Submitted to *Journal TBD*
manuscript (xxxxxxxx)

A Markov Chain Model of Military Personnel Dynamics

Personnel retention is one of the most significant challenges faced by the U.S. Army. Central to the problem is understanding the incentives of the stay-or-leave decision for military personnel. Using three years of data from the US Department of Defense, we construct and estimate a Markov chain model of military personnel. Unlike traditional classification approaches, such as logistic regression models, the Markov chain model allows us to describe military personnel dynamics over time and answer a number of managerially relevant questions. Building on the Markov chain model, we construct a finite horizon stochastic dynamic programming model to study the monetary incentives of stay-or-leave decisions. The dynamic programming model computes the expected payoff of staying versus leaving at different stages of the career of military personnel, depending on employment opportunities in the civilian sector. We show that the stay-or-leave decisions from the dynamic programming model possess surprisingly strong predictive power, without requiring personal characteristics that are typically employed in classification approaches. Furthermore, the results of the dynamic programming model can be used as input in classification methods and lead to more accurate predictions. Overall, our work presents an interesting alternative to classification methods and paves the way for further investigations on personnel retention incentives.

Key words: Markov chain estimation, stochastic dynamic programming, reenlistment propensity

1 INTRODUCTION

The size of the U.S. Army is often a function of competing budget constraints and mission requirements. Moreover, the United States experiences fluctuating periods of economic growth and decline combined with varying sentiments towards a propensity to serve. All of these affect how our nation assesses new recruits into the armed forces. A similar scenario is evident when the Army takes action to influence the retention (reenlistment) of enlisted soldiers. The Army uses forecasting methods to predict the number of reenlistment eligible soldiers it must retain each fiscal year. When projections without incentives fall below requirements, natural questions persist. Can we target incentive programs and policies to specific categories of individuals? How can we identify those most likely to leave?

The U.S. Army cannot adequately answer the more precise question related to incentive size unless it uses the most relevant data and techniques to predict personnel behavior. Current U.S. Army incentive programs frequently target specific skills or military occupation specialties without rigorous analysis that determines propensities to serve. There is also a serious lack of understanding of how individuals make stay-or-leave decisions. There exists a research gap for separating individual retention propensity and for explaining individual retention decisions.

We model the career progression of enlisted personnel as a discrete-time homogeneous Markov chain and estimate state transition probabilities using three years of reenlistment data from the US Department of Defense. The estimated transition probabilities offer at least two benefits: (1) it provides an easily constructible method for evaluating the probabilistic progression of enlisted personnel through different career states, and (2) it can be used in a subsequent stochastic dynamic programming model to understand reenlistment behavior.

The Markov chain model allows us to answer managerial questions from policy makers. For example, it allows easy computation of various statistics at both individual and aggregate levels. At the individual level, it can be used to describe the probabilistic progression for military personnel at a given career stage. At the aggregate level, it can be used to derive information on overall continuation rates and separation behavior, which are critical inputs in developing retention programs.

Our estimation methodology only requires a data set with a limited time range rather than the complete career time horizon of personnel cohorts. This should be viewed as an advantage, as the military is constantly fluctuating in size, structure and policies. Our Markov chain estimation approach is not only more adaptable and easier to implement than tracking cohort data, but also captures transition probabilities while omitting obsolete data.

Building on our Markov chain estimation result, we construct a finite horizon stochastic dynamic programming model to understand the stay-or-leave decision of a military personnel at different career stages. Our model assumes that an individual is forward-looking and can rationally evaluate his expected payoff of staying in the military, taking into account future career progressions and retirement benefits. The military individual leaves the army if the expected payoff in the civilian sector dominates the expected payoff of staying in the army. Therefore, our model takes into account data on both military pay and civilian pay. Taken together, the dynamic programming model allows us to calculate the reenlistment propensity of military personnel at different career stages and therefore can be a useful tool for policy makers.

The dynamic programming model is a viable alternative to more traditional approaches based on statistical analysis. A popular approach to predicting the stay-or-leave decision is to use a binary logistic regression model. We show that the dynamic programming approach can predict retention behavior with similar accuracy as a well-constructed logistic regression model. However, unlike the logistic regression model, the dynamic programming model requires far fewer predictive variables. Our dynamic programming model uses only grade and time-in-grade variables, while the logistic regression model uses many variables on personal attributes. Using personal attributes for retention incentives often leads to perceived discrimination and therefore can be problematic. The strength of the dynamic programming models is that it takes into account expected future compensation, which is not incorporated in the logistic regression model. We also consider a hybrid logistic regression model that incorporates the output from the dynamic programming model and shows that the predictive power can be dramatically improved.

As a by-product of the dynamic programming model, we obtain the difference in expected future compensations between the military and civilian sectors based on a military personnel's career status. This allows a tailored approach for retention incentives based on variables such as grade and time-in-grade, and has the potential to generate substantial savings for the military.

The balance of the paper is organized as follows. Section 2 discusses some relevant literature. Section 3 presents the Markov chain model and its estimation from data. Section 4 presents an aggregated Markov chain model based on grouped states. This section also presents functional manipulations of the transition matrix that provide insightful steady state interpretations. In Section 5, we present a dynamic programming model for the stay-or-leave decision and then in Section 6 we integrate the results from dynamic programming with logistic regression results to assess the impact on predictive modeling. Section 7 concludes.

2 RELATED LITERATURE

Strategic workforce planning is well documented in many domains, to include the defense sector. Literature on the topic spans an array of disciplines and organizations such as the healthcare industry, the Federal Bureau of Investigation, and other government agencies. Within government, the Government Accountability Office studies principles for effective strategic workforce planning across selected agencies in order to align human capital with strategic goals while the Department of Defense publishes a strategic workforce plan on an annual basis (Government Accountability Office, 2003; Department of Defense, 2014). Most research related to strategic workforce planning is policy driven and focuses on process improvements. Department of Defense (2014) states that the technical solution for assessing total force capability versus manpower requirements does not currently exist; our research improves upon the common analytical gap in strategic workforce planning, particularly with predicting individual retention behavior.

The desire to attract, retain and develop talent is an increasing priority for businesses and organizations. April et al. (2014) developed a simulation-optimization approach to strategic workforce planning called OptForce, which "offers a broad range of predictive analytics capabilities including sophisticated workforce demand planning, fine grained employee retention models, and agent-based simulation forecasting." As a decision making tool, OptForce goes beyond traditional approaches based on static assumptions and is flexible enough for any human capital process. It is effective in measuring the impact of policies and practices on retention by skill populations and can model the likely impact of broad compensation strategies; however, our approach aims to build the framework for retention incentives that are more individually tailored by attributes.

Military workforce planning inherently has an additional level of complexity than other environments. An important distinction from strategic workforce planning in the general space is increased precision required. For example, the lack of lateral entry paths is just one of many unique aspects to the military manpower system. Wang (2005) classifies military workforce planning techniques into four categories: Markov chain models, computer simulation models, optimization models and supply chain management through Systems Dynamics. Our approach addresses limitations described by Wang of Markov chain models in workforce planning. By incorporating a hybrid Markov chain and dynamic programming framework, we include mathematical programming techniques that lends itself to optimization of outcomes such as minimizing cost. Moreover, the size of the military produces adequate sample sizes, which Wang (2005) highlights as a limiting aspect of many workforce environments.

Discrete-time Markov chains are frequently used to model and evaluate the long-term behavior of individuals in a variety of disciplines. Describing a process as a Markov chain and understanding its properties enables informed decision making and policy guidance. Constructing and evaluating Markov chains has a precedence in the medical industry, particularly in describing the progression and treatment of chronic diseases and illness. A Markov chain model can be used to study long-run behavior even when only data with limited time range is available. This property is important for certain applications, including the one in the current paper.

Beck and Pauker (1983) describe how a Markov chain model can replace a analytical methods such as decision trees within the medical decision making process. They demonstrate the utility of the fundamental matrix, especially coupled with simulation. Craig and Sendi (2002) use discrete-time Markov chains to evaluate treatment programs and health care protocols for chronic diseases. They present different types of situations that make matrix estimation techniques unique. In particular, they describe a discrete-time homogeneous Markov model in which the estimation techniques use observation intervals that coincide with the cycle length. This situation happens to align fairly accurately with the structure of the enlisted Army career network we describe in this paper.

Markov models are also applied to the financial industry as described by Cyert et al. (1962). They developed a Markov model describing the behavior of accounts receivable balances and used matrix properties, including the fundamental matrix, to make a variety of interpretations about the behavior of accounts at different stages of a time horizon. While the context of the Cyert et al. (1962) research is much different than military personnel dynamics, we demonstrate the same utility of an implementable method for interpreting a variety of system related questions by extending the use of the fundamental matrix for interpreting military retention behavior.

Pfeifer and Carraway (2000) introduce a Markov chain modeling approach to model customer relationships with a firm or business in which customer retention pertains to the potential for future customer-relationships. They highlight the usefulness of Markov chain modeling for decision making and retention situations that do not have algebraic solutions. We leverage their approach of using Markov chain models for one-one-one direct marketing to individual customers rather than broad cohorts and extend the idea for improving decision making and policies related to individuals in the military. Markov chain models are applicable to almost any domain that involves predicting choice decisions, as evidenced by Kvan and Sokol

(2006) in predicting National Collegiate Athletic Association (NCAA) basketball tournament outcomes. Kvan and Sokol (2006) present a logistic regression/Markov chain approach in which a logistic regression model is used to calculate transition probabilities. Our paper demonstrates a similar approach to combining methods; however we use a Markov Chain model as the starting point and also implement dynamic programming.

A small subset of retention research uses dynamic programming, which can be described as an approach for sequential decision making (Puterman, 2005). Using a modeling approach similar to Hall (2009) and Gotz and McCall (1983), the enlisted retention model can be structured as a manpower network. Hall models Army officer retirement, while Gotz and McCall model a distinct proportion of the Air Force officer. Gotz and McCall assert that, between the 10 and 20 year marks, retirement pay is the most important factor for officer retention with officers making decisions in an optimal sequential fashion. In contrast to the aforementioned, enlisted personnel behave differently than commissioned officers in the military. The enlisted career path and military experience is much different than that of officers, making it unreasonable to conclude that retirement pay is as influential without a separate analysis.

The ability of current technology to accommodate more computationally intensive models suggests that more complex dynamic programming methods are suitable for modeling optimal reenlistment behavior than the simplified Annualized Cost of Leaving model produced by Warner and Goldberg (1984), which focus on first and second terms of service and do not incorporate future uncertainty. Dissimilar to Asch et al. (2008), we do not model the retention decision of all services or incorporate the potential transition from active duty to the reserves. Asch et al. estimate model parameters that affect the decision to stay on active duty or leave, and a parameter related to the variance of the stochastic shocks affecting the alternatives of being a civilian or a reservist. Whereas Asch et al. (2008) estimate model parameters associated with the means, variances, and covariance of the preference for active and reserve service, we computationally construct parameters, specifically pertaining to transitions rates, from historical data sets.

