NPS-HR-20-001

ACQUISITION RESEARCH PROGRAM SPONSORED REPORT SERIES

Retention Analysis Modeling for the Acquisition Workforce

4 November 2019

Dr. Sae Young (Tom) Ahn, Assistant Professor

Dr. Amilcar A. Menichini, Associate Professor

Graduate School of Defense Management

Naval Postgraduate School

The research presented in this report was supported by the Acquisition Research Program of the Graduate School of Defense Management at the Naval Postgraduate School.

To request defense acquisition research, to become a research sponsor, or to print additional copies of reports, please contact the Acquisition Research Program (ARP) via email, arp@nps.edu or at 831-656-3793.

Executive Summary

The Department of Navy (DoN) and Department of Defense (DoD) Acquisition Workforce (AWF) Strategic Plans describe the plan and vision to strategically restore and strengthen the civilian AWF after more than 20 years of cyclical contraction. The DoN Strategic Plan emphasizes the changing “size, composition, and skill” needs of the workforce in “parallel with technology advances and global trends.” Combined with recommendations from the Section 809 Panel, the mandate of the leadership is to reform and reshape the workforce to improve acquisition and delivery of world-class warfighting capabilities for the U.S. military.

This technical research report provides a cutting-edge modeling and simulation tool that leverages the increase in availability of AWF data and the exponential increases in computing power in the last three decades. The tool is designed to support the continuing mission of recruiting, hiring, and developing a diverse, agile, highly qualified, and motivated workforce of acquisition professionals. This report accomplishes two foundational tasks that will support the long-term effort to build, estimate, and simulate a custom-designed dynamic retention model (DRM) for the AWF that uses the principles of a powerful mathematical/econometric technique called dynamic programming.

The first step in building the DRM is understanding the observed behavior of the civilian AWF employees by studying the available data. The summary statistics, long-run trend analysis, and survival analysis via a Cox proportional hazard model of a subset of the data suggest that:

- The average career length is 12 years, with little differences observed across gender or race: the AWF does well in retaining diversity. However, the workforce remains predominantly white due to racial imbalance at the point of hiring.
- One of the strongest predictors of career longevity is prior military service. The leadership should augment recruiting from active duty to seamlessly transition them into the civilian workforce.
- Employees with more education have longer careers. The leadership should be ambitious in recruiting advanced degree holders.

- In fact, employees who get a master's degree mid- career are observed to have the longest careers. Then, the leadership may invest in the workforce by encouraging/subsidizing education without worrying about brain drain.

Subsequent to the empirical analysis, we develop the proof-of-concept DRM that could eventually assist the leadership in reshaping the entire force as desired. This model takes a complex, multi-period problem and breaks down the solution into a series of much simpler, one-period sub-problems in a recursive manner. We use this model to:

- Simulate the long-run stay-or-leave decision of a representative employee and aggregate the individual up to the entire workforce.

- Introduce compensation changes (ex. pay increase, bonus) to simulate the response of the workforce.

- Preview the ability of the model to forecast and actively shape the AWF through dynamic simulations.

Overall, we find that the proposed model, though simple at this stage, already replicates the main patterns of the observed employee retention behavior. Complementing this analysis, we study how the out-of-balance AWF would evolve over time with no active and forward-looking intervention to shore up the workforce. The critical lack of experienced professionals that is predicted to arise due to neglecting to construct a deep bench of seasoned professionals many years prior to the crisis hitting full force highlights the need for data-driven decision making where manpower planning and talent management should look decades ahead.

The analysis and models in this report will assist the leadership in making such long run plans to allow the DoN to "build, develop, and sustain a balanced workforce to compete and win."

NPS-HR-20-001

ACQUISITION RESEARCH PROGRAM
SPONSORED REPORT SERIES

Retention Analysis Modeling for the Acquisition Workforce

4 November 2019

Dr. Sae Young (Tom) Ahn, Assistant Professor

Dr. Amilcar A. Menichini, Associate Professor

Graduate School of Defense Management

Naval Postgraduate School

THIS PAGE LEFT INTENTIONALLY BLANK

Table of Contents

THIS PAGE LEFT INTENTIONALLY BLANK

Introduction

Highly educated, skilled, and experienced government acquisition professionals are vital now and in the future, to provide warfighters the products they need.

–DoN Acquisition Workforce FY 19-24 Strategic Plan

All this relies on our most important asset, our people, and the approaches we take to recruit, train, and retain the workforce we need to compete and win in support of our national defense strategy.

–The Hon. James F. Geurts, ASN (RD & A)

As the defense acquisition system continues its rapid transformation, it is more vital than ever to have a talented, experienced, and well-qualified civilian workforce with a deep enough bench to have the capability and flexibility to support the ever-changing needs and increasing demands of the modern warfighters. As part of this broad effort, the Department of Navy (DoN) Acquisition Workforce (AWF) FY 19-24 Strategic Plan has clarified the need to continually reshape the workforce by developing metrics to track the overall "health" of the AWF and predict how different policy choices in talent management can affect its effectiveness. To increase these capabilities, the report recommends modernizing analytical tools, leveraging data, and building workforce forecasting tools.

Section 809 Panel has proposed changes to the Department of Defense's (DoD) career management framework to continue to improve and develop the workforce. In addition, the DoD AWF Strategic Plan—FY 2016–FY 2021 has previously highlighted efforts undertaken since 2010 to strategically restore and reshape the AWF after a period of roughly 20 years of continuous shrinkage.

Figure 1, borrowed from the DoD Strategic Plan, shows that a severe reduction in the AWF resulted in a "hollowing out" of the workforce, especially in the mid-career segment. The long-run prospects of neglecting the shape of the workforce as it stood in Fiscal Year (FY) 2008 was dire. Beyond the simple lack of manpower due to the culling of the workforce, the distribution of the remaining workforce was problematic. As the workforce aged, the "bathtub" depression would simply travel forward on the horizontal axis, leading to potentially serious shortfall in experienced professionals after one to two decades.

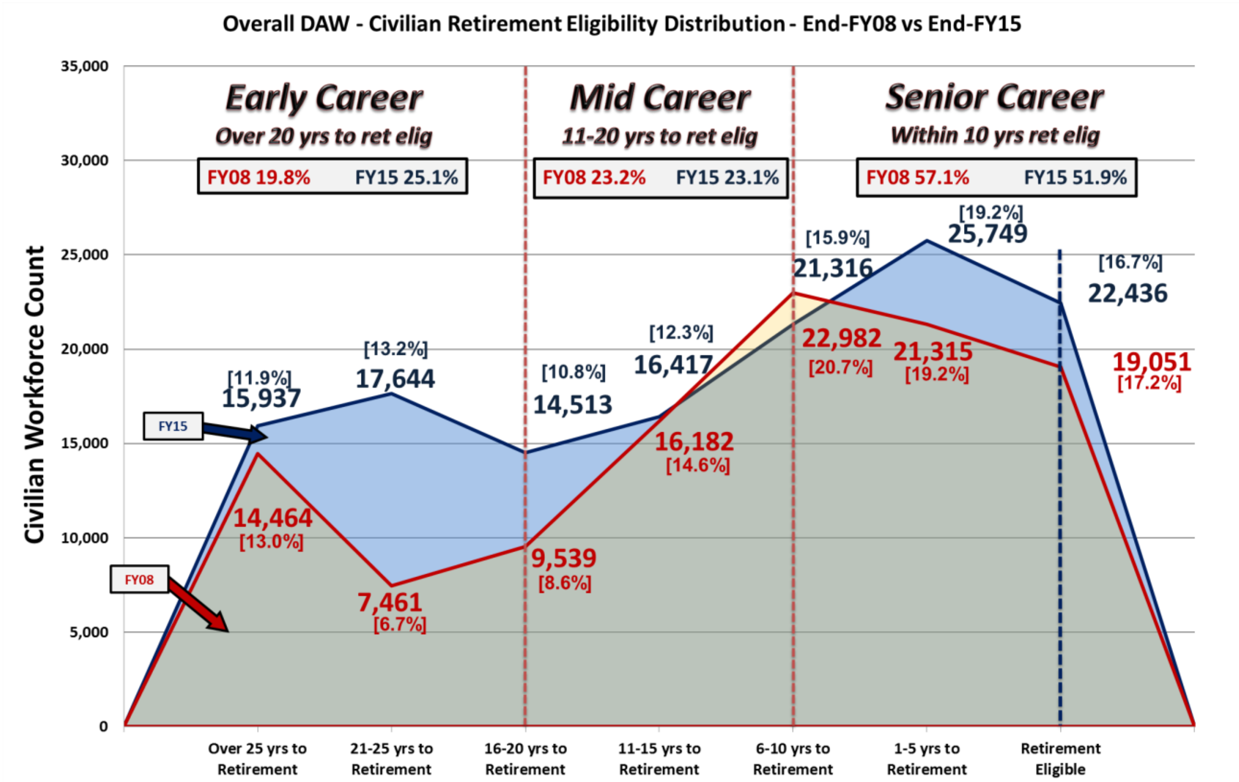

Figure 1. Civilian AWF Retirement Eligibility Distribution. Source: Department of Defense (DoD; 2015).

