NPS-AM-26-291

ACQUISITION RESEARCH PROGRAM SPONSORED REPORT SERIES

Aviation MOSterpiece: Markov-Based Forecasting of Aviation MOS Inventories During Platform Divestment

March 2026

Maj Victoria L. Cannon, USMC

Thesis Advisors: Dr. Chad W. Seagren, Senior Lecturer

Department of Acquisition, Finance and Manpower

Naval Postgraduate School

The research presented in this report was supported by the Acquisition Research Program of the Department of Acquisition, Finance and Manpower Naval Postgraduate School.

ABSTRACT

Platform divestment introduces complex personnel transitions that challenge traditional manpower forecasting models developed for steady-state environments. This thesis examines how enlisted aviation manpower inventories evolve during the sundown of the Marine Corps F/A-18 community. Using personnel data from the Marine Corps Total Force System covering Fiscal Year (FY) 2009–FY2025, this study develops and validates Markov models to forecast transitions across five F/A-18 maintenance military occupational specialties (MOSs). The analysis evaluates two model structures: an annual fiscal year–to–fiscal year model and a six-month transition model, each implemented with MOS-level and MOS-by-rank state definitions. The analysis uses rolling validation windows to estimate transition matrices and assess predictive accuracy using Weighted Absolute Percentage Error (WAPE) and aggregate bias metrics. Results show that a six-month transition model incorporating rank structure and using a longer historical training window provided the best predictive performance. Forward projections through FY34 indicate that personnel inventories decline more slowly than institutional planning targets assume during the F/A-18 sundown, producing temporary overages in most MOS and grade combinations while persistent shortfalls emerge in the E-5 cohort. These findings demonstrate that Markov transition modeling can improve visibility into manpower inventory dynamics during platform divestment.

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

First and foremost, I would like to thank my husband Nico. This thesis began while I was pregnant with our daughter and continued through her arrival five weeks early. Throughout the process he was my emotional rock, whether that meant reassuring me that the thesis would eventually make sense, holding our newborn so I could finish writing, or distracting our two year old son long enough for me to get a few uninterrupted paragraphs on paper. This thesis would not exist without his patience, support, and ability to manage chaos while I tried to finish "just one more thing."

I would also like to thank my advisor, Chad, for his steady encouragement and guidance as I navigated the thesis process, often somewhat clumsily. His willingness to answer questions and gently steer me back on course was invaluable. I am equally grateful to Jake, who first introduced me to the building blocks of Marine Corps manpower and helped me understand the foundations that ultimately shaped this research.

A special thank you goes to my neighbor Larry, who generously allowed me to verbally dump half formed ideas in his office while I tried to work through this project. In the process, he also provided invaluable insight into M&RA processes and frequently pointed me toward the right people to talk to when I inevitably realized I did not know something.

Finally, I would like to thank the lunch table group. Their daily humor and quips provided much needed comic relief and a safe space throughout my time at NPS and especially during the thesis writing process.

THIS PAGE INTENTIONALLY LEFT BLANK

NPS-AM-26-291

ACQUISITION RESEARCH PROGRAM
SPONSORED REPORT SERIES

Aviation MOSterpiece: Markov-Based Forecasting of Aviation MOS Inventories During Platform Divestment

March 2026

Maj Victoria L. Cannon, USMC

Thesis Advisors: Dr. Chad W. Seagren, Senior Lecturer

Department of Acquisition, Finance and Manpower

Naval Postgraduate School

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

LIST OF FIGURES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

ASR	Authorized Strength Report
CDI	Collateral Duty Inspector
CDQAR	Collateral Duty Quality Assurance Representative
EAS	End of Active Service
F/A-18	McDonnell Douglas Hornet
FY	Fiscal Year
GAR	Grade Adjusted Recapitulation
H1/H2	Half Year Time Steps (Oct-Mar / Apr-Sep)
ID	Identification Number
M&RA	Manpower and Reserve Affairs
MCTFS	Marine Corps Total Force System
MMEA	Manpower Management Enlisted Assignments
MMIB	Manpower Management Integration Branch
MOS	Military Occupational Specialty
MOSR	Military Occupational Specialty and Rank
MPP	Manpower Plan, Programs, and Budget Branch
NATTC	Naval Air Technical Training Center
NCO	Non-Commissioned Officer
NPS	Naval Postgraduate School
P2T2	Patient, Prisoner, Transient, or Trainee
PII	Personally Identifiable Information
RUC	Reporting Unit Code
SMCR	Selected Marine Corps Reserve
TFDW	Total Force Data Warehouse
UAV	Unmanned Aerial Vehicle
USMC	United States Marine Corps
VMFA	Marine Fighter Attack Squadron
WAPE	Weighted Absolute Percentage Error

THIS PAGE INTENTIONALLY LEFT BLANK

AI DISCLOSURE

With the permission of my advisor, I used generative artificial intelligence (AI) tools to support the preparation of this thesis. Specifically, I used OpenAI's ChatGPT to assist with spelling, grammar, punctuation, tense consistency, clarity, and flow. I also used the tool to assist with writing, troubleshooting, and debugging portions of the R code used for data preparation and model implementation.

I used the tool by inputting text or code that I had written and asking for recommendations to improve readability, grammatical accuracy, or code functionality. For coding, the tool was used to help identify errors and improve code structure. AI was not used to generate research findings or interpret the results presented in this thesis.

Since generative AI tools may introduce inaccuracies in meaning or code, I reviewed all suggested edits and code recommendations before incorporating them. I verified all code produced the intended outputs and that all written revisions preserved the original meaning of the text. I remain fully responsible for the content, analysis, code, and conclusions presented in this thesis.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION AND SIGNIFICANCE OF STUDY

Military platform divestments represent a recurring feature of force restructuring, which the Marine Corps undertakes to rebalance resources, modernize capabilities, or align the force with evolving strategic priorities. Planners often evaluate force changes through the lens of platforms and budgets. However, these decisions also initiate complex personnel transitions that unfold over multiple years. The effects on enlisted manpower are particularly complex and far-reaching. Enlisted Marines assigned to sundowning squadrons do not exit the force uniformly; instead, they navigate a variety of pathways, including lateral transfer to other Military Occupational Specialties (MOSs) or separation from the service. These transitions occur across every rank and over staggered timelines, creating a dynamic manpower environment that deviates significantly from standard steady-state assumptions. Together, these interdependent aviation specialties form an *Aviation MOSterpiece* of technical expertise and personnel dynamics that sustain Marine aviation capability.

Unlike steady-state environments, platform divestments introduce non-stationary transition behavior as Marines respond to shifting incentives and training availability. Transition rates vary across time, grade, and MOS, particularly as unit deactivations approach. These dynamics create challenges for manpower planners, who must forecast inventories during periods of structural change.

I examine how transition modeling can be used to forecast enlisted aviation manpower during platform divestment, specifically analyzing the sundowning of F/A-18 squadrons. By leveraging historical personnel data and Markov chain models, this research seeks to provide a more responsive and analytically grounded approach to manpower forecasting in non-steady-state environments.

A. PROBLEM

The Marine Corps must simultaneously execute platform divestments while maintaining enlisted end strength across remaining MOSs and grades. The sundowning of

F/A-18 squadrons introduces reductions in manpower that ripple throughout the Marine Corps as Marines transition into other communities or exit the service. The transitions affect not only total inventory but also grade distribution and accession requirements.

This thesis examines how to more accurately forecast enlisted manpower transitions during the sundowning of a legacy platform so that I can better align projected personnel inventories with institutional manpower requirements. Researchers developed existing approaches for steady-state environments, and it remains unclear whether they maintain predictive accuracy during periods of force reduction. Without this evaluation, the Marine Corps risks basing manpower planning decisions on forecasts that do not fully capture transition behavior during periods of restructuring.

B. IMPORTANCE OF THE STUDY

This research holds both operational and institutional importance. More accurate manpower forecasting during platform divestment has implications for day-to-day readiness as well as long-term manpower policy decisions.

1. Operational Importance

Operational readiness within the Marine Corps depends on maintaining the correct mix of enlisted Marines by MOS and grade. As F/A-18 squadrons divest, the Marine Corps must absorb transitioning personnel without exceeding training capacity or creating grade imbalances. Inaccurate forecasts risk understaffing MOSs or producing surpluses that strain receiving communities.

Surpluses at specific grades can be as problematic as shortages, particularly when training pipelines or promotion timing cannot accommodate excess personnel. More accurate forecasting provides planners with earlier visibility into emerging imbalances, enabling more deliberate retention planning and the exploration of lateral transfer opportunities. Improved forecast accuracy supports decisions that preserve readiness and expand assignment options for planners.

2. Institutional and Policy Importance

At the institutional level, this research supports Marine Corps manpower initiatives focused on data driven decision making and talent management. As the aviation community transitions away from legacy platforms, planners must balance end strength constraints with evolving requirements across the force. Forecasting tools that highlight how adjustments occur provide critical insight into the policy levers decision makers can use.

The analysis also addresses the broader challenge of aligning forecasted inventories with strength management mechanisms, such as the Grade Adjusted Recapitulation (GAR). While the GAR establishes bi-annual constraints on inventory by grade and specialty, it does not prescribe how planners should achieve those constraints. By identifying periods of divergence between projected inventories and planning targets, this analysis highlights where and when planners may require redistribution strategies.

3. Aviation Maintenance Training Pipeline

Producing a qualified maintainer requires a significant investment in training, time, and institutional resources. Marines who enter aircraft maintenance MOSs complete boot camp and Marine combat training before attending their MOS specific instruction at Naval Air Technical Training Center (NATTC) schools. Depending on the MOS, the initial training pipeline can take several months to more than a year before a Marine reports to an operational squadron.

After arriving at a squadron, maintainers continue developing technical proficiency through structured on-the-job training and qualification programs tied to the training and readiness program, (United States Marine Corps [USMC], 2021). As Marines gain experience, they assume increasing levels of responsibility within the maintenance organization. Experienced maintainers often serve as work center supervisors and qualify for critical quality assurance roles such as Collateral Duty Inspector (CDI) and, at more advanced levels, Collateral Duty Quality Assurance Representative (CQAR). These positions require additional training, certifications, and follow-on schooling to ensure that Marines can inspect and certify aircraft maintenance actions in accordance with naval aviation safety standards (Department of the Navy, 2022).

