



INSTITUTE FOR DEFENSE ANALYSES

Forecasting Demand for Air National Guard Enlisted Initial and Technical Schooling

Akshay Jain
John W. Dennis
Julie Lockwood
Minerva Song
Nathaniel Latshaw
Erin Eifert
Joseph King

March 2022

Approved for public release;
distribution unlimited.

IDA Paper P-32920

Log: H 21-000488

INSTITUTE FOR DEFENSE ANALYSES
730 East Glebe Road
Alexandria, Virginia 22305



The Institute for Defense Analyses is a non-profit corporation that operates three federally funded research and development centers to provide objective analyses of national security issues, particularly those requiring scientific and technical expertise, and conduct related research on other national challenges.

About this Publication

This work was conducted by the Institute for Defense Analyses under contract HQ0034-19-D-0001, project DZ-6-4720, "Forecasting Demand for Air National Guard (ANG) BMT and Tech School Slots to Improve Readiness," for the Chief of the National Guard Bureau (CNGB) and the Director of the Air National Guard (DANG). The views, opinions, and findings should not be construed as representing the official position of either the Department of Defense or the sponsoring organization.

For More Information:

Dr. John W. Dennis, Project Leader

jdennis@ida.org, (703) 845-2166

ADM John C. Harvey, Jr., USN (ret) Director, SFRD

jharvey@ida.org, 703-575-4530

Acknowledgment

The authors wish to thank the following contributors, Dr. David M. Tate, Dr. Nicholas A. Wagner, Mr. Jenns A. Robertson, Mr. Michael J. Norton

Copyright

© 2022 Institute for Defense Analyses

730 East Glebe Road,

Alexandria, Virginia 22305 • (703) 845-2000

This material may be reproduced by or for the U.S. Government pursuant to the copyright license under the clause at DFARS 252.227-7013 (Feb. 2014).

INSTITUTE FOR DEFENSE ANALYSES

IDA Paper P-32920

Forecasting Demand for Air National Guard Enlisted Initial and Technical Schooling

Akshay Jain
John W. Dennis
Julie Lockwood
Minerva Song
Nathaniel Latshaw
Erin Eifert
Joseph King

This page is intentionally blank.

Executive Summary

Air National Guard (ANG) airmen frequently experience delays of up to a year to enter training, which harms unit readiness and contributes to inefficient resource allocation. These delays may be caused, at least in part, by inaccurate forecasts in earlier years of training needs: the demand for basic military training (BMT) and for technical (or “tech”) school slots. To satisfy planning and budgeting requirements, the ANG must request BMT and tech school training slots up to five years in advance. Currently, individual ANG Career Field Managers (CFMs) collect data through various methods and forecast school slot demand generated by airmen who exit the ANG or switch occupations, and by end strength adjustments to particular career fields. More precise and reliable ANG school slot demand forecasts may help to reduce entry delays.

To assist in developing ANG school slot demand forecasts, the Institute for Defense Analysis (IDA) produces annual “exit” forecasts for each career field subdivision, commonly referred to as three-digit Air Force Specialty Code (AFSC), for each June from 2022 through 2026. We employ the Retention Prediction Model (RPM), a machine learning (ML) capability developed by IDA, to forecast five-year retention probabilities for each airman in service in June of 2021. Since individuals can switch occupations during their time in service, we layer a Markov transition model over the raw RPM exit forecasts to predict the AFSC from which enlisted personnel will exit the ANG in the future. We then use both of these probabilities to forecast the expected number of airmen who will exit the ANG from each AFSC. This method does not account for individuals not yet in ANG service who will both join and exit during the forecasting window of July 1, 2022 to June 30, 2026.

To illustrate the performance of this forecasting method, we compare the forecasts this method would have produced for each June from 2017 through 2021 to actual exits during the same period. This exercise indicates an average mean absolute error (by exit share) of 7% across all years, representing a 70% improvement from a baseline extrapolation model using historic attrition rates at the three-digit AFSC level. While these exit forecasts represent only one component of the demand generating process, this method represents a substantial improvement in ANG school slot and general attrition forecasting capabilities.

This page is intentionally blank.

Contents

1.	Introduction	1
2.	Background and Scope	3
	A. Background	3
	B. Components of Training Demand	4
	C. Scope of the Analysis	5
3.	Methodology.....	7
	A. Mapping AFSC-level Personnel Flows.....	7
	B. Estimating Total Enlisted School Slot Demand	8
	C. Estimating Exits from the ANG from each AFSC	9
	1. Method 1: Forecasting Remaining in Service	10
	2. Method 2: Forecasting Transitions Between AFSCs	11
4.	Data and Descriptive Statistics	13
5.	Performance.....	17
	A. Retention Prediction Model Individual-Level Performance	17
	B. AFSC Exit Prediction Performance	21
6.	Conclusion.....	25
	A. Synopsis of Findings	25
	B. Future Direction	25
	Appendix A. Selecting Transition Model Samples.....	A-1
	Appendix B. Data Preparation Steps	B-1
	Appendix C. References	C-1
	Appendix D. Illustrations.....	D-1
	Appendix E. Abbreviations.....	E-1

This page is intentionally blank.

1. Introduction

Air National Guard (ANG) airmen frequently experience matriculation delays of up to a year to enter training, which harms unit readiness and contributes to inefficient resource allocation. Training consists of two parts: basic military training (BMT), which lasts approximately eight weeks and provides individuals with basic skills to become an airman, and technical (or “tech”) schools, which prepare airmen to enter their specific occupations, enumerated by Air Force Specialty Codes (AFSCs). ANG leadership relies on the active component of the U.S. Air Force (AF) to provide enlisted entrants with basic and technical skills training, and due to planning and budgeting requirements, the ANG must account for BMT and tech school training slots up to five years in advance. Currently, individual ANG Career Field Managers (CFMs) collect data through various methods to forecast school slot demand generated by airmen who exit the ANG or switch occupations, and end strength adjustments to particular career fields. More precise and reliable ANG school slot demand forecasts may help to reduce these delays.

Demand for BMT and tech school enrollments flow from many potential sources related to airmen’s career paths to-date. Direct enlisted ANG accessions (recruits without prior military service) must complete both BMT and tech school trainings, while prior non-AF service members must generally only complete tech school. Depending on their previous AFSC and technical training, prior active/reserve enlisted Air Force service members may be required to complete tech school. Current ANG airmen changing to another AFSC may also need retraining, depending on the similarity between their previous training and their future AFSC. This analysis addresses one piece of the school slot demand-generating process—exits from service—and provides the complete concept for estimating school demand.

Using the Retention Prediction Model (RPM-ANG), a machine learning (ML) capability developed by the Institute for Defense Analyses (IDA), this analysis provides the ANG with enhanced exit forecasts for each career field subdivision, expressed as a three-digit Duty AFSC. In particular, we use a combination of gradient-boosted trees and Markov transition models to estimate exits from ANG service (accounting for personnel movement between AFSCs) annually from June 2022 through June 2026. This document provides detailed technical information on the data and methodologies used to produce these forecasts, which have been provided to the ANG separately.

This paper is organized as follows: Chapter 2 provides background information and discusses the scope of this work. Chapter 3 provides the methodology used to produce the

forecasts. Chapter 4 overviews the data used to produce the forecasts. Chapter 5 provides details and discusses the performance of the forecasting method. Chapter 6 concludes with a synopsis of our results and discusses future directions for research that can benefit the ANG by way of improved forecasts.