The objectives as stated in the Strategic Plan were to increase the overall size of the force as well as to reshape it "by deliberate and targeted growth of specific career fields" (DoD, 2015). Concurrently, the DoD made substantive efforts to improve the quality of the civilian acquisition employees. Figure 1 also shows that the leadership had strived to "fill in" the hollowed-out portion of the workforce by concerted hiring, such that by FY 2015, the overall size of the workforce had increased and the relative depth of the bathtub had become much shallower. While the overall trends in reshaping the AWF has been admirable, it is difficult to argue that shrinking the workforce in such a way as to create the bathtub in the first place was optimal.

In addition, it is unclear whether the workforce has been "re-grown" in an efficient and effective manner. While it is impossible to definitively critique the new shape of the

AWF based on Figure 1 alone, the workforce now seems very “top-heavy,” with senior employees more numerous than mid- and early-career professionals combined. In most organizations, the preferred distribution of employees by experience-level is more bottom-heavy, with many more junior and mid-level employees and fewer senior-level managers. The imbalance in workforce composition may lead to wasteful use of effort, with senior-level employees engaged in work that does not make optimal use of their abilities or experience.

As these senior professionals reach retirement age in the next five years, we may again see a sharp (natural and unavoidable) shrinkage of the workforce. Along with the loss in numbers, the experience lost may create a vacuum in leadership that is difficult to replace immediately. The leadership may then be forced to “play defense” again, shoring up the workforce by hiring in greater numbers. To prevent such crises in the future, hiring of employees in re-growing the workforce should not be done indiscriminately, but with an eye toward re-building a deep bench of acquisition professionals.

The DoN and DoD AWF Strategic Plans also describes the future goals to sustain the progress made in restoring its workforce, making further changes in the force structure to accommodate the required skillsets, improve employee quality and professionalism, and recruit and retain a more diverse workforce to maximize effectiveness. These are ambitious goals that are further complicated by the need to attract and retain (and perhaps terminate) employees differentiated along multiple skill and characteristic dimensions. A centralized plan is required to keep track of the myriad of different types of employees and how they may all have different career profiles and objectives.

Many of the goals described in the strategic plans are aspirational in nature without a long-term plan for specifics. How many employees need to be hired at each period? What are the transitional paths toward the final shape of the workforce? What is the final distribution of quality, experience, education, and other socio-demographic characteristics? In addition, the ambitious goals of each sub-branch of the AWF may be more easily achievable with a clearer understanding of the overall long-term plan. Efficiency gains may also be achievable by avoiding duplicates in recruiting and hiring efforts.

To assist the DoN and the DoD in these goals, and directly address the recommendations put forth in the DoN FY 19-24 strategic plan, this project will develop a Dynamic Retention Model (DRM) specifically designed from the ground-up for the acquisition workforce. A simpler version of the model being proposed has been used by the DoD to assess the potential impact of various talent management policy changes. The DRM and its extensions have been the workhorse of personnel analysis in the military for the past 30 years, yielding valuable insights into the labor market decisions of officers and enlisted personnel.

Our DRM uses a powerful mathematical/econometric technique called dynamic programming. It takes a complex, multi-period problem (such as the lifetime labor market decisions of an acquisition field employee) and breaks it down into simpler, one-period sub-problems in a recursive manner. Solving a single-period problem "nests" the future decisions that the employee will make, allowing the estimation and prediction of complex behavior in a surprisingly manageable framework. The set-up of the model allows researchers, once estimation of the model is complete with data, to simulate how various modifications in policy, such as changes in salaries, retirement, and bonuses, would have affected the labor market decisions of the workforce. In doing so, the model can help the leadership in achieving the desired workforce size and structure, as described by the two strategic plans.

To understand the relevant features of the workforce, in this report, we conduct the following analysis:

1) We do a deep dive into a sub-group of the AWF. We construct a panel (longitudinal) dataset that tracks employees from the start of their careers to, in many cases, the end. We examine how demographic characteristics such as gender and race affect the length of their career. We zero-in on education as a particularly important characteristic to examine. We find that employees with different amounts of education have sharply different career experiences. Higher levels of education are associated with longer careers. Surprisingly, we also find that employees who attain an additional degree while working stay substantially longer at their jobs compared to employees who started with the higher credentials. This has strong implications for how the leadership should recruit and invest in its workforce over the long run. We conduct both graphical trend analysis as well as survival regression analysis using a Cox proportional-hazard model to examine this issue further.

2) We create a proof-of-concept model that provides an intuitive way to examine and experiment with the long-term labor market behavior of the employees. An individual employee's career decisions (whether to stay in or separate) at every point in time in his or her career is modeled using insights from dynamic programming. Once we create a model that allows for a logical and time-consistent decision-making behavior by one employee, we aggregate up to the entire workforce. Then, once we can examine the behavior of an entire workforce, we simulate changes in monetary factors that can affect an individual's career decisions and follow how the sum of these individual decisions leads to large changes in the shape of the AWF. Thus, we successfully demonstrate the utility of the proof-of-concept model, although we do note that the model does not use any of the empirical data yet.

The survival analysis using the data in this study can be used to guide the construction of the empirical dynamic programming model, building on the proof-of-concept model. We can then allow for other compensation factors and differences in individual characteristics and abilities, incorporate additional career fields, and continue the process of empirically estimating the DRM model. In addition, we can explore how large, centralized hiring and firing of employees can be used to substantially reshape the workforce in a certain time horizon.

The over-arching long run aim of this research project is to use the principles of dynamic programming, update the technique with recent innovations from academia, and build a model specifically tailored for the civilian AWF. With the estimation results from the model, we plan to run simulations to:

1) Better predict size/composition of the workforce due to changes in monetary compensation.
2) Evaluate the effects of various non-pay policies.
3) Assess the impact on employee quality.
4) Incorporate the effects of the state of the economy.
5) Expand analysis sample to additional career fields.

The final product will allow analysis of personnel systems and policies to aid in data-driven decision-making by leaders. One of the primary benefits of the simulation capabilities of the model is to allow for the creation of easy to use (and understand) analytics and visualizations to help experiment, understand, and share/distribute insights and recommendations for various personnel policies. Simulations will allow ranking of the efficacy of various proposed policies using consistent performance metrics and allow robust management of (sometimes large) risks associated with changes in recruitment, retention, and separation policies.

The next section introduces the data we use to analyze the AWF. We then describe the data using simple summary statistics and graphical long-run trend analysis. Our primary empirical model is presented and analyzed in the next section. After the empirical section is concluded, we switch gears to describe dynamic programming in more detail, both with a general and technical overview. We then run a proof-of-concept model to show off some of the capabilities of the model.

Data

We received a data extract (covering September 1987 through December 2018) from Defense Manpower Data Center (DMDC) of the 1,000-byte Appropriated Funds (APF) Civilian Personnel Master File. The file was transferred after encryption and anonymization to our secured workstation. The physical file is in flat ASCII format, at 0.98 GB (approximately 1,000 MB). The data file was converted into Stata Data File format (.dta) file for analysis using StatTransfer 12. All empirical analysis in this technical report is conducted on Stata 16.0 MP–Parallel Edition.

The full list of variables in the extract follows in Table 1. For this technical report, we use the variables that are bolded. In future reports, we plan on analyzing and using the italicized variables to examine promotion and geography in more detail.

