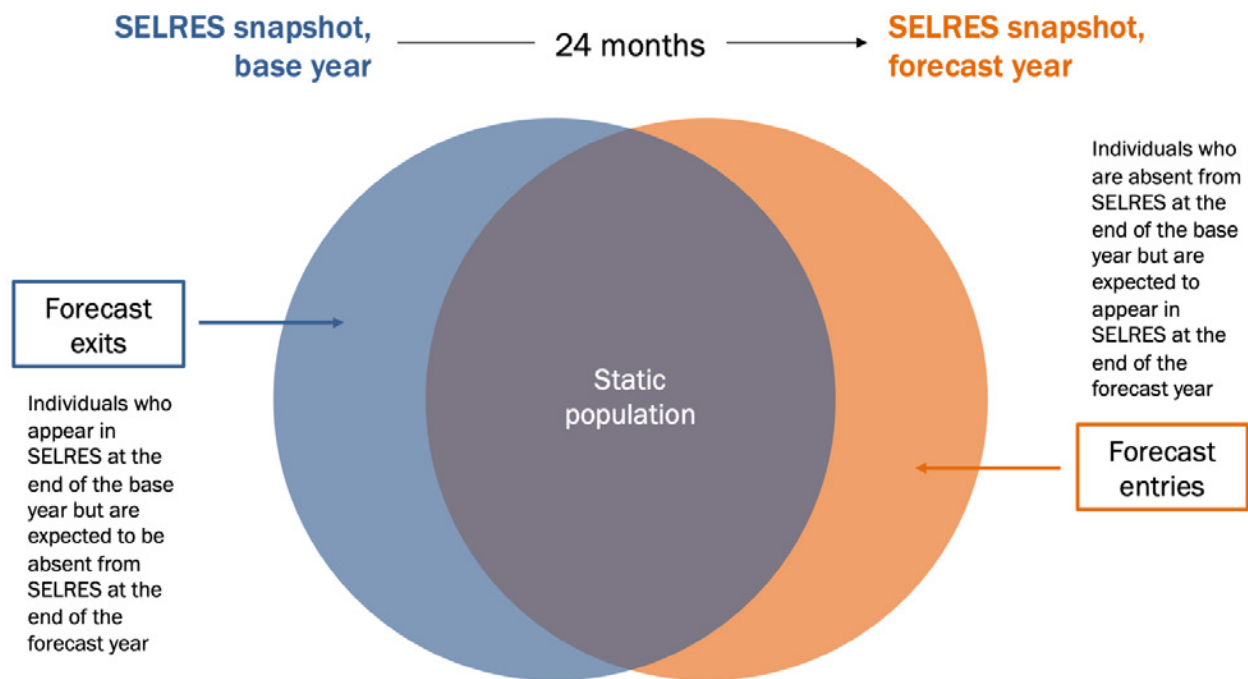
We begin with the current, or *base year*, population and partition it to align with the modeling tools that generate retention forecasts. As explained in the next section, we aggregate these forecasts to estimate the fraction of the base year population that will exit SELRES over the next 24 months. We subtract the forecast exits, represented by the blue crescent on the far left, from the base year population to obtain an estimate of the *static population*: individuals who appear in SELRES in the base year and are predicted to remain in SELRES over the next 24 months.

We then turn our attention to forecasting *retained entrants*: individuals who enter SELRES during the 24-month period and are retained long enough to be present in SELRES at the end of the forecast year. This population is represented by the orange crescent on the far right of

Figure 4.1. As explained later in this chapter, we partition the set of potential entrants to align with the modeling tools that generate accession forecasts, pairing each of the accession tools with a retention tool that applies to the same subpopulation. We aggregate the paired forecasts across subpopulations to estimate the number of retained entrants. By adding the forecast entries to the static population, we obtain the desired 24-month forecast of USAR end strength, represented by the orange circle.

The remainder of this chapter details how we partitioned the base year population and set of potential entrants to permit assignment of the modeling tools reviewed in Chapter 3. It also provides a quantitative estimate of the relative prevalence of each tool within the IMC by computing the share of exits and share of retained entrants that would have been covered by each tool had the base year been FY 2016. We include a numerical example that uses the actual personnel flows during the period FY 2016–FY 2018 to demonstrate how the IMC would aggregate outputs from the various modeling tools to generate a 24-month end strength forecast.

Summary of Populations Covered by Each Modeling Tool

Partitioning the base year population and set of potential entrants in a manner that permits the assignment of modeling tools requires a clear and detailed understanding of the populations covered by each tool. Drawing from the information presented in Chapter 3, Table 4.1 provides a structured account of these populations.

Each model covers the populations indicated in Table 4.1 only to the extent that the requisite data are available. For example, the DRM can predict the retention behavior of an enlisted reservist who began his service in the Army AC, did not return to the AC having separated, and has no history of service outside the Army—but only if the available data contain a complete record of the reservist's service history. Similarly, the ML approach can forecast accessions by reservists with prior service, but only for those whose prior service is documented in the available data.

The model assignments shown in the next two sections track the populations described in Table 4.1. In some cases, populations overlap, meaning that two or more modeling tools cover some of the same exits or entrants. In those cases we choose the modeling tool with the greater capability to assess the effects of policy changes but note that an alternative choice exists.[1]

---

[1] Capability to assess the effects of policy changes is not the only criterion one could use to select among modeling tools. Other criteria include the accuracy of the predictions generated by the tool and the cost of operating the tool. In this study, we chose to apply the policy analysis criterion because doing so was feasible and of interest to OCAR. Because the RRRM and R-RST were under development while this study was underway, and operating the DRM was outside the study's scope, application of an accuracy or cost criterion was not an option. However, these alternatives should be considered in any future implementation of the IMC.

Table 4.1. Populations Covered by Each Modeling Tool

| Modeling Tool	Retention	Accession
RRRM	• None	• Enlisted personnel with no prior service of any kind
R-RST	• Enlisted personnel who have no prior USAR service; join USAR directly from either civilian life, the ARNG, or an RC outside the Army; and are in their first term with USAR	• None
DRM	• Enlisted personnel with a history of USAR service only	• Enlisted personnel with prior service in USAR only
	• Enlisted personnel who began their service in the Army AC, did not return to the AC having separated, and have no history of service outside the Army	• Enlisted personnel who began their service in the Army AC, did not return to the AC having separated, and have no history of service outside the Army
	• Officers who began their service in the Army AC; did not return to the AC having separated; have no history of service outside the Army; have never been enlisted; and are not in the medical, legal, or chaplaincy career fields	• Officers who began their service in the Army AC; did not return to the AC having separated; have no history of service outside the Army; have never been enlisted; and are not in the medical, legal, or chaplaincy career fields
ML	• Reservists of any type (enlisted, officer, or warrant officer)	• Reservists of any type (enlisted, officer, or warrant officer) who have prior service

NOTES: The table omits the RDI because the RRRM provides an estimate that is more directly relevant to the flow of enlisted personnel into USAR (see Chapter 3 for details). While the R-RST includes first-term attrition and reenlistment by prior service-civil life gains (as indicated in the table), the IMC pairs the R-RST with the RRRM and, in doing so, applies the R-RST only to enlisted recruits with no prior service of any kind. See Figure 4.3 and the accompanying text for additional details. There exists a version of the DRM that covers accessions by officers with prior service in USAR only, but that version combines ARNG and USAR officer populations such that the parameter estimates are likely driven by the behavior of ARNG officers. For additional details, see Michael G. Mattock, Beth J. Asch, and James Hosek, *Making the Reserve Retirement System Similar to the Active System: Retention and Cost Estimates*, Santa Monica, Calif.: RAND Corporation, RR-530-A, 2014.

Forecasting Exits

Figure 4.2 depicts the partition on the base year population that permits assignment of the modeling tools that predict retention—namely, the DRM and the ML approach. The R-RST does not lend itself to predicting exits from the base year population because its capability is limited to first-term attrition.

The initial node at the top of Figure 4.2 represents forecast exits from SELRES, coinciding with the blue crescent in Figure 4.1. In a sequence of six steps, Figure 4.2 partitions the population of forecast exits to align with the modeling tool populations described in Table 4.1. Each terminal node bears one of three colors: purple to indicate a DRM assignment, green to indicate an ML assignment, and red to indicate no assignment. The subpopulations represented by the terminal nodes are mutually exclusive and jointly exhaustive; that is, every forecast exit belongs to one, and only one, terminal node.

Figure 4.2. The Model Integration Plan for Forecasting Exits
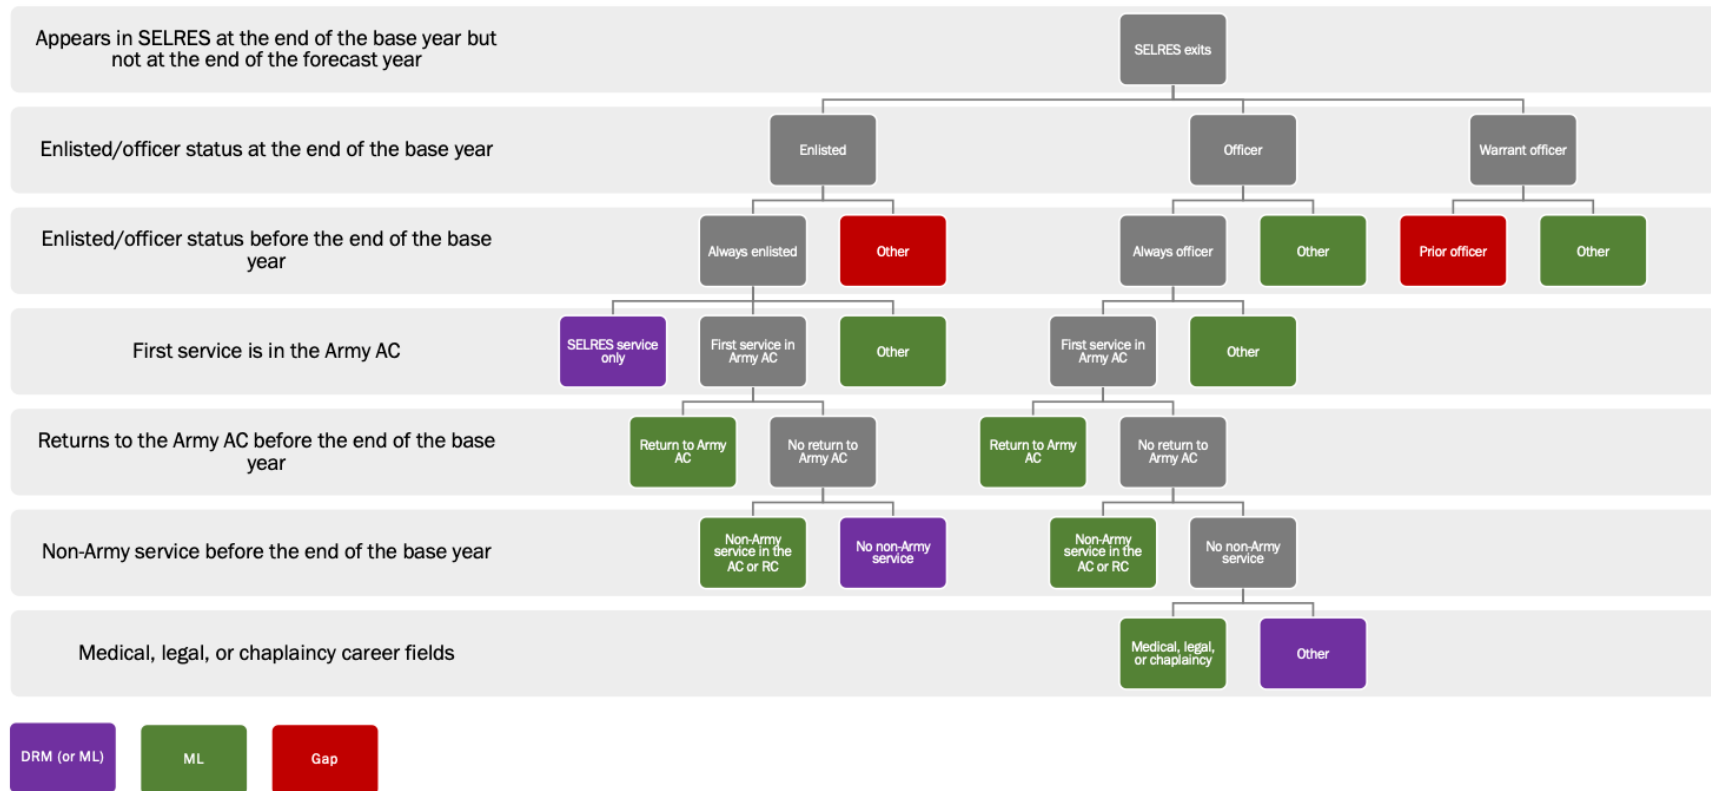
NOTE: Gray nodes are not terminal, meaning that the populations they represent will be divided further.