This progression represents a substantial institutional investment. Developing an experienced maintainer requires years of training and operational experience before a Marine reaches billets responsible for supervising maintenance or certifying aircraft safety. As a result, aviation maintenance personnel represent highly specialized and costly human capital. If manpower forecasts produce excess inventory during a platform divestment, the Marine Corps risks losing trained maintainers through early separation or unnecessary lateral transfers before the service fully realizes the return on that investment. More accurate forecasting therefore helps manpower planners preserve trained expertise while managing inventory reductions associated with platform transitions.

C. RESEARCH OBJECTIVES

The primary objective of this study is to develop and validate a forecasting model that captures enlisted aviation manpower transitions during the sundowning of F/A-18 squadrons. To accomplish this, the study constructs Markov models using historical enlisted data to represent movement across MOS, grades, and separation outcomes. It further evaluates whether shorter time-step models better capture manpower dynamics during platform divestment more effectively than traditional fiscal year approaches. The analysis assesses model performance using validation metrics, with particular attention to periods of divergence between forecasted inventories and strength targets.

Validation results show that the six-month MOS model using a twenty-six-span training window provides the strongest predictive performance, achieving approximately 4.5 percent WAPE with minimal aggregate bias. Forward projections through FY34 indicate that enlisted inventories decline more slowly than institutional planning targets assume during the F/A-18 divestment period, producing temporary overages across most grades. However, the projections consistently show persistent shortfalls in the E-5 cohort, highlighting a structural gap in the mid-grade inventory that manpower planners may need to address through targeted retention, promotion timing, or redistribution strategies.

D. MANPOWER MANAGEMENT TOOLS

When a platform is divested, Marine Corps manpower planners employ several mechanisms to manage the associated personnel transitions. One approach is to adjust reenlistment opportunities within affected MOS communities. By limiting reenlistment approvals or adjusting retention incentives, planners can gradually reduce inventory as Marines reach the end of their contracts.

Another common tool is lateral transfer, which allows Marines to retrain into MOSs experiencing manpower shortages elsewhere in the force. Lateral moves preserve experienced Marines within the service while redistributing skills to areas of higher demand. In addition, planners may temporarily assign Marines to B-billets, training pipelines, or institutional assignments while the aviation force structure adjusts.

Each of these mechanisms allows the Marine Corps to manage personnel reductions over time rather than through immediate separations. However, effective use of these tools depends on accurate forecasts of how inventories will evolve during platform divestment. Without reliable projections, planners risk either losing valuable trained personnel prematurely or retaining excess inventory that cannot be absorbed elsewhere in the force.

E. SCOPE, CONSTRAINTS, AND ANALYTICAL BOUNDARIES

This study focuses on enlisted aviation Marines within the F/A-18 MOS community to isolate the manpower effects of a single, clearly defined platform divestment. I conduct the analysis at an aggregate level; the objective is to understand overall manpower flows and inventory behavior during a transition period, rather than to optimize individual career outcomes. The analysis leverages historical personnel data to estimate transition behavior, assuming that observed patterns remain informative of near-term future outcomes. The Markov model developed here focuses on forecasting manpower flows and inventory trajectories to support high level planning and policy decisions, rather than predicting specific individual outcomes.

F. ORGANIZATION OF THE THESIS

I organize this thesis as follows. Chapter II provides institutional context for enlisted aviation manpower management. Chapter III reviews relevant literature on military manpower planning, Markov modeling, and prior studies addressing force restructuring. Chapter IV describes the data sources, modeling framework, transition definitions, and validation methodology employed in this research. Chapter V presents model validation results and forward projections used to assess forecasting performance during the F/A-18 sundown period. Finally, Chapter VI discusses the implications of these results for enlisted manpower planning and policy and presents conclusions and recommendations for both practitioners and future research.

II. INSTITUTIONAL AND MANPOWER CONTEXT

This chapter provides institutional context for enlisted aviation manpower management and the mechanisms planners use to manage manpower transitions during platform divestment.

A. MARINE AVIATION FORCE RESTRUCTURING

Marine Aviation has undergone substantial restructuring as part of broader force modernization efforts (USMC, 2020). Shifts in national defense priorities drive these changes and the need to reallocate finite resources toward capabilities more relevant to future operating concepts. This restructuring affects not only aircraft inventories and squadron alignments but also the enlisted manpower communities responsible for maintaining, operating, and supporting these platforms.

Marine Aviation restructuring differs from routine force adjustments by deliberately reducing legacy platforms rather than incrementally changing force structure. As the Marine Corps divests aircraft and deactivates squadrons, the associated enlisted manpower does not disappear simultaneously. Instead, Marines transition through a range of pathways, including reassignment to other aviation platforms, lateral transfer to non-aviation specialties, or separation from service. These dynamics introduce complexity into manpower planning that extends well beyond traditional steady-state assumptions.

B. F/A-18 PLATFORM DIVESTMENT

As part of Marine aviation restructuring, the Marine Corps is divesting the F/A-18 platform through phased deactivation or transition of multiple squadrons. Beginning in 2012 with the transition of VMFA-121, the F/A-18 community has steadily shrunk as additional squadrons transitioned to new platforms or deactivated entirely, including: VMFA-251 (2017), VMFA-314 (2019), VMFA-312 (2020), VMFA-242 (2020), VMFA-225 (2021), VMFA-533 (2023), VMFA-115 (2023), VMFAT-101 (2023), and VMFA-224 (2025). The two remaining squadrons, VMFA-232 and VMFA-323, will deactivate in FY 2028. The Naval Aviation Plan (USMC, 2026) documents the retirement of legacy

F/A-18 aircraft and the transition to the F-35 platform. Prior to the start of these transitions, the F/A-18 maintenance community included approximately 2,884 enlisted Marines across the associated MOSs in 2010. As squadrons divest, the requirement for F/A-18 specific maintenance personnel declines accordingly.

The enlisted Marines associated with F/A-18 maintenance do not exit the force in a uniform manner. Instead, divestment unfolds over several fiscal years, with manpower effects occurring as squadrons sundown and maintenance MOS accessions decrease. This staggered timeline complicates forecasting because transition behavior shifts as divestment milestones approach.

C. ENLISTED AVIATION MANPOWER MANAGEMENT

The Marine Corps organizes enlisted aviation manpower management around MOSs that define technical skill sets, training pipelines, and career progression pathways (USMC, 2018). Because many aviation MOS communities are tied to specific aircraft platforms, they are particularly sensitive to platform-level force structure decisions. When the Marine Corps divests a platform, the requirement for the associated maintenance specialties declines, creating immediate inventory pressure within those MOS communities.

Planners must determine how to redistribute Marines from divesting squadrons while maintaining appropriate grade structure, promotion flow, and training pipeline capacity across the aviation enterprise. Planners cannot simply remove Marines whose MOS demand declines. Instead, they reassign Marines within the community, retrain them for other specialties, move them into institutional billets, or manage them through normal separation timelines. As a result, platform divestment creates a period in which manpower supply and billet demand may diverge, requiring deliberate inventory management.

D. MANPOWER REQUIREMENTS AND INVENTORY MANAGEMENT

Within enlisted manpower management, the Marine Corps categorizes billets as A-billets or B-billets based on their alignment with a Marine's primary MOS (USMC, 2025). A-billets correspond to requirements directly associated with a Marine's primary occupational specialty. B-billets represent requirements outside a Marine's primary MOS and typically include assignments such as recruiting duty, drill instructor duty, or other special assignments that support institutional needs.

During platform divestment, the distribution of Marines across A- and B-billets often shifts. As the Marine Corps eliminates operational A-billets tied to a divesting platform, some Marines transition into B-billet assignments while they await retraining, reassignment, or separation. These shifts can obscure the true availability of personnel within a given MOS and complicate how planners interpret manpower inventories during the transition period.

The Marine Corps releases the Authorized Strength Report (ASR) bi-annually to establish congressionally approved end strength levels by grade and specialty and provide the baseline for inventory management. The ASR reflects institutional decisions regarding force size and composition and serves as the primary reference point for manpower planning across the Marine Corps. While the ASR defines target inventory levels, it does not specify how adjustments should occur over time. During platform divestment, inventories often temporarily diverge from ASR targets as transitions unfold, challenging planners who must reconcile observed manpower flows with authorized levels.

The Grade Adjusted Recapitulation (GAR) translates ASR targets into strength management controls by grade and specialty. The GAR adjusts inventory levels to account for institutional requirements such as A-billets and B-billets and serves as the primary tool for managing enlisted inventory relative to force structure requirements. However, because these adjustments occur on a bi-annual cycle, they may be misaligned with the timing of manpower transitions during periods of platform divestment.

While the ASR and GAR establish the billet requirements the Marine Corps intends to maintain, they do not describe how personnel inventories actually evolve over time. In practice, manpower planners must continuously reconcile the difference between required billets and the available inventory of Marines qualified to fill them. Forecasting tools therefore play an important role in identifying when inventories are likely to diverge from GAR targets and where adjustments to accessions, assignments, or retention policies may be required.

In addition to billet-based planning mechanisms such as the ASR and GAR, Marine Corps manpower planners rely on inventory forecasting to inform accession decisions. Organizations within Manpower and Reserve Affairs (M&RA), including Manpower Plans, Programs, and Budget Branch (MPP)-20 and Manpower Management Enlisted Assignments (MMEA), play key roles in this process. MMEA maintains visibility over current personnel inventories, assignment timelines, and anticipated manpower movements within MOS communities. These inputs inform broader manpower planning efforts that MPP-20 conducts.

MPP-20 uses inventory projections and force structure requirements to estimate future manpower needs and determine accession targets. This process commonly uses linear forecasting models that incorporate expected separations, promotions, and inventory levels alongside billet requirements the GAR establishes. Based on these inputs, the model estimates the number of new accessions required to maintain desired inventory levels across MOS communities.

While these models provide a structured method for aligning accessions with force structure requirements, they are typically calibrated using steady-state assumptions about personnel flows. During periods of structural change, such as platform divestment, historical transition behavior may shift as Marines respond to changing assignment opportunities and career incentives. As a result, forecasting methods that explicitly model transition behavior may provide additional insight into how inventories evolve during these periods (Bartholomew, Forbes, & McClean, 1991).

The Markov modeling framework developed in this study complements existing planning tools by estimating how personnel transitions occur across MOS communities over time. By capturing promotion, lateral transfer, and separation behavior directly from historical personnel data, the model provides an alternative method for projecting inventories during the F/A-18 divestment period.

With the operational context established, the next chapter reviews the academic and professional literature relevant to manpower forecasting. It examines how researchers have applied Markov models to workforce and military manpower problems and identifies the methodological gaps that motivate the modeling approach used in this study.