Duala and Moffitt (1995) also used a stochastic dynamic programming model to estimate the effect of reenlistment incentives on military retention rates using panel data. They model decision points every 4 years after reenlistment up to twentieth year, but did not analyze the relevance of military occupation specialties. The modeling approach proposed in this paper can be used in aggregate or decomposed by skill. We also incorporate the entire career horizon and different parameter estimation techniques. Rather intuitively, in the presence of unobserved heterogeneity, Duala and Moffitt (1995) also conclude that a higher military-civilian pay difference significantly affects reenlistment behavior.

A distinct difference between modeling officer and enlisted retention decision lies in the gap between career expectations. Identical retirement systems may shape intuition about similar officer and enlisted retirement rate; however, the reality is that the probability of an enlisted personnel reaching retirement is significantly less than officers. Recent statistics show that less than 15% of enlisted personnel, compared to 46% of officers, will become eligible for retirement annuity (Stewart and White, 2004; Henning, 2011). This implies the relative importance of effectively modeling the enlisted retention decision prior to retirement eligibility. The retention model proposed in Section 5 assumes that the goal of enlisted personnel is to maximize pay and entitlements.

Expected civilian compensation is a key input regarding the decision to reenlist. Generally, officers are provided equal or better compensation in the civilian workforce when they retire or separate from the military. Due to lower average civilian education levels and variance in skill levels, enlisted trends do not necessarily mirror the officer domain. Therefore, one of Hall (2009)'s primary assumptions, that an officer's initial civilian pay is equivalent to their final military pay, is not reasonable for our model. Uncertainty in future civilian pay compared to current military compensation adds to model complexity but cannot be sacrificed for simplicity.

Duala and Moffitt (1995) acknowledged that ``military service is not completely substitutable for civilian work" and addressed civilian compensation by using estimates from Internal Revenue Service data on the post-service civilian earnings of veterans, allowing the civilian wage profile to depend on time spent in the military. Similarly, Asch et al. (2008) used data to map civilian wages with respect to total years of experience using a civilian median-wage profile based on population surveys of corresponding education ranges. Neither of the aforementioned models account for uncertainty of civilian wages. The method we use to estimate expected civilian pay entitlements is described in Section 5.1.

Perhaps, our research can most accurately be described as a compliment to the Dynamic Retention Model presented in the aforementioned Asch et al. (2008) and further extended in Asch et al. (2013). The Dynamic Retention Model has a proven record for influencing overarching military policies and modeling the effects of compensation decisions such as retirement reform. However, it is limited to analyzing the behavior of individuals in the aggregate, whereas our model seeks to provide a method for influencing the behavior of specific individuals based on certain attributes. We contend that the ability to target retention incentives by skill and attribute can contribute to significant budget efficiencies.

3 A MARKOV CHAIN MODEL FOR ENLISTED MILITARY PERSONNEL

This section introduce a discrete-time Markov chain model for enlisted Military Personnel. We use three years of historical data to estimate parameters for the Markov Chain.

3.1. Model Description

Our model considers individuals in the rank of Private (E1) through Sergeant Major (E9). However, since grades E1 through E4 belong to the same skill level category and have similar behavior patterns and decentralized. States refer to a combination of grade and years-in-grade. The state space is augmented by voluntary and involuntary separations, which are modeled as two separate absorbing states. Our model therefore ignores the possibility of returning to the military after voluntary and involuntary separations.

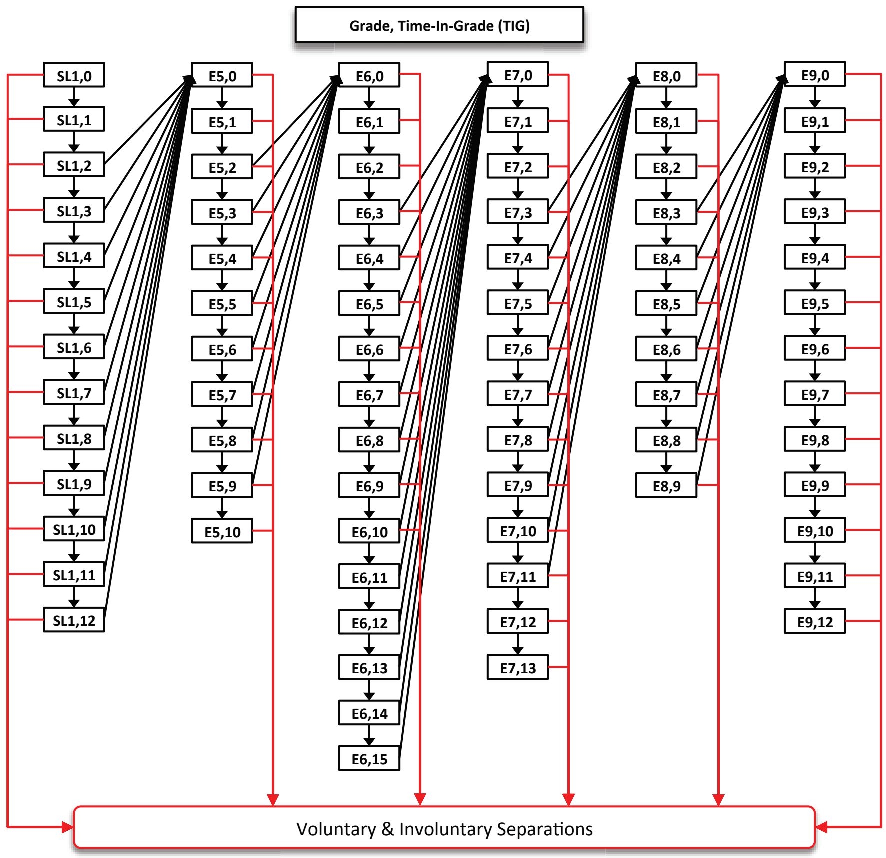

Figure 1 **State Transition Diagram of the Markov Chain Model**

The state definition in our model should be contrasted with some previous work where state is defined as an individual's pay grade and/or time-in-service (see, e.g., Hall, 2009). Note, however, that Hall (2009) models officers who are typically promoted in cohort. Compared with enlisted personnel, the variation in state transition is much smaller. For enlisted promotion criterion, which is relevant in our context, a minimum time-in-service requirement is typically stipulated. To incorporate time-in-service promotion criterion, we impose minimum time-in-grade requirement in the state transition. For example, an SL1 is not eligible for promotion to E5 with less than two years in grade and a E5 is not eligible for promotion to E6 with less than two years in grade, etc. The minimum time-in-grade captures the observations from the data quite well, as there are rarely violations. It also supersedes time-in-service requirements.

Figure 1 shows the state transition diagram for our model. In total, there are 79 states, including voluntary and involuntary separations as two absorbing states. Note that for each grade, there is a maximum time in grade. For example, for grade E5, the maximum time-in-grade is 11 years.

3.2. Data Description

In order to estimate parameters for the Markov chain, we analyze three years of historical data for enlisted personnel obtained from Headquarters, U.S. Army, Deputy Chief of Staff, Army G-1 (Personnel). Our data covers the period from October 2007 to September 2009. Our main data sources are personnel inventory and transaction files in Total Army Personnel Database (TAPDB). The personnel inventory file contains end of month snapshots of the active Army enlisted force each month. It includes a large number of personal attribute fields for each enlisted personnel. The transaction file is compiled at the end of each month and includes a record for any major personnel transaction occurred during the month. We merge inventory and transaction files by social security number (SSN) and date. Table 1 describes sample variables from TAPDB inventory and transaction files.

Table 1 Description of inventory and transaction data variables

Extract	Variable	Definition
Transactions	SSN	Social Security Number pertaining to the Soldier who executed the transaction
Transactions	Transaction Date	Actual date of transaction (DDMMYYYY)
Transactions	Transaction Category	Major category type of transaction (Gain, Loss, Extension, Promotion, Demotion)
Transactions	Gain Type	Description of specific type of gain to active duty (Prior Service, Non-Prior Service, Immediate Reenlistment
Transactions	Loss Type	Description of specific type of loss from active duty (Expiration of Term of Service, Retirement, Misconduct, Physical Disability, Dropped From Rolls, Entry Level Separation, Hardship or Parenthood, Pregnancy, Unfit, Unsatisfactory Performance, Reduction In Force, Early Release, Early Retirement, Immediate Reenlistment, Other)
Inventory	Pay Grade	Pay grade scale of individual (E1 to E9)
Inventory	Time-in-Grade	Months of service at current grade (converted to years)
Inventory	Time-in-Service	Total months of active service (converted to years)
Inventory	AFQT	Armed Forces Qualification Test score
Inventory	Education	Level of education centered on high school graduate
Inventory	Marital Status	Marital status of married, divorced or single
Inventory	Dependents	Number of minor dependents
Inventory	Race	Racial category of white, black, hispanic or asian/other
Inventory	Reenlistment Quantity	Number of previous reenlistments
Inventory	Deployed	Binary category equals 1 if previously deployed or 0 if never deployed
Constructed	MSD	Months since previous deployment
Constructed	Age	Age centered on 18 years
Constructed	Employment	Employment rate of change (12 month moving average)

Descriptive statistics of the data used for predictive modeling are provided in Appendix A. Not included are statistics for variables in Table 1 which are not used for prediction, such as gains, promotions and involuntary losses. Additionally, education level statistics are not presented due to the unique coding system adopted by the Army. However, 98.6% of individuals completed a level of schooling equivalent to a high school degree and 9.2% of enlisted individuals have some amount of eduction beyond high school.

3.3. Estimation Procedure

We use the historical data described in Section 3.2 to estimate the transition matrix for the discrete-time Markov chain model. Our estimation procedure follows Craig and Sendi (2002). Even though our data are collected on a monthly interval, we use a cycle length of one year. This choice allows us to accommodate seasonality in promotions observed in the data. Since we have complete observations for three full years, we use a version of the estimation procedure where the observation intervals coincide with cycle length (Craig and Sendi, 2002).

An important assumption we make is that the transition probabilities are stationary over time.This assumption is not unreasonable as military operations are fairly consistent during the three year observational period. With the stationarity assumption, the observed transitions in all three years can be pooled together to form an observed one-year transition count matrix.

We also manipulate the data to more honestly reflect involuntary separation. Recall that in our state definition, there is a maximum time-in-grade allowed for each grade, which is also called a retention control point (RCP). When a RCP is reached for an enlisted individual, all our original records show voluntary separation. In reality, the voluntary separation decision is made with known information about promotion potential and a barrier for continued transition within the current grade. Therefore, data for individuals who reached an RCP are treated as involuntary losses even if a record depicts a voluntary loss.

For notational simplicity, let m denote the number of states. For our model, m is the constant 79. The first step in our estimation procedure is to produce the transition count matrix $C = [c_{ij}]$, which is an $m \times m$ matrix. The entry c_{ij} represents the total number of yearly transitions from state i to state j. We assume without loss of generality that the states are ordered such that the last two states are voluntary and involuntary losses, which are also the absorbing states.