30

The first step in the figure separates individuals based on enlisted member/officer/warrant officer status at the end of the base year. This division is necessary because the DRM treats enlisted personnel and officers separately and excludes warrant officers altogether. The "Enlisted," "Officer," and "Warrant officer" nodes are gray to indicate that the populations will be divided further.

In the second step, the three populations are subdivided based on their prior status. Records showing enlisted reservists with prior experience as officers or warrant officers are considered data errors, with the corresponding node colored red to indicate that neither the DRM nor ML can predict these exits. Exits by officers with prior enlisted or warrant officer experience, shown in green, will be estimated using ML because the DRM excludes this population. Records showing warrant officers with prior experience as officers are considered data errors and colored red; all other warrant officer exits are colored green to indicate they will be estimated using ML.
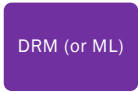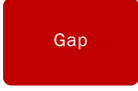
The remaining steps subdivide the two populations represented by the gray nodes: enlisted reservists who have always been enlisted and reservist officers who have always been officers. The divisions track the DRM populations described in Table 4.1. We assign the DRM wherever possible to retain the model's capacity for estimating the effects on retention of changes to compensation policy, including changes that have not occurred in the past, such as was done prior to implementation of the Blended Retirement System. The purple node marked "SELRES service only" represents exits by enlisted personnel with a history of USAR service only. The purple node marked "No non-Army Service" represents enlisted personnel who began their service in the Army AC; did not return to the AC having separated; and have no history of service outside the Army. The purple node marked "Other" represents officers who began their service in the Army AC; did not return to the AC having separated; have no history of service outside the Army; have never been enlisted; and are not in the medical, legal, or chaplaincy career fields. We assign the ML approach to all other terminal nodes.

The DRM requires that the user make an assumption about the future of the military compensation system. The user may assume that the system currently in place will simply persist, but, regardless, an assumption must be made in order for the DRM to generate retention estimates. The ML approach does not require any such assumption, at least not explicitly; the model implicitly assumes that the conditions, including the compensation system(s), that prevailed during the time period covered by the data used to train the neural network will remain going forward. In addition, the ML approach has the potential to generate predictions in a more economical fashion, as discussed in Chapter 3. Hence, if the effects of compensation changes on USAR end strength are not of primary interest, substituting the ML approach for the DRM should be considered.

The Prevalence of Each Modeling Tool

Table 4.2 uses the SELRES populations in September 2016 and September 2018 to compute the share of exits that would have been covered by each modeling tool had FY 2016 been the

Table 4.2. Share of Exits Covered by Each Modeling Tool, FY 2016–FY 2018

Modeling Tool	Population	Exit Count	Exit Share (%)
DRM (or ML)	Enlisted	37,195	74.6
	Officer	728	1.5
	Warrant officer	0	0.0
	Total	**37,923**	**76.1**
ML	Enlisted	6,118	12.3
	Officer	5,269	10.6
	Warrant officer	443	0.9
	Total	**11,830**	**23.7**
Gap	Enlisted	77	0.2
	Officer	0	0.0
	Warrant officer	28	0.1
	Total	**105**	**0.2**
Total exit count and share		**49,858**	**100.0**

base year. The counts shown are actual numbers, not estimates generated by the modeling tools. The DRM covers over 76 percent of the exits that occurred during the 24-month period. However, the DRM's prevalence is due largely to its capacity to forecast retention among enlisted reservists; the DRM covers only 12 percent of officer exits and none of the warrant officer exits. A very small fraction of exits from the base year population, 0.2 percent, cannot be addressed by either the DRM or the ML approach.

Forecasting Retained Entrants

This section explains how we partitioned the set of potential entrants to permit assignment of the modeling tools that generate accession forecasts. As before, the divisions track the modeling tool populations described in Table 4.1. We begin with enlisted entrants and continue with reservists who join SELRES as officers and warrant officers.

Enlisted Entrants

The integration plan for forecasting retained entry of enlisted reservists uses all four modeling tools. The RRRM and R-RST are paired to forecast the accession and attrition, respectively, of enlisted personnel with no prior service of any kind. This assignment is represented by the blue node in the fourth row of Figure 4.3. While the R-RST can predict first-term attrition of both new recruits and prior service–civil life gains (as indicated in Table 4.1), the pairing of the R-RST with the RRRM restricts the population to which the R-RST is applied.

Figure 4.3. The Model Integration Plan for Forecasting Retained Entry
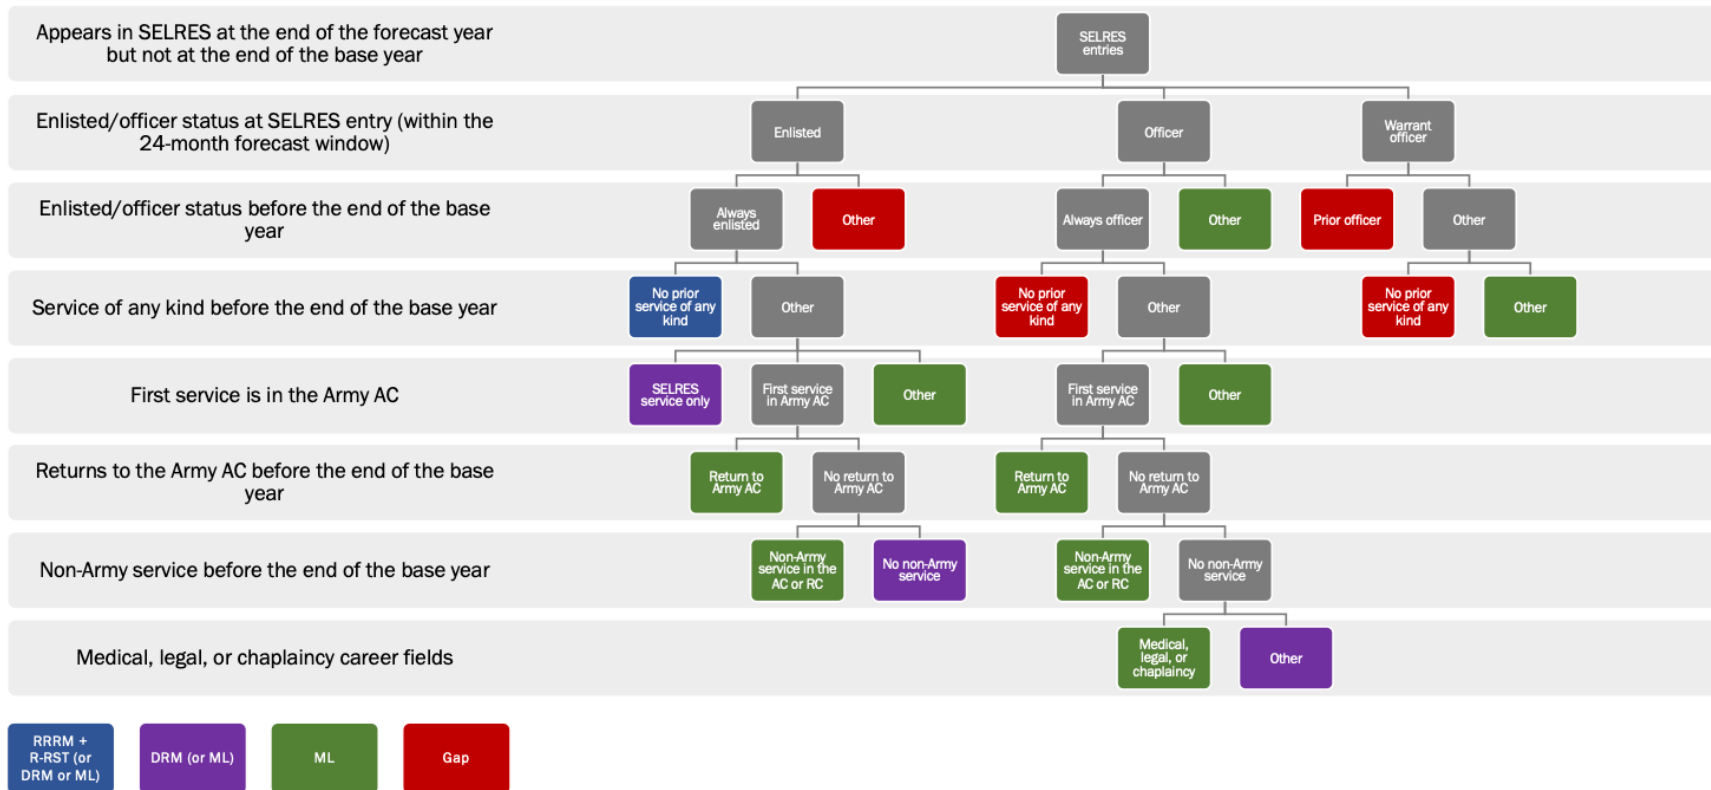
NOTE: Gray nodes are not terminal, meaning that the populations they represent will be divided further.

Either the DRM or ML can be used in lieu of the R-RST to estimate retention of the entrants predicted by the RRRM. The choice of modeling tool depends on the policies of primary interest. If changes to recruit eligibility policies—such as medical and conduct waivers, quality targets, and enlistment incentives—are the focus, the R-RST is the better choice. If changes to compensation policy in either the reserves, national guard, or AC are at the forefront, the DRM may be the more appropriate choice. If policy changes are not at issue, ML may be the preferred route.[2]

As is shown in Figure 4.3, accession and retention of enlisted personnel with prior service are covered by the DRM and the ML approach, with the DRM assigned whenever possible to retain the model's capacity for estimating the effects of changes to compensation policy. The DRM assignments track the populations described in Table 4.1. The purple node marked "SELRES service only" represents retained entry by enlisted personnel with prior service in USAR only, and the purple node marked "No non-Army service" represents enlisted personnel who began their service in the Army AC, did not return to the AC having separated, and have no history of service outside the Army. The ML approach is assigned to all other terminal nodes.