THIS PAGE INTENTIONALLY LEFT BLANK

III. LITERATURE REVIEW AND CONCEPTUAL FRAMEWORK

This chapter reviews research that applies Markov models to military manpower problems, including forecasting inventories, modeling promotion and attrition behavior, and assessing how personnel communities evolve over time. This analysis focuses on how these tools have been applied to former and current manpower issues, what those studies found, and what gaps remain.

A. MARKOV MODELS AND MANPOWER

Researchers have applied Markov models to workforce planning problems across a wide range of non-military settings. McClean (1991) explains how manpower planning models use Markov and semi-Markov processes to forecast personnel flows such as promotion, recruitment, transfers, and attrition in structured employment systems. Because Markov models represent movement between discrete career states over time, they allow analysts to estimate how workforces evolve based on empirically observed transition probabilities. These models therefore provide organizations with a systematic method for forecasting workforce composition and evaluating the long-term effects of staffing policies.

Scholars have applied this approach to several professional environments characterized by hierarchical career structures. Stewman (1978) used Markov and renewal models to analyze personnel dynamics within a state police system and demonstrated that transition-based models can represent staffing behavior within public-sector organizations. Dimitriou, Georgiou, and Tsantas (2013) extended this approach by modeling manpower movement within and between departments, incorporating recruitment, transfers, and attrition into a multivariate Markov framework. Across these applications, researchers consistently show that Markov models provide a structured and transparent method for forecasting workforce dynamics when individuals transition between defined career states over time.

These characteristics make Markov models particularly well suited for analyzing military manpower systems, where personnel move through structured career pathways and organizations must balance inventories across multiple ranks and occupational specialties. Researchers have applied Markov models to several military manpower problems, including forecasting personnel inventories, modeling promotion and attrition behavior, and evaluating the long-term sustainability of personnel communities. The following studies illustrate how these models examine manpower dynamics within military organizations and provide the methodological foundation for the analysis conducted in this thesis.

Taylor's (2020) study of the Marine Corps cyber officer community exemplifies how Markov modeling can support a nascent and rapidly evolving manpower structure. Because the cyber community was in its early developmental stages in 2020, historical data was scarce. To address this, Taylor (2020) used existing cyber officer data to establish starting inventories and frame plausible scenarios. He then utilized the unmanned aerial vehicle (UAV) officer community as a proxy for estimating promotion, retention, and attrition behavior, using its longer personnel history to inform the transition probabilities in the cyber models. His model projected the accessions, promotions, and lateral moves required to grow the cyber community into a sustainable force.

Taylor's (2020) forecasts, even with partial observations, show which year groups were at risk of collapsing and which ranks needed improved accession pipelines. His model demonstrates that without targeted increases in accessions, the cyber officer inventory would fall short of manpower objectives for senior Captains and Majors. In several scenarios he shows that the community would experience manpower gaps of 10 to 20 percent within a few years unless there was immediate intervention. His findings reinforce that early investment in new communities is essential because gaps take years to correct.

Taylor (2020) further validates his model through *backcasting*, which compares the rank distributions generated by his model against observed historical distributions, showing that the model closely reproduces actual community patterns. This comparison provides evidence that the model reproduces the promotion and retention behavior of the population, even in a data-limited environment. Taylor's (2020) results are based on several key assumptions, including stationary transition probabilities, a fixed inventory structure, and the use of a proxy community to approximate missing transitions. His findings demonstrate that Markov modeling can still produce credible results for manpower analysis when data is incomplete.

Rai's (2019) thesis on the Navy Nurse Corps demonstrates how Markov chains support manpower planning, particularly for communities where readiness depends on balancing subspecialties. Using historical personnel data, Rai (2019) estimated transition probabilities across both ranks and subspecialties. He then applied fixed-inventory and steady-state Markov models to forecast future personnel distributions under peacetime and wartime conditions.

One of Rai's (2019) major findings was that the Nurse Corps faced significant risks in its wartime surge specialties. His results indicated that the Navy would fall short of required wartime staffing levels for trauma and critical care nursing by approximately 14 to 30 percent. Rai (2019) subsequently adjusted continuation rates and target training throughputs to demonstrate how these shortfalls could be mitigated, providing clear courses of action for Navy manpower planners. Using multiple *backcasts* and Sales' (1971) goodness-of-fit method, he compared model-generated inventories to historically observed end strength data, showing that the model reproduced past Nurse Corps inventories within roughly 1 to 2 percent at the aggregate level.

Erhardt (2012) contributes a critical perspective to the literature by highlighting the consequences of failed Markov assumptions. He attempted to build a Markov model for forecasting continuation rates for both Prior Service and Non-Prior Service enlisted Marines. His key finding was that the Selected Marine Corps Reserve (SMCR) continuation process is not stationary. Erhardt (2012) showed, through validation, that transition rates calculated using annual aggregate monthly data failed the Markov assumption of homogeneity. His testing revealed significant seasonality in attrition behavior; the probability of a Marine staying or leaving depended heavily on the time of year. Consequently, his model produced inaccurate forecasts, especially for short-term projections. Erhardt (2012) recommended shifting to monthly, season-specific transition matrices rather than a single annual matrix. His study reinforces the necessity of testing underlying transition patterns for stability before applying a Markov model.

Licari's (2013) thesis further examines the impact of time-step selection and data sufficiency by developing models to predict end strength for SMCR officers. Using Total Force Data Warehouse personnel data, Licari constructed monthly transition matrices for both Prior Service and Non-Prior Service officer populations, comparing models based on aggregate monthly transition rates to those using unique monthly rates. His validation results showed that aggregate monthly models slightly outperformed unique-rate models in predicting end strength for Prior Service officers. However, Licari was unable to validate models for Non-Prior Service officers due to limited observations and high variance in transition estimates. His findings highlight the importance of sufficient sample sizes, appropriate time-step selection, and empirical validation.

This review of prior research reveals several themes regarding how Markov models should be constructed, validated, and applied to military manpower. All of the above authors use Markov chains to forecast inventories or continuation behavior, but each confronts a different structural challenge. Taylor (2020) deals with a new, rapidly developing community with limited historical data. Rai (2019) models a complex multi skill type community where certain subspecialties dominate operational readiness. Erhardt (2012) evaluates a reserve population with distinct continuation patterns. Additionally, Licari (2013) examines how time-step selection and data sufficiency affect model performance when forecasting reserve officer end strength.

Taylor (2020) and Rai (2019) both produce validated Markov models that closely approximate historical behavior when *backcasted*, with Taylor (2020) identifying accession-driven bottlenecks and Rai (2019) highlighting 14–30 percent wartime shortfall. In contrast, Erhardt's (2012) model did not hold for the SMCR. His discovery demonstrates a critical point often overlooked in manpower forecasting: the Markov framework only works if transitions are stable. Where Taylor (2020) and Rai (2019) focus on structural changes or subspecialty demand, Erhardt (2012) reveals that time-varying behavior can undermine a model's accuracy. Licari's (2013) findings reinforce this point from a different perspective, showing that even when monthly transition behavior is relatively stable, insufficient observations per state can undermine model validity, particularly for smaller populations where transition rates exhibit high variance.

From a methodological standpoint, this implies that any manpower environment with cyclic or seasonal patterns may require time-specific transition matrices rather than annual averages. For this thesis, Erhardt's (2012) and Licari's (2013) results serve as cautionary precedents. Before building the main forecasting model, it is necessary to validate whether F/A-18 MOS transition rates are stable and whether sufficient observations exist at the chosen level of disaggregation, or if shorter time intervals are required to accurately capture observed behavior.

Even though Markov modeling is widely used in manpower research, the literature reveals several clear gaps that are directly relevant to my study. First, most Markov theses focus on communities growing, maturing, or maintaining steady-state strength, rather than communities that are downsizing or deactivating. There is a lack of research on billet removal or planned redistribution, which are both core aspects of the F/A-18s sundown. Second, while Erhardt (2012) and Licari (2013) explicitly test the limits of Markov assumptions through non-stationarity and data sufficiency, many studies assume stable annual transitions without examining whether those assumptions hold. Taylor (2020) and Rai (2019) both assume stable annual transitions, which works for their communities. Aircraft maintenance communities may behave more like the SMCR than like the steady-state medical or cyber personnel. Third and finally, most existing studies focus on stable or growing communities and provide limited insight into manpower behavior during periods of deliberate downsizing or platform divestment. These gaps form the core methodological plan where my thesis may contribute.

B. CONCEPTUAL FRAMEWORK

Markov models are commonly used to represent systems in which individuals transition between a finite set of discrete states over time, where future states depend only on the current state and not on the full history of prior transitions. In a manpower context, these states typically represent rank, MOS, or separation outcomes, while transitions reflect promotion, lateral movement, or attrition. As described by Ross (2014), discrete-time Markov chains are particularly well-suited for population systems where individual behavior is uncertain but aggregate transition patterns are relatively stable.

Bartholomew, Forbes, and McClean (1991) further formalize the use of Markov models for human populations, distinguishing between fixed-inventory models, which forecast population changes over a finite horizon, and steady-state models, which examine long run equilibrium behavior. In manpower applications, fixed-inventory formulations are more appropriate when end strength, accessions, or force structure are constrained. These foundational works provide the theoretical basis for applying Markov modeling to USMC personnel systems and frame the assumptions that must be tested when forecasting manpower transitions.

C. CONTRIBUTION

Prior research demonstrates that Markov models are effective tools for forecasting military manpower inventories when transition behavior is carefully tested and validated. Studies by Taylor (2020), Rai (2019), Erhardt (2012), and Licari (2013) show that Markov modeling can identify structural weaknesses, readiness risks, and assumption failures across a range of military communities. However, existing work largely examines stable or growing populations, providing limited guidance for communities undergoing deliberate platform divestment.

This chapter examined how researchers have applied Markov models to military manpower problems. Prior studies demonstrated how transition-based models forecast personnel inventories, identify structural risks, and evaluate manpower policy decisions across several communities. At the same time, the literature revealed a consistent limitation: most existing work focuses on stable or growing personnel populations rather than communities undergoing deliberate downsizing or platform divestment.

This thesis addresses that gap by applying a Markov framework to an aviation community experiencing planned reductions associated with the sundown of the F/A-18 platform. Chapter IV translates this conceptual framework into a practical modeling approach. It describes the dataset, data preparation procedures, model structure, and validation methodology used to estimate transition probabilities and evaluate forecasting performance.