Let the $m \times m$ matrix $M = [\theta_{ij}]$ denote the transition matrix, where θ_{ij} is the probability of moving from state i to state j by the end of a cycle. With the observed count matrix C, the maximum likelihood estimate of the transition probability θ_{ij} is the row proportions of the counts

$$\theta_{ij} = \frac{c_{ij}}{\sum_{j=1}^{m} c_{ij}}, \quad \forall \; i, j. \tag{1}$$

We also have $\theta_{m-1,m-1} = \theta_{m,m} = 1$, because the last two states are absorbing states. Since we assume there is no demotion, entries below the diagonal in M are equal to zero. Moreover, the transition probabilities in the diagonal for non-absorbing states are also equal to zeros since time-in-grade uses the same measurement as the cycle length.

3.4. The Transition Matrix

The state progression of an enlisted personnel follows the Markov chain described in Figure 1 with transition matrix M. There are four distinct types of transitions:

- Remain in the same grade (time-in-grade increases by 1)
- Promote to the next grade (time-in-grade equals 0)
- Exit the system due to involuntary separation
- Exit the system due to voluntary separation

Since the estimated transition matrix M is a 79×79 matrix, we choose not to report the full details in the paper. Instead, we report slices of the estimates to give a flavor of the result.

3.4.1 Continuation within Grade

The calculation for continuation within the same grade requires tracking individual records from one year to the next. Moreover, individuals are only included in the continuation counts if they are in the initial data set at time 0. Figure 2 shows the probability of an individual continuing within the same grade from one year to the next. For example, a SL1 individual has an 87.8% probability of beginning and ending their second year in grade without advancing in grade due to promotion or leaving the Army. The probability decreases to 43.4% during the fourth year in SL1.

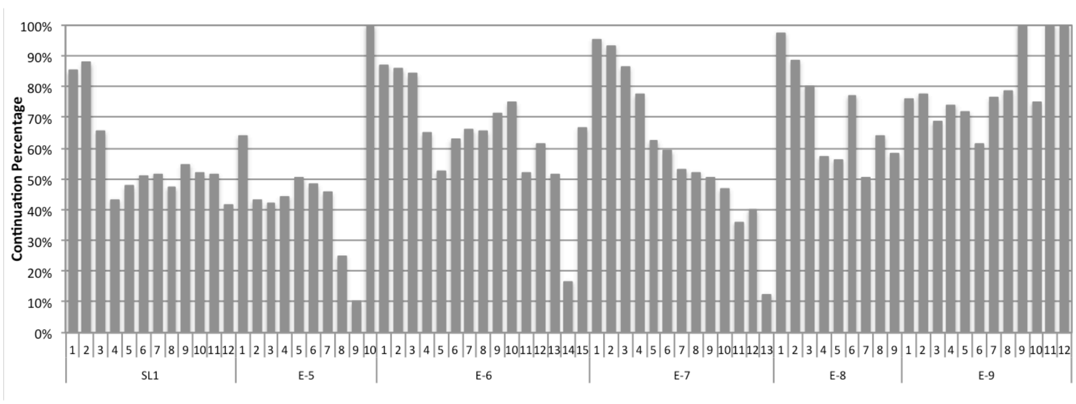

Figure 2 Continuation probabilities within grade

3.4.2 Promotions

Promotion counts from state to state are calculated in the same manner as voluntary and involuntary losses. The Army applies retention control points to each grade, which force an individual to exit the military when reached. Retention control points are enforced based on an individual's time-in-service and vary by grade. Figure 3 shows the decrease in promotion probability by grade (to E6 and above) and Figure 4 shows the probability of promotion for each year of service for semi-centralized or centralized Army promotion systems to E6 through E8. The time-in-grade axis reflects the number of years spent in the previous pay grade.

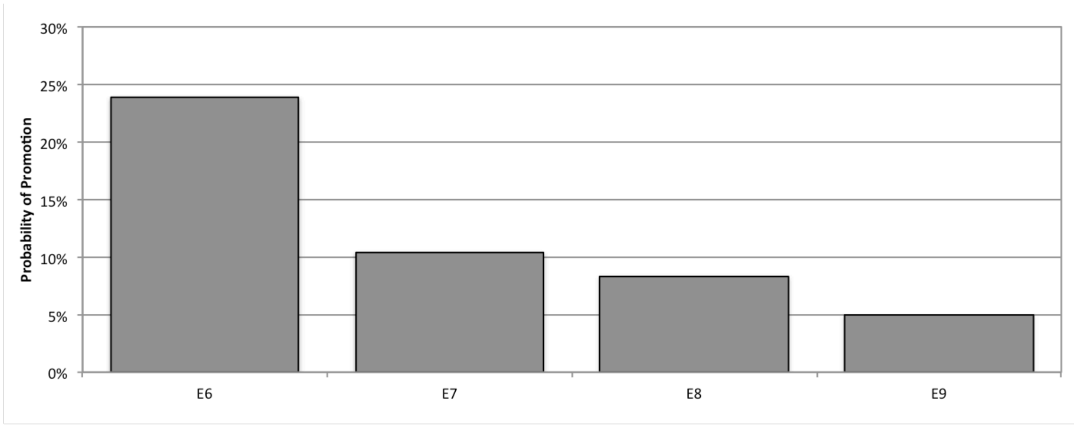

Figure 3 **Promotion probabilities by grade**

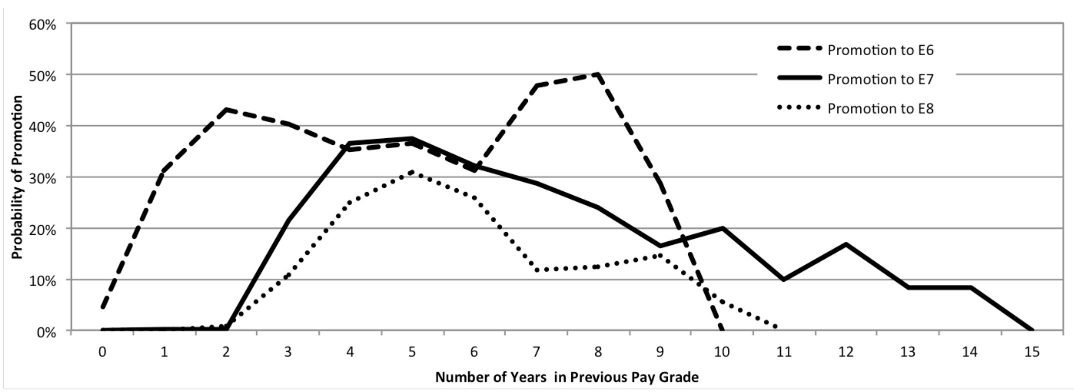

Figure 4 **Probability of promotion to grade (by time-in-grade of previous grade)**

3.4.3 Involuntary Losses

Involuntary losses include all administrative and adverse losses and can be calculated from the loss transactions for the relevant periods. Involuntary losses occur primarily during the grades of E5 and lower. Involuntary loss rates from fiscal years 2006 through 2009 are shown in Table 2.

3.4.4 Voluntary Losses

Voluntary losses are composed of two loss sub-categories: non-disability retirement (NDR) and expiration of term of service (ETS). Voluntary losses occur at all grade levels with the rates increasing at higher grade and higher time-in-grade combinations. This intuitive result is depicted in Figure 5 where the highest voluntary loss rates occur at grade E7 with greater than 10 years in grade. Voluntary loss counts are generated in the same manner as involuntary losses as transaction counts are merged with inventory data. It can be seen from Figure 6 that the probability of an individual choosing to leave voluntarily outweighs the probability of any type of involuntary loss at grade E5 and beyond. It is also evident that loss probabilities from both major categories are noticeably low ($< 5\%$) at the grade of E6; particularly likely for voluntary separations due to the associated time-in-service requirements and the proximity to retirement eligibility.

Table 2 CMF 11 Average Involuntary Loss Rates FY06-09

TIG	E4	E5	E6	E7	E8	E9
0	27.6%	4.7%	2.4%	2.3%	0.9%	0.0%
1	11.6%	7.2%	3.6%	2.2%	0.8%	0.4%
2	10.0%	10.2%	4.0%	2.5%	0.9%	1.4%
3	14.7%	14.2%	4.0%	1.8%	0.7%	1.3%
4	18.0%	12.9%	4.9%	1.2%	1.0%	6.1%
5	20.2%	10.9%	4.0%	0.9%	0.8%	0.0%
6	18.0%	19.4%	3.5%	1.4%	2.3%	0.0%
7	20.9%	8.4%	2.5%	2.6%	0.0%	0.0%
8	22.5%	20.1%	0.7%	2.2%	0.0%	0.0%
9	21.1%	0.0%	0.9%	1.3%	RCP	0%
10	33.1%	RCP	0.7%	7.6%	-	0%
11	39.5%	-	2.3%	5.0%	-	RCP
12	RCP	-	2.9%	0.0%	-	-
13	-	-	12.5%	RCP	-	-
14	-	-	0.0%	-	-	-
15	-	-	RCP	-	-	-