2. Background and Scope

A. Background

Each ANG airman is assigned to a particular occupation, designated by an AFSC, and while all direct accession ANG recruits with no prior service must attend BMT, the demand for technical schools is driven by specific occupations. Some occupations may require more recruits in a given year than others and may not exhibit consistent annual demand patterns. While BMT provides skills useful for all airmen regardless of their occupation (and all service members, in general), technical school training is tailored to the responsibilities that each airman will be performing. We do not possess specific data regarding which technical schools are offered, nor do we possess data regarding which AFSCs require the same technical schools; however, we are able to provide forecasts on an AFSC level, which we believe will provide ANG leadership with flexibility.

An AFSC is a five-digit alpha-numeric code where each character represents a hierarchical ordering of an airman's occupation. For enlisted AFSCs, the first digit corresponds to the career group, the first two digits together correspond to the career field, and the first three digits correspond to the career field subdivision. The fourth digit corresponds to the skill level, and all five digits correspond to the specific AFSC (Air Force Manual 2021). At the direction of ANG leadership, we forecast exits for each career field subdivision, hereafter referred to as a three-digit AFSC.

Each airman can have up to five AFSCs, which can vary or can overlap: primary, secondary, tertiary, duty, and control. Primary AFSC corresponds to the specialty in which the airman is best qualified and has the highest skill level. Secondary and tertiary AFSCs correspond to the specialty in which the airman is second- and third-best qualified, respectively. Control AFSC corresponds to the specialty that helps determine training requirements, promotions, and other administrative matters. Duty AFSC corresponds to the specialty in which the airman is actually serving (Air Force Manual 2021). At the direction of ANG leadership, we forecast exits using the duty AFSC. This is more useful to the ANG, as training demand forecasts would then correspond to actual labor needs rather than administrative assignments and historical training information.

Not only does the demand of new entrants vary by AFSC in any given year, but the demand of new entrants for a particular AFSC may vary by year. In other words, a particular occupation may have different labor needs depending on the year, which can be a result of many factors, including the current geo-political climate or military policies. For example, the 2013 budget cuts that led the Air Force to make significant drawbacks on its

maintenance fields illustrates a structural shift that resulted in significant changes in the need for certain occupations (Mehta 2018). Other temporal effects and changes can dramatically shift training needs and illustrates the inability to directly forecast demand using historical data and subsequently, highlights the need for an individually driven demand model.

B. Components of Training Demand

As is common in many professions, new individuals must be hired each year to account for individuals who left the occupation and to account for structural resizing. We use the term AFSC demand to refer to this accounting. AFSC demand is driven by three factors:

- Exits from the ANG from each specific AFSC by individuals who leave ANG service altogether;
- Transfers from one AFSC to another by individuals who change their occupation to a different AFSC; and
- Structural adjustments for the AFSC sought by ANG leadership to resize to a level that ANG leadership desires.

These three components make up the total number of individuals who must be added to that occupation to maintain the necessary labor levels.

Similarly, various mechanisms are used to fill AFSC demand. We use the term AFSC supply to refer to the number of individuals who are allocated to fill AFSC demand. There are five sources of AFSC supply:

- Non-prior service entrants – individuals who enter the AFSC and the ANG from civilian life;
- Pre-trained prior service entrants – individuals who enter the AFSC and the ANG from another service and who already have the corresponding technical school training;
- Non-trained prior service entrants – individuals who enter the AFSC and the ANG from another service and who do not have technical training;
- Pre-trained transfers into the AFSC – individuals already in the ANG who change their occupation from another AFSC and who already have the corresponding technical school training; and
- Non-trained transfers into the AFSC – individuals already in the ANG who change their occupation from another AFSC and who do not have technical school training.

Table 1 details which of the five supply sources require BMT and technical school training prior to entering the AFSC.

Table 1: AFSC Supply Sources

Supply Source	Requires BMT	Requires Technical School Training
Non-Prior Service Entrants	Yes	Yes
Pre-Trained Prior Service Entrants	Already Satisfied	Already Satisfied
Non-Trained Prior Service Entrants	Already Satisfied	Yes
Pre-Trained Transfers Into AFSC	Already Satisfied	Already Satisfied
Non-Trained Transfers Into AFSC	Already Satisfied	Yes

Since only some entrants require training prior to entering an AFSC, training demand is equal to the number of untrained individuals who fill the AFSC demand. From this definition, the total training demand can be calculated only after the supply variables are determined.

There are three stages to determine the number of AFSC technical school training and BMT slots required within a given time period. For each AFSC: First, estimate the demand; second, determine how to satisfy demand using individuals from the five supply sources; and third, based on the supply source allocations for each AFSC, calculate both the total number of AFSC entrants who will require technical school training, and the total number of new AFSC entrants who will require BMT across all AFSCs.

C. Scope of the Analysis

While we are interested in forecasting training demand for BMT and technical school training, collectively referred to as training demand, training demand is difficult to forecast, as doing so requires knowledge regarding many components previously discussed. We deconstruct this problem into a series of steps that can provide reasonable estimates for training demand. However, due to limitations in data availability, we only focus on estimating one piece of the overall training demand: exits from service from each AFSC.

We do not estimate adjustments to the force structure desired by the ANG, as the Air Force determines those adjustments. Further, estimating the supply of individuals available to satisfy AFSC demand requires understanding whether the ANG can control the flow through each available source (combinations of new ANG members, prior service ANG entrants, and ANG members switching AFSCs).¹ For example, the ANG may control whether individuals can transfer into or out of AFSCs, or how many individuals enter from prior service. To adjust the exit forecasts, we assume that transfers into and out of AFSCs

¹ Some of these sources may be controlled-but-bounded, in that the ANG can choose how many to select but are faced with a limited pool.

are uncontrolled by the ANG; however, we lack data on which transfers require retraining prior to entering an AFSC. Similarly, information regarding the flow of prior service entrants to the ANG (and their associated AFSCs) is unavailable to this analysis.

Due to these data limitations, this paper forecasts exits from ANG service. As we illustrate in Chapter 3, ANG exits are a large driver of—but do not exclusively determine—AFSC school slots demand. Chapter 3 presents a concept for how to estimate total ANG training demand once information is available on controlled supply sources, retraining requirements, and former service entrant flows.

3. Methodology

We focus on forecasting AFSC demand generated by exits from the ANG from each AFSC. Sections 3.A and 3.B provide the relevant context and technical framework for forecasting technical school demand and explain how our forecasts fit into the larger picture regarding overall training demand. Section 3.C details the specific methodology used in this analysis to forecast exits from the ANG from each AFSC.

A. Mapping AFSC-level Personnel Flows

To support estimation of ANG school slots demand, we focus on personnel flows into and out of duty AFSCs. These flows, which are discussed in section 2.B, are illustrated in Figure 1 for a generic AFSC denoted by j .