IV. DATA, MODEL IMPLEMENTATION, METHODOLOGY, AND VALIDATION

This study used data from the Marine Corps Total Force Data Warehouse (TFDW) via the Department of Defense Advana analytics platform. Data extraction and processing occurred within Advana's Databricks environment using the `jup-nps-manpower`cluster. The primary source table, 'mctfs_restricted.mil_mctfs_tfdq_marine1' contains individual-level personnel records derived from the Marine Corps Total Force System (MCTFS). This table includes longitudinal snapshot data capturing Marines' rank, primary MOS, billet MOS, and Reporting Unit Code (RUC). I filtered the dataset to relevant months and primary MOSs to construct the career histories and transition states for the analysis.

A. DATA

The dataset contains monthly snapshots of all enlisted F/A-18 manpower on active duty from the end of Fiscal Year (FY) 2009 through the end of FY 2025. The final observation set includes 6,228 unique Marines. I used R Studio to calculate historical transition rates and then exported them for analysis. The key variables utilized in the dataset are defined below:

1. Marine Identification (ID)

Before exporting the data from Databricks, the process assigned each individual a unique Marine ID. This allowed for the accurate tracking of the pathways for each Marine throughout their career, while maintaining Personally Identifiable Information (PII) protections through removal of social security numbers.

2. Sequence Number

The dataset assigns each month a chronological sequence number, starting with 247 representing September 2009 and continuing until 439 representing September 2025.

3. Present RUC

A four-digit numeric code representing a Marine's assigned unit during the corresponding sequence number.

4. Present Grade Code

The ranks range from E1 to E8, with E1 and E2 grouped together to align with the GAR reporting format.

5. Primary MOS Code

Per NAVMC 1200.1E Military Occupational Specialties Manual (USMC, 2025), this thesis uses five F/A-18 specific MOSs:

- MOS 6217, Fixed-Wing Aircraft Mechanic, F/A-18
- MOS 6257, Fixed-Wing Aircraft Airframe Mechanic, F/A-18
- MOS 6287, Fixed-Wing Aircraft Safety Equipment Mechanic, F/A-18
- MOS 6317, Aircraft Communications/Navigation/Radar Systems Technician, F/A-18
- MOS 6337, Aircraft Electrical Systems Technician, F/A-18

6. Billet MOS Code

The billet MOS code identifies the MOS associated with a Marine's assigned billet at the time of observation. When the billet MOS matches the primary MOS, the Marine is serving in an A-billet; when the billet MOS differs from the primary MOS, the Marine is classified as serving in a B-billet or in a P2T2 status.

B. DATA CLEANING AND PREPARATION

I imported the raw F/A-18 personnel dataset into R Studio to serve as the foundational input for modeling. The cleaning process removed records with missing or invalid RUCs to ensure data integrity; specifically, the process excluded observations where the present RUC was coded as zero or null, retaining only valid unit assignments.

The analysis then restricted the dataset to Marines who entered the Marine Corps as one of the target MOSs: 6217, 6257, 6287, 6317, or 6337. Identifying the first observation for each Marine in a target MOS allowed the model to remove all prior records. This restriction ensures that each Marine's observed career history begins at entry into the target MOS community and prevents pre-community service from skewing estimated transition probabilities. The analysis excluded Marines who did not enter a target MOS from the dataset.

The analysis treated lateral transfers out of the target MOS community systematically. Once a Marine's Primary MOS transitioned away from a target MOS, the analysis coded that transition point as a lateral transfer. The model coded all subsequent observations for that Marine as lateral transfers. This ensured that post-transfer career history did not re-enter the target MOS states, maintaining a consistent representation of permanent exits from the community.

Multiple validation checks occurred throughout the cleaning process. These checks verified that no invalid RUCs remained, that each Marine's first retained observation corresponded to a target MOS, and that lateral transfers were consistently applied and persisted. The process then exported the resulting cleaned dataset to a new Excel file to serve as the input for model estimation, validation, and forecasting.

C. MARKOV MODELS

1. Markov Models in Manpower Analysis

A Markov model is a framework used in this thesis to represent how enlisted Marines in F/A-18 communities transition among a finite set of personnel states over time. In this application, states are defined by MOS or by a combination of MOS and rank. Transitions capture promotion, lateral MOS movement, and attrition. The defining characteristic of a Markov process is the Markovian property, which assumes that the probability of transitioning to a future state depends only on a Marine's current state and not on prior career history. Rather than predicting individual career outcomes, the model estimates aggregate flows between states over a defined time step, allowing future inventories to be forecast from an observed initial population. Figure 1 illustrates the MOS only Markov model structure used to capture occupational flows and attrition behavior, while Figure 2 depicts the expanded MOS-and-Rank (MOSR) structure that incorporates promotion-dependent transitions and career-stage effects.

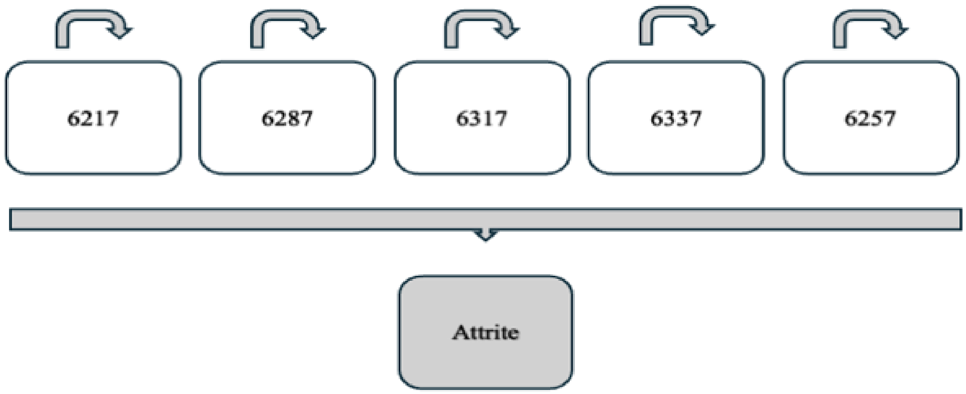

Figure 1. MOS Only Markov State Transition Model

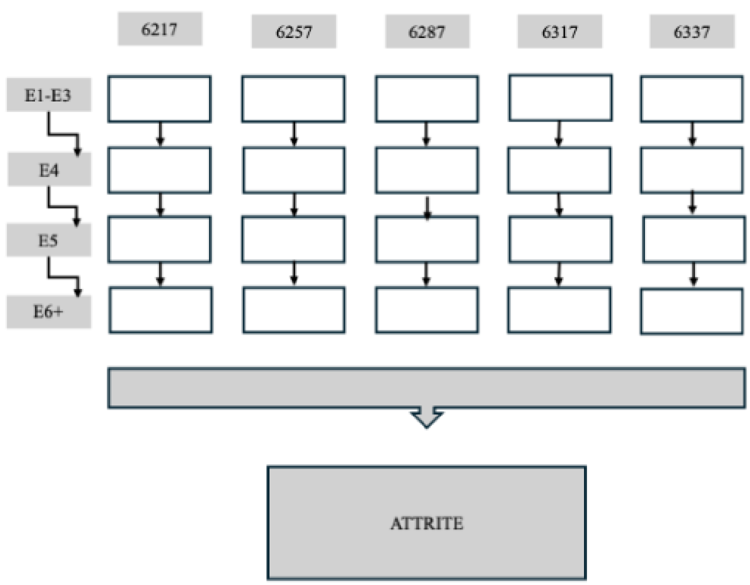

Figure 2. MOS and Rank Markov State Transition Model

2. Model Variants and Time Step Design

This thesis implements two Markov manpower models that differ in state definition and time step. In both cases, the underlying idea is the same: Marines are observed at regular snapshot intervals, each Marine is assigned to a "state" at time t, and historical transitions from t to $t + 1$ are used to estimate a transition probability matrix P. The models focus on the five target enlisted F/A-18 MOSs and treat movement outside that community and separation as explicit modeled outcomes. This structure follows standard Markov manpower modeling practice, where observed historical flows are converted into an empirical transition matrix and then iterated forward to project inventories (e.g., Bartholomew et al., 1991; Ross, 2014).

Two state-space variants appear throughout the analysis: a MOS-only model and an MOSR model. In the MOS-only model, the state at time t is simply the Marine's primary MOS. In the MOSR model, the state augments MOS with an enlisted rank grouping where rank groups are bucketed as E1-E2, E3, E4, E5, E6, and E7. Using the MOSR version makes the state space more detailed, so it can capture how transitions differ by grade. The tradeoff is that the data becomes thinner in some of the MOSR combinations, so to reduce noise, the analysis excludes records with unknown rank groups when constructing MOSR states.

Across both variants, the model assigns each observed transition to one of three "next-state" categories. If a Marine does not appear in the following snapshot, the model codes the transition as End of Active Service (EAS). If the Marine remains in one of the five target MOSs at the next snapshot, the model codes the transition as a within-community transition to the appropriate MOS (MOS model) or MOS|RANK Group (MOSR model). If the Marine appears at the next snapshot but with a primary MOS outside the target set, the model codes the transition as a lateral transfer. In the forecasting matrices, EAS and lateral transfer function as absorbing states, ensuring that once a Marine exits the target community in the observed data structure, the model keeps that Marine outside of it in projected inventories.

Time-step selection is the second major design choice. The FY-to-FY model uses one transition per fiscal year, built from September (end of FY) snapshots. Each Marine contributes a single annual transition from FY t to FY $t + 1$.

The six-month model uses a finer time step by retaining March and September snapshots interpreted as two half-steps per FY: H1 (Oct–Mar) and H2 (Apr–Sep)**.** The model creates a half-year time index $HALF\ ORDER$ to enforce strict ordering.

$$HALF\ ORDER = 2 \cdot FY + \mathbb{1}\,(\text{HALF} = H2),$$

so transitions occur from $HALF\ ORDER = h$ to $h + 1$. This design supports two goals: (1) capturing shorter-term churn masked by annual aggregation, and (2) aligning more naturally with within-year timing effects while still producing FY-level outputs by reporting the projected inventory at each FY's H2 (September) snapshot.

3. Transition Probability Estimation and Forecasting Mechanics

Each model converts observed transitions into empirical transition probabilities. Let N_{ij} be the number of observed transitions from state i to state j in the training set, and let $N_i = \sum_j N_{ij}$. The maximum-likelihood estimate of the one-step transition probability is

$$\hat{p}_{ij} = \frac{N_{ij}}{N_i}, \text{ with } \sum_j \hat{p}_{ij} = 1 \text{ for each origin state } i.$$

These $\hat{p}_{ij}$values populate the transition matrix P, with absorbing rows enforced for EAS and lateral transfers.