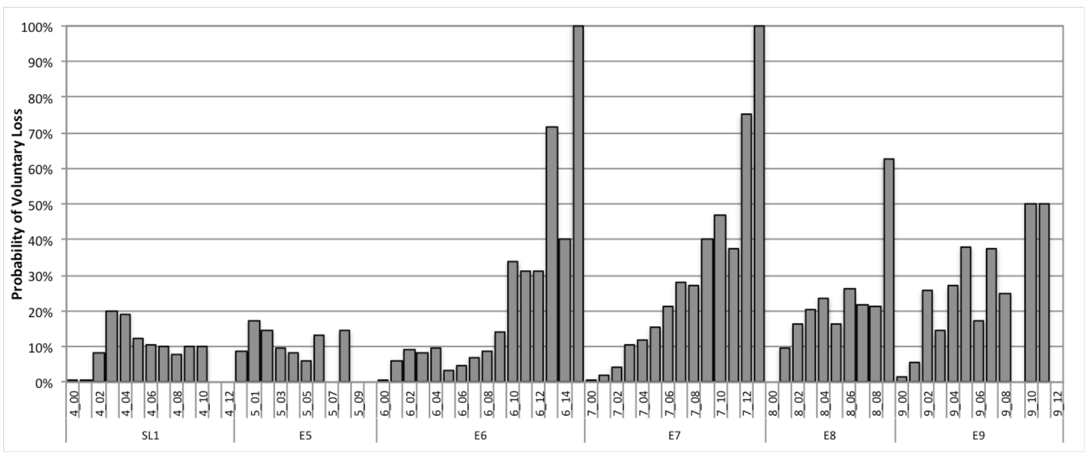

Figure 5 Probabilities of voluntary loss from each state in a period

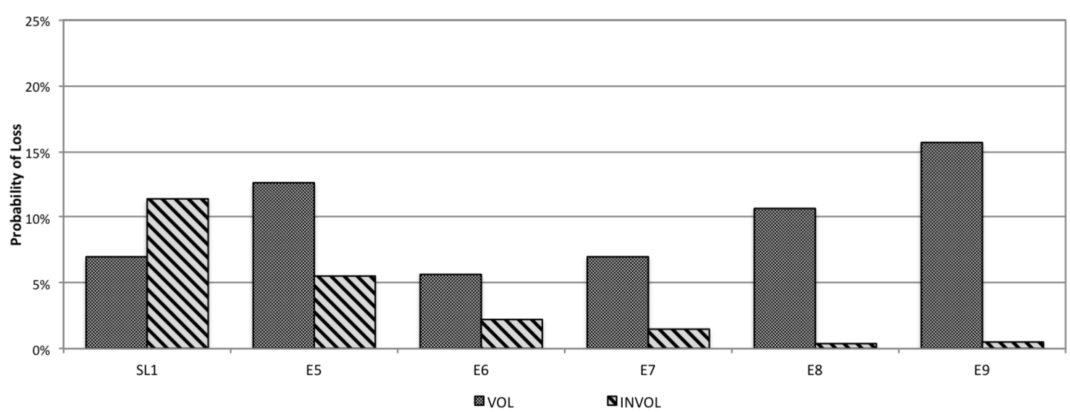

Figure 6 Comparison of loss probabilities in a single period by grade

4 AN AGGREGATED MARKOV CHAIN MODEL

The Markov chain model presented in Section 3 defines state as grade and time-in-grade combinations. As we showed in Section 3, the model allows us to make many detailed observations. This section considers an aggregated model where the states are grouped to the grade level. Grade is arguably the most widely used unit of analysis for military personnel planning. The aggregated model can be used to answer many managerially relevant questions for different grades. We also use the aggregated Markov chain to validate our model using a chi-squared goodness-of-fit test. We return to the expanded Markov chain in Figure 1 in subsequent section when we incorporate a dynamic programming model for military personnel retention.

4.1. State Aggregation by Grade

We aggregate states by grade to form a new Markov chain. The states for each grade level forms a state in the new model. For example, $(SL1\_0, SL1\_1, \ldots, SL1\_12)$ makes up the state s_1. In this fashion, The six Grade categories (Skill Level 1 through E9) make up states s_1 to s_6. s_7 and s_8 represent voluntary and involuntary loss states, respectively. Figure 7 shows the state transition diagram for the aggregated Markov chain.

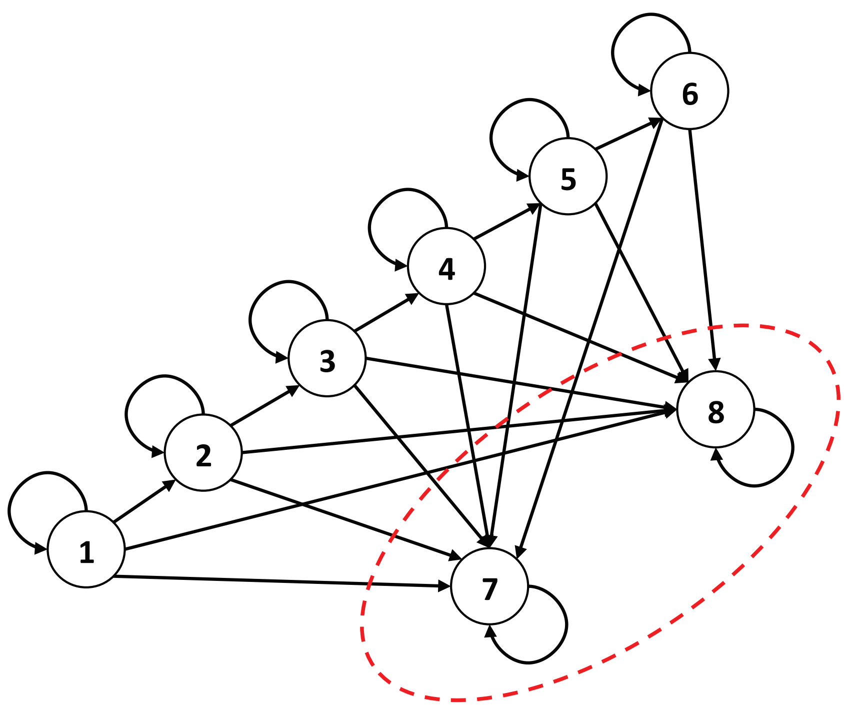

Figure 7 State Transition Diagram for the Aggregated Markov Chain Model

We use an 8×8 matrix $P = [p_{ij}]$ to represent the transition matrix of the aggregated Markov chain, where p_{ij} denotes the transition probability from state s_i to state s_j. The transition matrix P can be generated from the transition matrix M. The idea is that transitions within the same grade level are considered transitions back to the same state in the aggregated model, while only promotions or voluntary/involuntary separations are considered transitions between states. Clearly, we have $p_{77} = p_{88} = 1$ because states s_7 and s_8 are absorbing states. The resulting transition matrix P is given by

$$P = \begin{bmatrix} 0.6691 & 0.1588 & 0 & 0 & 0 & 0 & 0.0646 & 0.1075 \\ 0 & 0.5335 & 0.2542 & 0 & 0 & 0 & 0.1495 & 0.0627 \\ 0 & 0 & 0.6992 & 0.1110 & 0 & 0 & 0.1408 & 0.0490 \\ 0 & 0 & 0 & 0.8100 & 0.0800 & 0 & 0.1000 & 0.0100 \\ 0 & 0 & 0 & 0 & 0.8150 & 0.0499 & 0.1282 & 0.0069 \\ 0 & 0 & 0 & 0 & 0 & 0.8380 & 0.1554 & 0.0066 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1.0000 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1.0000 \end{bmatrix}. \tag{2}$$

4.2. Validation and Bootstrap Confidence Intervals for P

The transition matrix for the Markov chain is validated using a Chi-square goodness-of-fit test applied to our sample data (Montgomery and Runger, 2003). The Chi-square goodness-of-fit test provides a useful measure of fit to our data as described by Anderson and Goodman (1957). A test of the difference between the observed counts at each state in an initial period 0 and final period 3 are evaluated using the transition probability matrix P.

Let O be the vector of initial observed counts and O' be the vector of counts after three period. Under the estimated transition matrix P, the vector of expected count is given by $E = O^T P^3$, where E_i denotes the expected count in state i. The test statistic is given by

$$\chi_0^2 = \sum_{i=1}^{8} \frac{(O_i' - E_i)^2}{E_i} \tag{3}$$

Table 3 lists the values of vectors O, O', and E. The test statistic $\chi_0^2 = 6.7604$, which is less than $\chi^2_{0.05,7} = 14.067$. Hence the test is insignificant at P-value 0.05, implying that the transition matrix is a good fit.

Table 3 Expected and observed counts for the goodness-of-fit test

Period	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8
O	7752	1869	1692	779	258	54	0	0
O'	2259	1690	1762	862	285	67	3060	2419
E	2322.1	1624.7	1718.2	877.9	302.2	66.6	3064.6	2427.0

Next, we use the bootstrap method to assess the uncertainty of the maximum likelihood estimate and to construct confidence intervals for the transition matrix. Let $\tilde{C}$ be the aggregated count matrix. Let t_r denote the total number of transitions for row r. We bootstrap row r by sampling t_r transitions with replacement from the observed t_r transitions. For example, since there are 69,835 transitions in the first row of $\tilde{C}$, 69,835 draws are taken from the corresponding empirical distribution to form a new set of transition counts for row 1. This process is repeated for each row until a new count matrix C^* is formed for each bootstrap sample. The new set of transition counts for row r of the count matrix is used to construct a new set of transition probabilities and a new transition probability matrix, P^*. The process is repeated for a number of bootstrap samples until a confidence interval can be constructed for the function.

Table 4 reports the confidence interval for the transition matrix based on 500 bootstrap samples. The confidence intervals are created from a collection of bootstrapped transition matrices, which approximate sampling distributions. The goal of constructing bootstrap confidence interval is to calculate dependable confidence limits for a parameter of interest θ from the bootstrap distribution of an estimator $\hat{\theta}$. Having chosen to use a particular $\hat{\theta}$, bootstrapping is a general methodology to measure how accurate $\hat{\theta}$ is as an estimator of θ.

We can use bootstrapping to assess the uncertainty of each entry in the transition matrix as well as any function of the transition matrix. While sensitivity analysis is a very helpful technique to investigate the behavior of a Markov model, Craig and Sendi (2002) claim that it should not be used to construct a confidence interval because it does not take into account model restrictions and complex dependency of all the transition probabilities. To assess the precision of this estimate, it substitutes considerable amounts of computation in place of theoretical analysis and is more accurate than standard errors when dealing with nonparametric confidence intervals (DiCiccio and Efron, 1996).

Table 4 Confidence intervals for transition probabilities based on 500 bootstrap samples

Initial State	s_1	s_2	s_3	s_4
s_1	(0.66899, 0.66929)	(0.15871, 0.15896)	-	-
s_2	-	(0.53333, 0.53390)	(0.25387, 0.25438)	-
s_3	-	-	(0.69906, 0.69961)	(0.11078, 0.11119)
s_4	-	-	-	(0.80959, 0.81028)
s_5	-	-	-	-
s_6	-	-	-	-

Initial State	s_5	s_6	s_7	s_8
s_1	-	-	(0.06452, 0.06468)	(0.10733, 0.10753)
s_2	-	-	(0.14929, 0.14971)	(0.06262, 0.06290)
s_3	-	-	(0.14046, 0.14088)	(0.04888, 0.04915)
s_4	(0.07978, 0.08027)	-	(0.09979, 0.10033)	(0.00990, 0.01007)
s_5	(0.81405, 0.81525)	(0.04975, 0.05042)	(0.12776, 0.12888)	(0.00682, 0.00707)
s_5	-	(0.83741, 0.84006)	(0.15339, 0.15603)	(0.00627, 0.00684)

4.3. Interpretations

One of the benefits of the Markov chain model is that it allows us to draw conclusions on long-run statistics, even though our model is only estimated from data with a limited time range. In this section, we expect some interesting interpretations from our model.