Table 1. Full List of Variables in the DMDC Extract. Source: Defense Manpower Data Center (DMDC; 2019)

Variables
Unique ID
Date of Birth
Gender
U.S. Citizenship Status
Race Code
Education Level
Year Degree or Certificate Attained
Instructional Program
Pay Plan
Grade, Level, Class, Rank, or Pay Band
Step or Rate
Work Schedule
Tenure
Pay Basis
Agency-Subelement
Organizational Component
Unit Identification Code
Duty State
Duty Country–FIPS
Locality Pay Area
Core Based Statistical Area

Combined Statistical Area
Duty Station Zip Code
Duty Station Zip Code Extension
Occupation
DoDOCC
Occupational Category Code
Functional Classification
Position Title Description
Rating of Record (Level)
Rating of Record (Period)
Service Computation Date (Retirement)
Service Computation Date (Special Retirement)
Creditable Years of Military Service
Frozen Service Years
Retirement Plan
Retirement Eligibility
Annuitant Indicator
FEHB–Health Plan
FEGLI–Life Insurance
Position Sensitivity
Disability
Targeted Disability Category
Date Overseas Tour Expires
Prior Military Experience
Supervisory Status
Basic Pay
Locality Adjustment
Adjusted Basic Pay
Total Salary
Retention Incentive
Special Pay Table Identifier
Administratively Uncontrollable Overtime (AUO)
Drawdown Action Indicator
Award
Oracle Date and Time Stamp from DCPDS
Nature of Action (1)
Nature of Action (2)
Reason for Separation
Effective Date of Personnel Action
File As of Date

We note that our current analytical sample does not constitute the whole of the AWF. For this year's analysis we were restricted to analyzing civilian employees who were ever in the Contracting, Industrial Property Management, or Purchasing fields (Occupation Codes 1102, 1103, and 1105) across the entire DoD.[1] That is, we also include employees who entered from, or left to work in other DoD departments, even outside of the AWF as a whole, if they were ever in Contracting, Industrial Property Management, or Purchasing. Restricting analysis to employees who begin and finish their careers in Contracting significantly cuts down on the sample. We only keep employees born after January 1, 1950 and before December 31, 1980. Many employees born before 1950 would have retired by 1990, and their labor market experience (primarily concentrated in the 1970s to 1980s) would be less relevant for forecasting career pathways. Employees born after 1980 would be at the very early stages of their careers and would not provide much information on long-term career trajectories.

Even after restricting the sample, we have more than 2 million observations, as the data tracks employees monthly. More than 13,000 DoD employees are tracked throughout their careers in this current sample. Potentially including other career fields within the AWF is expected to significantly increase the sample size. Table 2 presents some summary statistics for our sample.

Table 2. Summary Statistics for the DoD Acquisition Workforce. Source: DMDC (2019).

Variables	Mean (Std. Dev) [Min/Max]
Female	0.632
White	0.776
African-American	0.222
Hispanic	0.045
Asian	0.081
Native American / Native Alaskan	0.011
Has Identified Disability	0.202
Prior Military Service	0.619
Has Bachelor's Degree	0.547
Has Post-graduate Degree	0.332
Gained Additional Education in AWF	0.441
Career Length in AWF (in months)	143.6 (103.8) [1 / 309]

[1] We have the capacity to restrict the analysis to Navy AWF employees. For the analysis in this report, all DoD employees are used.

Age at Entry	33.0 (8.2) [15 / 65]
Age at Exit	48.2 (10.55) [20 / 68]
Position Type: Professional	0.657
(Ever Held) Technical	0.245
Blue-Collar	0.018
White-Collar	0.297
Ever Ranked Not Fully Satisfactory	0.575
Highest Salary	95,143.67 (30,410.74) [27,397 / 189,600]
Observations	13,590

The workforce is majority white and female. The majority of the workforce is highly educated with more than half of the sample having a college degree or above, at some point in their careers. Interestingly, 44% of the sample acquire additional education during their careers. While the bulk of these are employees who enter with a bachelor’s degree obtaining a post-graduate degree, a small number of employees who start with a high school degree successfully finish college while in the AWF.

On average, employees start their career at 33, which means that the first time we observe them in the data set is usually not their first job. In fact, more than 60% of the workforce has had prior military experience. Most employees stably match within the DoD for a long period of time. Average tenure is around 12 years, which is close to double the average tenure in private sector occupations. For the purposes of this analysis, inter-departmental transfers *within* the DoD are not treated as separation. Therefore, at least from the summary statistics, the AWF does not seem to have major issues holding on to employees.

Most of the employees we examine in this year’s technical report hold professional or technical positions. This is reflected in their education level, as described above, as well as the level of the maximum salary, at almost $100,000, they attain in their career. That said, a deeper dive into performance ratings of these employee show that more than half the workforce has received a performance rating below fully satisfactory at some point in their DoD careers. A deeper dive to examine possible predictors of poor performance may be valuable.

Descriptives and Trend Analysis

Examining the fraction of employees who leave the workforce by months of employment, shows some surprising trends. Figure 2 shows the differences in attrition rates by gender. The lines on the graph should be interpreted in the following manner: among newly hired Acquisition employees, 100% are still with the job after “zero” months. After about 100 months (approximately eight to nine years), about 70% of the original employees remain. By about 300 months (approximately 25 years of service), roughly one-quarter of the employees are left. An additional feature worth noticing is that there does not exist substantive differences in attrition rates across gender. In fact, women on average tend to have slightly longer careers, especially beyond 200 months.

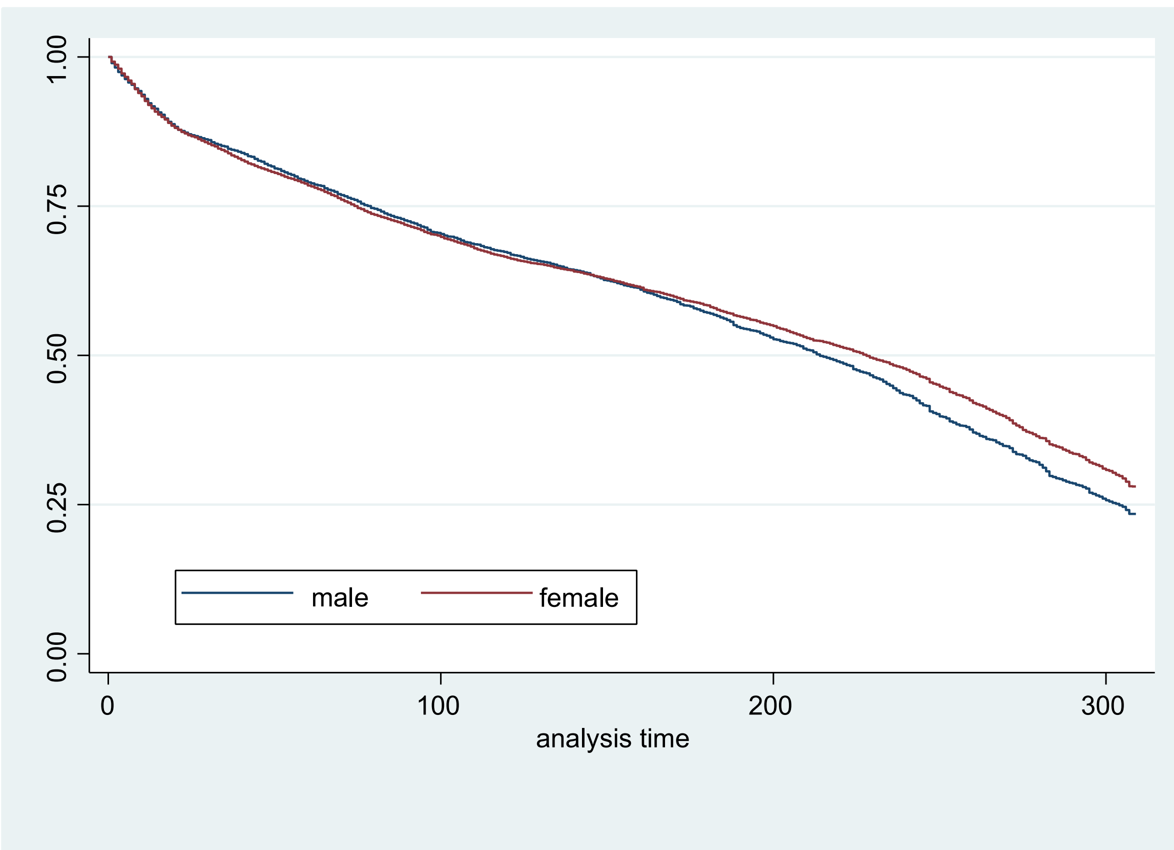

Figure 2. Career Trajectory of Employees by Gender

This trend is in stark contrast to what we observe among naval officers. Figure 3, from “Retention Analysis Model (RAM) for Navy Manpower and Personnel Analysis,” (Ahn et al.,2019), shows very early and widening gaps in career longevity across gender. By 55 quarters (i.e., around 165 months) of active service from the date of commissioning, approximately half of the male officers are still in the Navy, while more than 70% of female officers have left. At similar career points in the civilian AWF, male and female attrition rates both hover around 60%.