The model carries out forecasting using the standard expected-value Markov inventory update. Let v_t denote the row vector of state counts at time t, and let P denote the transition probability matrix estimated from historical data. The expected inventory one step ahead is given by

$$v_{t+1} = v_t P + a_t$$

where $v_t P$ represents transitions among Marines already in the system and a_t represents new accessions entering the model during period t.

In this study, the model incorporates accessions as an external inflow applied to the entry states of the system. Based on information provided by MPP-20 within M&RA, a simple accession rule is applied by adding a fixed increment each fiscal year: +1 to each MOS in the MOS model and +1 to each MOS|E1–E2 state in the MOSR model. In the six-month model, the model applies this increment once per fiscal year when entering H1, ensuring that accessions occur annually while allowing transitions to occur within the fiscal year.

Rather than relying solely on the expected-value update, the model employs simulation to capture uncertainty in future inventories. In this approach, the number of Marines in each origin state is treated as a pool randomly redistributed to next states based on the observed transition probabilities. For a given state i with n_i Marines, the model assigns individuals to next states according to a multinomial draw using the transition probabilities for that state. The simulated counts are then combined across all origin states to produce a simulated inventory for the next time step.

The simulation repeats this process iteratively to generate a distribution of possible inventory outcomes by fiscal year and state. The model reports the median projection and a central prediction interval from this distribution. For example, a 95 percent prediction interval is defined by the 2.5th and 97.5th percentiles of the simulated outcomes.

4. Model Validation

The analysis performs model validation by repeatedly training the model on past data and testing it on the subsequent period, mirroring practical application. For the FY-to-FY model, each window trains on fiscal years t through $t + \text{span} - 1$ and tests on fiscal year $t + \text{span}$, predicting transitions from $t + \text{span}$ to $t + \text{span} + 1$. Here, "span" defines the number of consecutive fiscal years of historical transitions included in the training window. For example, a span of two estimates transition probabilities using two years of data, while a span of five uses five years of transitions. The same logic applies to the half-year model on the $HALF\ ORDER$ axis. By varying the span and the starting point of the training window, the validation procedure generates multiple train/test splits, allowing for direct comparison of how results change with the length of the training window and the chosen time step. For the six-month model, the span corresponds to the number of consecutive half-year transitions included in the training window (e.g., a span of four uses two full years of half-year transitions).

Within each test window, the model performs the validation at the state-to-state transition level. For an origin state i, the test set provides observed counts O_{ij} to each next state j, and the training matrix provides probabilities $\hat{p}_{ij}$. The predicted expected counts are

$$E_{ij} = O_i \, \hat{p}_{ij}, \text{ where } O_i = \sum_j O_{ij}.$$

These E_{ij} values are compared against O_{ij} to produce both overall and diagnostic metrics.

The primary accuracy metric used for ranking model windows is Weighted Absolute Percentage Error (WAPE). The model computes WAPE over all cells in the test transition table:

$$\text{WAPE} = \frac{\sum_{i,j} | O_{ij} - E_{ij} |}{\sum_{i,j} O_{ij}}.$$

WAPE is particularly suitable in this setting because it is scale-normalized, easy to interpret as an overall percent error, and naturally places more weight on high-volume transitions than sparse ones. In addition to WAPE, the analysis reports an aggregate bias to identify whether the model systematically over- or under-predicts total transition volume where negative values indicate underprediction, and positive values indicate overprediction.

Together, these validation outputs support two key decisions: (1) selecting the preferred training span for each state-space variant, and (2) comparing whether the annual FY-to-FY step or the half-year step provides better predictive performance.

Because all windows are scored using identical metrics and out-of-sample structures, the ranked-window tables provide a transparent, reproducible basis for choosing the model structure used in the FY26–FY34 forecasts. This framework establishes a consistent basis for the model selection and evaluation presented in the following chapter.

D. WORKED EXAMPLE OF THE VALIDATION PROCESS

This section illustrates the validation procedure using a specific training and test window from the MOSR six-month model with a training span of 26 half-year transitions.

In this example, the model trains on observed transitions from 2009 H1 through 2021 H2 and then evaluates its predictions using the next observed half-step, 2022 H1 to 2022 H2. Consider the origin state 6217|E3, which represents Marines in MOS 6217 (Fixed-Wing Aircraft Mechanic) at the E-3 rank. Using the transitions observed in the training window, the model estimates transition probabilities for each possible next state. These probabilities form the Markov transition matrix used for prediction.

The test data then provide the number of Marines who began the next half-step in the same origin state. The model multiplies this test-period inventory by the transition probabilities estimated from the training window to generate expected transition counts for each possible next state. Table 1 presents the observed and expected transitions for the selected state during the test window.

For example, the model expected approximately 19.90 Marines to remain in state 6217|E3, while the observed data recorded 25 Marines, producing an absolute error of 5.10. Similarly, the model expected 8.82 Marines to transition to 6217|E4, while the observed data recorded 6 Marines, resulting in an absolute error of 2.82. Smaller expected flows also appear in other states, such as 0.60 Marines expected to transition to 6217|E1–E2 and 1.57 Marines expected to separate through EAS, though no observed transitions occurred in those cells during the test period.

Comparing observed and expected transitions produces the cell-level absolute errors used in the validation process. These errors contribute to the aggregate error metrics calculated for the full transition table, including WAPE. This worked example demonstrates how the validation framework compares predicted and observed personnel flows for each state transition within a given test window.

Table 1. Observed and Expected Transitions for State 6217|E3 During the 2009 H1-2021 H2 Validation Window

Observed and Expected Transitions

MOSR Model Validation Example (Span 26)

Origin State	Next State	Observed	Expected	Absolute Error
6217|E3	6217|E3	25	19.90	5.10
6217|E3	6217|E4	6	8.82	2.82
6217|E3	6217|E1-E2	0	0.60	0.60
6217|E3	LAT_TRX	0	0.08	0.08
6217|E3	6257|E3	0	0.02	0.02
6217|E3	6337|E3	0	0.01	0.01
6217|E3	EAS	0	1.57	1.57

With the modeling framework established, the analysis now turns to empirical results. Chapter V evaluates the validation outcomes for each model configuration, compares the performance of annual and six-month transition structures, and applies the selected model to generate forward projections of aviation MOS inventories through FY34.

V. MODEL VALIDATION AND FORWARD PROJECTION

Chapter IV described the development of the Markov modeling framework used to estimate transition probabilities across the F/A-18 enlisted community, including state definitions, time-step construction, and the rolling validation methodology. With that structure established, this chapter presents the empirical validation results and applies the selected training span to generate forward projections through FY34. The analysis evaluates two time-step designs, annual FY-to-FY transitions and six-month transitions, as well as two state-space definitions: MOS only and MOSR. The analysis assesses each configuration using identical procedures to ensure that differences in performance reflect modeling choices rather than methodological variation.

The purpose of this chapter is to determine which model most accurately reproduces observed transition behavior during the F/A-18 divestment period and to assess whether shorter transition intervals improve forecast accuracy in a non-steady-state environment. After comparing validation performance across all variants, the chapter identifies the operational forecasting model and presents forward projections, including simulation-based prediction intervals and comparison to GAR targets. The implications of these projected outcomes for enlisted manpower planning are discussed in Chapter VI.

A. ANNUAL (FY-TO-FY) MODEL VALIDATION

This section evaluates the performance of the annual FY-to-FY transition model using the rolling validation framework described in Chapter IV.

The FY-to-FY model estimated one transition per fiscal year using September snapshots. The analysis conducted validation using rolling train-test windows, where each span represents the number of consecutive fiscal years included in the training matrix. The analysis evaluated performance using WAPE and aggregate bias.

1. MOS Only FY-to-FY Model

Table 2 presents the ranked span performance for the FY-to-FY MOS only model.

Table 2. Ranked Results for FY-to-FY Model (MOS Only)

Rank	Span	Test transitions (N)	Windows (count)	WAPE (%)	Bias (%)	Bias direction
1	2	25,949	14	5.88	-1.75×10^{-15}	Under-forecast
2	1	28,434	15	5.96	-8.00×10^{-16}	Under-forecast
3	3	23,579	13	6.16	1.93×10^{-15}	Over-forecast
4	5	19,120	11	6.28	-5.95×10^{-15}	Under-forecast
5	4	21,331	12	6.29	-1.07×10^{-15}	Under-forecast
6	6	16,924	10	6.60	0.00	No bias
7	8	12,463	8	6.61	-3.65×10^{-15}	Under-forecast
8	7	14,692	9	6.69	-4.64×10^{-15}	Under-forecast
9	9	10,351	7	6.87	0.00	No bias
10	11	6,578	5	7.47	-6.91×10^{-15}	Under-forecast

The annual MOS-only model achieved its lowest WAPE using a two-year training span, with the one-year span performing nearly as well. This pattern suggests that shorter historical windows better capture recent transition dynamics during divestment. As span length increased, predictive performance declined, indicating that longer averaging windows diluted structural shifts occurring within the F/A-18 community.

2. MOSR FY-to-FY Model

The MOSR FY model expands the state space to include rank. Table 3 summarizes ranked spans.

Table 3. Ranked Results for FY-to-FY Model (MOS and Rank)

Rank	Span	Test transitions (N)	Windows (count)	WAPE (%)	Bias (%)	Bias direction
1	2	21,443	14	3.11	3.63×10^{-1}	Over-forecast
2	1	23,565	15	3.13	4.48×10^{-2}	Over-forecast
3	3	19,493	13	3.30	2.11×10^{-1}	Over-forecast
4	5	15,866	11	3.32	-9.64×10^{-2}	Under-forecast
5	4	17,664	12	3.33	6.76×10^{-2}	Over-forecast
6	8	10,270	8	3.41	9.03×10^{-1}	Over-forecast
7	7	12,106	9	3.46	9.28×10^{-1}	Over-forecast
8	6	14,015	10	3.49	2.71×10^{-1}	Over-forecast
9	13	3,045	3	3.57	−2.93	Under-forecast
10	11	5,487	5	3.64	-2.58×10^{-1}	Under-forecast

The best performing FY-to-FY MOSR model used a two-year span, producing the lowest WAPE of 3.11 percent across 14 validation windows. Shorter training windows appear to better capture the transition dynamics observed during the divestment period, likely because they allow the model to respond more quickly to structural changes in the inventory.