By partitioning the matrix P as

$$P = \left[\begin{array}{c|c} Q & R \\ \hline 0 & I \end{array}\right], \tag{4}$$

we expose a useful concept in the analysis of Markov chains, the fundamental matrix. From standard Markov chain theory, we can compute the *fundamental matrix* as

$$N = \begin{bmatrix} 3.0221 & 1.0287 & 0.8694 & 0.5079 & 0.2196 & 0.0677 \\ 0 & 2.1436 & 1.8115 & 1.0583 & 0.4577 & 0.1410 \\ 0 & 0 & 3.3245 & 1.9422 & 0.8399 & 0.2587 \\ 0 & 0 & 0 & 5.2632 & 2.2760 & 0.7011 \\ 0 & 0 & 0 & 0 & 5.4054 & 1.6650 \\ 0 & 0 & 0 & 0 & 0 & 6.1728 \end{bmatrix}. \tag{5}$$

The (i, j)-th entry of N, n_{ij}, is the expected number of transitions (visits) to state j before being absorbed, starting in state i. The fundamental matrix extends the interpretation of our data. For example, the matrix NR provides probability estimates of being absorbed into each of the absorbing states, s_7 and s_8. By definition, the fundamental matrix assumes an infinite number of transitions. The career path of an enlisted personnel has a finite timeline, typically 30 years. Therefore, we expect that computations based on the fundamental matrix has certain level of error. In Appendix B, we show that the error from using an infinite summation is very small. Similarly, we formulate the variance of N using Markov chain theory, which is given by

$$V = \begin{bmatrix} 6.1108 & 2.3234 & 4.1552 & 4.5804 & 2.1065 & 0.7630 \\ 0 & 2.4515 & 6.9516 & 8.9618 & 4.2805 & 1.5795 \\ 0 & 0 & 7.7276 & 14.7298 & 7.5344 & 2.8682 \\ 0 & 0 & 0 & 22.4377 & 17.1490 & 7.4624 \\ 0 & 0 & 0 & 0 & 23.8130 & 16.1183 \\ 0 & 0 & 0 & 0 & 0 & 31.9311 \end{bmatrix}. \tag{6}$$

where each element of V provides the variance for the number of periods one expects to be in state j given a start in state j.

Figure 8 shows the probability of remaining in the Army each year over a horizon of 30 years when starting in state s_1 (Skill Level 1). This graph is interesting because it allows us to see how changes to short-term behavior can affect long-term continuation rates. From Figure 8, we can see that there is only a 45% probability that an enlisted individual will remain in the Army beyond 4 years. Only 17% can be expected to remain in the Army for 10 years. Whether these rates are acceptable is a strategic decision; however, it is useful to know whether short term policies shift the curve in a desirable direction.

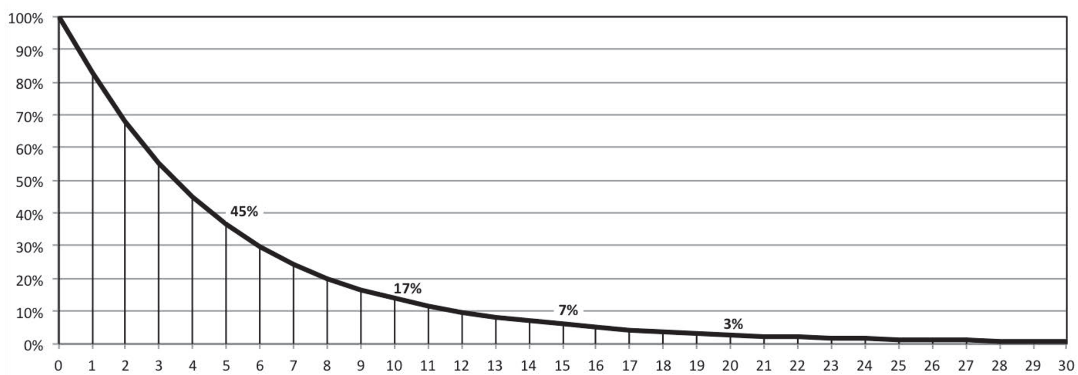

Figure 8 **Continuation rates over 30 year time horizon**

From the fundamental matrix (5), the expected amount of time an individual spends in skill level 1 is 3.02 years. The expected amount of time an individual spends as an E7 is 0.51 years. Even though this information is interesting, it is not intuitive because it incorporates the probability of not reaching the state. A somewhat related, but perhaps more interesting question, is the expected amount time someone spends in grade E7 given he/she reaches the grade. Manipulations of the fundamental matrix N can be used to answer a number of such questions. Such an approach is taken by many authors, including Doyle and Snell (2000) and Cyert et al. (1962).

First, the row sum of matrix N,

$$T = \begin{bmatrix} 5.7153 \\ 5.6121 \\ 6.3652 \\ 8.2402 \\ 7.0704 \\ 6.1728 \end{bmatrix} \tag{7}$$

gives the expected number of periods before being absorbed from each state. The interpretation of T is quite intuitive. The expected number of cycles or periods before being absorbed into a loss state is 5.7 years.

This expectation increases with an increase in grade due to the increased probability of remaining until the retirement threshold. The peak is upon reaching the grade of E7 when an individual is expected to remain for 8.2 years until being absorbed, presumably as a voluntary loss category of retirement. Beyond the grade of E7, the expected time until absorption decreases due to the finite time horizon of mandatory retirement or reduced incentive for delaying retirement.

Second, the probability of moving from each of the transient states to an absorbing state is given by the matrix NR, which is shown below:

$$NR = \begin{bmatrix} 0.5609 & 0.4390 \\ 0.7619 & 0.2378 \\ 0.8102 & 0.1898 \\ 0.9270 & 0.0730 \\ 0.9517 & 0.0483 \\ 0.9593 & 0.0407 \end{bmatrix}. \tag{8}$$

The elements in NR has interesting interpretations. For example, the probability of eventually becoming a voluntary loss 0.5609 is higher than the probability of becoming an involuntary loss 0.4390 when starting from the initial state, s_1. Despite the significantly higher involuntary loss rates of skill level 1 individuals, these probabilities account for expected future transitions to other states. We can see that the disparity between absorbing into the two loss states diverges significantly as the grade level progresses.

It is possible to use bootstrap method to estimate the confidence intervals for the absorption times from different starting states. Consider state s_1, which has expected absorption time of 5.7153 years. The results of 500 bootstrap samples are shown in Figure 9. Using an two-tailed confidence interval, the 95% confidence interval for the expected absorption time is $[5.715, 5.721]$.

The probability of ever making a transition into state j starting in state i is given by

$$f_{ij} = \frac{n_{ij}}{n_{jj}}$$

where n_{ij} is the (i,j)-entry of N (Ross, 1996). Let $F = [f_{ij}]$. We have

$$F = \begin{bmatrix} 0.6691 & 0.4799 & 0.2615 & 0.0965 & 0.0406 & 0.0110 \\ 0 & 0.5335 & 0.5449 & 0.2011 & 0.0847 & 0.0228 \\ 0 & 0 & 0.6992 & 0.3690 & 0.1554 & 0.0419 \\ 0 & 0 & 0 & 0.8100 & 0.4211 & 0.1136 \\ 0 & 0 & 0 & 0 & 0.8150 & 0.2697 \\ 0 & 0 & 0 & 0 & 0 & 0.8380 \end{bmatrix}. \tag{9}$$

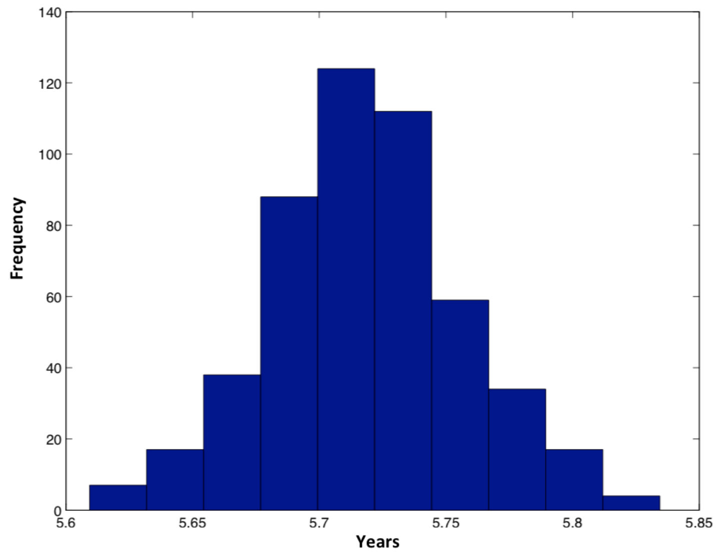

Figure 9 **500 bootstrap samples of absorption time from state 1**

So, what is the probability of an enlisted personnel ever reaching the grade of E9? Based on F, the probability is close to 1%. In fact, an individual has less than a 5% probability of reaching the grade E8. Perhaps equally indicative of the likelihood of reaching the grade of E9 is that the probability of reaching this state is still less than 5% (0.0419) when an individual reaches the grade E6.

5 A DYNAMIC PROGRAMMING MODEL FOR MILITARY PERSONNEL RETENTION

Based on the Markov chain estimated in Section 3, we construct a Markov decision process (MDP) model for a military personnel's stay-or-leave decision at each point of his/her career. The model captures career path dynamics, including military pay and career advance opportunities. We assume that in making a stay-or-leave decision, a military personnel weighs the expected military pay, including possible retirement benefits, against the expected pay from civilian job opportunities. Even though not explored in the present paper, our model can be used to shape and target incentives to military personnel.

The model is formulated as a finite-horizon MDP using an optimal stopping framework (Puterman, 2005). The state i is the grade and time-in-grade combination in each period. Important assumptions are made in order to develop the model, such as the assumption that when an individual decides to leave the military, they cannot return to the military at later periods in the time horizon. We assume individuals do not change their skills during the time horizon and all military service is conducted while on active duty. The reclassification rate is historically less than 1%, even during periods of volatile force shaping drawdowns of enlisted strength. Although promotions, involuntary losses and retention decisions occur at monthly intervals, following Section 3, a simplifying assumption is made to define the period length as one year. The objective of the model is to maximize total expected payoff over the finite horizon.

Since leaving the military is modeled as a decision, the transition probabilities in our model follows the transition probabilities estimated in Section 3, conditional on not being a voluntary loss. Let q_{ij} denotes the probability of transitioning from state i to state j in the following period, we have

$$q_{ij} = \frac{\theta_{ij}}{1 - \sum_{j \neq \mathbf{VL}} \theta_{ij}}, \qquad \forall i, j. \tag{10}$$

The state space includes all states for the Markov chain estimated in Section 3 except the voluntary loss state, which we denote by VL. Appendix C provides a notation dictionary.

5.1. Estimating Civilian Pay

A key input the MDP model is the expected civilian pay in the event of choosing to leave the military. As mentioned earlier, it is not reasonable to assume initial civilian pay is equal to the last military pay. We assume the MDP model is constructed for a specific military skill, which is cross-mapped to comparable civilian occupations using the Occupational Information Network (O*NET).

Table 5 illustrates a distribution of potential civilian employment for an enlisted personnel with a military occupation of Infantry. The second column list the comparable civilian occupations for an infantryman. The third column reports the number of workers employed in each of the civilian occupation. The empirical distribution of the different occupations is reported in the second-to-last column, where we use α_x to denote the proportion of workers employed in occupation x. The compensation for each occupation is reported as a multiple, δ_x of the military compensation $m_{i,t}$. By using a multiplier to military compensation rather than a constant value (i.e median of civilian occupation x), we account for an increase in expected pay as military experience increases. It is reasonable to assume that two individuals with the same military occupation specialty and expected civilian occupation can expect different levels of pay attributed to technical acumen and leadership experience.