The simple gender trends observed in the civilian AWF presents an interesting contrast with the most recent data from the Current Population Survey (CPS). In general, while average job tenure in the private sector among men has fallen from 8.3 years to 7.4 years in the span of 30 years (1983 to 2012), women’s average job tenure has risen from 5.8 years to 6.9 years. While job tenure is observed to be on a path to convergence across gender, it should still be noted that average career length for men is still longer. Our sample shows a more extreme shift, with women on average having longer careers compared to men.

The patterns observed in the CPS sample can be attributed to ongoing changes in the labor market that allow women to more easily avoid long spells of employment disruptions due to childbirth along with better and earlier options for childcare. Viewed from this perspective, a stable job in the DoD AWF with generous benefits may then be particularly attractive for females who desire to have long careers.

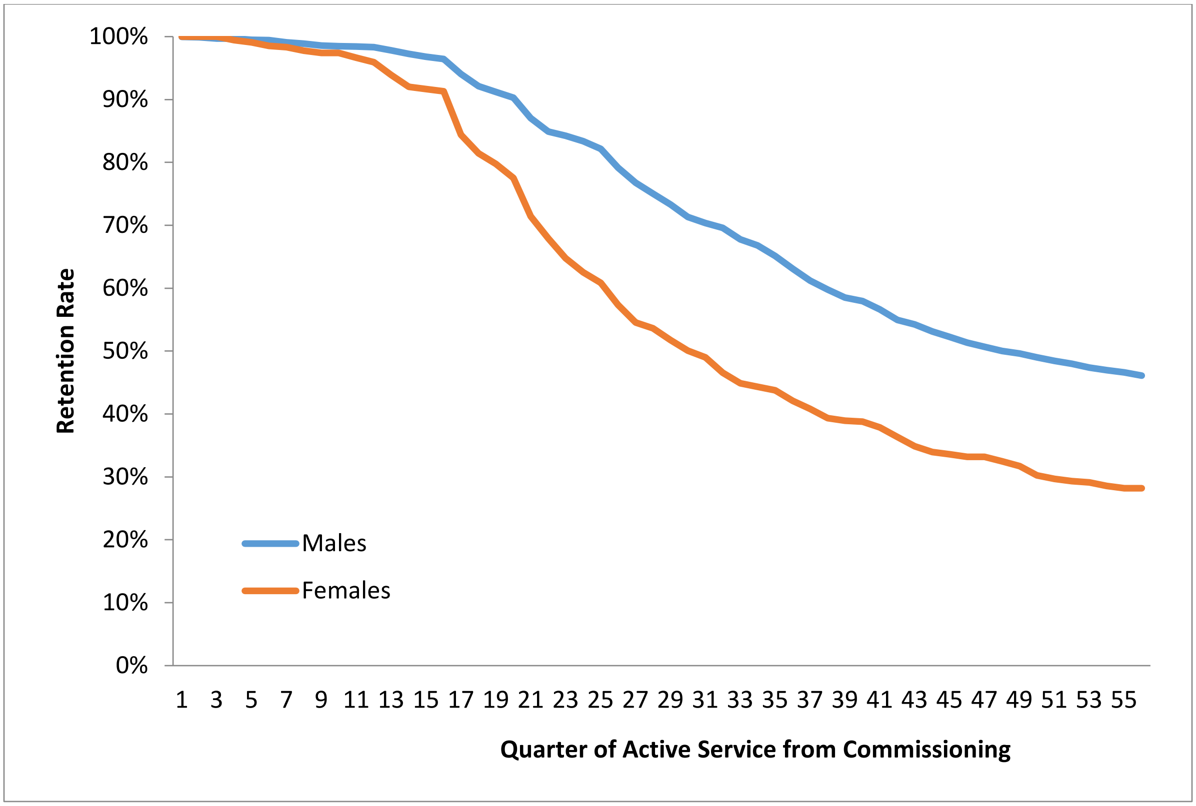

Figure 3. Quarterly Retention Rates Among Naval Officers Commissioned in FY1999, by Gender. Source: Ahn, et al. (2019).

Figure 4 shows differences in career lengths by ethnicity. African-American and Hispanic employees are combined into the "minority" category, and White and Asian employees are combined into the "white" category. Similar to observed trends across gender, no large differences are observed across race. Minority employees attrite slightly faster during mid-career (between seven and 15 years) and slightly slower near the end of their careers. We were also unable to find substantive differences in attrition across citizenship status or duty station. Unfortunately, we do not have access to marital status or number/age of dependents for the civilian data.

While a diversified workforce is important in its own right, we are also interested in the effectiveness and productivity of the workforce. To that end, we examine the impact of education level on career longevity in Figure 5. We take the employee's highest degree obtained (whether education was obtained prior to entry or acquired during his or her career).

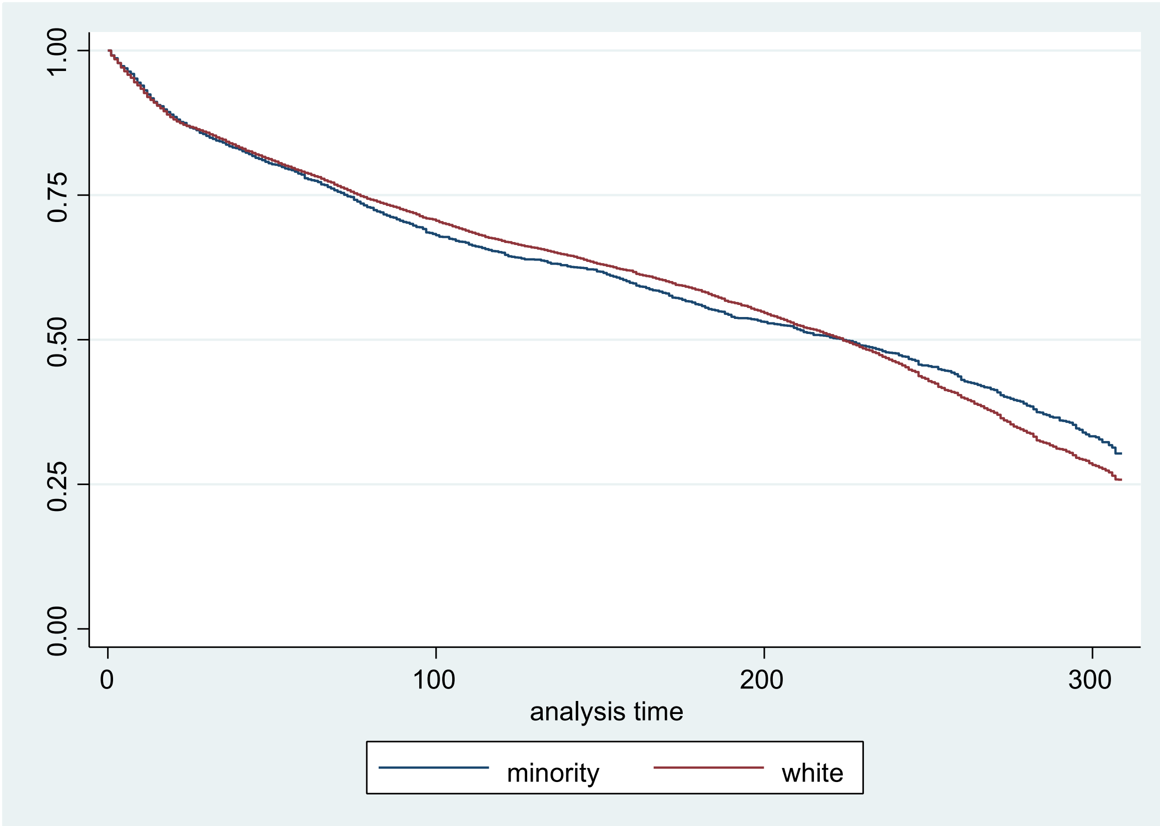

Figure 4. Career Trajectory of Employees by Minority Status

Encouragingly, there is a strong positive correlation between education level and longevity. While those with a high school degree or less attrite such that less than half the of the original cohort remains after 15 or so years, the most qualified employees, identified as those with post-undergraduate degrees, are retained at more than 50% well beyond 20 years. Attrition rate does increase rapidly beyond 20 years for these employees, as these employees are reaching retirement age more quickly than employees without graduate degrees, due to entering the workforce at least two to three years later. It is possible that employees with lower levels of education leave the AWF to return to school, so we should not automatically assume that lower productivity employees are leaving the workforce en masse.

We do emphasize that while the attrition rate for the average employee looks encouraging, there are several important caveats that warrant the reader's attention.