Figure 1: ANG Enlisted Personnel Flows into and out of AFSC j Positions

Enlisted personnel moving into and out of AFSC j potentially trigger demand for BMT and/or technical school training slots. Total enlisted school slot demand for BMT and technical schools equals the number of billets in an AFSC to be filled by untrained individuals. That is, technical school slots are needed when ANG personnel flow via

pathways (2), (8), and (15), and BMT slots are needed when ANG personnel flow via pathway (1). By contrast, school slots are not required for pathways (7) and (13) as those arrows reflect individuals entering AFSC j who are already trained. Depending on whether a vacated billet is to be filled, vacant billets arise from pathways (9), (10), and (11). For AFSC j , personnel flows via pathways (3), (4), (5), (12), (14), and (16) trigger neither school slot demand nor billet vacancies.² We exclude from the flowchart those who drop out prior to AFSC assignment.

To formalize this process, let

- $\Delta_{j,t}$ = Desired change in steady state for AFSC j in year t (e.g., planned resizing), which is not shown in Figure 1
- $\varphi_{j,t}$ = Exits from ANG from AFSC j between $t-1$ and t , i.e., pathways (11)
- $\lambda_{j,t}$ = Switches to other AFSCs from AFSC j between $t-1$ and t , i.e., pathways (9) and (10)
- $\gamma_{j,t}$ = AFSC switch-ins to AFSC j between $t-1$ and t who need tech school training, i.e., pathway (8)
- $\theta_{j,t}$ = AFSC switch-ins to AFSC j between $t-1$ and t who do not need tech school training, i.e., pathway (7)
- $\mu_{j,t}$ = Prior service entrants to AFSC j between $t-1$ and t who need tech school training, i.e., pathway (15)
- $\rho_{j,t}$ = Prior Service Entrants to AFSC j between $t-1$ and t who do not need tech school training, i.e., pathway (13)
- $\omega_{j,t}$ = Non-prior service entrants to AFSC j between $t-1$ and t who need tech school training, i.e., pathway (2)
- d_t = Drop outs between BMT tech school training between $t-1$ and t , i.e., pathway (1) minus the sum of pathways (2) and (3)³

These parameters are represented as first-differences, that is, the numerical difference between the current and the previous period. As we use annual data in the context of this work, each model period corresponds to one calendar year.

B. Estimating Total Enlisted School Slot Demand

The two outcome variables that we seek to estimate are:

- B_t = Basic military training slots to offer in year t
- $D_{j,t}$ = Technical school slots to offer for AFSC j in year t

² Pathway (6) simply denotes fully trained personnel migrating into an AFSC.

³ We lack data on dropouts from BMT and tech schools; however, we only exclude dropouts from technical schools as we suspect that the number of dropouts from technical schools is small and the number of dropouts after or during BMT is much higher.

As demand for the service as a whole is driven by demand at an AFSC level, we first focus on supply and demand for an AFSC where billet demand (on the left side) in AFSC j equals the supply of ANG enlisted personnel.

$$\Delta_{j,t} + \varphi_{j,t} + \lambda_{j,t} = \gamma_{j,t} + \theta_{j,t} + \mu_{j,t} + \rho_{j,t} + \omega_{j,t}$$

Letting A be the set of all AFSCs in the forecasting year, demand for BMT school slots in time t (B_t) is then written as:

$$\begin{aligned} B_t &= \sum_{j \in A} (\omega_{j,t}) + d_t \\ &= \sum_{j \in A} \max(0, \Delta_{j,t} + \varphi_{j,t} + \lambda_{j,t} - \gamma_{j,t} - \theta_{j,t} - \mu_{j,t} - \rho_{j,t}) + d_t \end{aligned} \quad (1)$$

Demand for slots for each technical school in time t ($D_{j,t}$) can then be written as:

$$\begin{aligned} D_{j,t} &= \gamma_{j,t} + \mu_{j,t} + \omega_{j,t} \\ &= \max(0, \Delta_{j,t} + \varphi_{j,t} + \lambda_{j,t} - \theta_{j,t} - \rho_{j,t}). \end{aligned} \quad (2)$$

Both demand models use a lower bound of 0 to account for the possibility that existing supply may already exceed demand, resulting in no new desired training slots.

C. Estimating Exits from the ANG from each AFSC

We forecast the number of ANG airmen observed in period t who exit the ANG from each AFSC over each of the next 5 years. This is not the same as forecasting *total* exits from an AFSC for each of the next 5 years, as we are unable to forecast future exits of those not yet in service using this framework. For example, if an individual joins service two years from now and exits service from a particular AFSC three years from now, then this framework will not accommodate prediction of that exit, as it is not able to account for entry of that individual. For this reason, we expect all predictions of $\varphi_{j,t+h}$ after the first forecast year ($h > 1$) to be lower than the observed total exits from the ANG from each AFSC, since the latter includes new entrants during the prediction window. Given this limitation, we then estimate (3).

$$\varphi_{j,t+h} = \sum_{k=1}^J \left(\sum_{i \in X_{k,t}} (S_{t+h-1}(i) - S_{t+h}(i)) (I_{j,t+h-1}(i)) \right) \quad (3)$$

where

J = the number of AFSCs in the forecasting year

$h = 1, \dots, 5$ is the forecast horizon

$X_{j,t+h}$ = individuals in AFSC j in $t + h$.

$S_{t+h}(i) = 1$ if i is in service in $t + h$ and 0 otherwise

$I_{j,t+h}(i) = 1$ if i is in service and in AFSC j in $t + h$ and 0 otherwise

Our methodology combines two distinct methods to obtain an estimate for (3). The first method estimates the probability that an individual will remain in service in $t + h$ (the probability associated with $S_{t+h}(i)$). The second method estimates the probability that an individual will appear in AFSC j in period $t + h - 1$ (the probability associated with $I_{j,t+h-1}(i)$).

1. Method 1: Forecasting Remaining in Service

We denote $P_{t+h}(i)$ to be the probability that individual i remains in service in $t + h$. To estimate this probability (and thus the probability they will exit service between $t + h - 1$ and $t + h$), we use the Finite Interval Forecasting Engine (FIFE). FIFE is a free and open source machine learning package developed by IDA and sponsored by the Office of the Under Secretary of Defense for Personnel and Readiness to use machine learning for discrete-time survival analysis on panel datasets (Institute for Defense Analyses 2021).⁴ Specifically, we use the gradient boosted tree implementation to forecast survival. We refer to the resulting estimated model as a retention prediction model, which uses a collection of numeric and categorical features containing information on each individual to predict the probability that each individual will still remain in service for each of the next 5 years.

We estimate two separate retention prediction models: an evaluation model and a future forecasting model. The evaluation model is used to evaluate the performance of the method in forecasting future survival. This model is trained using 75% of the individuals in the data between 2005 and 2015, reserving the remaining 25% as a validation set,⁵ and it is then applied to all 2016 individuals to forecast their exit probabilities for 2017 through 2021. Note that, since this method does not account for new entrants, individuals entering the data after 2016 are excluded from the evaluation forecast window. Forecasting performance is then evaluated by comparing predictions for this time period to the actual observed exit occurrences.

The future forecasting model is used to produce the forecast interest for 2022 through 2026. It is trained using 75% of the individuals in the data between 2005 and 2020, reserving the remaining 25% as a validation set, and it is then applied to all 2021 individuals to forecast their exit probabilities for 2022 through 2026.