Officer Entrants

The integration plan for forecasting retained entry of officers uses only two modeling tools: the DRM and the ML approach. Neither the RRRM nor the R-RST considers recruits who join SELRES as officers. As is indicated in Table 4.1, the DRM covers only those officers who began their service in the Army AC; did not return to the AC having separated; have no history of service outside the Army; have never been enlisted; and are not in the medical, legal, or chaplaincy career fields. These restrictions leave the purple node in the bottom row of Figure 4.3. The IMC applies the ML approach to all other officers as long as they have service history. A gap remains, as indicated by the red node marked "No prior service of any kind." Between September 2016 and September 2018, there were 631 such recruits—approximately 11 percent of all reservists entering SELRES as officers during this period.[3]

As before, ML can be substituted for the DRM. The choice of whether to do so depends on the policies of interest. The effects of changes in compensation policy are handled best by the DRM, but ML may be a more efficient choice when there is no need to consider these policy changes.

---

[2] As was explained in Chapter 3, the RRRM estimates are drawn from the contract production submodel, which predicts execution year accessions of enlisted personnel with no prior service *given a specified resourcing plan*. This fact introduces two complications: (1) it requires the user to make an assumption about the allocation of recruiting resources that will prevail during the 24-month forecasting period; and (2) it requires that additional care be taken when pairing the RRRM with the R-RST, DRM, or ML approach. Because the R-RST was designed to work with the RRRM (see Figure 3.1), aligning the assumptions of the two modeling tools is likely to be straightforward. The same might not be true when pairing the RRRM with either the DRM or the ML approach.

[3] To mitigate this gap, one might consider augmenting the ML approach by drawing in data on Army Reserve Officer Training Corps members. We leave this and any other attempts to mitigate the gaps identified in this report to future work.

Warrant Officer Entrants

Figure 4.3 shows that the integration plan for forecasting retained entry of warrant officers uses ML only. None of the other modeling tools cover this small population of reservists. Records showing warrant officers with prior experience as officers are considered data errors, with the corresponding node colored red to indicate a gap in coverage. Warrant officers with no prior service of any kind are also left uncovered because the ML approach cannot forecast accessions in the absence of a service history.

The Prevalence of Each Modeling Tool

Table 4.3 uses the SELRES populations in September 2016 and September 2018 to compute the share of retained entry that would have been covered by each modeling tool had FY 2016 been the base year. The counts shown are actual numbers, not estimates generated by the modeling tools.

Table 4.3. Share of Retained Entry Covered by Each Modeling Tool, FY 2016–FY 2018

Modeling Tool	Population	Retained Entry Count	Retained Entry Share (%)
RRRM + R-RST (or DRM or ML)	Enlisted	20,220	50.1
	Officer	0	0.0
	Warrant officer	0	0.0
	Total	**20,220**	**50.1**
DRM (or ML)	Enlisted	10,172	25.2
	Officer	565	1.4
	Warrant officer	0	0.0
	Total	**10,737**	**26.6**
ML	Enlisted	3,871	9.6
	Officer	4,513	11.2
	Warrant officer	198	0.5
	Total	**8,582**	**21.3**
Gap	Enlisted	157	0.4
	Officer	631	1.6
	Warrant officer	2	0.0
	Total	**790**	**2.0**
Total exit count and share		**40,329**	**100.0**

The RRRM paired with the R-RST covers about half of the reservists who entered and were retained during the 24-month period. Among enlisted reservists, the RRRM and R-RST cover nearly 60 percent. Retained entry among officers is covered largely by the ML approach. Between September 2016 and September 2018, approximately 5,700 officers entered SELRES and were

retained through the end of the 24-month period. Of these, 79 percent would have been covered by the ML approach, and 10 percent would have been covered by the DRM, with 11 percent remaining as a gap. The gap is associated with the 631 officers with no prior service of any kind.

Numerical Example

We conclude this chapter with a numerical example that uses the personnel flows reported in Tables 4.2 and 4.3 to demonstrate how the IMC would aggregate outputs from the various modeling tools to generate a 24-month end strength forecast. Recall that the counts reported in the two tables are actual numbers, not estimates generated by the modeling tools.

On September 30, 2016, there were 198,273 individuals in SELRES, of which 160,998 were enlisted, 33,868 were officers, and 3,407 were warrant officers. To obtain a forecast of USAR end strength on September 30, 2018, we would

1. subtract the exits estimated by the DRM (37,923) and the ML approach (11,830), with an ad hoc deduction to account for the gap (105)
2. add the retained entries estimated by the RRRM and R-RST (20,220), the DRM (10,737), and the ML approach (8,582), with an ad hoc addition to account for the gap (790).

After subtracting the total estimated exits (49,858) and adding the total estimated retained entries (40,329), we obtain an estimate of USAR end strength on September 30, 2018, of 188,744. Again, this estimate was not constructed with outputs from the modeling tools; actual population numbers were used as proxies for the outputs.

The same process can be used to construct 12-month forecasts or to generate forecasts by subpopulation. For example, there were 33,868 officers in SELRES on September 30, 2016. After subtracting the officer exits estimated by the DRM (728) and the ML approach (5,269), we are left with 27,871 officers. We then add the officer entries estimated by the DRM (565) and the ML approach (4,513) and make an ad hoc adjustment to account for the gap (631) to obtain an estimate of USAR officer end strength on September 30, 2018, of 33,580.

We note that the subpopulation forecasts do not account for transitions between subpopulations. For example, the officer forecast described above does not account for reservists who transition from enlisted member to officer or from warrant officer to officer during the 24-month time frame. Consequently, the estimate of USAR officer end strength (33,580) falls short of the actual count (34,302) by 722 reservists. More generally, we caution against applying the process described in this chapter to subpopulations that are small; the accuracy of such forecasts is likely to be low, particularly if transitions into and out of the subpopulation are prevalent or if the subpopulation includes gaps that are managed in an ad hoc fashion.

Chapter Summary

This chapter has described

1. how one can partition the current (base year) SELRES population to align with the modeling tools that predict retention
2. how one can partition the set of potential entrants to align with the modeling tools predict accession
3. how the mapping of subpopulations to modeling tools can be used to integrate outputs from the tools and construct 24-month forecasts of SELRES end strength.

The IMC relies heavily on the DRM to predict retention among enlisted reservists and on ML to predict retention among reservist officers and warrant officers. Only 0.2 percent of separations from the SELRES cannot be addressed by any of the modeling tools. Enlisted accessions are estimated by the RRRM and R-RST and, to a lesser extent, by the DRM and ML; officer and warrant officer accessions are estimated almost entirely by ML. Two percent of SELRES accessions cannot be addressed by any of the modeling tools.

The IMC performs best when forecasting total SELRES end strength. Because the modeling tools do not predict changes in rank, end strength forecasts for subpopulations, such as enlisted reservists or reservist officers, are likely to be less accurate than the total end strength forecast.

5. Conclusion

While providing an accurate prediction of 24-month end strength for USAR may seem like a simple requirement, the many paths into and out of the reserves outlined in this report illustrate how complex the task really is. Several modeling tools with relevant capabilities are available, but none was designed to forecast end strength. Most were developed as policy analysis tools; they predict accession and/or retention outcomes under alternative recruiting resource allocations (RRRM), recruit selection criteria (R-RST), and compensation systems (DRM). Nevertheless, with the appropriate care and caution, the estimates generated by these tools can be combined to construct the desired end strength forecasts. A detailed plan for doing so, the IMC, is the primary product of this project. Because some of the tools were under development while this project was underway, we did not implement the IMC—that is, we did not operate the tools or integrate their outputs.

The complexity of the IMC required to fully articulate 24-month end strength depends on a variety of factors, such as the level of granularity required and the level of interest in modeling effects of policy decisions that may be under consideration. It may be possible to forecast overall end strength with greater accuracy than, for example, estimates for a specific Military Occupational Specialty (MOS), rank, or demographic group. Econometric models that enable exploration of the effects of changes to compensation policies or recruiting bonuses may be more complex than those that perform only projections based on assumed continuation of the status quo.

The IMC presented in this report represents one possible approach to solving this problem, and aims to find a balance among complexity, versatility, and applicability. Two key considerations guided the selection of modeling tools applied in the IMC. First, the IMC aims to leverage prior work that was sponsored by USAR and the broader Army by applying existing modeling tools whenever possible. Second, the IMC preferentially includes tools that support exploration of alternative policy decisions, such as those related to alternative recruiting or compensation policies, when such options are readily available. These decisions were made to maximize the versatility and future applicability of the IMC to support counterfactual analysis and planning into the future.

Although the approach presented here offers an optimal capability, it is not the only alternative for modeling the personnel flows required to make 24-month end strength predictions. Even limiting alternatives to those that use the models described in Chapter 3, multiple approaches with different strengths and weaknesses could be considered. The IMC described in this report relies heavily on the DRM for retention modeling because of its expansive capabilities and long history of successful use. An alternative modeling concept might replace this approach with one that more heavily favors data-driven ML models, possibly reducing

38

requirements for data collection and complex parameter estimation. However, such a choice would also limit the applications of the IMC and would not support the same kinds of counterfactual analyses.

The predictive capacity of any data-driven ML model is limited by the data used to train the algorithm. If a particular scenario or compensation scheme has not occurred in the past, the algorithm has no basis on which to formulate a retention forecast. The DRM is different in that it estimates the parameters of a structural model, which can then be used to forecast the effects of changes in compensation, regardless of whether they have occurred previously. This distinction between the two approaches suggests that the ML approach may be the better option when compensation (for both the AC and RC) is not expected to change in the foreseeable future or is the same as a compensation scheme that appears in the training data. The DRM may be the better option when compensation has changed very recently or is expected to change in a way that has not occurred in the past.

Identification of the "best" modeling concept thus depends heavily on the ways in which the concept will be used to support policy decisions. The IMC presented in this report leans toward the DRM because changes to compensation policy may be of interest in current conditions. For example, the effects on recruiting and retention of the Blended Retirement System, implemented in January 2018, may be of interest. In order to successfully train an ML model to capture these effects, one would need about eight years of data on individual retention decisions under the Blended Retirement System.

Future implementation of the IMC should begin with a careful evaluation of the applications and requirements for IMC outputs throughout OCAR's planning processes, especially as these may have evolved since the writing of this report. The thorough mapping of personnel flows to modeling tools presented in this study provides some flexibility in terms of supporting alternate tool assignments to best suit future needs, identifying some of the reasons why different choices might be made. The work of developing the IMC as described here has resulted in a deep understanding and careful documentation of the applicability of various modeling tools to different personnel flows and can thus provide a foundation for future implementation that supports a variety of needs.

Appendix A. The Decomposition of Personnel Flows

This appendix provides additional detail about Figure 2.2, which shows the various sources from which individuals entered USAR and the various destinations to which the individuals went upon leaving USAR. Both Figure 2.2 and the figures in this appendix cover the 24-month period between September 30, 2016, and September 30, 2018.