Table 5 Cross-Mapping of Military to Civilian Skills (Enlisted Infantryman)

i	Civilian Occupation	Workers (1,000)	α_x (%)	δ_x
1	Training and Development Specialists	226	3.5	2.10
2	First-Line Supervisors of Police and Detectives	88	1.4	2.13
3	Correctional Officers and Jailers	312	4.8	1.32
4	Police Patrol Officers	578	8.9	1.77
5	Security Guards	572	8.8	0.95
6	First-Line Supervisor of Construction Trades and Extraction Workers	517	8.0	1.75
7	Construction Laborers	937	14.4	1.08
8	Operating Engineers and Other Construction Equipment Operators	315	4.9	1.36
9	First-Line Supervisor of Production and Operating Workers	574	8.8	1.63
10	Heavy and Tractor-Trailer Truck Drivers	2,368	36.5	1.25

5.2. Military Compensation and Retirement pay

Individuals have the choice to stay in the military and continue to receive military pay or leave the military before or after retirement eligibility. Enlisted personnel are compensated based on a military pay scale, which increases by pay-grade and years-of-service. Individuals receive additional compensation pay that includes a housing compensation based on local civilian housing markets and a basic allowance intended to pay for food. Individuals who leave the Army prior to retirement eligibility do not receive any pension compensation. However, after reaching retirement eligibility, individuals can leave the military and receive pension.

Upon retirement, a military personnel receives a proportional amount of the so-called high-3 average retirement compensations (Hall, 2009). Let $m_{i,t}$ denote the base military compensation for an individual in state x in period t. The high-3 average retirement compensations is calculated as

$$\widehat{m}_{i,t} = \frac{m_{i'',t-3} + m_{i',t-2} + m_{i,t-1}}{3}, \tag{11}$$

where i' and i'' are the states in two years preceding retirement. Note that the the calculation of retirement compensation does not violate the Markovian property because it is reasonable to assume that an individual can be promoted with less than or equal to one year of time-in-grade, but not in two consecutive years. Therefore, the pay scale can be backed out from the state upon retirement.

The proportion of the high-3 average compensations received by an individual upon retirement depends on the total number of years in service. Let w_t denote the proportion received by an individual with total number of years of service t. Since military retirement pay is only earned if a Soldier serves a total of 20 years of active duty, $w_t = 0$ for $t < 20$. The standard retirement format (Mattock et al., 2012), an individual receives a base 50% plus 2.5% for each additional year served beyond 20 years. Hence w_t is defined as

$$w_t = \begin{cases} 0, & \text{if } t < 20, \\ 0.5 + 0.025(t-20), & \text{if } t \geq 20. \end{cases} \tag{12}$$

5.3. Value Function

First, consider the value of remaining in the Army for one more year. The expected compensation in period t for an individual in state x is given by

$$m_{i,t} + c_{i,t} + \beta \sum_{j=1}^{S} P_{i,j} V_{t+1}(j), \tag{13}$$

where S represents the total number of states and β denotes a personal discount factor such that $\beta = \frac{1}{1+\gamma}$ and γ is the personal discount rate.

In case the decision is to leave the army, we can use (14) to represents an individual's future total cash flow from that decision

$$w_t \widehat{m}_{i,t} \sum_{k=t}^{\tau-\eta} \beta^{k-t} + \delta_x m_{i,t} \sum_{k=t}^{T-1} \beta^{k-t}(1 + \xi_{k-t}) \tag{14}$$

The first term in the above cash flow equation represents an individual's expected retirement pay at $\tau - \eta$, where τ represents life expectancy and η represents an individual's age upon entering the Army (we assume 21 years). To simplify the model, we assume an individual's life expectancy to be the same; however, the model may be refined with the use of applicable data pertaining to life expectancies by selected attribute to better represent the military retirement time horizon. The second part of equation takes into account individual's initial civilian sector pay plus civilian pay raises ξ_{k-t} and time T is considered the ceiling for full retirement benefits.

Now, we are ready to construct the value functions. In the terminal period T, the only viable option is to retire, and no civilian pay will be received. By simplifying the first term in (14), we obtain the value function

$$V_T(i) = w_T \widehat{m}_{i,T} \frac{1-\beta^{\tau-\eta+1}}{1-\beta}. \tag{15}$$

For $t < T$, the decision is whether to stay in the army. We also incorporate the possibility of becoming an involuntary loss. For notational purposes, let IL denote the state of involuntary loss. Then $q_{i,\mathrm{IL}}$ is the probability of becoming an involuntary loss in state x. The optimality equations are given by

$$\begin{aligned} V_t(i) = (1-q_{i,\mathrm{IL}}) E \Bigg[\max \Bigg\{ & m_{i,t} + c_{i,t} + \beta \sum_{j \neq \mathrm{IL}} \frac{q_{ij} V_{t+1}(j)}{(1-q_{i,\mathrm{IL}})}, \\ & w_t \widehat{m}_{i,t} \sum_{k=t}^{\tau-\eta} \beta^{k-t} + \delta_x m_{i,t} \sum_{k=t}^{T-1} \beta^{k-t}(1+\xi_{k-t}) \Bigg\} \Bigg] \\ & + q_{i,\mathrm{IL}} \left(w_t \widehat{m}_{i,t} \sum_{k=t}^{\tau-\eta} \beta^{k-t} + \delta_x m_{i,t} \sum_{k=t}^{T-1} \beta^{k-t}(1+\xi_{k-t}) \right), \qquad \forall i, t < T. \end{aligned} \tag{16}$$

The right-hand side of equation (16) is comprised of two parts. The first part considers the probability of not being an involuntary loss multiplied by the expected maximum value of staying in the military or leaving the military. The second part considers the probability of an involuntary loss multiplied by the value of leaving the Army. The value of of staying in the military, $m_{i,t} + c_{i,t} + \beta \sum_{j \neq \mathrm{IL}} \frac{q_{ij} V_{t+1}(j)}{(1-q_{i,\mathrm{IL}})}$, consists of military compensation plus additional compensation plus the conditional value function in the next period. The value of leaving the Army, $w_t \widehat{m}_{i,t} \sum_{k=t}^{\tau-\eta} \beta^{k-t} + \delta_x m_{i,t} \sum_{k=t}^{T-1} \beta^{k-t}(1+\xi_{k-t})$, consists of the value of future retirement pay plus the value of future civilian pay.

5.4. Optimal Policy Index

One method of evaluating the stay-or-leave decision between individuals is to compare the expected values of compensation for each decision at different states and time periods. The total expected value of compensation will change in respect to each variable and it is worth noting the scale in which the values change. For the following computational experiments, we use a life expectancy of 75 years. We assume that the average age of individuals entering the Army is 22 years, the total time an individual works in the military and civilian sector is 40 years, and the civilian pay raise per period is 2.5%.

Applying the assumptions above to the value function, the following time parameters are used for computation: $\tau = 75$, $\eta = 22$, and $T = 40$. The parameters $\beta = 0.8696$ and $\xi = 0.025$ relate to the personal discount factor and civilian pay raise. Probability matrices are used for q and P. Because of variations across occupations in civilian pay, the stay-or-leave decision depends on the realization of civilian occupation.

We use a policy index to represent the propensity to leave. The optimal policy index is derived by calculating the values of leaving and staying in the military for each cross-mapped civilian occupation at each state and time period using equation (16). We use an empirical distribution for civilian compensation parameters. The index represents the proportion of instances that the value of leaving exceeds the value of staying given potential comparable civilian occupations.

For example, an optimal policy index of 0.3 would indicate a 30% probability that the expected compensation value of leaving is greater than the value of staying, given an empirical distribution of expected civilian pay. A depiction of the indices for each state-time combination within the Infantry skill are shown in Figure 10. Darker shading levels represent higher index values and a greater propensity to leave. The shading clearly depicts the impact of retirement benefits after 20 years as well as the forced decisions at each retention control point.

Individuals are characterized by different attributes and preferences which lead to a variance in personal discount rates. The policy indices in Figure 10 are calculated using a constant personal discount rate of 15%. This personal discount rate is consistent with the average rate simulated by Asch et al. (2008), which they found to be the best fit and most reasonable approximation of enlisted Army behavior. It is worth noting that a personal discount rate of 15% is significantly lower than the mean nominal enlisted discount rates reported by Warner and Pleeter (2001). Moreover, Warner and Pleeter (2001) suggest that there is a significant difference between officer and enlisted discount rates that is significantly attributable to observable demographic and characteristics differences.

Table 6 shows how changes to the personal discount factor affect the average index values for a selection of grade and YOS combinations. The average index is calculated over all possible time-in-grade values. Intuitively, higher discount rates result in increased compensation values for leaving. Note that an individual with grade E-5 and 7 years-of-service has an average index value of 0.090 when $\beta = 0.99$. Therefore, they have strong propensity for continued service. When the personal discount factor decreases to 0.75, the index increases to 0.373 and the propensity decreases considerably.

Figure 11 shows the change in expected value for each combination of state and years-of-service for two different discount factors. The impact of the personal discount rate is most evident in the states associated with lower pay grades.

6 INTEGRATING DYNAMIC PROGRAMMING RESULTS IN PREDICTIVE MODELING

The dynamic programming model in the previous section tries to rationalize the stay-or-leave decision for military personnel by taking into account long-term payoffs from stay-or-leave decisions. There are two