The two models are different, as the underlying training data changes. The first model allows us to determine the accuracy of our forecasts if we had attempted to answer the question of interest 5 years ago, which serves as a lens through which we can consider our actual forecasting results from the second model. While we could simply apply the first evaluation model to 2021 individuals to obtain future forecasts, doing so would not

⁴ FIFE is written in Python and available via the Python Package Index (PyPI), or via GitHub at <https://github.com/IDA-HumanCapital/fife>.

⁵ The validation set is used to optimize the hyper-parameters in the model; this is a common method for selecting the best model in a way that avoids data snooping bias.

consider data between 2016 and 2020, resulting in less than optimal performance for future forecasts.

2. Method 2: Forecasting Transitioning Between AFSCs

To forecast the probability of an individual residing in a given AFSC in a future time period, we use a discrete-time Markov chain model. We use this method to estimate the aggregate transition probability for ANG enlisted personnel between AFSCs (“states”) based only on their current AFSC.⁶ Because these transition probabilities are aggregate, they do not account for individual-level variation in transition probabilities. We may expand this capability in the future to account for individual-level heterogeneity in transitions; however, results from the primary forecast algorithm will be aggregated at the three-digit AFSC level, so we will rely on these AFSC-level predictions for forecasting these transitions.

Mirroring the two retention-prediction evaluation and forecast models, we estimate two Markov transition models – one for evaluation and a second for forecasting over the period of interest. For each model, we define M to be the one-period transition matrix containing the probabilities of an individual transitioning between AFSCs and $\hat{M}$ to be its estimate. M is a $J \times J$ matrix whose k, j th element gives the probability that an individual will transition from AFSC k to AFSC j one period in the future.

M is estimated using counts of switching of the individuals among AFSCs. For example, a specific transition probability between AFSC A and B is estimated by dividing the number of individuals who switched from A to B by the total number of individuals who started in A. Estimation of the probability of transitioning more than one period in the future is accomplished by multiplication of the matrix M ; that is, the h -period transition probabilities are given by M^h .

We can forecast the quantity of interest, $\varphi_{j,t+h}$, defined in (3) by accounting for the probability that each individual will exit service between $t + h - 1$ and $t + h$:

$$\hat{\varphi}_{j,t+h} = \sum_{k=1}^J \hat{M}_{k,j}^{h-1} \sum_{i \in X_{k,t}} \left(\hat{P}_{t+h-1}(i) - \hat{P}_{t+h}(i) \right) \quad (4)$$

We restrict the transition model training sets to years 2014-2015 for the evaluation model and to years 2013-2020 for the future forecast model, as examination of individuals in the validation set revealed sensitivity to the validation window, indicating potential non-stationarity (see Appendix A for more information). The evaluation forecast and future forecast windows remain the same as with the two respective retention prediction models. For the same reason mentioned in the discussion of the evaluation of the retention prediction model (this method does not account for new entrants), individuals entering the

⁶ This is called the Markov property. This property implies that the transition probability between states depends on past states only through the most recent state (Hamilton 1994).

data after 2016 are excluded from the evaluation forecast window (2017-2021) for the evaluation model.

4. Data and Descriptive Statistics

To generate the required three-digit AFSC-level school slot demand forecasts, this project uses administrative military personnel data from the Defense Manpower Data Center (DMDC). IDA receives this data on a quarterly basis as part of an institutional data sharing agreement and houses it in IDA’s Personally Identifiable Information Enclave. Measured monthly, the data currently spans January 2000 to June 2021. IDA performs routine cleaning and regularization of these data assets to ensure their consistency. We construct our analytic dataset using DMDC Reserve, Pay, Title 10 Activations, Defense Enrollment Eligibility Reporting System (DEERS),⁷ Deployment, and Casualties files. Additionally, we use ANG data corresponding to Extension Course Institute (ECI) training, provided to us by the National Guard Bureau/Training Resources and Programming Branch (NGB-A1DU). We also include external labor market information from the Current Population Survey (CPS) by matching an airmen’s current domestic mailing zip code.

Table 2: Data Sources

Dataset	Description	Years Used	Source
Master	Administrative and Personal for Servicemembers	2005-2021	DMDC
Pay	Pay/Bonus Information for Servicemembers	2005-2021	DMDC
Activations	Title 10 Activations for Servicemembers	2005-2021	DMDC
Deployments	Deployments for Servicemembers	2005-2021	DMDC
Family	Family information for servicemembers	2005-2021	DMDC
Casualties	Casualties for individuals or in UIC	2005-2021	DMDC
ANG Training	ECI courses completed by ANG personnel	2005-2020 ⁸	NGB-A1DU
Labor Market	Unemployment Rates by Occupation and Location	2005-2021	CPS

⁷ We augment the historic DMDC Family files with information from the Defense Enrollment Eligibility Reporting System (DEERS).

⁸ Due to data reception limitations, we only have this data through April of 2020. For this reason, our model uses lagged values to correspond with training completed between 26 and 14 months prior to observation.

We restrict the analysis to all enlisted ANG personnel present in our DMDC/ANG data between June 2005 and June 2021. Due to data quality concerns impacting years 2000 through 2004, we restricted our subset to calendar year 2005 and thereafter. For computational reasons, as well as limited month-to-month variation in the data, the final analysis dataset is a panel consisting of annual snapshots as of June 30 of each year reflecting the then-current status of each individual, as well as summarized information pertaining to the previous 12 months.

If an enlisted airman commissions as an officer within our analytic period, we retain observations prior to commissioning. If an individual exits the ANG and subsequently joins another service or component while in our dataset, we keep only the ANG observations. We consider an individual to have exited enlisted ANG service (or simply “exited service”) when they no longer appear in our ANG analytic dataset; a given exit could represent an individual leaving military service altogether, commissioning, or joining another service. Due to a lack of data, we are unable to account for individuals who have “dropped out” of ANG service, for example, by failing to attend drill. Because individuals are observed annually, statements like “ X individuals exited in period t ” should be interpreted as X individuals having exited the dataset at some point between periods $t-1$ and t . Several additional steps were taken to transform these data sources into a usable feature space for modeling. More information on these steps can be found in Appendix B. Table 3 provides descriptive statistics for the data used in these analyses.

Table 3: Descriptive Statistics

	Evaluation Model	Forecasting Model
Training Set Years	2005-2015	2005-2020
Validation Set Years	2005-2015	2005-2020
Test Set Years	2016	2021
Forecasting Period	2017-2021	2022-2026
Training and Validation Sets (Combined)		
Number of Total Observations	1,011,831	1,463,888
Number of Distinct Individuals	176,289	215,895
Average Number of Individuals Per Year	91,985	91,493
Number of Distinct Duty AFSCs	386	403
Average Number of Duty AFSCs Per Year	192.1	175.4
% Male	81.2%	80.5%
% White	84.2%	83.9%

Test Set		
Number of Total Observations	89,630	91,144
Number of Distinct Individuals	89,630	91,144
Average Number of Individuals Per Year	89,630	91,144
Number of Distinct Duty AFSCs	148	113
Average Number of Duty AFSCs Per Year	148	113
% Male	76.7%	77.8%
% White	83.4%	83.1%
Number of Total Features	633	633
Number of Numeric Features	419	419
Number of Categorical Features	214	214

Figure 2 presents the Kaplan-Meier survival curve, which illustrates the retention probability after completing each year of service in the ANG. The retention probability drops substantially after the sixth year of ANG service, corresponding to the completion of the initial contract for individuals with no prior military service (ANG 2021).