Flows into the Selected Reserve

This section provides a decomposition of the five inflow sources shown in Figure 2.2. Figure A.1 breaks down the largest source of entry: no prior service in the previous six months. The overwhelming majority of the reservists who were civilians during the six months prior to joining USAR—95 percent—had no prior service ever. Among the remaining 5 percent, nearly 75 percent had prior service in the Army specifically—either in the AC, the ARNG, the IRR, or SELRES.
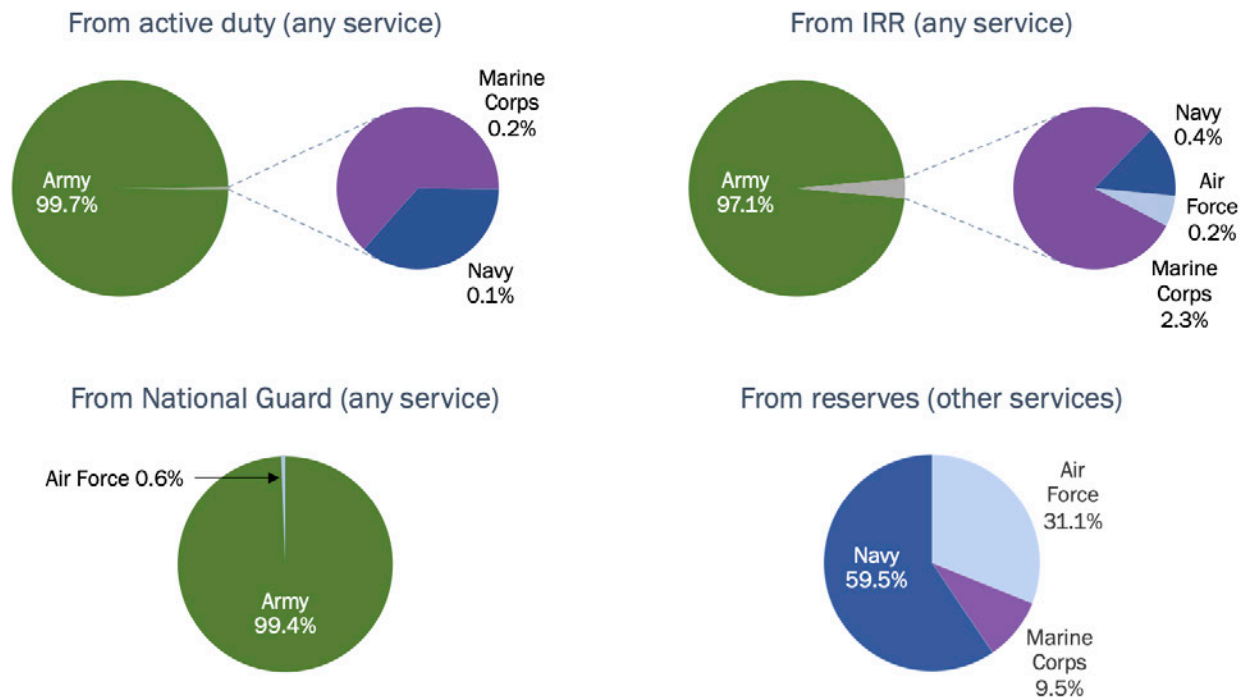
Figure A.1. Decomposition of No Service in the Six Months Prior to Joining
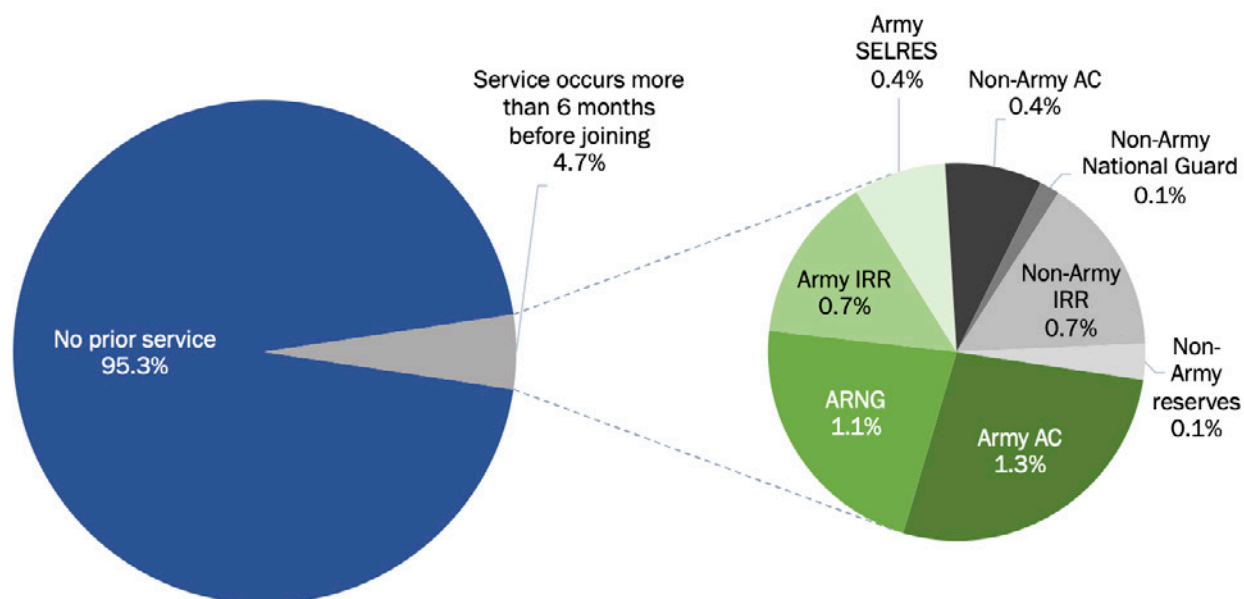
NOTE: Because of rounding, the sum of the percentages shown in the chart on the right does not equal precisely 4.7 percent.

Figure A.2 breaks down the remaining four sources of entry by branch of service: Air Force, Army, Marine Corps, and Navy. To prevent double counting, each individual was assigned to only the source that reflects his or her most recent service. Individuals entering SELRES directly from active duty, the IRR, or the National Guard were overwhelmingly from the Army. The small number entering directly from the reserves of another service branch were largely from the Navy.

Figure A.2. Decomposition of Other Inflow Sources

From active duty (any service)

From IRR (any service)

From National Guard (any service)

From reserves (other services)

Flows out of the Selected Reserve

This section provides a decomposition of the five outflow destinations shown in Figure 2.2. Figure A.3 breaks down the largest source: no service in the six months following separation. Nearly all of these reservists, 99.4 percent, had no subsequent military service. Among the remainder, over 85 percent returned to Army service specifically. However, these results were likely affected by censoring: because the separations occurred between September 30, 2016, and September 30, 2018, we were able to observe only up to two years following separation.

Figure A.4 breaks down the remaining four destinations by branch of service. As before, we prevented double counting by assigning each separating reservist to his immediate destination only. As is shown in the figure, individuals who left SELRES for active duty, the IRR, or the National Guard overwhelmingly remained within the Army. The small number that transitioned to the reserves of another service branch tended to join the Air Force.

Figure A.3. Decomposition of No Service in the Six Months Following Separation
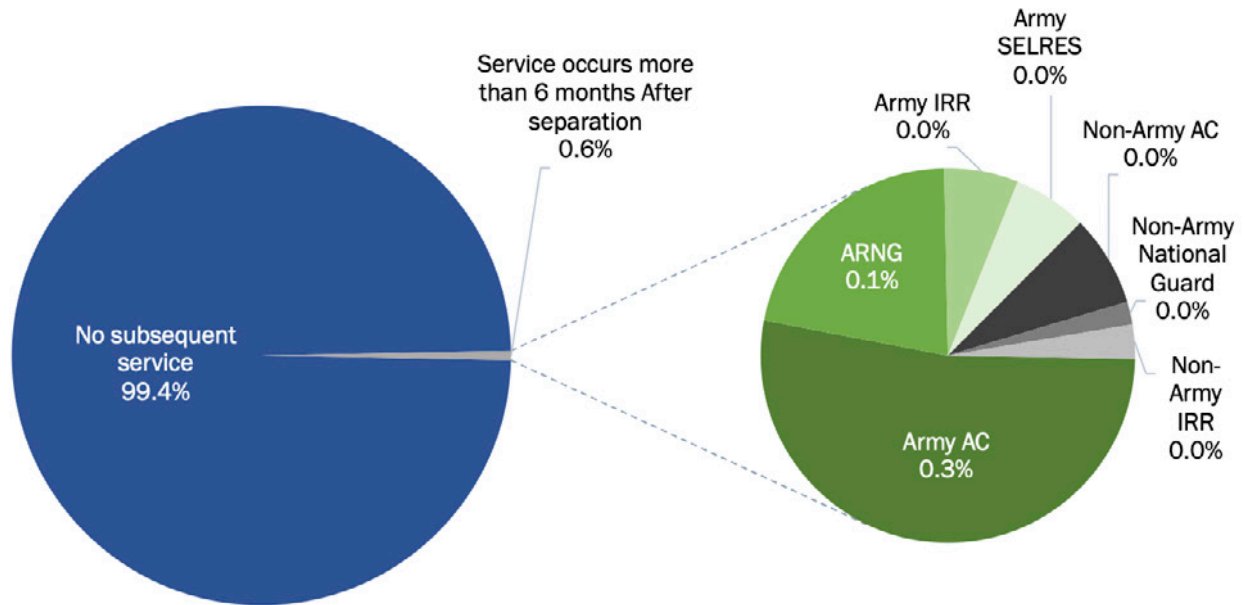
NOTE: Because of rounding, the sum of the percentages shown in the chart on the right might not equal precisely 0.6 percent.

Figure A.4. Decomposition of Other Outflow Destinations

42
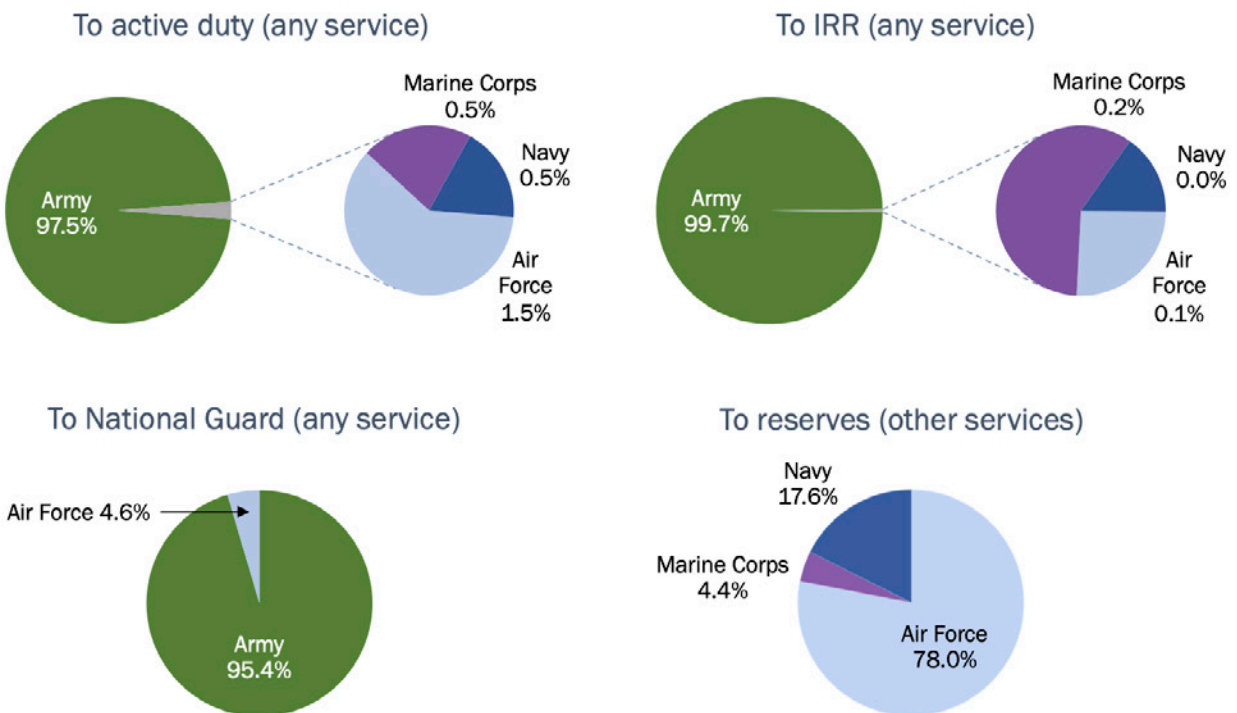
Appendix B. The Clustering Approach to Understanding SELRES Career Paths

As is described in Chapter 2, our population of SELRES personnel is large and heterogeneous. The data set contains 762,549 individuals, with 3,639 unique data points on each individual (most of which are monthly time series entries). While the focus of this report is on the inflows and outflows of this population to and from SELRES, this data also contains a wealth of information about individuals' behaviors *during* their time in SELRES. In order to better understand behavior patterns during service in SELRES, we apply techniques similar to those used in the predictive ML approach described at the end of Chapter 3. The goal of this exercise is to identify groups of individuals whose behavior and path during time in SELRES are similar, and then to describe the characteristics of these groupings.