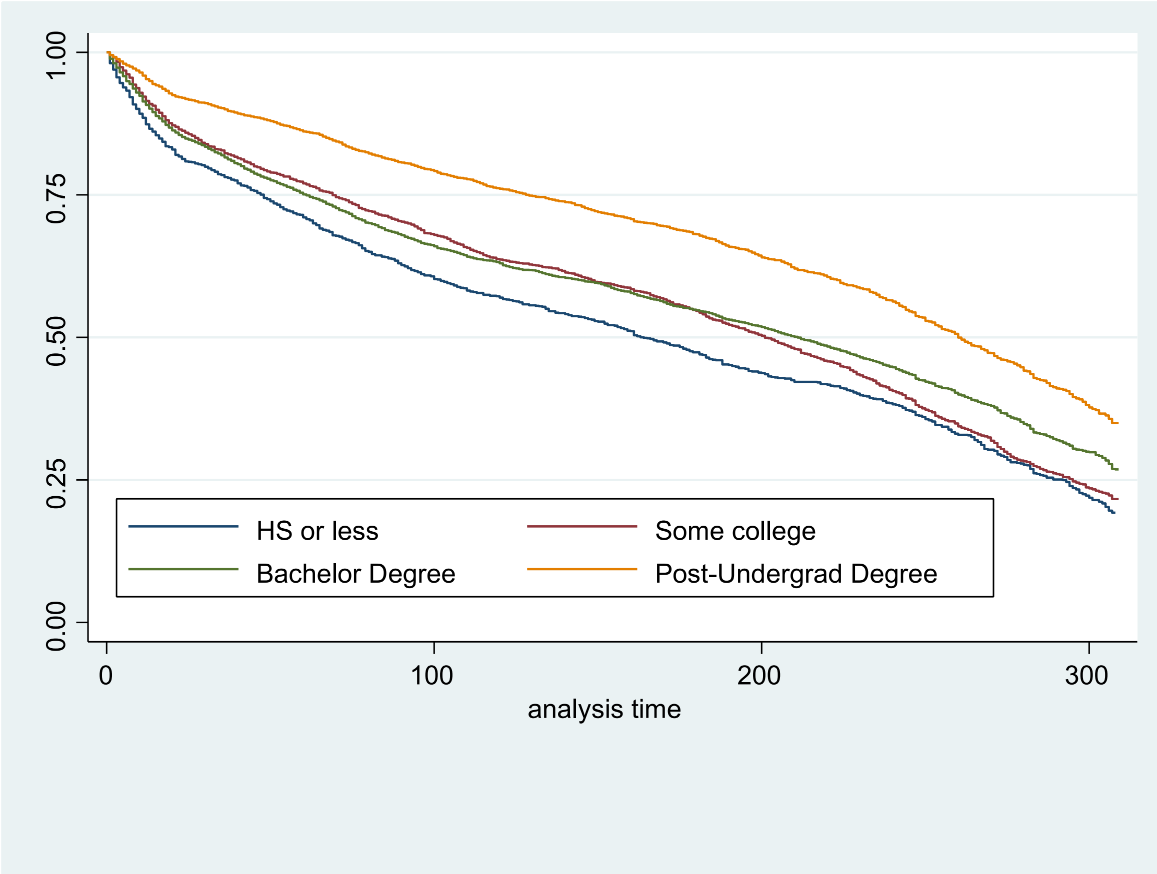

Figure 5. Career Trajectory of Employees by Education Level

First, the graphs of the attrition rates do not imply anything about the overall composition of the workforce. From the summary statistics, we saw that white and female employees comprise the vast majority of the AWF. We see that women have lower attrition rates compared to men, especially late in their careers. As for race, even though attrition may look largely the same, if whites comprise the majority of the initial cohort, the aggregate gap will not close. Table 3 shows this in detail.

Table 3. Compositional Change of AWF Cohort Through Time

	Percent female	Percent white
At entry	0.632	0.778
At 1 year	0.632	0.779
At 10 years	0.681	0.801
At 20 years	0.739	0.785

If the leadership wishes to close the racial gap, it will have to change the demographic distribution drastically at recruitment, or by retaining higher percentages of minority employees during their careers. Once estimation and simulations are completed, the demographic implication of changes to hiring and promotion policies can be fully examined in our dynamic programming model.

Second, while the trends are informative, they do not reveal compositional differences within categories. For example, even though females comprise roughly 63% of the employees at entry, narrowing the labor pool to those holding a post-undergraduate degree shrinks this number to 54%. Similarly, white composition increases in 77.8% to 79.8% when we restrict the sample to only those who hold post-undergraduate degrees. This implies that the "lower rungs" of the career ladder tend to be more heavily female and minority, compared to the rest of the workforce.

Third, trend analysis cannot reveal *why* behavioral differences across groups may or may not arise. To a certain extent, extrapolation of intent from the data becomes unavoidable. However, we must be cautious in interpretation of such data, as policy based on speculation may lead us astray. Figure 6 demonstrates this. Because we found education level to be a strong differentiator in career longevity, we examined employees who attained additional education during their careers, compared to those who chose not to acquire additional human capital. The results are quite striking. Employees who attained an additional degree have much longer careers. This result is somewhat puzzling, as one would assume that additional education would make the employee more valuable in the private sector. That these ambitious employees tend to stay much longer in the AWF could be an encouraging sign.

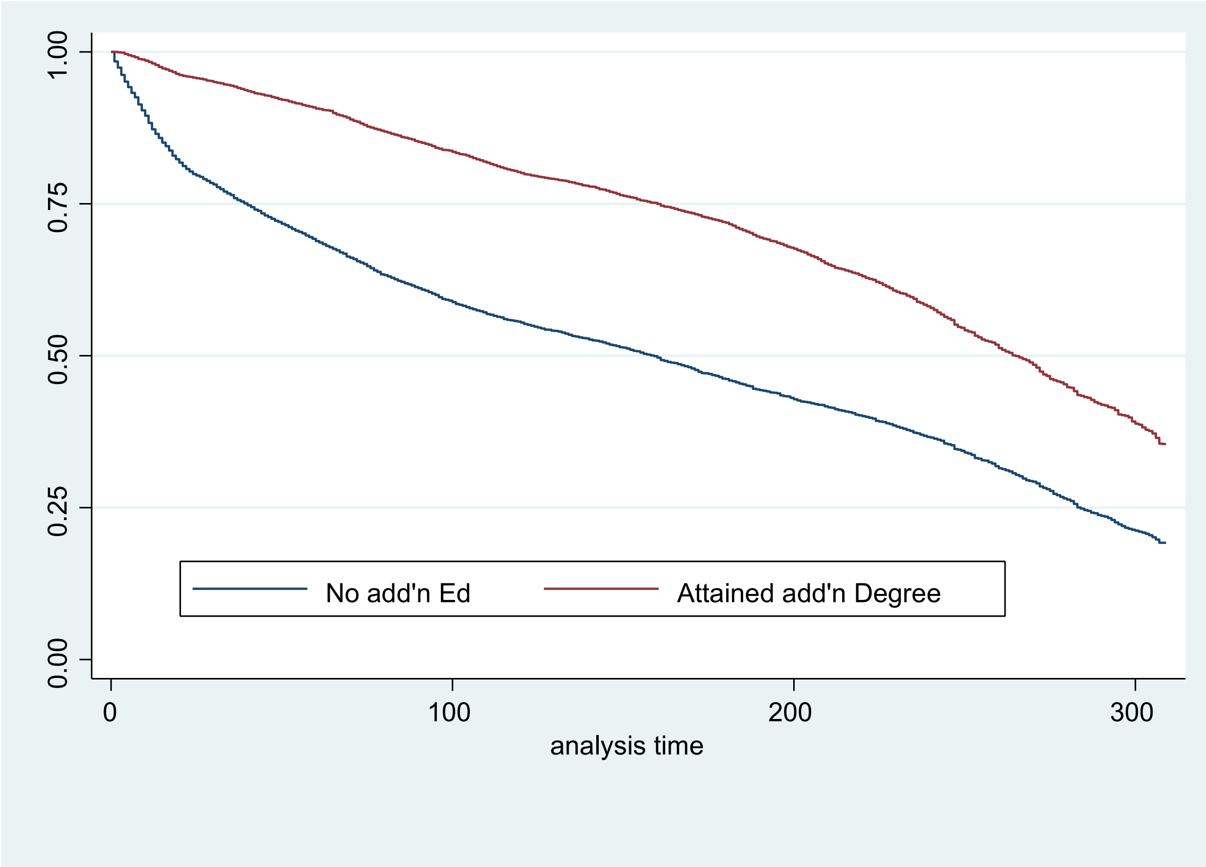

Figure 6. Career Trajectory of Employees by Those Who Do/Do Not Obtain Additional Education

However, we must be cautious about drawing too many conclusions from these trends. Figure 7 charts the career of those who hold post-undergraduate degrees. We split this sample into those who entered the workforce with the degree initially and those who acquired the degree during their career. Although both groups have the same terminal degree, those who acquire the degree mid-career have much longer careers in the AWF.