Figure 2: Kaplan-Meier Retention Probability by Year of Service

This page is intentionally blank.

5. Performance

This chapter documents the performance of the evaluation model. The evaluation model predicts exits between 2017-2021, and although it employs methods identical to the future forecast model, differences in the underlying training data make it a different model. As such, these results should be interpreted as illustrating how the forecasting method would have performed if applied five years ago. We then compare the results of the evaluation model to what actually occurred during the evaluation period, and to a set of loss predictions generated by extrapolating historic attrition rates at the three-digit AFSC level (approximating current practice).

A. Retention Prediction Model Individual-Level Performance

We document the performance of the individual-level exit model within the evaluation time period. We report individual-level performance in addition to AFSC-level performance as the ANG Retention Prediction Model (RPM-ANG) is a core development of this analysis.

From an analytic perspective, the data is “imbalanced” in the sense that the vast majority of individuals choose to stay in any given time period. This makes both our prediction and our evaluation tasks more challenging. Many evaluation metrics have difficulty capturing performance in a satisfactory manner in the presence of this type of imbalance in the outcome, so care must be taken in choosing appropriate evaluation criteria. For example, if a population faced 95% survival in one period, a one-period survival analysis model that predicted “stay” for everyone, would obtain a 95% accuracy.⁹ To account for this issue, we focus on precision and recall metrics corresponding to exits rather than retentions (Saito 2015).¹⁰ In general, one should note that no single evaluation criteria can perfectly capture model performance.

To operationalize our metrics, we must select the classification threshold, α , for the probability of survival, $P_{t+h}(i)$, above which any individual i is said to survive for analytical purposes. A common method for identification of the optimal threshold value

⁹ Accuracy represents the fraction of correctly predicted observations: $\frac{TP+TN}{TP+TN+FP+FN}$, where TP is true positives, TN is true negatives, FP is false positives, and FN is false negatives.

¹⁰ Precision represents the fraction of individuals who were predicted to exit who were observed to have exited: $\frac{TP}{TP+FP}$. Recall represents the fraction of individuals who were observed to have exited who were correctly predicted as exiting: $\frac{TP}{TP+FN}$.

involves choosing the threshold associated with the maximum F1 score; we use this method.¹¹ Table 4 presents key performance metrics for each period, and table 5 provides confusion matrices for each period based on the optimal classification threshold value.

Table 4: Performance Metrics by Period for the Evaluation Method

Period	Classification Threshold	Precision ¹²	Recall	Accuracy	F1 Score
1	0.85	0.41	0.51	89.24%	0.45
2	0.75	0.47	0.55	83.59%	0.51
3	0.77	0.49	0.70	74.51%	0.57
4	0.75	0.51	0.78	68.36%	0.61
5	0.71	0.53	0.79	64.81%	0.64
Pooled	0.75	0.50	0.72	76.09%	0.59

Table 5: Classification Performance for Each Period

Period	Classification Threshold	True Positives (TP)	False Positives (FP)	False Negatives (FN)	True Negatives (TN)
1	0.85	4,016	5,829	3,811	75,974
2	0.75	7,505	8,614	6,091	67,420
3	0.77	15,357	16,250	6,594	51,429
4	0.75	22,649	22,101	6,257	38,623
5	0.71	27,607	24,236	7,307	30,480
Pooled	0.75	77,377	77,329	29,817	263,627

Figure 3 provides the precision-recall curve (left) and precision-recall gain curve (right) for pooled observations across the five-year window. Precision-recall curves offer a visual representation of model performance based on binary (stay/exit) classification using various threshold values ranging from 0 to 1. The area under the precision-recall curve and precision-recall gain curve can be used as consolidated measures of performance and can assist with identifying the optimal threshold value. A traditional precision-recall curve must be interpreted relative to some baseline measurement, whereas precision-recall gain curve incorporates the baseline measurement, making it easier to interpret (Flach and Kull 2015).

¹¹ The F1 score is the harmonic mean of precision and recall: $\frac{2*Precision*Recall}{Precision+Recall}$. There are alternatives to the F1 score when precision and recall should not be equally weighted.

¹² A precision of 0.50 indicates that 50% of those predicted to exit actually exit. While this number seems exceptionally low on its own, it is a noticeable improvement on simple guessing (0.24), which is due largely to class imbalance. Conversely, precision for retentions is 0.90, which indicates that 90% of those predicted to stay actually stay.

We pool observations from all periods to create these plots for general illustration of model performance. In general, one may wish to examine precision-recall curves for each period in the forecasting window.




Figure 3: Precision Recall, Precision-Recall Gain Curves for all Five Years Pooled.

When interpreting the precision-recall curve, it is important to note that the baseline for evaluation shifts based on the prevalence of the positive class. In this particular case, the baseline, or “Everyone Exits” classification rate results in a precision of 0.24 (if every observation is predicted to exit, 24% of those observations would be correctly predicted). This baseline is not intended to reflect a current method used by the ANG; but rather, to illustrate the performance of a rudimentary model. The area under the curve reveals that the RPM performs substantially better than the baseline, as the area under the precision-recall curve is 0.63 (compared to 0.24 for the baseline).¹³ The area under the precision-recall gain curve is 0.81 (compared to a baseline of 0.5).

Figure 4 illustrates the distribution of predicted retention probabilities separated by outcome for pooled observations across the five-year window. The red line marks the classification threshold value which maximizes the F1 score and separates predicted exit from predicted retention.¹⁴ The axes are scaled differently to account for the class imbalance.

¹³ Given the imbalance in the prevalence of stay and exit outcomes, one may wonder how the model performs in comparison to an “Everyone Stays” baseline. In this case, the area under the precision-recall curve for the baseline measure is 0.761, and the area under the precision-recall curve for the associated model using “stay” as the positive class is 0.925, yielding a positive performance difference of 0.164. This appears to offer lower performance gains relative to the Everyone Exits comparison ($0.634 - 0.239 = 0.395$); however, these metrics can be thought of as different sides of the same coin. In both cases, the developed forecasting model outperforms the baseline.

¹⁴ If we let the positive class be “exits,” then the top left represents true positives, the top right represents false negatives, the bottom left represents false positives, and the bottom right represents true negatives.




Figure 4: Predicted Retention Probabilities Separated by Observed Outcome

Figure 4 suggests two observations. First, the model is excellent at determining when someone will stay in service, likely because certain data features are highly correlated with retention. Second, the model is less adept at predicting exits compared to its ability to predict retentions. We observe that for those actually exiting, the distribution of predicted retention probabilities is much wider and centered farther from the truth (0). For comparison, a perfect model would produce a probability distribution for true exits with a point mass at 0. Plausible reasons for this include a potential lack of features correlating highly with exit and the imbalance in the data resulting in many more true non-exits from which the model learns. Future analysis could potentially improve on this limitation by improving the feature space with inclusion of additional drivers of attrition. One such source would include information regarding civilian careers, which are often linked to attrition as service may present an obstacle to airmen's civilian careers.¹⁵

¹⁵ This data was not included in this analysis as IDA was informed that quality data on civilian careers did not exist.