Before we can proceed with an automated clustering approach, we must adjust our data set to make it more amenable to such techniques. In addition to being large, the data on SELRES careers over time is also quite sparse,[1] especially for service members whose careers are short. Clustering algorithms tend to perform poorly in such high-dimensional sparse spaces. To address this problem without applying prior knowledge about which variables are interesting to consider, we perform dimensionality reduction using a neural network autoencoder. We then use the dimensionality-reducing portion of the autoencoder to preprocess our data prior to applying clustering techniques. These two steps are described in more detail below. Here and in Appendix C, we build our neural nets using Keras 2.2.4 with a TensorFlow 1.14.0 back end.

The Neural Network Autoencoder for Dimensionality Reduction

Neural networks are general-purpose ML tools that, after being trained on known examples, can predict a particular output given a set of inputs. In the canonical case, neural networks digest large, complex data sets to produce simple outputs (as was the case with our prediction task), but other paradigms are also possible. One example is a neural network autoencoder: a neural network that is trained to produce the same output as its original input (thus, it is a *self-encoder*). Mapping an input vector onto itself to produce an output vector is a trivial problem; what makes an autoencoder useful is the presence of a bottleneck in its network topology that forces the network to learn a more compact way of representing the input data. In other words, an autoencoder can be thought of as a type of compression algorithm. Indeed, autoencoders can be used for image compression, though they are more often used in natural language processing to build language

---

[1] The data set is "sparse" because many of the entries in the data table are zero, indicating, for example, that the individual is not currently in SELRES.

models. Because our time series data can be considered to be analogous to a sentence (with each month of data corresponding to a word), we believed an autoencoder could be an appropriate tool to build a compact representation of each service member's data history.

A representation of our autoencoder is provided in Figure B.1. It is built as a deep feedforward network composed of multilayer perceptrons. The autoencoder has a network topology that narrows at the middle, beginning with 3,639 neurons that capture the full data set completely,[2] shrinking to only eight neurons at its narrowest point and widening out again to reproduce the original data using only the information that was contained in the narrow eight-neuron encoding. Each of the autoencoder's dense, fully connected layers (the boxes in Figure B.1) is followed by a batch normalization layer (the arrows in Figure B.1).[3]

Figure B.1. Notional Representation of the Dimensionality Reduction Neural Network
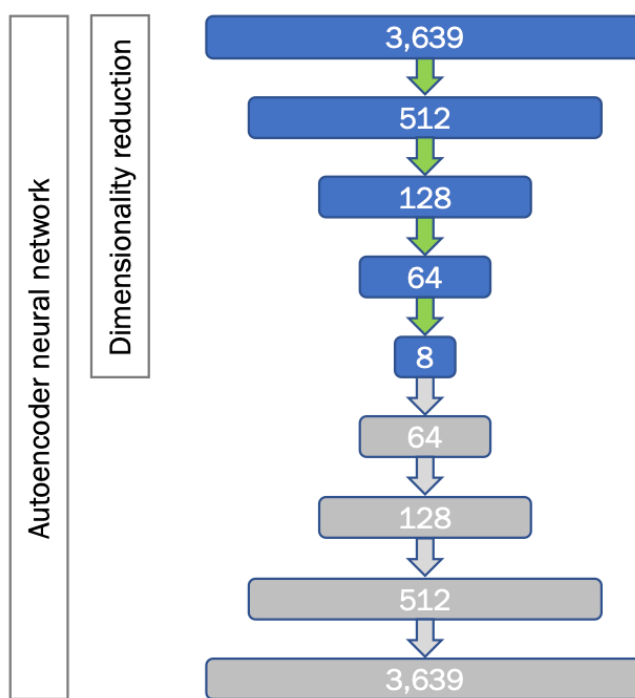
NOTE: The layer sizes of the autoencoder are described in the text. After testing a variety of network sizes, we found that a bottleneck of eight neurons was sufficient to capture most of the detail about a service member's history.

---

[2] Of the variables, 3,600 correspond to monthly time series data over 30 years of history (individuals with fewer than 30 years of history have the missing entries filled with zeros), and 39 variables correspond to the static categorical variables for each individual (e.g. race, sex, and age). More details about the exact variables used can be found in Appendix C.

[3] The neurons have exponential linear unit activation. See Djork-Arne Clevert, Thomas Uterthiner, and Seep Hochreiter, "Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)," arXiv, last revised February 22, 2016.

As with the neural network described in Chapter 3, the autoencoder neural network learns to make predictions by adjusting the weights of the links between its neurons to optimize the value of a *loss function* during training. Intuitively, the loss function is a representation of model error, which is to be minimized. Since the goal of an autoencoder is to reproduce the input signal, we selected as our loss function the mean-squared error between the input and output arrays. We train our autoencoder by feeding in batches of 256 individuals in five epochs,[4] using adaptive moment estimation (Adam optimization).[5] Manual inspection of the input and output vectors confirmed that the autoencoder was capturing the relevant information for each individual.

Clustering of Reduced Data

Once the autoencoder is trained, we simply remove the second half of the network (the gray boxes in Figure B.1), resulting in an encoder that converts the high-dimensional sparse space of raw data into an eight-dimensional dense space.[6] We then apply a standard clustering algorithm to the dense representation of our collection of individuals.

The data is encoded prior to being fed into the network. We have ten data fields for each of the 360 months considered: four are one-hot encodings for the CMF in that month (each CMF is considered as one of Force Sustainment, Operations Support, Maneuvers & Fires, or Other); four are related to the component for that month (Active, SELRES, Civilian, or Other); one is grade (which is divided by ten to yield a real number between zero and one); and one describes activation and deployment (a zero for not activated or deployed, and a one otherwise). Our static variables (race, age, sex, education level, and prior service) are also one-hot encoded for each category in each variable, yielding a vector of length 39.

We pass the complete data for each enlisted individual (762,549 SELRES members) in our data set through the encoder network,[7] and then apply K-means clustering to the transformed, dense distribution of data.[8] Because we do not have an a priori understanding of the number of clusters we expect to see, we produced two sets of K-means clustering results, for six and 12 clusters. Comparing the two results, we find that many of the clusters in the 12-cluster

---

[4] Batching is a common practice that balances the extremely manual process of training the model on each individual observation with the large computing demands of using the entire sample at once.

[5] Adam optimization is a popular extension of the standard gradient descent algorithm. See Diederik P. Kingman and Jimmy Lei Ba, "Adam: A Method for Stochastic Optimization," arXiv, last revised January 30, 2017.

[6] Inspection of the value distributions for the eight-dimensional space confirmed that it was dense; using a larger number of neurons tended to result in some of the dimensions being underutilized.

[7] In this case, we "flatten" the time series data fields for each individual and concatenate them end to end with the static data fields, yielding a single vector with about 3,600 entries. The variables considered are the same ones used in the predictive analysis of Chapter 4.

[8] Trevor Hastie, Robert Tibshirani, and Jerome Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed., New York: Springer, 2009, pp. 509–514.

exercise are further subdivisions of those of the six-cluster exercise and do not yield significant additional insights. Therefore, we focus our analysis on the six-cluster results.

The results presented in this report are of interest primarily because they confirm that this approach is successful in identifying groups that may be useful for understanding career paths. Automated clustering does not work well in all applications, but this proof-of-concept exercise suggests that it may indeed work well for detecting groupings of similar SELRES careers. As discussed later in this appendix, the current implementation does not tackle issues related to censoring, correlation between variables, inclusion or exclusion of different potential predictors, or other data issues. Such exploration is reserved for future work that focuses more explicitly on career trajectories *during* time in SELRES.

Clustering Results

As is described in Chapter 2, and specifically Table 2.1, the clustering exercise identifies six groups within the SELRES sample analyzed, with clusters differing in terms of career length, prior service status, and, to a lesser extent, individual job category and demographic characteristics. We hypothesized that the prevalence of career length and YOS in cluster membership definition may be partially due to censoring issues—that is, a short career might indicate that the service member left USAR after a short time, or it could simply mean that not enough time has elapsed for that member to record a longer history. To test this, we performed a small experiment to study the effect of censoring: we repeated the autoencoder training and clustering exercise but restricted the data sample to reservists who joined USAR prior to 2007. This reduced significantly the number of service members who were still in service at the end of the time period considered (which ended in 2017). The new clusters based on this subset of data showed a very similar pattern as the original clusters: in-cluster membership was driven strongly by career length. We therefore concluded that our clustering approach was sufficient for our stated goal of examining similar career trajectories.

Figure B.2 shows the breakdown of total YOS by cluster (left panel) and the entry cohort to which individuals belong (right panel). We see, for example, that the "Satisfied reservists" cluster (in purple) comprises mostly individuals who began their SELRES careers between 1990 and 1995 and contains no individuals who began their service in 2016. Clearly, we cannot observe whether an individual who began his or her career toward the end of our sample period will become a long-term SELRES member. However, further analysis could investigate whether automated assignment to the "Satisfied reservist" cluster for those who began their careers in 2011–2015 is based on similarities between the early parts of these reservists' careers and the early parts of the careers of those who decided to remain in SELRES for a long time. Similar issues that are not explicitly addressed in this proof-of-concept analysis include the fact that age-related factors are necessarily correlated with other variables such as career length, YOS, and fiscal years in which the individuals joined SELRES; prior service; or that prior service data may be less complete for those who joined SELRES early during our period of observation.