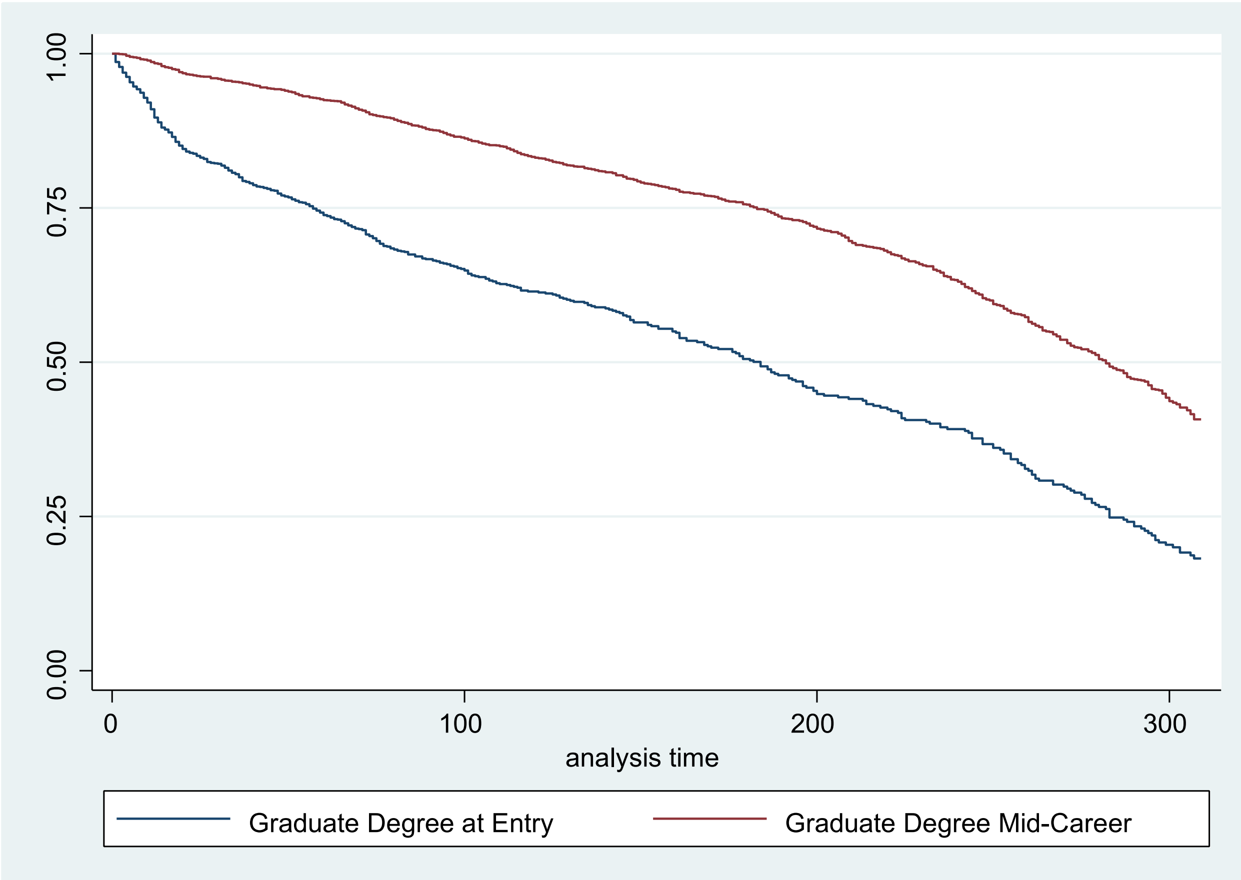

Figure 7. Career Trajectory of Employees Who Obtain Master's Degrees at Different Points in Career

One possible interpretation of this figure is that those who attain the post-undergraduate degree are employees who are most enthusiastic about their government careers, and the degree is a signal to the leadership that indicates the willingness of these employees to improve at their craft. In contrast, those who entered the field with a graduate degree may have had the technical skills required for the job, but not the drive, leading them to exit prematurely. If this is the correct description of the workforce, leadership should discourage hiring of those with master's degrees and above, concentrate on those with bachelor degrees, and actively promote going to graduate school to advance careers in the AWF.

However, another interpretation is that there is a quality difference in the graduate degree obtained by those mid-career compared to those who already had the degree at entry. The civilian employees who stayed are not choosing to do so, but are not highly sought after by private industry. Indeed, even if there is no gap in "quality," if the field of study that these mid-career students choose to pursue is tightly and narrowly connected to their tasks in their current position, this may not boost their desirability to the private sector, even with the advanced degree.

To examine this issue further, Figure 8 graphs the career trajectory of three separate sub-samples: those who enter with a bachelor's degree and never pursue an additional degree, those who enter with a bachelor's degree and go on to obtain a post-undergraduate degree, and those who enter with a post-undergraduate degree.

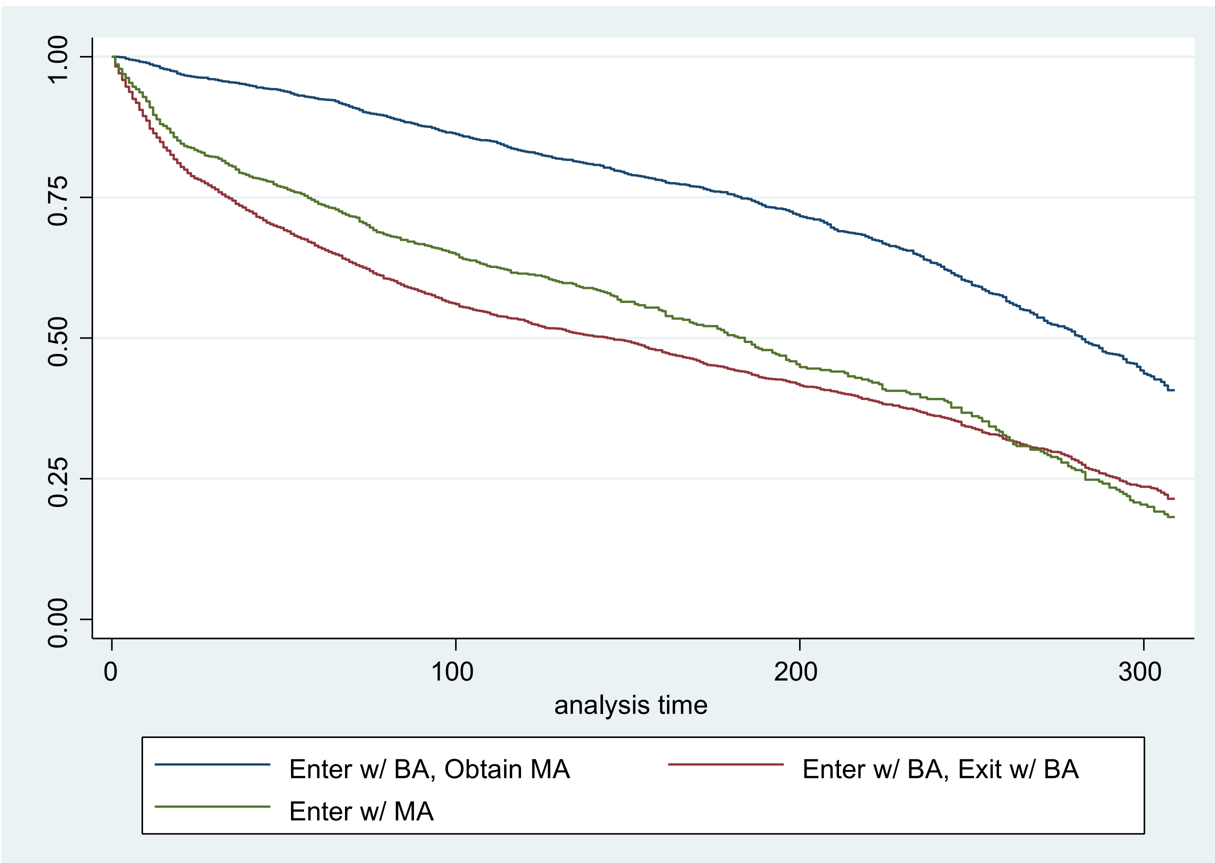

Figure 8. Career Trajectory of Employees Who Enter the AWF with a Bachelor's, Master's, or Obtain a Master's Degree Mid-Career

The relative proximity of the lines for those who enter/exit with only a bachelor's degree or master's degree, compared to those who make additional investments in education, lends support to the assertion that the human capital investments civilian employees make are directly related to a desire for a longer career in the DoD.