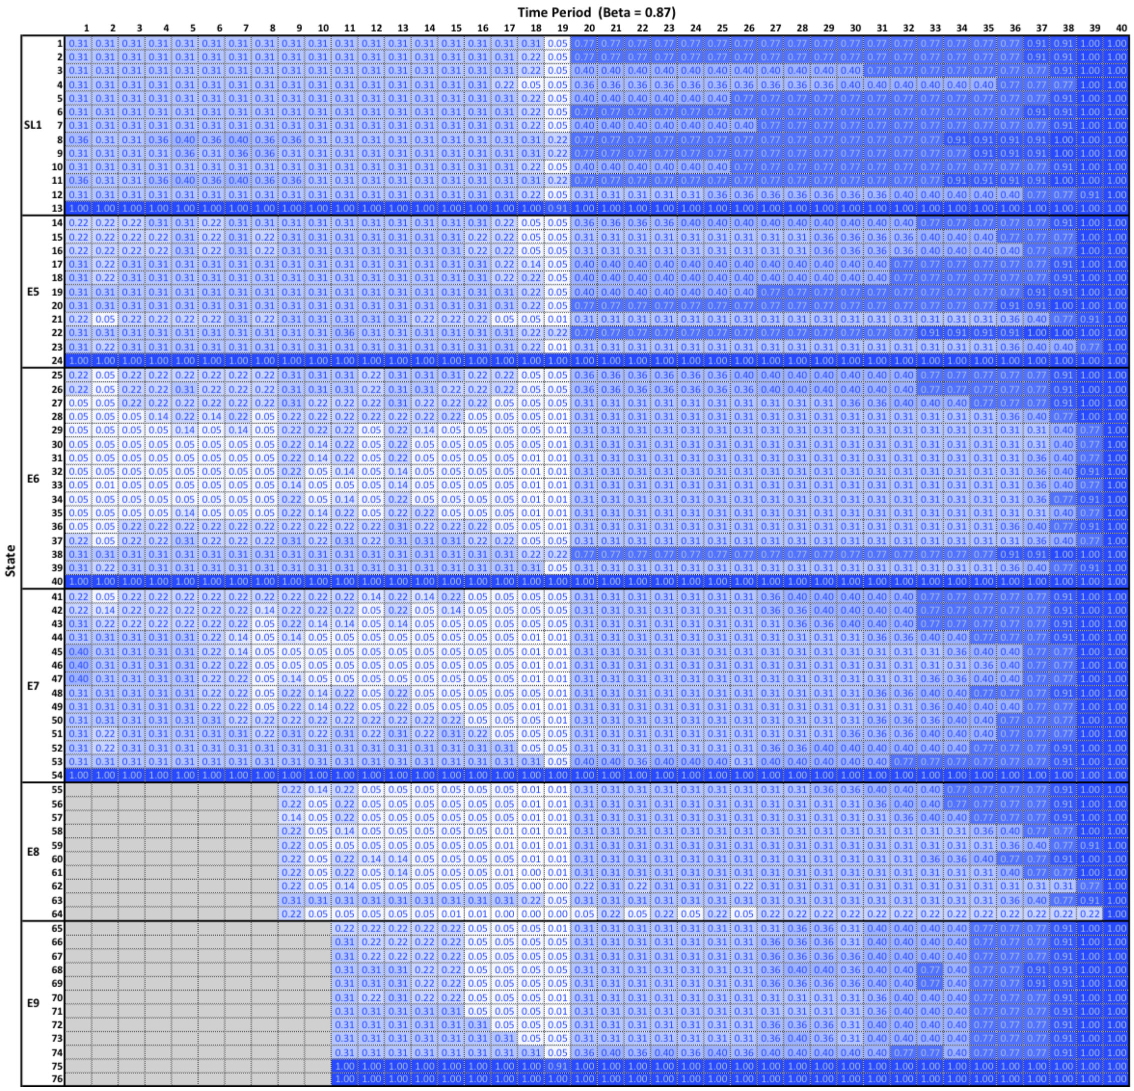

Figure 10 Shading represents the propensity to leave with darker shading and higher values indicating a greater value for leaving.

Table 6 Average Index values for all Time-in-Grade.

Grade	YOS	Personal Discount Factors (β)					
		0.99	0.95	0.90	0.85	0.80	0.75
E4	5	0.077	0.085	0.246	0.277	0.315	0.415
E5	7	0.090	0.090	0.191	0.273	0.282	0.373
E5	9	0.090	0.090	0.218	0.273	0.336	0.445
E6	12	0.088	0.113	0.206	0.275	0.331	0.419
E7	15	0.179	0.186	0.271	0.336	0.450	0.485

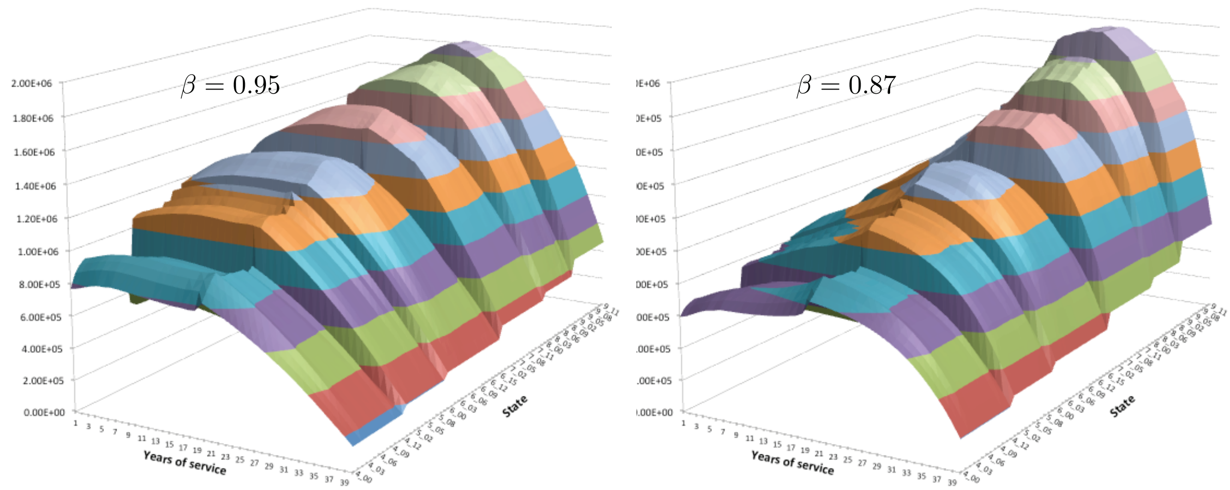

Figure 11 **Value by State and Time Period**

major differences between the model and more commonly used classification approach. First, the model captures long-run personnel dynamics through its use of the estimated transition probabilities. Second, it does not use demographic variables typically employed in classification approaches. Because of these difference, we believe the result from the model can be used to supplement popular predictive modeling approaches. In this section, we discuss how to integrate the results of the dynamic programming model with one of the most popular binary classification method, the logistic regression.

We first discuss a logistic regression model based on the same data as the Markov chain model in Section 3. The variables used in the binary logistic regression equation are described earlier in Table 1. In addition to variables included in the original database we construct additional endogenous and exogenous variables such as *age*, *months since last deployment*, and *employment rate of change*. The employment exogenous variable is constructed from data obtained from the Bureau of Labor Statistics by calculating a twelve month moving average of the monthly change in employment rate. Table 7 shows the Logit results. Most of the results are aligned with tuition, including the impact of experience and grade being more influential than marital status or number of dependents. The table only includes the racial category of *Black* because there is not a significant difference in behavior between *White*, *Hispanic* and *Asian/Other*.

There are two approaches to improve the classification accuracy of the logistic regression model. The first one is to incorporate the policy indices from the dynamic programming model as an additional variable in the logit model. The second one is to use a combination of the DP policy indices and the logistic regression classification as a classifier. That is, we use a weighted index $\alpha I + (1-\alpha)L$ as a classifier, where I is the DP policy index, L is the logit index, and α is the weight assigned to the DP policy index I.

Table 8 shows the cross-validation results of classifying the stay-or-leave decision using binary logistic regression. The overall results provide a marginal lift in comparison to predicting a stay decision for all individuals. Moreover, the majority of the error is a result of misclassifying the observed leave decisions.

Table 7 Estimated Logistic Regression Coefficients

	Coefficients	
Variables	Logit	Logit w/ Index
Constant	-3.025***	-4.073***
	(0.053)	(0.087)
MINOR_DEP	-0.413***	-0.417***
	(0.018)	(0.018)
E5	-0.093*	-
	(0.031)	-
E6	-2.134***	-2.003***
	(0.061)	(0.055)
E7	-2.386***	-2.573***
	(0.113)	(0.118)
E8	-1.081***	-1.392***
	(0.286)	(0.307)
E9	-2.120*	-2.560***
	(0.690)	(0.683)
AFQT	-0.001*	-0.001
	(0.001)	(0.001)
REENLIST_QY	-2.212***	-2.177***
	(0.034)	(0.034)
TIME_IN_SERVICE	0.809***	0.769***
	(0.011)	(0.011)
AGE_CENTER	0.130***	0.163***
	(0.008)	(0.008)
AGE_CENTER^2	-0.009***	-0.011***
	(0.000)	(0.000)
MARRIED	-0.564***	-0.556***
	(0.030)	(0.030)
DIVORCED	-0.686***	-0.670***
	(0.090)	(0.091)
BLACK	-0.430***	-0.440***
	(0.057)	(0.058)
EDUCATION_CENTER	0.079***	0.090***
	(0.017)	(0.017)
EMPLOYMENT	-0.005***	-0.005***
	(0.001)	(0.001)
MONTHS_SINCE_DEPLOY*DEPLOY	0.004***	0.004***
	(0.001)	(0.001)
INDEX	-	0.563***
	-	(0.035)

Standard errors in parentheses; $*\ p < .05$, $**\ p < .01$, $***\ p < .001$

The best fit logit model accurately predicts the stay decision in 89.9% of observations and classifies 58.7% of the leave decisions correctly. This finding is significant, since the focus of incentive programs should is directed at targeting individuals who are predicted to leave. Table 8 also shows that the additional parameter does not provide noticeable improvement to the overall classification accuracy. The misclassification error is not significantly improved when predicting the leave decision; however, it continues to be the primary source of error. When performing logit with the Index parameter and a cut for the decision at 0.3, the overall classification percentage slightly decreases. However, the classification accuracy for the leave decision increases to 84.5%.

Table 8 Classification Results of Stay or Leave Decisions (FY 06-09)

Method	Cut	Overall %	Stay %	Leave %
Logit	0.5	79.6%	89.9%	58.7%
Logit w/ Index	0.5	79.6%	89.8%	59.0%
Logit w/ Index	0.3	76.9%	73.1%	84.5%
Index	0.3	71.1%	99.9%	12.9%
Index	0.5	70.3%	99.9%	10.5%

Table 9 shows the forecasting accuracy for different parameter combinations using the weighted approach, where we also experimented with different cut-off value z. The best overall accuracy with a personal discount factor of 87% is achieved with weights of $\alpha = 0.40$ for the retention index value and $1 - \alpha = 0.60$ for the logit value. The best cut line between stay and leave for the weighted index is $z = 0.33$. The two scenarios shown below the top line show the accuracy levels when z and α are fixed at 0.5 respectively. A slightly better accuracy specifically regarding the leave decision is achieved with $\alpha = 0.5$ and $z = 0.30$.

Table 9 Classification Accuracy for Different Combinations of α and z Values

β	Overall Accuracy%	Stay%	Leave%	α	z
0.87	80.8%	85.1%	72.2%	0.40	0.33
0.87	79.6%	89.9%	58.7%	0.00	0.50
0.87	80.7%	84.1%	73.9%	0.50	0.30
0.95	80.8%	85.8%	70.6%	0.40	0.26
0.95	80.7%	85.5%	71.1%	0.50	0.30

Predicting a leave decision based on a dynamic programming approach filters a portion of leave predictions from the hybrid model. However, the smaller isolated set is predicted purely on expected compensation values. Predicted leaves decisions in states and time combinations with a positive value for stay versus leave compensation are not as likely to be influenced by monetary incentive.

The results of this model are significant beyond the integration of both modeling techniques because, individually, we do not significantly sacrifice predictive capability using the dynamic programming model.

In fact, we can achieve slightly better results while requiring only three factors: grade, TIG and TIS. This eliminates the need for collection and analysis of specific information about individuals. More importantly, it significantly reduces the potential perception of individual bias or discrimination that would exist if incentives are tailored by more descriptive factors. This result is important because of the nature of policy making. Many policies, particularly regarding incentive options, are naturally broad due to an aversion to policies that are based on demographic characteristics. Targeting individuals based on grade and time-in-grade states, present a framework in which previously unpalatable policy frameworks involving explanatory demographic characteristic are replaced by feasible state criterion.