Figure B.2. Censoring Considerations in the Preliminary Clustering Exercise
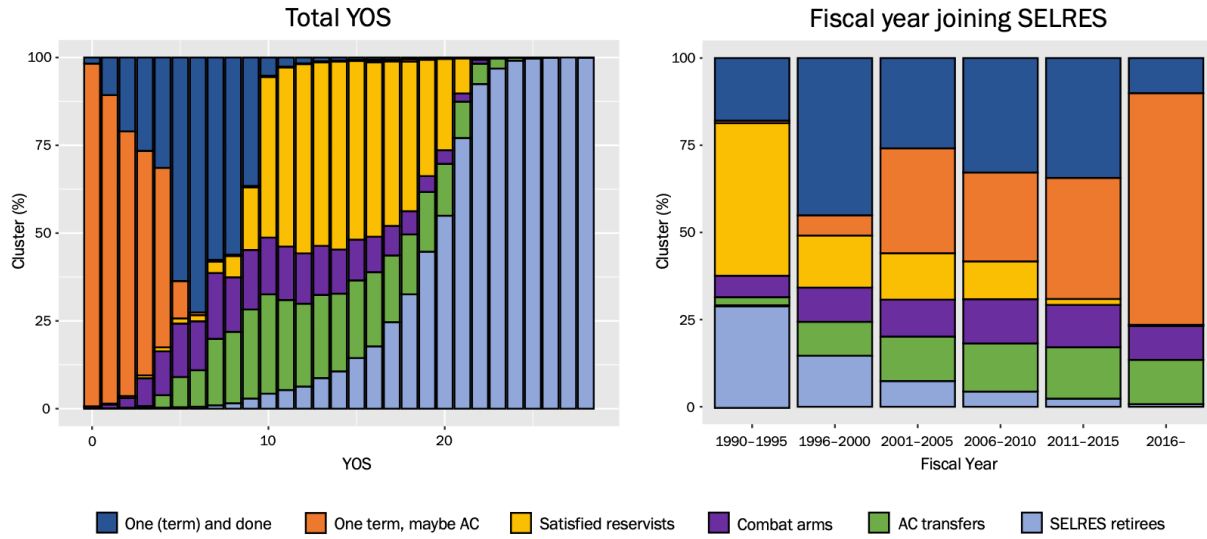
Figures B.3 and B.4 capture how clusters differ in terms of career characteristics and personnel flows, while Figure B.5 shows how clusters differ across demographic characteristics. These figures are included to make the distinctions between clusters highlighted in Table 2.1 more concrete, as well as to allow the reader to detect other patterns or trends in this purely descriptive analysis. Figure B.3 shows how cluster populations differ by prior service status and job specialty. We see that three clusters primarily include individuals with no prior service, while the remaining three clusters primarily include individuals with at least some prior service—consistent with the

Figure B.3. Career Characteristics, by Cluster
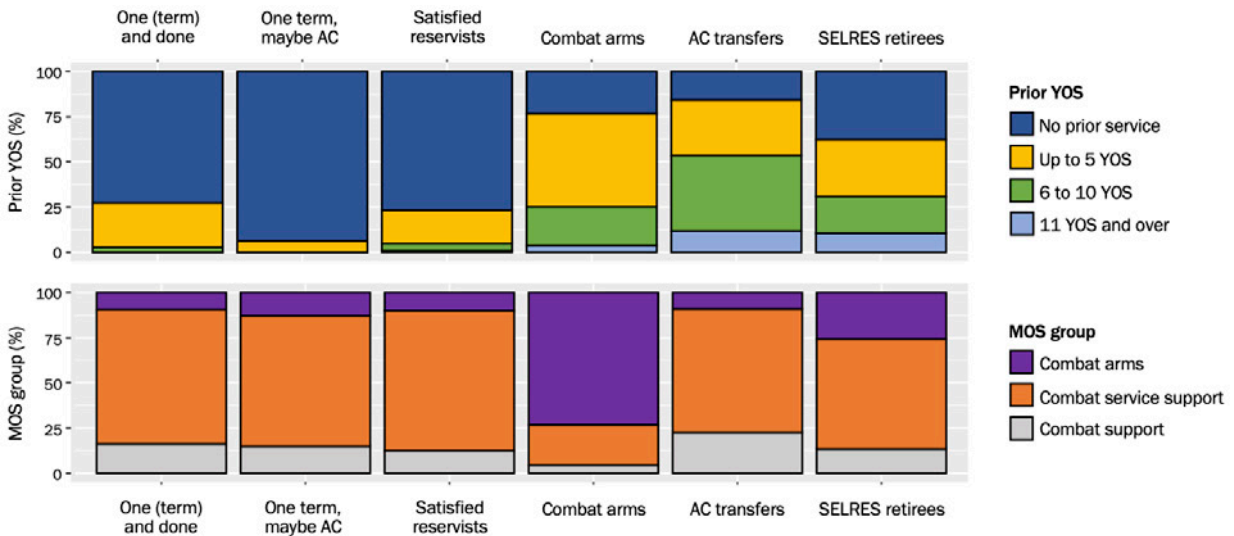
NOTE: MOS groups are defined as follows: "Combat arms" includes service as a combat engineer or in air defense artillery, armor, aviation, field artillery, infantry, or the Special Forces; "Combat support" includes service as a chemical, biological, radiological, and nuclear specialist or in audio-visual signaling, military intelligence, or the Military Police; and "Combat service support" includes other groups.

cluster definitions detailed in Table 2.1. The breakdown of individuals by job category is similar across clusters, except for the "Combat arms" cluster, which has a higher proportion of individuals with prior service and specifically includes a larger proportion of individuals in career fields with roles in combat arms (as a combat engineer or in air defense artillery, armor, aviation, field artillery, infantry, or Special Forces).

Figure B.4 shows how cluster populations differ by gender, race, educational attainment, and age. We do not see large differences in demographic composition across clusters, with the exception of age and, to a lesser extent, educational attainment, both of which are correlated with overall career length. The one exception may be that combat arms cluster members tend to be more male and more white than members of other clusters.

Figure B.4. Demographic Characteristics, by Cluster
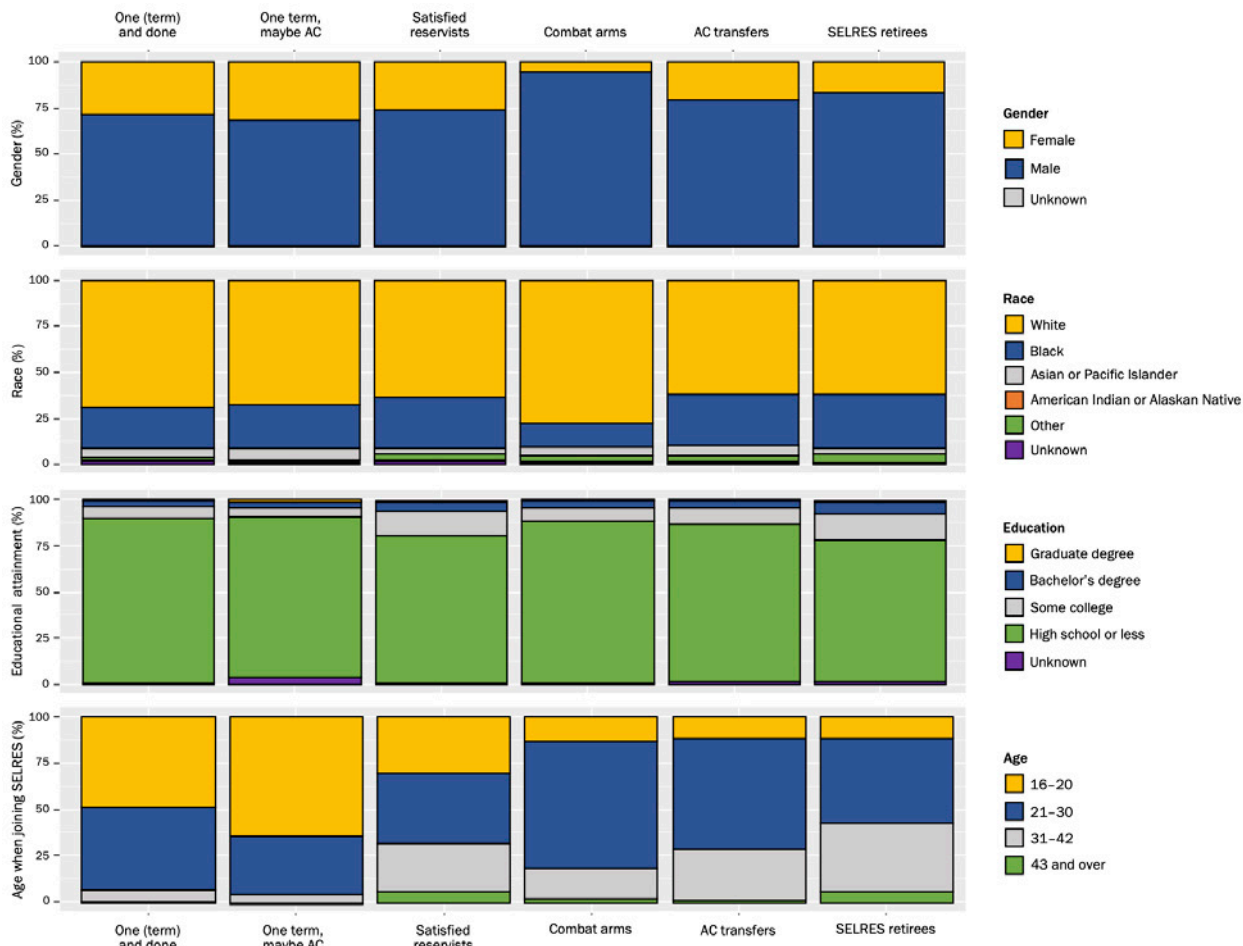
A summary of cluster differences in terms of personnel flows and the paths individuals in each cluster take into and out of SELRES is the portion of the clustering analysis that is most relevant to the IMC presented in this report. A description of personnel flows across clusters is included in Chapter 2 and Figure 2.4.

Appendix C. The Machine Learning Model

In this appendix we describe a newly developed proof-of-concept neural network model designed to predict individual-level retention behavior. The data-driven proof-of-concept model developed for this study is similar to analysis conducted by the Institute for Defense Analyses in 2019 but makes adjustments to tailor it to the specific needs of this project.[1]

Input Data

The data set used to train and test the neural-network-based model contains about 1.3 million individuals, about half of whom are removed from our sample due to missing data (as described below). Each individual is associated with two types of data: (1) monthly time series data of his or her military history, and (2) static demographic data. The time series data fields are grade; CMF (binned into one of four categories: force sustainment, maneuvers and fires, operations support, and other); deployment/activation; and component (AC, National Guard, SELRES, etc.). Static fields are age, sex, race, education level, and date of SELRES entry.

Neural networks are flexible and generally appropriate for the kind of mixed data available on SELRES individuals: a combination of time series and static fields, with both categorical and numerical entries. However, our heterogeneous data set does have one particular disadvantage for use in a neural network: many of the entries (individuals) are missing one or more values in their data fields. Of the 1.3 million entries in our data, only 762,549 contain all the entries needed, with most of the missing entries being in the static fields (such as education level). Missing data fields for inputs into a neural network can be handled in one of two main ways: (1) imputing the missing data (with the mean or mode of the data field across the entire set, for example), or (2) simply discarding the entries with missing data. We take the latter approach here for simplicity, as our model is a proof of concept to show how one could make predictions about USAR careers. Discarding data does run the risk of biasing the results if there is a correlation between the presence of missing data and other fields; a more advanced algorithm could be employed in future iterations.[2]

After discarding entries with missing data fields, we perform several data preprocessing steps. CMF variables, which represent SELRES job types, are binned into one of four categories: force sustainment, maneuvers and fires, operations support, and other. These and other categorical variables (education level, race, etc.) are then encoded into a sparse representation. We also

---

[1] Pechacek et al., 2019.

[2] Another approach is to perform a variable importance study on the individuals without missing fields by training a model only on them and then perturbing the inputs to see which variables have the highest impact on accuracy. We leave such a study for future iterations of the analysis.

apply scaling factors to bring all numeric values to be between zero and one so the neural network can operate on variables with standardized magnitudes.

Model Architecture

The architecture of our neural network has been adapted from the Institute for Defense Analyses model to suit the needs of this study. Figure C.1 provides a simplified visual overview of the model architecture, which has both recurrent and feedforward components. Our priority is determining an end strength forecast at 24 months from a particular date (the *observation date*), though we are also interested in a 12-month forecast. With this in mind, we design our neural network to make one of three predictions: whether a given individual will exit USAR in less than one year, between one and two years, or more than two years. Given these predictions, it is straightforward to calculate the end strength at 12 and 24 months. We take as our input all static variables and the 100 most recent months of time series data prior to the observation date. For service members with less than 100 months of history, the missing entries are filled with zeros.[3] Static variables are fed directly into the predictive neural network (the "Fully connected neural network" in Figure C.1), while time series data first undergo dimensionality reduction.