B. AFSC Exit Prediction Performance

In this section, we present performance metrics for the joint RPM-Transition model. We examine mean absolute error to understand how well the model predicts exits from each AFSC during the evaluation period. To control for differences in AFSC sizes, we consider the predicted share exiting at each forecast horizon. That is, we divide the expected number of individuals exiting by the initial AFSC headcount to obtain an expected share exiting, which we then compare to the observed exit share at that point in time. This normalizes accuracy at the AFSC level.

Table 6 contains the mean absolute error for the RPM-Transition model within each year of the evaluation period.¹⁶ To provide a comparison metric, we include an “AFSC-Average” model, which is akin to a simple trend extrapolation mechanism. For each forecast horizon, the AFSC-Average model takes historical data and calculates the average exit share from that AFSC for that window. For example, for the second period in the forecast, for each year between 2005 and 2014, the model will take each two-year window (i.e., 2005 to 2007, 2006 to 2008, etc.), obtain the observed exit shares for each AFSC, and then average those shares to obtain a prediction for the two-year forecasted exit share (corresponding to 2017 for the evaluation model).

The AFSC-Average model slightly outperforms the RPM-Transition model in the first period, but the RPM-Transition model outperforms in all other years. Averaged across periods, the RPM-Transition model is inaccurate by approximately 7% of exits for any given AFSC, whereas the AFSC-Average model is off by approximately 23% of exits for any given AFSC: an average improvement of 70% over the 5-year horizon. We believe the difference in performance in the first year can likely be explained by the inability of individuals to switch AFSCs before exiting in the first year. As a result, the AFSC-Average model predicts the historical average one-year exit rate from an AFSC, which remains relatively constant for a given year.


Table 6: Mean Absolute Error of Exit Share (RPM-Transition vs. AFSC-Average)¹⁷

Year	Forecast Horizon	RPM-Transition	AFSC-Average	% Improvement using RPM-Transition
2016	$t+1$	0.03	0.03	-6.85%
2017	$t+2$	0.04	0.12	66.79%
2018	$t+3$	0.06	0.21	71.39%
2019	$t+4$	0.09	0.33	73.91%
2020	$t+5$	0.12	0.43	73.29%
Average		0.07	0.23	70.28%

¹⁶ In this evaluation, we unified the data in the case of AFSC renaming.

¹⁷ Values may not add up within table due to rounding

Figure 5 displays model performance for the 20 largest AFSCs. Model performance in this figure is measured as the percentage difference of predicted exits from actual exits per AFSC $\left(\frac{\text{Predicted}-\text{Observed}}{\text{Observed}} * 100\right)$, where we average across periods. AFSC size in this figure is calculated according to the number of airmen in 2016. Positive (red) bars for a given AFSC denote overestimates, while negative (blue) bars correspond to underestimates. The model appears to underestimate exits more frequently than overestimate them, which could be a result of some combination of underestimating transfers into an AFSC, underestimating exits on an individual level, or just randomness in the data. None of these suggestions indicate that future forecasts will necessarily also underestimate exits; rather, they serve to illustrate a potential source of future investigation.¹⁸



**Figure 5: Forecast Error among the 20 Largest Duty AFSCs
(by Assigned Personnel in 2016)**

In addition to the model performance caveats described thus far, exits from some AFSCs may be more challenging to forecast than others. This may be due to observable differences between individuals across AFSCs as represented in their feature values or unobserved or partially observed effects (e.g., policy changes, restructuring, specific economic conditions) that can impact exits from service and AFSC transfers. While the

¹⁸ For example, we suspect that AFSC 3A1 was recoded to AFSC 3F5; however, it did not meet our threshold for a recode (AFSC 3A1 appeared in our data in 2019 and less than 90% of individuals in 3A1 moved to 3F5 after 2018). A list of actual recodes would likely result in a noticeable increase in model performance.

retention prediction model captures some individual heterogeneity through observable information, the current version of the transition model cannot, indicating that such information could be used to improve forecasts of the transitions. Further, some unobserved information simply cannot be included in either model due to its unavailability (for example, information regarding comparable civilian careers).

Table 7 documents the 10 AFSCs with the smallest forecast error and the 10 AFSCs with the largest forecast error, where “Mean Error” corresponds to headcounts. Observe that larger AFSCs result in lower absolute forecast error on a percentage basis than smaller AFSCs. This could be due to the difference in magnitude (the same size error would result in a greater percentage error for smaller AFSCs); however, this reasoning is confounded by the observation that, on average, the magnitude of the mean error is smaller for smaller AFSCs.¹⁹ Additionally, it appears that AFSCs corresponding to systems, air crews, and other operations fields experience lower absolute forecast error on a percentage basis than occupations like recruiters and mission support. This could be related to the size of the AFSCs or that some AFSCs are potentially more sensitive to shifting policy and other temporal effects. These hypotheses could be investigated in future analyses.

Table 7: AFSCs with the Highest and Lowest Average Forecast Error

Model Performance	AFSC	Description	Population in 2016	Mean Error (Std)	Mean Percent Difference (Std)
Best	1A2	Aircraft Loadmaster	1,088	5.83 (9.31)	0.29 (7.33)
	2A9	Bomber	73	-1.38 (3.26)	0.33 (23.9)
	5R0	Chaplain Assistant	285	1.64 (6.99)	-0.62 (17.99)
	2A6	Aerospace Propulsion	6,133	22.19 (60.01)	-0.64 (10.19)
	2W1	Aircraft Armament Systems	1,808	9.75 (19.44)	0.87 (10.01)
	8U0	Unit Deployment Manager	56	-1.55 (1.95)	-0.95 (18.83)
	3F0 ²⁰	Personnel	2,709	14.0 (39.47)	-1.0 (8.83)
	2A3	Avionics Systems	3,771	16.2 (48.38)	-1.06 (10.87)
	2A8	Mobility Air Forces	1,046	6.45 (20.59)	-1.36 (16.44)
	2S0	Materiel Management	2,764	17.11 (49.78)	-1.45 (13.21)
Worst	3A1	Administration	1,164	64.35 (87.59)	29.3 (43.86)
	1N2	Signals Intelligence Analyst	121	-7.81 (2.41)	-29.96 (17.5)
		Enlisted Accessions			
	8R0	Recruiter	318	9.8 (4.54)	32.81 (8.94)
	3F4	Equal Opportunity	117	-18.0 (8.18)	-35.92 (9.05)
	8I0	IG Superintendent	129	21.41 (21.37)	48.38 (39.5)
	8R2	Second-Tier Recruiter	217	32.01 (32.89)	50.83 (42.46)

¹⁹ The Pearson correlation coefficient was 0.50 between the size of an AFSC in 2016 and the absolute value of the mean error for all AFSCs.

²⁰ We believe that 3S0 was recoded to 3F0 in 2018; however, for evaluation purposes, we compare the predictions for 3S0 to the observed exits from 3S0 in 2017 and the observed exits from 3F0 in 2018 through 2021.

Model Performance	AFSC	Description	Population in 2016	Mean Error (Std)	Mean Percent Difference (Std)
	8T0	Professional Military Educ Airborne Cryptologic	24	2.56 (2.76)	52.67 (57.52)
	1A8	Linguist	23	-3.02 (2.02)	-59.18 (17.18)
	3F3	Manpower	13	1.49 (1.52)	64.6 (66.77)
	1C8 ²¹	Radar	691	207.79 (202.51)	104.52 (84.35)

²¹ AFSC 1C8 first appeared in 2015, meaning that the error is inflated because the model was not well equipped to predict the average number of retentions and switch-outs. To our knowledge, no parallel problem appears to exist in 2020, so we have reason to believe that this issue will not appear with the future forecast model.