In contrast to the MOS-only specification, incorporating rank detail improved predictive accuracy in the annual model. The best MOS-only FY model produced a WAPE of 5.88 percent, compared to 3.11 percent for the MOSR model. This reduction in forecast error suggests that including grade structure provides additional information about transition behavior at the annual time scale. Because promotions, separations, and lateral movements occur through rank pathways, modeling the system at the MOSR level better captures the personnel dynamics driving inventory change during the divestment period.

B. SIX-MONTH MODEL VALIDATION

This section applies the same validation framework to the six-month transition model in order to evaluate whether shorter time steps improve predictive accuracy. The six-month model estimates transitions at half-year intervals defined as H1 (Oct–Mar) and

H2 (Apr–Sep). This structure captures within year transition behavior that may be obscured by annual aggregation.

1. MOS Only Six-Month Model

Table 4 presents ranked span performance for the MOS-only half-year model.

Table 4. Ranked Results for Six-Month Model (MOS Only)

Rank	Span	Test transitions (N)	Windows (count)	WAPE (%)	Bias (%)	Bias direction
1	4	51,237	28	4.50	4.44×10^{-16}	Over-forecast
2	2	56,138	30	4.51	2.03×10^{-15}	Over-forecast
3	6	46,540	26	4.65	-4.89×10^{-16}	Under-forecast
4	8	42,107	24	4.74	-2.16×10^{-15}	Under-forecast
5	5	48,867	27	4.79	1.40×10^{-15}	Over-forecast
6	3	53,653	29	4.80	8.48×10^{-16}	Over-forecast
7	10	37,702	22	4.85	-1.81×10^{-15}	Under-forecast
8	7	44,292	25	4.90	1.54×10^{-15}	Over-forecast
9	9	39,896	23	4.91	3.99×10^{-15}	Over-forecast
10	11	35,506	21	5.11	8.97×10^{-15}	Over-forecast

The best performing six-month MOS-only model used a span of four half-year transitions, producing a WAPE of 4.50 percent across 28 validation windows. This configuration provided the lowest error among the MOS-only specifications and demonstrated stable aggregate bias across test windows. Shorter spans exhibited slightly higher variance, while longer spans produced gradual increases in forecast error. Overall, the six-month MOS-only model improved upon the FY-to-FY specifications but remained less accurate than the MOSR six-month model.

2. MOSR Six-Month Model

Table 5 presents ranked span performance for the MOSR half-year model.

Table 5. Ranked Results for Six-Month Model (MOS and Rank)

Rank	Span	Test transitions (N)	Windows (count)	WAPE (%)	Bias (%)	Bias direction
1	26	6,459	6	1.74	−1.25	Under-forecast
2	29	3,090	3	1.75	−1.24	Under-forecast
3	25	7,638	7	1.76	-9.71×10^{-1}	Under-forecast
4	27	5,292	5	1.78	−1.15	Under-forecast
5	28	4,189	4	1.87	−1.51	Under-forecast
6	24	8,906	8	1.98	-4.13×10^{-1}	Under-forecast
7	30	2,039	2	2.03	−1.71	Under-forecast
8	2	51,116	30	2.06	4.92×10^{-2}	Over-forecast
9	22	11,683	10	2.07	-2.80×10^{-2}	Under-forecast
10	4	46,588	28	2.07	1.90×10^{-1}	Over-forecast

The MOSR six-month model achieved the best predictive performance in the validation exercise. The top-ranked configuration used a span of 26 half-year transitions and produced a WAPE of 1.74 percent across six validation windows. Incorporating rank detail improved predictive accuracy in the six-month framework, suggesting that grade-specific transition dynamics provide meaningful signal when modeled at shorter time intervals. While the MOSR model increases the dimensionality of the state space, the longer training span provides sufficient observations to stabilize the transition probabilities and reduce forecast error relative to the MOS-only specification.

C. TIME STEP COMPARISON

After evaluating each model independently, the analysis compares the results across time-step structures to determine which approach provides the most accurate forecasts. Figure 3 compares the best performing configurations across all models.

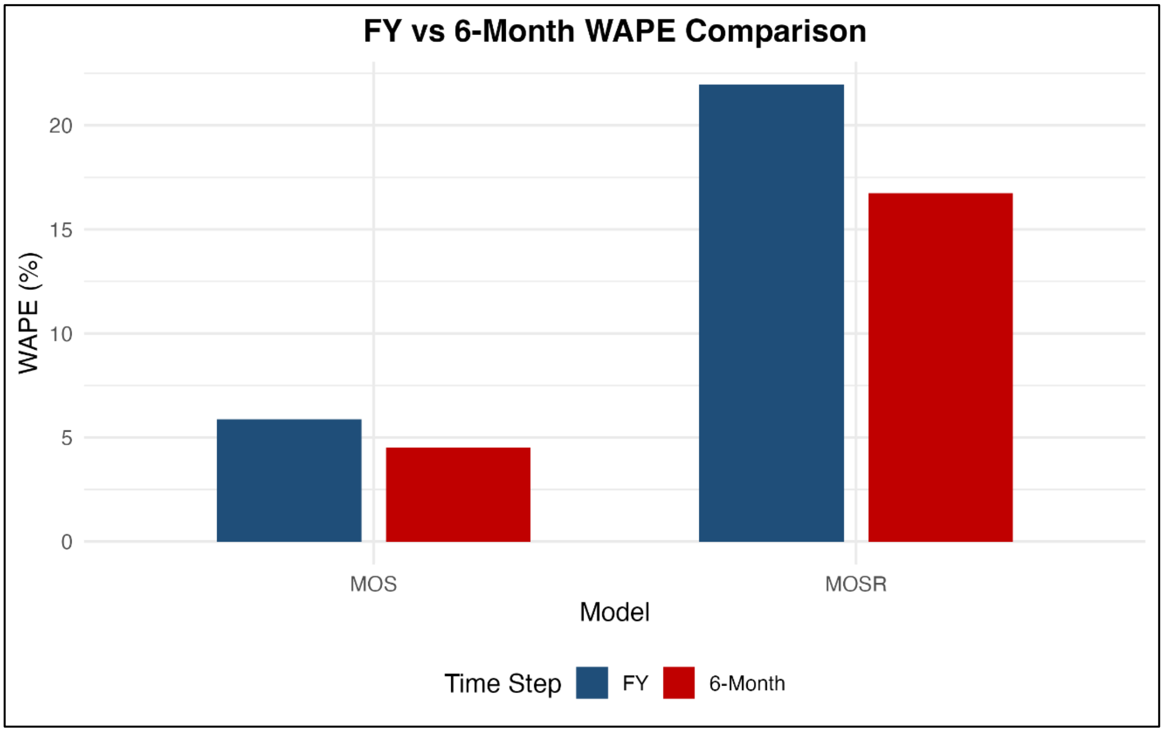

Figure 3. FY vs. Six-Month WAPE Comparison

The results demonstrate three key findings. First, the six-month time step outperforms the FY-to-FY model. Figure 3 shows that the six-month models produce lower forecast error. This indicates that shorter transition intervals better capture manpower churn during platform divestment by reflecting the more frequent movement of personnel across states.

Second, incorporating grade detail improves predictive accuracy when combined with the six-month transition structure. The MOSR six-month model produces the lowest WAPE observed in the validation exercise, outperforming both the MOS-only six-month specification and the FY-to-FY models. This suggests that rank-specific transition dynamics become more informative when transitions are modeled at shorter time intervals.

Third, span sensitivity differs across model structures. The MOS-only six-month model achieves its best performance with a relatively short span of four half-year transitions, while the MOSR six-month model requires a substantially longer span to stabilize the larger transition matrix. This pattern reflects the tradeoff between state detail and data availability. Models with greater state granularity require longer training windows to accumulate sufficient transition observations. Overall, the results indicate that both the choice of time step and the level of state detail influence forecast accuracy during periods of structural manpower change.

D. FORWARD PROJECTION FRAMEWORK

Forward projections use the twenty six-month span, six-month transition matrices for both the MOS and MOSR state spaces. For each state-space variant, the analysis initializes projections using the observed FY25 H2 inventory calculated from the cleaned personnel dataset. The model then applies the estimated half-year transition matrices forward to project inventory trajectories through FY34.

For each projection, the model applies the estimated half-year transition matrix (span four for MOS and span twenty-six for MOSR) from FY26 through FY34. In the MOS-only specification, transitions occur across primary MOS states, with lateral transfer and EAS treated as absorbing states. In the MOSR specification, transitions occur across MOS-rank group combinations, preserving grade-sensitive promotion and attrition behavior.

E. PROJECTION RESULTS

The model calculated projected inventories at each fiscal year H2 snapshot through FY34.

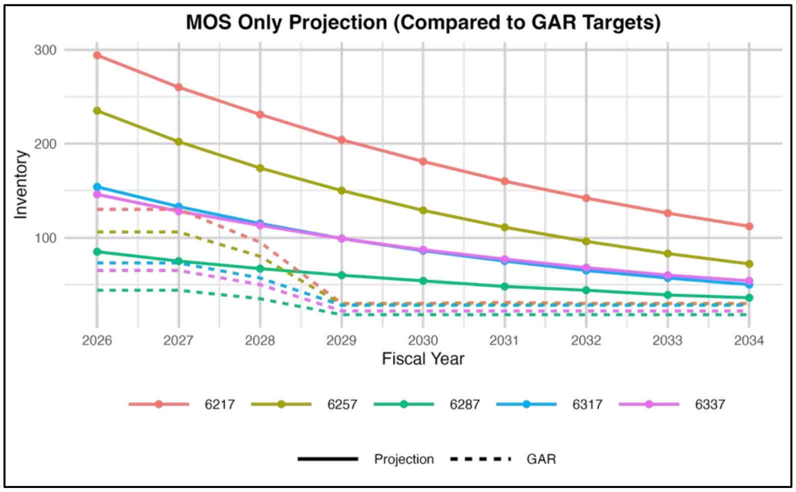

Figure 4. MOS Only Projected Results Versus GAR Target Numbers

Figure 4 illustrates projected enlisted inventory trends for the five F/A-18 maintenance MOS communities from FY2026 through FY2034. Across all MOSs, inventories decline steadily over the projection period as the F/A-18 platform continues its planned divestment. Higher-population MOSs begin the period with larger inventories but follow a similar downward trajectory as smaller communities. Although the rate of decline varies slightly by MOS, the overall pattern indicates a gradual drawdown rather than an abrupt reduction in personnel. By the end of the projection period, inventories across all MOSs converge toward lower levels, reflecting the long-term reduction in maintenance requirements associated with the sundown of the platform.