7 CONCLUSIONS

We represent the enlisted career network as a discrete-time homogeneous Markov chain and estimate its transition probabilities using three years of personnel data from the US Department of Defense. The model allows us to answer a wide variety of questions related to personnel behavior. Understanding the progression of individuals throughout a career timeline can shape policies that influence the composition of the personnel inventory such as promotion criteria, retention control points, and forced reduction measures. The estimated parameters for the model are used to construct a stochastic dynamic programming model to understand individual stay-or-leave decisions.

The stochastic dynamic programming model provides an alternative to classical classification approaches to evaluating the stay-or-leave decisions. This approach can predict retention behavior more accurately than some existing approaches, such as logistic regression, which omit expected future compensation. Another advantage is that the dynamic programming model does not use personal attributes, and therefore can avoid perceived discriminating in personnel actions. Future research can build on our work to optimize resource allocations and tailor retention incentives.

Manipulations of the fundamental matrix and data results from the model are validated with practitioners and historical data. Fundamental matrix results are cross-validated with test data from the same time periods and continuation rates are consistent with an even broader time-frame. However, this research can be strengthened in the future with additional validation against more recent data sets. A viable study of behavioral changes may support the use of this type of analysis in favor of longitudinal cohort data.

Given a budget constraint, and a set of skill requirements by pay grade, it is possible to determine an optimal retention policy that initially targets individuals with grade, time-in-grade and time-in-service combinations most likely to be influenced by financial incentives. Such an optimization method would avoid offering incentives to those who are likely to leave due to personal attributes despite a higher expected value of staying. Conversely, by removing the budget constraint, it is possible to determine a minimum cost required to achieve optimal retention goals. A case study of optimal incentives could provide an actionable extension to the research presented in this paper to inform decision-making.

Appendix A: Descriptive Statistics of Data

Table 10 Frequency percentiles of ordinal data

Marital Status	Percent	Race	Percent	Grade	Percent
Married	49.4%	White	73.7%	SL1	48.8%
Single	47.7%	Hispanic	12.7%	E-5	25.8%
Divorced	2.9%	Black	7.2%	E-6	19.0%
		Asian/Other	6.3%	E-7	4.8%
				E-8	1.3%
				E-9	0.3%
Prior Reenlist	**Percent**	**Transaction**	**Percent**	**Deployed**	**Percent**
0	58.6%	ETS	28.9%	Yes	88.2%
1	19.4%	Reenlist	66.9%	No	11.8%
2	11.4%	Retire	4.3%		
3	10.6%				

N = 43,242; Career Management Field = 11 (Infantryman)

Table 11 Descriptive Statistics of scaled data

Variable	Minimum	Maximum	Mean	Std. Dev.
Age	18	58	26.04	5.454
Time-in-Service (Years)	0.2	30	5.67	4.43
Time-in-Grade (Months)	1	250	20.32	19.22
Months Since Deployment	0	228	12.07	16.26
AFQT	0	99	57.31	20.16
Education Minor Dependents	0	10	0.73	1.097

N = 43,242; Career Management Field = 11 (Infantryman)

Appendix B: Approximation Error of the Fundamental Matrix N

The definition (5) of fundamental matrix assumes an infinite number of transitions. This definition gives a compact representation of N using Q. The career path of an enlisted personnel has a finite timeline, typically 30 years. Therefore, we expect that computations based on the fundamental matrix has certain level of error. For this reason, we would like to investigate the possible error from using an infinite summation in (5). To that end, we compare the fundamental matrix with a finite summation of the sequence $N_t = \sum_{k=0}^{30} Q^k$. The approximation error of using N instead of N_30 can be evaluated using the absolute and relative errors:

$$||N - N_{30}||_p, \quad \frac{||N - N_{30}||_p}{||N_{30}||_p} \times 100\%.$$

Here, $||\cdot||_p$ represents matrix norm for given constant p.

Table 12 presents the absolute and relative errors for three different matrix norms with $p=1, p=2$, and $p=\infty$. The table shows that the approximation error is quite small, with the largest relative error less than 3%.

Table 12 Absolute and Relative Errors of using N to approximate N_30

Norm	Absolute error	Relative error
1-Norm	0.2867	2.64%
2-Norm	0.1839	2.07%
∞-Norm	0.2292	2.22%

Appendix C: Glossary of Notations for Section 5

$m_{x,t}$: Military base pay in state x and period t

$c_{x,t}$: Additional compensation pay in state x and period t

$\widehat{m}_{x,t}$: High-3 average retirement pay in state x and period t

$q_{x,y}$: Probability of becoming an involuntary loss in state x

w_t: Proportional amount of $\widehat{m}_{x,t}$ expected to receive in retirement

δ_i: Multiplier of final military pay to calculate expected civilian pay

T: Length of military and civilian career time horizon

β: Personal discount factor

τ: Individual life expectancy

η: Age upon entering military

ξ: Expected civilian pay raises

References

Anderson, T.W., Leo A. Goodman. 1957. Statistical inference about markov chains. *The Annals of Mathematical Statistics* **28**(1).

April, Jay, Marco Better, Fred Glover, James P. Kelly, Gary Kochenberger. 2014. Strategic workforce optimization: Ensuring workforce readiness with optforce. *Annals of Optimization* .

Asch, Beth J., James Hosek, Michael Mattock, Christina Panis. 2008. Assessing compensation reform. Tech. rep., RAND National Defense Research Institute.

Asch, Beth J., Michael G. Mattock, James Hosek. 2013. A new tool for assessing workforce management policies over time. Tech. rep., RAND National Defense Research Institute.

Asch, Beth J., John T. Warner. 2001. A theory of compensation and personnel policy in hierarchical organizations with application to the united states military. *Journal of Labor Economics* .

Beck, Robert J., Stephen G. Pauker. 1983. The markov process in medical prognosis. *Medical Decision Making* **3**(4).

Bellman, Richard. 1997. *Introduction to Matrix Analysis*. Society for Industrial and Applied Mathematics.

Blackstone, Tanja F. 2007. Recruiting and retention of military personnel.pdf. Tech. rep., Navy Personnel Research, Studies, and Technology.

Buddin, Richard J. 2005. Success of first- term soldiers: The effects of recruiting practices and recruit characteristics. Tech. rep., Rand Corporation.

Coates, Hunter R., Teresa S. Silvernail, Lawrence V. Fulton, Lana Ivanitskaya. 2010. The effectiveness of the recent army captain retention program. *Armed Forces & Society* **37**(1).

Craig, Bruce A., Peter P. Sendi. 2002. Estimation of the transition matrix of a discrete-time markov chain. *Health Economics* **11**.

Cyert, R. M., H. J. Davidson, G. L. Thompson. 1962. Estimation of the allowance for doubtful accounts by markov chains. *Management Science* **8**(3).

Department of Defense. 2014. Fiscal years 2014-2019: Strategic workforce plan report. Tech. rep., U.S. Department of Defense.

Department of the Army. 2010. *Army Regulation 600-8-19: Enlisted Promotions and Reductions*.

Department of the Army. 2011. *Army Regulation 601-280: Army Retention Program*.

DeWolfe, D. D, J. G Stevens, R. K Wood. 1993. Setting military reenlistment bonuses. *Naval Research Logistics* **40**.

DiCiccio, Thomas J., Bradley Efron. 1996. Bootstrap confidence intervals. *Statistical Science* **11**(2).

Doyle, Peter G., Laurie J. Snell. 2000. *Random walks and electric networks*. Free Software Foundation.

Duala, Thomas, Robert Moffitt. 1995. Estimating dynamic models of quit behavior: The case of military reenlistment. *Journal of Labor Economics* **13**(3) 499--523.

Durrett, Rick. 1999. *Essentials of Stochastic Processes*. Springer Science.

Efron, B., Tibshirani. 1986. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science* **1**(1).

Gotz, Glenn A., John J. McCall. 1983. Sequential analysis of the stay/leave decision: U.S. air force officers. *Management Science* **29**(3).

Gotz, Glenn A., John J. McCall. 1984. A dynamic retention model for air force officers. Project air force report, Rand Corporation.

Government Accountability Office. 2003. Gao-04-39 human captial: Key principles for effective strategic workforce planning. Tech. rep., U.S. Government Accountability Office.

Hall, Andrew O. 2009. Simulating and optimizing: Military manpower modeling and mountain range options. Ph.D. thesis, University of Maryland.

Harrison, Glenn W., Morten I. Lau, Melonie B. Williams. 2002. Estimating individual discount rates in denmark: A field experiment. *The American Economic Review* .

Hattiangandi, Anita U., Gary Lee, Robert W. Shuford, Aline O. Quester. 2013. Retirement choices: 2013. Tech. rep., Center for Navy Analyses.

Henning, Charles A. 2008. Military retirement: Background and recent developments. Tech. rep., Congressional Research Service.

Henning, Charles A. 2011. Military retirement reform: A review of proposals and options for congress. Tech. rep., Congressional Research Service.

Hosek, James, Michael G. Mattock, Christine Fair, Jennifer Kavanagh, Jennifer Sharp, Mark Totten. 2004. Attracting the best: How the military competes for information technology personnel. Tech. rep., RAND National Defense Research Institute.

Kvan, Paul, Joel S. Sokol. 2006. A logistic regression/markov chain model for ncaa basketball. *Naval Research Logistics* .

Mattock, Michael, Jeremy Arkes. 2007. The dynamic retention model for air force officers: New estimates and policy simulations of the aviator continuation pay program. Tech. rep., Rand Corporation.

Mattock, Michael G., James Hosek, Beth J. Asch. 2012. Reserve participation and cost under a new approach to reserve compensation. Tech. rep., RAND National Defense Research Institute.

Montgomery, Douglas C., George C. Runger. 2003. *Applied Statistics and Probability for Engineers*. John Wiley & Sons, Inc.

Nirmala, S., M. Jeeva. 2010. A dynamic programming approach to optimal manpower recruitment and promotion policies for the two grade system. *African Jouranl of Mathematics and Computer Science Research* **3**(12) 297--301.

Pfeifer, Phillip E., Robert L. Carraway. 2000. Modeling customer relationships as markov chains. *Journal of Interactive Marketing* .

Puterman, Martin L. 2005. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc.

Rao, P. P. 1990. A dynamic programming approach to determine optimal manpower recruitment policies. *Journal of the Operational Research Society* 5.

Ross, Sheldon M. 1996. *Stochastic Processes*. John Wiley & Sons, Inc.

Stewart, Derek B., James R. White. 2004. Military personnel: Active duty compensation and its tax treatment. Tech. rep., U.S. Government Accountability Office.

Styan, George P.H. 1973. Hadamard products and multivariate statistical analysis. *Linear Algebra and Its Applicatioms* **6**.

Wang, Jun. 2005. A review of operations research applications in workforce planning and potential modelling of military training. *Defense Science and Technology Organisation* .

Warner, John T., Matthew S. Goldberg. 1984. The influence of non-pecuniary factors on labor supply: The case of navy enlisted personnel. *The Review of Economics and Statistics* .

Warner, John T., Saul Pleeter. 2001. The personal discount rate: Evidence from military downsizing programs. *The American Economic Review* .