6. Conclusion

This analysis addresses the challenge of forecasting the number of school slots needed by the ANG for each AFSC at the three-digit level over the next five years. We develop a general, reusable toolkit for forecasting the exit of individual ANG service members—the Retention Prediction Model-ANG (RPM-ANG)—and apply this tool with an extension to produce a reusable toolkit for forecasting the number of school slots needed on an ongoing basis. This analysis illustrates one of many potential applications of the RPM-ANG toolkit. The extension expands on the underlying RPM-ANG capability to account for AFSC transitions during the forecast window. We apply these tools to predict exits from ANG service among those in service as of June 30, 2021, by AFSC-designated occupation. While exits from service comprise only one part of future school slot demand, they provide a foundation for future analysis and represent a large share of total school slot demand.

A. Synopsis of Findings

We develop and employ the RPM-ANG to forecast individual ANG members' exits from service using a gradient-boosted tree specification designed for survival analysis. Because exits are significantly rarer than stay decisions in any given year, and because exits are relevant to school slot demand, we calibrate the model to optimize exit prediction. Under these parameterizations, the model correctly identifies 72% of all exits from service from 2017 to 2021.

Since individuals can switch occupations during their time in service, we layer a transition model over the raw exit forecasts to predict the AFSC from which individuals exit service. We use the resulting joint RPM-Transition model to obtain retention probabilities for each individual within each AFSC for each of the next five years. When forecasting exits from each AFSC from 2017 to 2021, the model has an average mean absolute error (by exit share) of 7% across all years, representing a 70% improvement from a baseline AFSC-average model during the same period.

B. Future Direction

This analysis considers exits among those currently in service, and thus represents only one component of the overall AFSC school slots demand-generating process. As illustrated in Chapter 2, AFSC switching itself can be a source of school slot training demand, as it can result in re-training for both the sending and the receiving AFSC. Though

we do not examine AFSC switching as a component of the demand-generating process, we develop a preliminary switching model to adjust the forecasts for exits from service. Future research might attempt to model individuals' switching behaviors using additional data. As not every individual who switches AFSCs requires retraining, such analysis would require information on available AFSC transfer options and which transfers do and do not require retraining. Further, a comprehensive model would consider the retraining requirements for prior service entrants. Given these limitations, the forecasts provided in the present analyses cannot be directly interpreted as school slot demand predictions; however, they do form one of the critical components used to calculate school slot demand.

This methodology has certain limitations. The first is its inability to forecast exits of those not currently in service (i.e., those who enter service after 2021 but exit prior to 2026). Second, while this analysis presents a substantial improvement over existing methods, the development of a more nuanced individual switching model may improve exit forecasts. In particular, the method for forecasting AFSC transitions can be improved to account for individual level heterogeneity by incorporation of individual level data.

As in any modeling exercise, performance could be improved with additional information, particularly regarding drivers of attrition. This data may include features on service members' civilian careers, or specific contract information beyond what is available from DMDC. Data that more clearly denotes constructive exit from service (e.g., failure to participate in drill weekends) would also likely improve model performance. Additional effort to identify and account for features whose meaning or statistical properties shift over time (non-stationarity) may also improve the exit predictions.

Finally, this work may yield additional returns for the ANG by integrating this effort with existing modeling of Air Force Active Component attritions. Such an effort could both identify a flow of potential ANG recruits exiting AC service and further assist career field managers in understanding expected personnel intake.

In sum, this effort provides a meaningful start to ANG's use of advanced predictive retention and exit modeling in its personnel management enterprise, and provides a solid foundation for future effort and operationalization of these techniques. This first application of RPM-ANG is now available for ANG leaders and analysts to use and build upon in future work, across a broad array of analytic and operational applications.

Appendix A. Selecting Transition Model Samples

It is critical that the occupational switching behavior during the training window closely matches that of the prediction window. Unfortunately, given the shifting nature of military duties and occupations, it is unlikely that data on occupational switching from the early 2000s will closely resemble occupational switching after 2021. As such, we implement a validation method for examining temporal shifts in the allocation processes.

We construct a transition matrix for each starting year of the training data, and we compare this matrix to the corresponding matrix from each starting year of the validation data by examination of the Frobenius norm of the difference of the two matrices.²² Figure A-1 provides heat maps of the norms for each combination of training and validation matrices. The left graph corresponds to the evaluation model and the right graph corresponds to the future model.

Two things are directly observable from the heatmaps below: first, there is a clear correlation between performance and the alignment of validation and training start years. Second, a potential structural break seems to be indicated in 2009 via a marked rise in norms. Given these facts and our prior knowledge that DMDC data collection efforts were re-defined between 2012 and 2013, we choose a window for each model that is closer to the current period. For the evaluation model, we limit our training set to 2014-2015, and for the future model, we limit our training set to 2013 to 2020.

²² The Frobenius Norm of a matrix M is the square root of the sum of the squares of the elements of M :

$$\|M\|_F = \sqrt{\sum_{i,j} |m_{i,j}|^2}.$$




Figure A-1: Transition Matrix Selections

Future analyses should consider robust external mechanisms for selecting the best validation and training windows. This is particularly challenging, as it results in a validation-of-the-validation sample problem.

Appendix B. Data Preparation Steps

Prior to modeling, we prepare the raw data for analysis. Many features in our data contain a large number of missing values. Where appropriate, we either fill the missing value with 0 or with “missing” to create a new category, and drop features that are more than 99.9% missing.²³ For the ECI training data, we impute missing course end dates by examining the duration of the same course for other individuals and adding that duration to the start date. Additionally, we drop features that are more than 99.99% constant. We transform features which have only two possible values—one missing and one non-missing—into a flag indicating whether the feature is populated. Furthermore, we transform time-based features into time-agnostic features. For example, we transform contract and commitment start dates and projected end dates to reflect the number of days since or until those events. There are almost certainly many features that exhibit non-stationarity which might adversely affect model performance in some situations; however, dealing with each of these is beyond the scope of this analysis.

In addition to including raw data features in model as appropriate, we use the raw data to produce new features. Table B-1 lists categories of features included in modeling, along with a description and the number of features in each category.

Table B-1: Feature Categories

Feature Category	Number of Features	Description
Activations	69	Frequency, tempo, duration and category of title 10 activations
Active Duty	12	Information pertaining to Active Duty service
Administrative	2	Administrative accounting and planning codes
Assignments	33	Specific assignments including occupations and units.
Benefits/ Retirement	47	Benefits, entitlement programs, retirement programs
Bonus	27	Bonus compensation
Career Hardship	17	Job difficulty and potential associated emotional or physical stress
Contracts and Commitment	28	Length of service and time since or until critical dates in their service (start date, projected end date, etc.)
Deployments	66	Frequency, tempo, duration and type of deployments

²³ While this is a higher threshold than many other studies consider for dropping missing values, the LightGBM algorithm allows for improved handling of missing values.