THIS PAGE LEFT INTENTIONALLY BLANK

Regression Analysis

The AWF sample is analyzed using duration (survival) analysis. Survival analysis has been widely used across many fields, including biostatistics, operations research, demography, actuarial studies, as well as economics and other social sciences. For analytical tractability, we assume that the observed career trajectories of the employees are of single-spell durations obtained by flow sampling. That is, we assume that once a employee chooses to leave the workforce, he or she does not return. The principal variable of interest is the "survival time" of a subject. In this case, survival time equates to the career length of the employee.

These categories of models attempt to estimate the probability that a duration (or spell) is less than some time period *t*. Then, this cumulative distribution function is defined as:

$$F(t) = \Pr(T \leq t) = \int_0^t f(s)ds$$

This simply means that we are concerned with identifying factors that increase or decrease the probability that the "duration" of an AWF career could end before a certain time period. In estimating these models, historically, we have been interested in estimating a "hazard function," which is the instantaneous probability of changing states (going from working to not working), conditional on having survived to time *t* (having worked until that time):

$$\lambda(t) = \lim_{\Delta t \to 0} \frac{\Pr[t \leq T \leq t + \Delta t \mid T \geq t]}{\Delta t}$$

This hazard function would answer a question such as: if an employee has "survived" for 10 years, what is the probability that he or she quits in the next month?

A curious reader may wonder why we resort to duration analysis, instead of the much simpler (and more easily interpretable) ordinary least squares (OLS). A primary reason for this is the presence of severe censoring in the data set. As our sample runs to 2018, the youngest employees in our sample will be approximately 38 years old. Given the average entry age into workforce is 33 years old, these employees will have had five to 10 years of work, with an average of two to seven more years to go until separation. Hence, a substantial number of employees will still be working as part of the AWF when our sample period terminates. Then, estimating the mean (average) duration of an employee's career using this dataset may lead to under-predicting the career longevity of the civilian employee as well as biasing the impacts of personal and professional characteristics of the employee on his or her career length.

Specifically, we estimate a Cox proportional hazards model. While the rigorous econometric justifications are outside the scope of this report, there are a few compelling reasons to use this estimator.[2] Other "fully parametric" duration models are simpler to estimate, but can produce inconsistent (read: incorrect) parameter estimates if any part of the model is mis-specified. Although the literature has introduced a myriad of ways to choose flexible parametric models that are robust to mis-specification, an alternative method is to use a semiparametric estimator that require less rigorous distributional assumptions. The Cox proportional hazard model has become the standard model for survival data due to this robust characteristic. For additional details on the model, refer to Cox (1972).

In a Cox proportional hazard model, the hazard function is defined as:

$$\lambda(t|x,\beta) = \lambda_0(t)\exp{(x'\beta)}$$

There is a "baseline hazard", $\lambda_0(t)$ which is a function of time alone. This can loosely be interpreted as how the probability of a representative (random) employee quitting in the next time period changes across time. For example, if an employee is in his or her first month of work, perhaps the probability that he or she quits within the next month is close to 0. However, an employee who has "survived" to, say, Year 30, may have a 10% chance that he or she quits in the next month.

Writing out the hazard function in this way makes the interpretation of estimated regression coefficients easy to interpret. If one of the *x* variables changes by one unit, say, for example, if we examine how changing the gender (from male to female) of an

[2] Interested readers are directed to Lancaster (1990).

otherwise identical employee affects the hazard rate, it turns out that the new hazard rate is simply $\exp(\beta_{female})$ times the original hazard rate:

$$\lambda(t|x_{new}, \beta) = \exp(\beta_{femlae})\lambda(t|c, \beta)$$

We estimate four different models, and present the coefficient estimates as well as the hazard ratios in Table 4. In general, a positive parameter estimate (and a hazard ratio greater than 1) indicates that the probability of transitioning to a different state (that is, going from working to leaving the workforce) increases. On the other hand, a negative parameter estimate (and a hazard ratio less than 1) means that the probability of leaving the workforce, conditional on surviving to this period, declines, compared to the base case. Note that the hazard ratio must be greater than zero.

To interpret, the parameter estimate on a female in Model 1 is -0.1866, and the hazard ratio is exp(-0.1866) = 0.8298. The negative number (and the hazard ratio less than 1) indicates that women, compared to men with identical attributes, are less likely to change status (leave), given that she survived to that period. As a concrete example, say that the probability that a male employee who just completed 10 years of employment at the AWF has a 5% chance of quitting within one month. Then, an otherwise identical female employee has about a 4.15% chance of quitting in the same period.

Model 1 contains the fewest number of independent variables. All variables in Model 1 are exogenous. That is, these variables are unquestionably predetermined prior to entry into the workforce, and thus, the regression estimates can reasonably assumed to be *causal*. We see that race has no effect on the career longevity of employees one way or another. Female and disability status yield moderately longer Acquisition careers, but the variable that has the greatest effect on longevity is whether the employee has prior military experience. In fact, the hazard ratio for prior military experience is roughly 1/20 in size compared to a similar individual without a military background. Employees who are sliding over from the military are almost guaranteed to have a long career in the acquisition field.

Looking forward a little bit, these results are robust across all four model specifications. Specifically, for the prior military experience variable, we may speculate that large impact on career length may be due to specific human capital match across the two jobs. It is most likely that these employees had complementary experience in the military and moved over to the civilian side to continue similar work. There may be little disruption in the types of day-to-day tasks performed, and most of the professional networks formed during the military career may be directly taken advantage of in the job. Given these circumstances, an employee with military experience would find the job less taxing compared to a civilian colleague who lacks similar experience, and it stands to reason that their performance would also stand out, leading perhaps, to faster promotions and more desirable assignments.

One other plausible explanation could be that the employee with prior military experience is over-specialized to the extent that the AWF (or comparable positions in the DoD) is the only place where he or she can receive compensation commensurate with his or her experience. In either case, it stands to reason that if the leadership hopes to maintain and increase the job tenure of the average employee, they should focus recruiting among soldiers and officers who are billeted to positions that deal with Acquisition issues.

Table 4. Cox Proportional Hazard Model Parameter and Hazard Ratio Estimates

	Model 1		Model 2		Model 3		Model 4	
	Coef.	Hazard	Coef.	Hazard	Coef.	Hazard	Coef.	Hazard
Female	-0.1866*	0.8298	-0.2292*	0.7952	-0.1619*	0.8505	-0.1126*	0.8935
	(0.0252)	(0.000)	(0.0259)	(0.000)	(0.0262)	(0.000)	(0.0261)	(0.000)
African-Am.	-0.0214	0.9789	-0.0250	0.9753	0.0008	1.0008	0.0573	1.0590
	(0.0291)	(0.463)	(0.0292)	(0.391)	(0.0292)	(0.978)	(0.0293)	(0.051)
Hispanic	-0.0492	0.9520	-0.0625	0.9394	-0.0247	0.9756	0.0352	1.0358
	(0.05461)	(0.368)	(0.0546)	(0.252)	(0.0547)	(0.652)	(0.0548)	(0.520)
Native Am.	-0.0414	0.9594	-0.0501	0.9511	0.0306	1.0311	-0.0090	0.9910
	(0.1178)	(0.725)	(0.1178)	(0.671)	(0.1179)	(0.795)	(0.1178)	(0.939)
Disability	-0.1331*	0.8754	-0.1312*	0.8771	-0.1154*	0.8910	-0.0723§	0.9303
	(0.0327)	(0.000)	(0.0327)	(0.000)	(0.0327)	(0.000)	(0.0328)	(0.028)
Prior Military	-3.0036*	0.0496	-2.9681*	0.0508	-2.9652*	0.0516	-3.0574*	0.0470
	(0.0358)	(0.000)	(0.0361)	(0.000)	(0.0364)	(0.000)	(0.0384)	(0.000)
BA degree	-	-	-0.1069*	0.8986	-0.0050	0.9950	0.0319	1.0324
			(0.0242)	(0.000)	(0.0275)	(0.841)	(0.0267)	(0.231)
Post-BA	-	-	-0.1598*	0.8523	-0.0051	0.9949	-0.0626§	0.9393
			(0.0282)	(0.000)	(0.0297)	(0.863)	(0.0314)	(0.046)
Add'n Degree	-	-	-	-	-0.4513*	0.6368	-0.3025*	0.7389

					(0.0272)	(0.000)	(0.0274)	(0.000)
Professional	-	-	-	-	-	-	-1.2607*	0.2835
							(0.0295)	(0.000)
Technical	-	-	-	-	-	-	-1.0919*	0.3356
							(0.0359)	(0.000)
Deficient Rank	-	-	-	-	-	-	-1.2102*	0.2981
							(0.0328)	(0.000)
Observations	1,951,719		1,951,719		1,951,719		1,951,719	
-ln L	63,297.701		58,795.086		58,652.802		57,393.441	

Note: §, * denote statistical significance at the 5% and 1% levels. For coefficient estimates, standard errors are in parentheses. For Hazard ratios, P-values are in parentheses.