The dotted lines represent the GAR targets provided by Manpower and Reserve M&RA for each MOS over the projection period. These values reflect the planned distribution of billet requirements associated with the gradual divestment of the F/A-18 platform. While GAR targets provide an institutional planning reference for desired inventory levels by MOS and grade, they represent billet requirements rather than executed staffing levels. In practice, assignment authorities such as MMIB and Manpower Management Enlisted Assignments MMEA retain flexibility to align Marines to billets based on operational needs. As a result, actual inventories may diverge from GAR targets as staffing adjustments occur during execution. Comparing projected inventories to GAR therefore provides a planning benchmark rather than a direct measure of forecast accuracy.

Because the MOSR model explicitly incorporates rank transitions, its primary analytical value lies not in visualizing total trajectories but in evaluating rank-level alignment relative to GAR requirements. Plotting MOSR projections as line graphs generates multiple overlapping series and obscures structural imbalance. Accordingly, the analysis presents MOSR results as Projection minus GAR to demonstrate the delta by rank and fiscal year.

Table 6. Difference Between Projected and GAR Values by MOS and Rank

MOSR Delta (Projection – GAR)								
	2027	2028	2029	2030	2031	2032	2033	2034
6217								
E7	84	80	78	73	69	65	61	57
E6	26	28	31	27	23	19	15	12
E5	5	7	14	8	4	0	−3	−5
E4	8	6	18	11	7	4	3	2
E3	−13	−10	6	3	1	1	1	1
E1-E2	1	1	0	0	0	0	0	0
6257								
E7	78	75	74	71	68	64	61	58
E6	25	24	27	23	20	16	14	11
E5	−4	−3	4	1	−1	−3	−5	−6
E4	2	1	10	6	3	1	0	−1
E3	−7	−5	8	5	3	2	2	1
E1-E2	1	1	0	0	0	0	0	0
6287								
E7	32	31	32	31	29	28	27	25
E6	8	7	7	6	4	3	2	1
E5	−5	−5	−3	−4	−5	−6	−6	−6
E4	−6	−4	1	0	0	−1	−1	−1
E3	−2	−1	2	2	1	1	1	1
E1-E2	0	0	0	0	0	0	0	0
6317								
E7	46	45	45	43	41	39	37	35
E6	15	13	12	9	6	4	2	0
E5	−5	−4	1	−2	−4	−6	−8	−8
E4	6	4	6	4	3	2	2	1
E3	−7	−4	4	3	2	2	2	2
E1-E2	1	1	0	0	0	0	0	0
6337								
E7	46	45	45	43	41	38	36	34
E6	18	15	14	11	8	6	4	2
E5	−3	−2	1	−1	−2	−3	−4	−5
E4	−3	−1	6	4	3	2	1	1
E3	−2	−1	5	4	3	2	2	2
E1-E2	1	1	0	0	0	0	0	0

Presenting MOSR outcomes in delta form directly quantifies projected inventory surpluses (positive values) and shortfalls (negative values) relative to GAR requirements by rank. This format highlights the fiscal year in which divergence begins and identifies the specific grade groups driving imbalance. Because the F/A-18 community is scheduled to sundown at the end of FY28, the projections illustrate how remaining personnel inventory evolves relative to requirements during the divestment period. Across most MOS and rank combinations, the model projects positive deltas following platform divestment, indicating that inventory begins to exceed GAR requirements as billets disappear faster than personnel transition out of the system. The primary exception occurs at the E-5 level, where projections remain negative across multiple MOSs and fiscal years. This pattern indicates that mid-grade maintainers represent the most persistent inventory shortfall during the transition, even as surpluses emerge in other ranks once the platform begins to divest.

F. SIMULATION AND 95% PREDICTION INTERVALS

The analysis conducted multinomial simulation for each projection path to capture uncertainty in transition behavior. For every origin state, the model redistributed individuals according to the estimated transition probabilities, and this process repeated across half-year steps.

From these simulated distributions, the analysis calculated the median projection and 95 percent prediction intervals using the 2.5th and 97.5th percentiles.

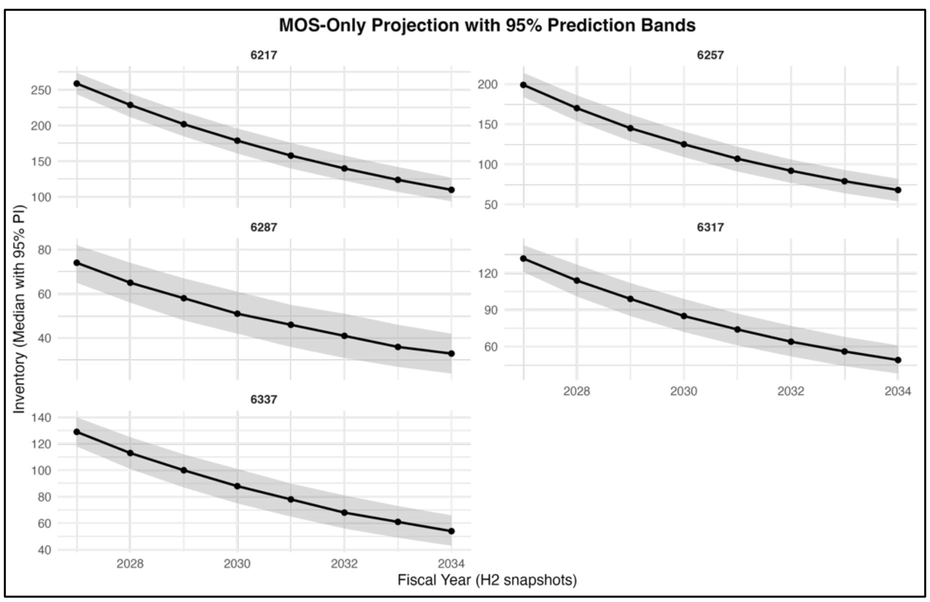

Figure 5. MOS Projection with 95% Bands

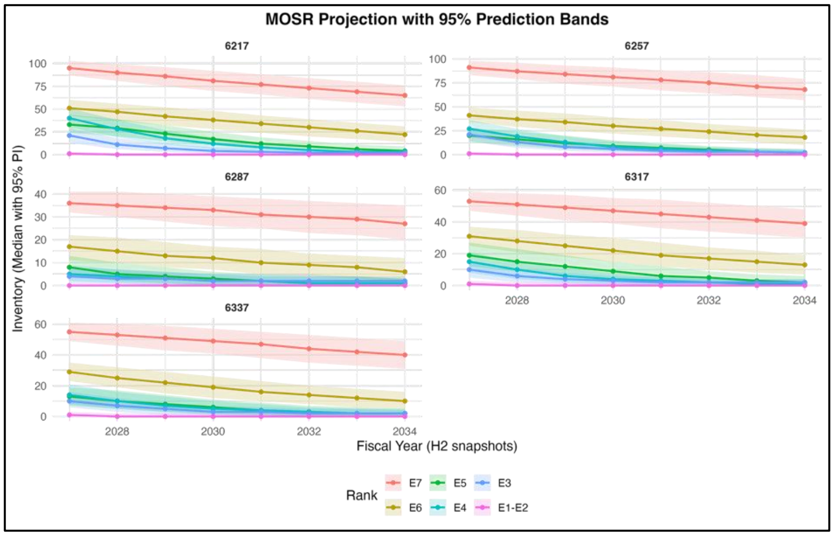

Figure 6. MOS and Rank Predictions with 95% Prediction Bands

Because projections apply the transition matrices sequentially across fiscal years, simulated uncertainty accumulates through each update. In the figures, the solid line represents the median projected inventory, while the shaded region represents the 95 percent prediction interval derived from the simulated outcomes. As projections move farther from the observed FY25 baseline, the shaded bands widen to reflect increasing uncertainty in future transition behavior.

G. MODEL CONSIDERATIONS

It is important to note that the projections use transition matrices estimated from historical data through FY25. As a result, projected transition behavior assumes that observed patterns remain informative of near-term future dynamics. If organizations implemented this framework for operational manpower planning, they would need to update transition matrices periodically as new personnel data becomes available.

Because platform divestment introduces structural change, transition probabilities may evolve as remaining squadrons deactivate and accession policies adjust. Accordingly, the forecasting framework presented here should be viewed as a dynamic tool rather than a static solution. Regular recalibration using updated transition observations would be necessary to maintain predictive accuracy.

The final chapter, Chapter VI, synthesizes these results and examines their implications for Marine Corps manpower planning. It summarizes the key findings of the study, evaluates the strengths and limitations of the modeling framework, and outlines potential directions for future research on manpower forecasting during platform transitions.

THIS PAGE INTENTIONALLY LEFT BLANK

VI. CONCLUSION AND RECOMMENDATIONS

This chapter summarizes the results of the analysis and discusses their implications for Marine Corps aviation manpower planning during platform divestment. It reviews the research objectives and key findings, evaluates the strengths and limitations of the modeling framework, and outlines potential directions for future research. The chapter concludes with observations regarding how the forecasting approach developed in this study may support manpower decision making as the F/A-18 community continues its transition.

A. SUMMARY OF RESEARCH OBJECTIVE AND APPROACH

This thesis set out to answer a practical problem facing Marine Corps aviation manpower planners: how do we forecast enlisted inventory behavior during a platform divestment in a way that is transparent, defensible, and grounded in historical transition behavior?

The sundown of the F/A-18 community under Force Design 2030 (USMC, 2020) is not simply a platform decision. It creates second and third-order manpower effects across multiple enlisted MOS communities. Marines do not leave uniformly. Some EAS, some laterally transfer, and others move into B-billets or remain within the MOS at different ranks. Without a structured forecasting framework, it becomes difficult to understand how quickly inventories will decline and where shortages will emerge.

To address this, I developed two Markov-based forecasting frameworks, a Fiscal FY-to-FY model and a six-month transition model. I apply both models at two levels, MOS and MOSR. Rather than selecting a model based on preference, this study evaluated and ranked rolling training windows using performance metrics, primarily WAPE and aggregate bias. This allowed model selection to be based on predictive accuracy rather than assumptions.

The analysis then used the selected models to generate forward projections through FY34 to assess expected inventory behavior during the divestment period.