Figure C.1. Summary of Machine Learning Model Architecture
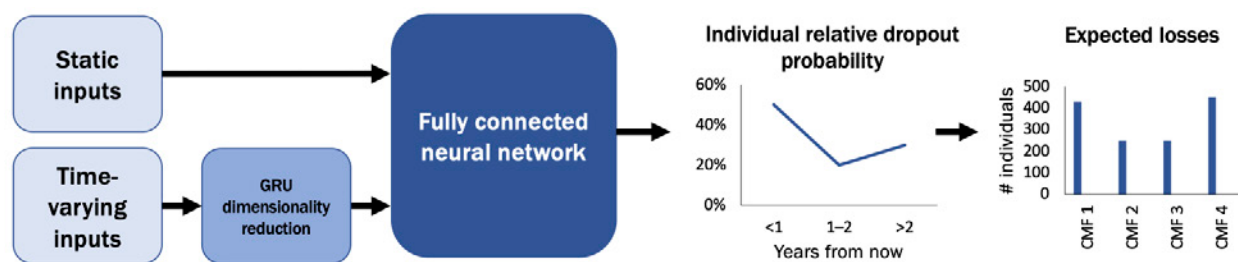
NOTES: GRU = gated recurrent unit. This is a simple visual representation of the ML model. Data on individual SELRES members enters on the left, with time series data fields passing through an extra filter of GRU neurons to capture long-term information. The main body of the neural network analyzes the inputs and returns the probability that each individual will exit service in the next two years. These probabilities can then be used to estimate the composition of the exiting population.

Specifically, the time series data is fed through a layer of 32 GRUs,[4] which can capture information on long time scales and encode the information contained in long time series data into a smaller number of dimensions. The main neural network employs three hidden layers of

---

[3] Because most of our variables are one-hot encoded, the vector containing each individual's data is sparse. We therefore believe the neural network should learn to ignore zeros and instead focus on nonzero entries. We find our autoencoder network (described in Appendix B) does exactly this, building a dense representation by ignoring zeros.

[4] Kyunghyun Cho, Bart van Merrienboer, Caglar Gulcehre, Dzimitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio, "Learning Phase Representations Using RNN Encoder-Decoder for Statistical Machine Translation," arXiv, last revised September 3, 2014.

128 neurons each that feed into a final output layer.[5] The output layer has three neurons whose values are passed through a softmax function,[6] which normalizes the results and yields values that can be interpreted as probabilities. In our case, the three values are the probabilities that a given individual will exit USAR in less than one year, one to two years, or more than two years. These bins are derived from our desire to forecast 12-month (the less than one-year bin) and 24-month (the one- to two-year bin) retention. In Figure C.1, this is represented by the line chart on the right, "Individual relative dropout probability." Given each individual's probability of dropping out, we can also calculate the characteristics of the population that needs to be replaced by USAR, represented by the chart of CMF losses on the far right of Figure C.1 ("Expected losses").

Our neural network learns to make predictions by adjusting the weights of the links between its neurons to optimize the value of a loss function during training. Intuitively, the loss function is a representation of model error, which is to be minimized. For this case, our loss function is Michael Gensheimer and Balasubramanian Narasimhan's so-called survival loss function:[7]

$$\mathcal{L} = h_j \prod_{i=0}^{j-1} (1 - h_i), \tag{1}$$

where h_j is the probability that an individual will depart in year j. In other words, Equation 1 is the likelihood of departure in year j, multiplied by the likelihood of survival during years zero through $j - 1$. In practice, we rewrite Equation 1 as

$$-\ln \mathcal{L} = -\ln h_j - \sum_{i=0}^{N} g_i, \tag{2}$$

where g_i is defined as

$$g_i = \begin{cases} 1 - h_i, i < j \\ 1, i \geq j \end{cases}. \tag{3}$$

Building the vector g is made difficult by the fact that the simplest way of coding it relies on nondifferentiable functions like $argmax()$. To make $\mathcal{L}$ differentiable, and therefore useful for training a neural network, we implement a soft $argmax()$ that closely approximates the $argmax()$ function. The entire network is trained using backpropagation.

---

[5] Final output layer uses rectified linear unit (ReLU) activation.

[6] The softmax function σ, defined by $\sigma(z_i) = \frac{e^{z_i}}{\sum_{j=1}^{N} e^{z_j}}$, where there are N elements in the vector of arbitrary real numbers z, transforms z into an array of probabilities that naturally sums to one.

[7] Michael F. Gensheimer and Balasubramanian Narasimhan, "A Scalable Discrete-Time Survival Model for Neural Networks," *PeerJ*, No. 7, January 25, 2019.

Training

Our neural network does not explicitly take into account the absolute date for any individual; instead, it simply considers 100 months of time series data (along with the static variables) and makes 12- and 24-month predictions. The neural network is therefore agnostic to the observation date itself, and training and using it to make a definite prediction requires some care.

On the observation date (or any given date, for that matter), we expect the USAR population to contain individuals with varying career lengths: some will have just joined, while others may have a decade or more of experience. However, because our data set covers several decades of history, most of the individuals in the set have completed their USAR careers. In other words, selecting individuals at random from our data set would result in a sample biased toward long-term careers. In order to make accurate predictions on both short- and long-term careers, we therefore need to train our neural network on a selection from our data set that oversamples short-term careers.

We draw samples from our data set in the following way: for individuals whose career ends before the observation date,[8] we select a random month from their history—subject to the condition that there is an equal chance of that month being less than one year, one to two years, or more than two years away from the date they exit USAR—and choose the 100 months of time series data that precedes it. For members still in service on the observation date, we require a minimum of two YOS (without which it is impossible to say which bin they will fall into); for those members with at least two years of USAR history, we select a random month that is at least two years prior to the observation date and choose the 100 months of time series data prior to that month.

We also keep track of how far each member is from exiting service (less than one year, one to two years, more than two years) at the selected month, and generate a vector with a one indicating which bin they fall into. For example, suppose we randomly select a member who entered USAR in June 2005 and exited in July 2013. Suppose then our algorithm randomly picks May 2012 as the date to consider. We grab all time series data from January 2004 to May 2012 (100 months), filling in zeros for the time period January 2004 to June 2005. This data becomes our input vector for the neural net, and the output vector is [0,1,0] (May 2012–July 2013 is more than one but less than two years). We train the neural network in batches of 1,024 individuals,[9] with epochs; our results indicate that more epochs did not appreciably improve the accuracy of the predictions.

---

[8] We treat members who are in a gap in their service at the observation date as if they had already completed their USAR career.

[9] Batching is a common practice that balances the extremely manual process of training the model on each individual observation with the large computing demands of using the entire sample at once.

Proof-of-Concept Application

In order to demonstrate predictive modeling using this neural network approach, we perform an illustrative exercise of predicting 12- and 24-month retention for the cohort of SELRES individuals who entered USAR between FY 2002 and FY 2013. The sample included in this illustrative application of the model consists of enlisted reservists with no prior service outside of SELRES. There are 78,082 such individuals with sufficiently populated data who were in SELRES in September 2015.

Using this sample, we predict what proportion of the population present on September 30, 2015 (the observation date) is still present on September 30, 2017. We train the neural network on individuals who entered USAR between FY 2002 and FY 2013 and use the trained network to predict the exit likelihood at 12 and 24 months for the individuals in our data set present on September 30, 2015. Because the neural network returns a probability (between zero and one) that an individual will exit, we sum the exit probabilities across all individuals to yield an expected value of the number of individuals still in SELRES.

The results of this illustrative analysis are summarized in Figures 4.1 and 4.2, suggesting that the ML approach predicts retention for existing reservists relatively well. The model underestimates retention over the 12-month period by 0.6 percent. Over the 24-month period, the model forecast overestimates retention by 7.3 percent. The accuracy of these forecasts can be more easily assessed using Figures C.2 and C.3. Along the vertical axis of these figures, we plot the ratio of the predicted number of retained reservists to the actual number of retained reservists. The dotted line marks a value of one (1) for this ratio, representing a prediction without error. We see that the model underestimates retention for individuals early in their careers (about zero to four YOS) and overestimates retention for individuals with about five to seven YOS.

Although our neural network was built for the purpose of making a total prediction of end strength, it is also interesting to consider how well it performs for individual reservists. For our particular experiment, our 12-month horizon yielded an ROC-AUC score of 0.54, and a C-index of 0.53, while our 24-month scores were very slightly lower. This indicates that our model was not particularly proficient at making individual-level predictions despite its ability to generate relatively accurate population-level predictions. On the other hand, it also indicates that there is ample room for improvement by future models.

While these observations apply only to the population analyzed, they provide some general insight into the applicability of these models to retention estimation and demonstrate the ML model's ability to provide satisfactory estimates in at least some applications.

Figure C.2. Error in Estimates of 12-Month Retention of Enlisted Reservists with No Prior Service
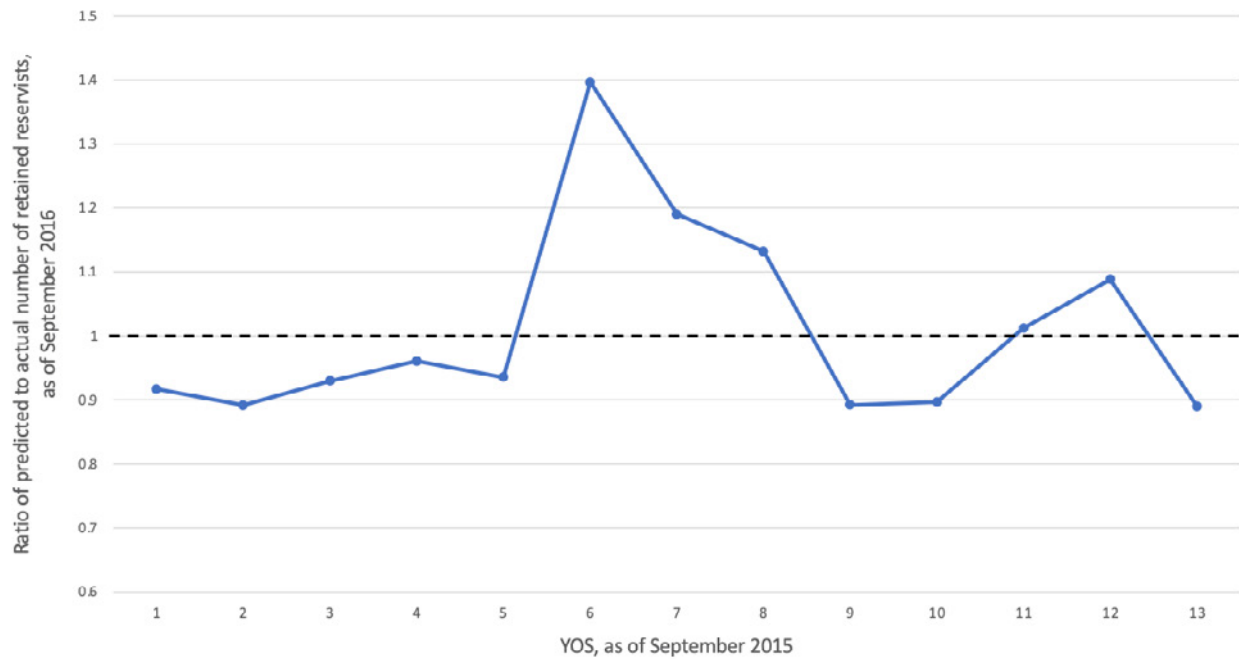
Figure C.3. Error in Estimates of 24-Month Retention of Enlisted Reservists with No Prior Service
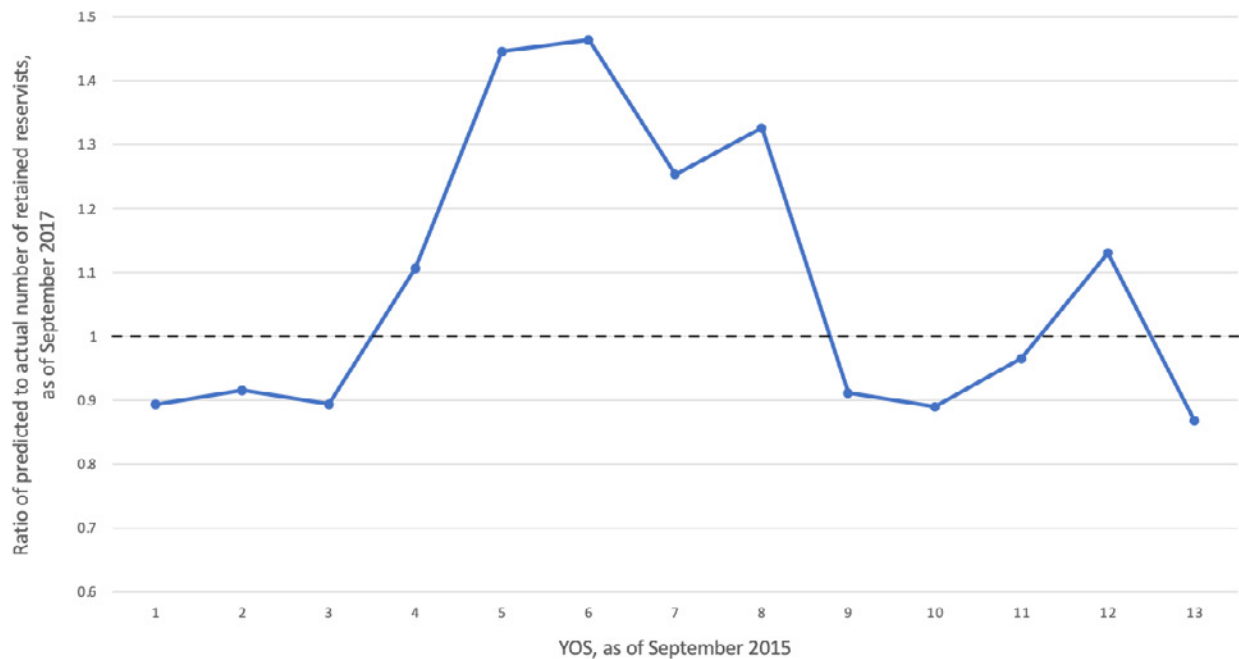
Further Applicability of the Machine Learning Approach