Feature Category	Number of Features	Description
Drilling	8	Information on drilling, specifically drilling pay
Education/Skills	9	Education level, programs, and test scores
Family	31	Information on family (spouses, dependents) and any changes in family status (newly married, divorce, new child, etc.)
Geography	16	Location of assignment, residence, and duty. Includes distance between the assignment and residence.
Housing	19	Housing arrangements, not including overseas payment information
Labor Market	3	Unemployment Rates for their occupation category and geographic area
Non-Standard Compensation	53	Special pay (i.e., pay not related to their salary or taxes)
Peer Effects	74	Qualities about other individuals in their UIC including family, education, and deployment features. Number of exits from their UICs and occupations each year and average tenure of those who served in their UIC or occupation
Personal Information	11	Demographic Information and other personal information
Prior Service	33	Information regarding any prior service in any service after 2000 (excluding service in ANG)
Rank	4	Information about their rank
Rank/Pay Mobility	9	Career mobility features reflecting changes in pay and rank
Special Positions/Duties	10	Non-typical or special jobs or duties
Standard Compensation	24	Compensation and related payment information
Training	10	Number of ECI courses taken (along with outcomes) and other training information
Traumatic Event Exposure	18	Whether person experienced a casualty themselves, or was exposed casualties in their UIC or family
Total Features	633	

Appendix C. References

- Air National Guard. “FAQ.” Air National Guard. <https://www.goang.com/faq.html>.
- Defense Manpower Data Center. Reserve Duty Master, Active Duty Master, Reserve Duty Pay, Active Duty Pay, Deployments, DEERs, Activations, and Casualty files. June 2005 – June 2021.
- Flach, Peter A, and Meelis Kull. “Precision-Recall-Gain Curves: PR Analysis Done Right.” *Advances in Neural Information Processing Systems* 1 (2015): 838–46.
- Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren and Michael Westberry. Integrated Public Use Microdata Series, Current Population Survey: Version 9.0 [dataset]. Minneapolis, MN: IPUMS, 2021. <https://doi.org/10.18128/D030.V9.0>
- Hamilton, James Douglas. 1994. Time Series Analysis. Princeton, NJ: Princeton Univ. Press.
- Institute for Defense Analyses. “FIFE: Finite Interval Forecasting Engine.” *FIFE*. Institute for Defense Analyses, 2021. <https://github.com/IDA-HumanCapital/fife>.
- Mehta, Aaron. 2018. “Former US Air Force head details decision to cut maintenance budgets in 2013.” *Military Times*. <https://www.militarytimes.com/news/your-military/aviation-in-crisis/2018/05/08/former-us-air-force-chief-details-decision-to-cut-maintenance-budgets-in-2013/>.
- National Guard Bureau/Training Resources and Programming Branch. ECI Training Data. June 2005 – April 2020.
- Saito, Takaya, Marc Rehmsmeier. “The Precision-Recall Plot Is More Informative than the ROC Plot When Evaluating Binary Classifiers on Imbalanced Datasets.” *PLoS ONE* 10(3): e0118432. <https://doi.org/10.1371/journal.pone.0118432>
- Secretary of the Air Force, Air Force Manual § (2021). https://static.e-publishing.af.mil/production/1/af_a1/publication/afman36-2100/afman36-2100.pdf

This page is intentionally blank.

Appendix D. Illustrations

Tables

Table 1: AFSC Supply Sources	5
Table 2: Data Sources	13
Table 3: Descriptive Statistics	14
Table 4: Performance Metrics by Period for the Evaluation Method	18
Table 5: Classification Performance for Each Period	18
Table 6: Mean Absolute Error of Exit Share (RPM-Transition vs. AFSC-Average)	21
Table 7: AFSCS with the Highest and Lowest Average Forecast Error	23
Table B-1: Feature Categories	B-1

Figures

Figure 1: ANG Enlisted Personnel Flows into and out of AFSC <i>j</i> Positions	7
Figure 2: Kaplan-Meier Retention Probability by Year of Service	15
Figure 3: Precision Recall, Precision-Recall Gain Curves for all Five Years Pooled	19
Figure 4: Predicted Retention Probabilities Separated by Observed Outcome	20
Figure 5: Forecast Error Among the 20 Largest Duty AFSCs	22
Figure A-1: Transition Matrix Selections	A-2

This page is intentionally blank.

Appendix E. Abbreviations

AF	Air Force
AFSC	Air Force Specialty Code
ANG	Air National Guard
AUC	Area Under the Curve
BMT	Basic Military Training
CFM	Career Field Manager
CPS	Current Population Survey
DEERS	Defense Enrollment Eligibility Reporting System
DMDC	Defense Manpower Data Center
ECI	Extension Course Institute
FIFE	Finite Interval Forecasting Engine
FN	False Negatives
FP	False Positives
IDA	Institute for Defense Analyses
ML	Machine Learning
NGB-A1DU	National Guard Bureau/Training Resources and Programming Branch
RPM	Retention Prediction Model
RPM-ANG	Retention Prediction Model for the Air National Guard
TN	True Negatives
TP	True Positives

This page is intentionally blank.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YY) XX-03-2022		2. REPORT TYPE Final		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE <i>Forecasting Demand for Air National Guard Enlisted Initial and Technical Schooling</i>				5a. CONTRACT NO.	
				5b. GRANT NO.	
				5c. PROGRAM ELEMENT NO(S).	
6. AUTHOR(S) Akshay Jain John W. Dennis Julie Lockwood Minerva Song Nathaniel Latshaw Erin Eifert Joseph King				5d. PROJECT NO.	
				5e. TASK NO. DC-6-4720	
				5f. WORK UNIT NO.	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Defense Analyses 4850 Mark Center Drive Alexandria, VA 22311-1882				8. PERFORMING ORGANIZATION REPORT NO. IDA Paper P-32920 Log: H 21-000488	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air National Guard Readiness Center Joint Base Andrews NAF, MD 20762				10. SPONSOR'S / MONITOR'S ACRONYM(S) ANG	
				11. SPONSOR'S / MONITOR'S REPORT NO(S).	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Expediently moving new Air National Guard (ANG) recruits through the initial training pipeline is critical to generating readiness and to the optimal use of personnel resources. In executing its training mission, the ANG faces challenges in accurately forecasting demand for basic military training, technical school slots and officer training school. Currently, the ANG collects data through various methods to establish future fiscal year requirements for tech training school seats. The ANG is experiencing significant wait times (more than a year) for critical career fields due to a lack of unit participation in the current data call process or gaps in manning. This uncertainty leads to delays in the deployability of new accessions, degraded readiness, and inefficient resource expenditure while new members await training. The ANG desires to build upon a more effective toolkit for predicting demand for scarce training resources to improve its readiness. The Institute for Defense Analyses (IDA) uses the Retention Prediction Model, a machine learning capability that IDA developed, to produce these forecasts.					
15. SUBJECT TERMS Air National Guard (ANG), training, Career Field Managers (CFMs), Basic Military Training (BMT), deployability, new accessions, degraded readiness, toolkit, forecasting, Tech School slots, Officer Training School (OTS)					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT U	18. NO. OF PAGES 48	19a. NAME OF RESPONSIBLE PERSON Mr. Joseph P. Torres
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include Area Code) (240) 612-9553

This page is intentionally blank.