B. SUMMARY OF KEY FINDINGS

Crucially, platform divestment does not immediately reduce manpower inventory. Across most MOS and grade combinations, projected inventories decline more gradually than the pace implied by platform divestment timelines. This indicates that historically observed transition behavior produces a slower drawdown of personnel than institutional planning targets assume. Marines assigned to divesting platforms do not disappear from the manpower system when squadrons deactivate; instead, they continue progressing through normal career pathways, including promotion, lateral transfer, and separation over time. Because the F/A-18 community is scheduled to sundown at the end of FY28, the projections illustrate how remaining personnel inventory evolves relative to requirements during the divestment period. Across most MOS and grade combinations, projected inventories eventually exceed GAR requirements after the platform begins to divest, indicating that personnel inventories decline more slowly than billet reductions. The primary exception occurs at the E-5 level, where projected inventories fall below GAR targets and remain negative across multiple MOSs later in the projection period. This pattern suggests that while most grades may experience temporary overages relative to requirements, the non-commissioned officer (NCO) cohort may remain constrained as the divestment progresses.

The forecasting framework developed in this thesis could complement existing manpower inventory tools that organizations such as MPP-20 and the Manpower Management Integration Branch (MMIB) use. Because the model calibrates directly on observed transition behavior, it provides short-term projections of MOS inventories during periods of structural change such as platform divestment. These projections illustrate how personnel inventories may diverge from GAR requirements over time. Used alongside existing inventory models, this approach could provide planners earlier visibility into emerging MOS or grade imbalances and support decisions such as adjusting accession targets, expanding retraining pipelines, or modifying lateral transfer opportunities.

Across all windows evaluated, the six-month MOSR span 26 model produced the lowest overall WAPE. The next best performer was the MOS six-month span 4 model, followed by the FY-to-FY MOS span 2 and FY-to-FY MOSR span 5. Several patterns emerged from this comparison.

First, shorter transition intervals (six-month) generally performed better than annual transitions. This suggests that enlisted manpower behavior during a divestment is more responsive than a full fiscal year model can capture.

Second, MOSR models performed best at the aggregate level in the six-month structure, indicating that incorporating rank detail improves predictive accuracy when the model captures transitions at shorter intervals. MOS-level models still provide useful aggregate insight, but they cannot show rank-specific inventory dynamics.

Finally, the validation framework demonstrated that span length matters. Models trained on too long of a window introduced lag and masked structural shifts. Models trained on too short of a window became overly sensitive to noise. The span 26 six-month model provided the best balance between stability and responsiveness.

C. STRENGTH AND LIMITATIONS OF THE MARKOV MODELING FRAMEWORK

The Markov modeling framework developed in this study offers several clear advantages for aviation manpower forecasting. First, it is transparent. Every projection is derived directly from observed transition probabilities within the historical data. There are no hidden coefficients, subjective weighting factors, or behavioral assumptions beyond what the data supports. This makes the results traceable and defensible.

Second, the framework is reproducible. The rolling-window validation structure allows future planners to update transition matrices, re-rank training spans, and regenerate projections using the same objective performance criteria. Model selection is therefore based on measured predictive accuracy rather than preference or convention.

Third, the model is flexible. Planners can adjust transition probabilities to simulate potential policy interventions, including retention incentives, lateral transfer restrictions, or modified separation assumptions. In this way, the framework functions not only as a forecasting tool, but also as a decision-support mechanism capable of evaluating alternative manpower strategies.

Finally, the model aligns with the Marine Corps planning cadence. Both the FY-based and six-month constructs mirror how manpower planners review and manage manpower in practice. This alignment increases the practical utility of the projections and facilitates integration into existing planning processes.

At the same time, several limitations define the boundaries of the model's application. Like all Markov systems, this framework assumes stationarity, that transition probabilities observed within the selected training window reasonably represent near-term future behavior. During periods of significant policy change or external shocks, that assumption may weaken.

In addition, the model operates at the aggregate transition level. It does not incorporate individual-level characteristics or behavioral drivers such as bonus structures, deployment tempo, or promotion opportunity. While this abstraction supports system-level forecasting, it limits behavioral specificity.

The scope of this study is also intentionally narrow. The analysis focuses exclusively on enlisted F/A-18 MOS communities. Officer manpower, cross-platform interaction effects, and enterprise-wide aviation implications fall outside the scope of this research.

Forward projections assume no major structural policy shifts beyond those reflected in the historical data. Significant retention initiatives, force-shaping tools, or changes in accession policy could alter observed transition behavior and therefore affect forecast accuracy.

Another limitation involves the use of GAR targets as the comparison baseline for projected inventories. While the GAR provides institutional planning targets derived from billet requirements, it does not necessarily reflect executed staffing levels. Assignment authorities such as MMIB and MMEA retain the flexibility to adjust staffing in response to operational conditions, meaning observed inventories may diverge from GAR targets even when manpower management is functioning as intended.

These limitations do not invalidate the modeling framework. Rather, they clarify the conditions under which planners should interpret and apply the projections. Within those bounds, the framework provides a structured and defensible method for forecasting aviation manpower during platform transition.

D. RECOMMENDATIONS FOR FUTURE RESEARCH

While this study focused on forecasting enlisted inventory behavior within the F/A-18 community, the modeling framework developed here provides a foundation for several important extensions.

One natural expansion would incorporate cross-platform interaction effects. Platform divestment rarely occurs in isolation, and future research could examine how reductions in one aviation community influence adjacent MOSs through lateral transfers, retraining pipelines, or enterprise redistribution. Integrating officer and enlisted models would also provide a more comprehensive view of aviation manpower dynamics, particularly in communities where rank structure and leadership flow significantly influence retention behavior.

In several projected fiscal years and grades, modeled inventory exceeds authorized or required GAR targets. Analysts should not treat excess inventory solely as a surplus problem. Marines who fall outside GAR constraints remain trained and experienced manpower assets. Rather than allowing excess inventory to dissipate through unmanaged separation or underutilization, planners could develop an optimization model to deliberately reallocate these Marines back into the broader aviation workforce.

Such a framework could incorporate constraints including available B-billet capacity, lateral transfer opportunities, retraining timelines and associated costs, promotion flow requirements, and enterprise-wide aviation maintenance demand. Integrating these constraints would allow planners not only to identify projected imbalances, but also to generate redistribution strategies that maximize workforce utilization while remaining within GAR and force structure limits.

Layering an optimization model on top of the validated Markov forecast would shift manpower planning from descriptive projection toward decision-oriented force design. This integration would provide commanders and manpower analysts with structured, actionable pathways for managing both shortages and excess inventory during platform transitions.

Future research could also evaluate model performance by comparing projected inventories against executed staffing decisions rather than relying solely on GAR targets. Because organizations such as MMIB and MMEA retain authority to adjust staffing assignments dynamically, comparing forecasted inventories with the actual adjustments made during execution could provide an additional measure of forecasting accuracy. Such an approach would allow researchers to assess not only how closely projections align with planning targets, but also how effectively they anticipate real-world manpower management actions.

E. FINAL CONCLUSION

The sundown of the F/A-18 community represents a structural shift in Marine Corps aviation. Managing that shift requires more than end-strength targets; it requires visibility into how Marines transition through the system over time.

This thesis demonstrates that a validated Markov modeling framework can provide that visibility. By combining rolling-window validation, objective span selection, and forward projection, the approach offers a structured and defensible way to forecast inventory behavior during divestment.

Most importantly, the methodology is adaptable. As Force Design 2030 (USMC, 2020) continues to reshape the Marine Corps, this framework can evolve alongside it, providing planners a tool for understanding and managing manpower change rather than reacting to it.

LIST OF REFERENCES

Bartholomew, D. J., Forbes, A. F., & McClean, S. I. (1991). *Statistical techniques for manpower planning* (2nd ed.). John Wiley & Sons.

United States Marine Corps. (2021). *Aviation Maintenance and Supply Training and Readiness Program* (NAVMC 4790.1D).

Department of the Navy. (2021). Naval aviation maintenance program (NAMP) (OPNAVINST 4790.2D). https://www.navair.navy.mil/sites/g/files/jejdrs536/files/2021-02/COMNAVFORINST%204790.2D%20NAMP.pdf

United States Marine Corps. (2026). *2026 Marine Corps Aviation Plan*. Headquarters Marine Corps. https://www.marines.mil/News/Press-Releases/Press-Release-Display/Article/4402473/2026-marine-corps-aviation-plan/

Dimitriou, V. A., Georgiou, A. C., & Tsantas, N. (2013). The multivariate non-homogeneous Markov manpower system in a departmental mobility framework. *European Journal of Operational Research, 228*(1), 112–121.

Erhardt, J. L. (2012). *Development of a Markov Model for Forecasting Continuation Rates for Enlisted Prior Service and Non-Prior Service in the Selective Marine Corps Reserve (SMCR)* (Master's thesis Naval Postgraduate School, Monterey, CA.) https://calhoun.nps.edu/entities/publication/04fff311-33bb-4927-9b7b-ea411446976b

Licari, J. P. (2013). *Developing a Markov Model for Forecasting End Strength of Selected Marine Corps (SMCR) Officers* (Master's thesis Naval Postgraduate School, Monterey, CA.) https://calhoun.nps.edu/search?query=licari&scope=8170278c-29e1-4af7-8cd6-c960e36516f9

McClean, S. I. (1991). Manpower planning models and their estimation. *European Journal of Operational Research, 51*(2), 179–187.

United States Marine Corps. (2015). *Total force structure process manual* (MCO 5311.1E). https://www.marines.mil/News/Publications/MCPEL/Electronic-Library-Display/Article/900533/mco-53111e-wadmin-ch-1/

United States Marine Corps. (2019). *Military occupational specialties manual* (NAVMC 1200.1E). https://www.marines.mil/News/Publications/MCPEL/Electronic-Library-Display/Article/4132818/navmc-12001l/

Rai, A. (2019). *Improving the Navy Nurse Corps' Wartime Surge Force Planning by Implementing a Markov Model* (Master's thesis Naval Postgraduate School, Monterey, CA.) https://calhoun.nps.edu/entities/publication/d162a626-c5b8-43bc-82fd-9149c96af766

Ross, S. M. (2014). *Introduction to probability models* (11th ed.). Academic Press.

Stewman, S. (1978). Markov and renewal models for total manpower system. *Omega, 6*(4), 341–351.

Taylor, J. A. (2020). *Direct Officer Accessions: Effecting the Timeline for Establishing a Cyber-Community Within the Marine Corps* (Master's thesis Naval Postgraduate School, Monterey, CA.) https://calhoun.nps.edu/entities/publication/4101286e-4c1f-42ff-9643-6c08a791964a

United States Marine Corps. (2020). *Force design 2030*. Headquarters Marine Corps. https://www.marines.mil/ForceDesign2030/

Acquisition Research Program
Naval Postgraduate School
555 Dyer Road, Ingersoll Hall
Monterey, CA 93943

www.acquisitionresearch.net