Although the framework described here only describes departures from USAR, it can be extended to cover certain categories of entrances—namely, the inflows from populations that have prior military experience, and thus available career history data. Instead of outputting a vector describing the probability of exit after less than one year, one to two years, and more than two years, the neural network could be modified to output a matrix describing both the probability of departure and where the individual goes to (for example, to civilian life, the AC, or the National Guard). We could then train clones of our neural network on AC and National Guard populations and use the trained networks to predict the number of members who will transfer to SELRES over the next 24 months. While application of similar ML approaches is possible for other portions of the SELRES end strength projection question, it is important to note that the test case considered here (only enlisted members whose data fields were complete) may be a best-case scenario for an ML model, since there are a large number of enlisted members with sufficient observable career histories on which to train. This may not be the case for other possible applications.

We also note that recent public codes have been developed that perform similar functions as our prediction framework, particularly the Finite-Interval Forecasting Engine package and the associated Retention Prediction Model built by the Institute for Defense Analyses.[10] We anticipate that utilizing the Finite-Interval Forecasting Engine framework in future iterations would reduce the amount of work required to build a simple model, allowing us to focus more effort on improving the model's performance.

---

[10] Python Software Foundation, "Finite-Interval Forecasting Engine: Machine Learning Models for Discrete-Time Survival Analysis and Multivariate Time Series Forecasting," webpage, last updated March 25, 2021.

References

Absher, Jim, "Guard and Reserve Retirement," Military.com, June 22, 2021. As of February 26, 2022:
https://www.military.com/benefits/military-pay/reserve-and-guard-pay/military-reserve-component-retirement-overview.html

Asch, Beth J., Michael G. Mattock, and James Hosek, *Effects of the Blended Retirement System on United States Army Reserve Participation and Cost*, Santa Monica, Calif.: RAND Corporation, RR-2591-A, 2019. As of January 31, 2022:
https://www.rand.org/pubs/research_reports/RR2591.html

Caldara, Dario, and Matteo Iacoviello, *Measuring Geopolitical Risk*, International Finance Discussion Papers 1222, Washington, D.C.: Board of Governors of the Federal Reserve System, February 2018.

Cho, Kyunghyun, Bart van Merrienboer, Caglar Gulcehre, Dzimitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio, "Learning Phase Representations Using RNN Encoder-Decoder for Statistical Machine Translation," arXiv, last revised September 3, 2014. As of February 26, 2022:
https://arxiv.org/abs/1406.1078

Clevert, Djork-Arne, Thomas Uterthiner, and Seep Hochreiter, "Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)," arXiv, last revised February 22, 2016. As of February 26, 2022:
https://arxiv.org/abs/1511.07289

Congressional Research Service, *Defense Primer: Reserve Forces*, fact sheet, Washington, D.C., last updated December 2021. As of February 26, 2022:
https://fas.org/sgp/crs/natsec/IF10540.pdf

Gensheimer, Michael F., and Balasubramanian Narasimhan, "A Scalable Discrete-Time Survival Model for Neural Networks," *PeerJ*, No. 7, January 25, 2019. As of February 25, 2022:
https://peerj.com/articles/6257/

Gotz, Glenn A., and John McCall, *A Dynamic Retention Model for Air Force Officers: Theory and Estimates*, Santa Monica, Calif.: RAND Corporation, R-3028-AF, 1984. As of January 31, 2022:
https://www.rand.org/pubs/reports/R3028.html

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed., New York: Springer, 2009.

Headquarters, Department of the Army, *Army Reserve: Assignments, Attachments, Details, and Transfers*, Washington, D.C.: Department of the Army, Army Regulation 140-10, July 16, 2021.

Kapp, Lawrence, and Barbara Torreon, *Reserve Component Personnel Issues: Questions and Answers*, Washington, D.C.: Congressional Research Service, October 4, 2018.

Kingman, Diederik P., and Jimmy Lei Ba, "Adam: A Method for Stochastic Optimization," arXiv, last revised January 30, 2017. As of February 25, 2022:
https://arxiv.org/abs/1412.6980

Knapp, David, Bruce R. Orvis, Christopher E. Maerzluft, and Tiffany Berglund, *Resources Required to Meet the U.S. Army's Enlisted Recruiting Requirements Under Alternative Recruiting Goals, Conditions, and Eligibility Policies*, Santa Monica, Calif.: RAND Corporation, RR-2364-A, 2018. As of January 31, 2022:
https://www.rand.org/pubs/research_reports/RR2364.html

Mattock, Michael G., Beth J. Asch, and James Hosek, *Making the Reserve Retirement System Similar to the Active System: Retention and Cost Estimates*, Santa Monica, Calif.: RAND Corporation, RR-530-A, 2014. As of January 31, 2022:
https://www.rand.org/pubs/research_reports/RR530.html

Neumark, David, "Datasets and Code Teaching," web database, University of California–Irvine, undated. As of February 6, 2022:
http://www.economics.uci.edu/~dneumark/datasets.html

Office of the Assistant Secretary of Defense, *Official Guard and Reserve Manpower Strengths and Statistics: FY 2007 Summary*, Washington, D.C., January 2007.

Office of the Under Secretary of Defense for Personnel and Readiness, Directorate of Compensation, *Selected Military Compensation Tables*, Washington, D.C., 2007.

Orvis, Bruce R., Craig A. Bond, Daniel Schwam, and Irineo Cabreros, *Resources Required to Meet the U.S. Army Reserve's Enlisted Recruiting Requirements Under Alternative Recruiting Goals, Conditions, and Eligibility Policies*, Santa Monica, Calif.: RAND Corporation, RR-A1304-1, 2022. As of July 14, 2022:
https://www.rand.org/pubs/research_reports/RRA1304-1.html

Orvis, Bruce R., Christopher E. Maerzluft, Sung-Bou Kim, Michael G. Shanley, and Heather Krull, *Prospective Outcome Assessment for Alternative Recruit Selection Policies*, Santa Monica, Calif.: RAND Corporation, RR-2267-A, 2018. As of January 31, 2022:
https://www.rand.org/pubs/research_reports/RR2267.html

Pechacek, Julie, Alan Gelder, Cullen Roberts, Joe King, James Bishop, Michael Guggisberg, and Yev Kirpichevsky, "A New Military Retention Prediction Model: Machine Learning for

High Fidelity Forecasting," paper presented at the 94th Annual Conference of the Western Economic Association International, San Francisco, June 2019.

Powers, Rod, "Military Reserves Federal Call Up Authority: From Reserves to Active Duty," The Balance Careers, last updated July 30, 2019. As of February 26, 2022: https://www.thebalancecareers.com/military-reserves-federal-call-up-authority-3332677

Python Software Foundation, "Finite-Interval Forecasting Engine: Machine Learning Models for Discrete-Time Survival Analysis and Multivariate Time Series Forecasting," webpage, last updated March 25, 2021. As of February 25, 2022: https://pypi.org/project/fife/

U.S. Army, "About the Army," webpage, undated. As of July 1, 2019: https://www.goarmy.com/about/serving-in-the-army/serve-your-way/army-reserve.html

U.S. Code, Title 10, Section 10104, Army Reserve: Composition, October 5, 1994.

U.S. Code, Title 32, Section 101, Definitions, 1956.

Wenger, Jeffrey B., David Knapp, Parag Mahajan, Bruce R. Orvis, and Tiffany Berglund, *Developing a National Recruiting Difficulty Index*, Santa Monica, Calif.: RAND Corporation, RR-2637-A, 2019. As of January 31, 2022: https://www.rand.org/pubs/research_reports/RR2637.html

Wenger, Jennie W., Bruce R. Orvis, David Stebbins, Eric Apaydin, and James Syme, *Strengthening Prior Service–Civil Life Gains and Continuum of Service Accessions into the Army's Reserve Components*, Santa Monica, Calif.: RAND Corporation, RR-1376-A, 2016. As of January 31, 2022: https://www.rand.org/pubs/research_reports/RR1376.html

T he U.S. Army Reserve (USAR) is an integral part of the U.S. Army and the country's national defense. Its mission is to provide trained individuals who can serve as active duty soldiers when the mission calls for it. Well-trained service members are central to the USAR mission, and personnel and career management are critical to building a well-trained force.

End strength forecasts are an important input to recruiting and retention policy decisions, as well as resourcing and planning discussions with other Army components. However, generating these forecasts is a complex task due to the many paths into and out of USAR. Currently, no single model is capable of providing such estimates.

The authors examined available modeling capability and identified a set of modeling tools that can estimate portions of the personnel flows into and out of USAR. Most of these tools were designed to support policy analysis, not forecast end strength. Nevertheless, with the appropriate care and caution, the estimates generated by these tools can be combined to construct the desired 24-month end strength forecasts. A detailed plan for doing so—the Integrated Modeling Concept—is the primary product of this study.

$29.50

www.rand.org