Model 2 adds education level to the analysis. In this case, we used the highest education level that the employee has attained, according to the administrative data. We should caution here that use of these specific variables “breaks” the exogeneity claim made in Model 1. That is, we can no longer state with certainty that having a bachelor’s degree or post-undergraduate degree *results* in a longer career for an employee.

The reason for this is that, as we saw in the summary statistics, a significant number of employees acquire an additional degree mid-way through their career. Since this choice is made after the employee has decided to join the AWF, and more importantly, since advancing one’s education makes a real difference in the employee’s productivity and thus his or her value to the organization, education level is no longer “outside the model.” Of course, we could have chosen to include education level at the point of entry into the workforce.

However, given that more than 40% of the workforce attain additional education while working in the AWF, it also seems somewhat disingenuous to discount this decision-making process. In fact, this is one of the important reasons to model the career of these employees using a dynamic programming framework. A more complex model than the proof-of-concept model we present later in this report can introduce additional decision-making capabilities of the employee merely beyond stay-or-leave the workforce. A third option of stay-and-acquire-education can be added. For this particular regression, we should interpret these parameter estimates more as correlative relationships. Employees who have advanced degrees are more likely to stay with the workforce for a longer period.

Model 3 takes the analysis from Model 2 one-step further, by adding an indicator variable for employees who attained additional education while in the workforce. When this variable is included, we see immediately that the parameter estimates for college and post-graduate degrees become much smaller in magnitude and statistically insignificant. Instead, the variable on additional education is large (in absolute magnitude) and statistically significant, indicating that employees who attain an additional degree are strongly associated with a longer career.

On the one hand, this is very encouraging. Additional education adds to human capital, which increases a given employee's productivity. That these employees, after acquiring more education, are shown to stay *longer* in the workforce may imply that the AWF is an attractive place to work, and highly skilled employees do not wish to leave. On the other hand, the magnitude and clarity of the association (by the small standard error), should lead us to treat these results with some degree of caution. The analysis of employees' choice in staying in the DoD may be overly optimistic for a number of reasons:

1) Employees who most desire staying with the AWF may be the ones earning the additional degree.

If the additional degree is purely a signal from the employees to the leadership of their desire to stay in the workforce, then the parameter estimates may only be reflective of the self-selection of these employees into getting additional education to show that they are valuable to the leadership. This additional investment may not lead to an increase in productivity.

2) Additional education may be a requirement for advancement.

The positive correlation between career longevity and acquisition of an additional degree may reflect a formal or informal professional requirement. As an example, if an employee without a college degree is precluded from advancing beyond a certain grade, he or she may be induced to obtain a degree. There are many examples in the private sector as well as the military where advanced degrees are either an explicit or an implicit pre-requisite for promotion. For example, teachers who aspire to become a principal must obtain an Ed.D. degree. Starting with those in the Year Group 2015, the Navy will require its officers to have a master's degree before taking major command (at the O-6 level)

(Burke, 2018). Outside of whether the additional degree adds productivity, there will be a functional positive relationship between career length and acquiring a degree.

As a simple illustrative example, assume that every employee enters the workforce with 16 years of education. After working 10 years in the AWF, they are promotable to a position that requires a master's degree (two years of additional education). Then, if we looked at a cross-tabulation of average years of education and career length, at the promotion point, we may observe something akin to the following:

Table 5. Example of Positive Functional Relationship Between Education and Career Length

	Before Promotion Point	After Promotion Point
Average Years of Education	15	16
Average Years in AWF	5	15

The relationship between education and career length now has a positive slope ((16-15)/(15-5) = +1/10), unrelated to whether education actually adds any productivity, purely based on the fact that it is an administrative requirement.

3) Employees may be adding human capital that is specific to the AWF.

Even if employees self-select into acquiring an additional degree to send a signal to management, the education itself may lead to a positive impact on employee productivity without leading to an increased exit for these newly productive employees if the additional degree adds specific human capital. Specific human capital refers to education that adds to the productivity of tasks that are directly related to the civilian acquisition workforce and it only.

This is in contrast to general human capital, which would be education or training that is valuable to other firms in the private sector or in government. As an example, learning to use a computer software that is custom-designed for acquisition work may make the employee more productive in his or her daily tasks within the current organization, but it is questionable whether this skill would transfer in a substantive fashion in a work environment where this software is not used. On the other hand, if an employee is trained in using Microsoft Excel (or other software suites that are used across fields and firms), this training should provide value to any firm that hires this employee.

While the employee becomes more productive, this productivity maximized in jobs within the AWF. Therefore, the value of these employees increase substantively for their increased capabilities in the acquisition field, but the private sector would not find these employees much more desirable than prior to obtaining the additional degree. In these circumstances, economic theory predicts that: a) employees themselves would not be interested in obtaining this degree without financial and professional incentives provided by the leadership, and b) employees should not expect to receive extraordinary compensation (commensurate with their new productivity) after attaining the degree. We plan to follow up on this analysis in subsequent years' reports.

We should note that these three potential reasons listed above are not necessarily negative factors for the AWF. As a signal or administrative requirement, the degree lets leadership know which employees are committed to a long career in the field. This can provide clarity on which employees to invest in for the long term. In addition, if indeed the additional education mostly adds specific human capital, the leadership should promote additional education mid-career and provide incentives to induce employees to take up training.

4) The additional degree does not add valuable *general* human capital.

While it is possible that the focus of the additional degree is in increasing specific human capital, it is somewhat difficult to envision a degree that does not add at least some amount of *general* human capital. Yet, the strong positive impact of additional education on career longevity points to the possibility that the quality of education pursued by these employees do not yield enough of a value to work *outside* the acquisition field to garner strong, competing interest from the private sector. If the degree does not add any useful human capital, the cost, in terms of direct tuition as well as the time and effort spent, is a waste.

We should note that this interpretation may co-exist with the signaling explanation. The point is not that signaling is not useful, but that if this additional education does not actually add productivity, this ends up being a very heavy price to pay just to let the senior leadership know that the employee is interested in a long-term employment relationship.

Ultimately, the summary statistics and long-run trend analysis lead us in tailoring our models in the following ways:

1) A simple static model will not be rich enough to capture the salient features of an employee in the AWF. While the decision to stay or leave the workforce is interesting and important, a primary differentiator across employees seems to be how they attain human capital. Since almost half of the sample acquire additional education mid-career, it is important for us to model this as well.
2) Demographic differences may not play as large a role in retention decisions compared to Navy officers. In prior research with Navy officer data, we had been struck by the stark differences in observed attrition dates across gender. We also found some differences across ethnicities. In contrast to the fundamental gender and racial imbalance in the Navy officers corps, we were surprised to find no such differences in the acquisition workforce. This is encouraging for the leadership in that the workforce, after intake of these employees, do not suffer from systematic over-attrition of any minority group. That said, the AWF is still a very white workforce, and more efforts at the point of recruitment may be necessary to achieve better demographic representation. For our modeling purposes, these factors imply that we may not need to be overly concerned with modeling gender and race separately throughout career. This should decrease the "state space," making the model more tractable.
3) More empirical exploration between the role of education/degrees and promotions may be necessary. While higher education (especially degree earned mid-career) have a positive effect on career longevity, it is not clear why career length is affected so strongly, when there are good theoretical reasons why the effect should be in the opposite direction.

We next describe how our full dynamic model will be constructed, starting with a non-technical description. The empirical insights in this section will inform the structure of the full model.

THIS PAGE LEFT INTENTIONALLY BLANK

General Description of Dynamic Programming

In this section, we provide a non-technical description of the dynamic programming model, showing how the simplicity of the per-period model can lead to a false solution if we myopically “solve” the problem without considering the dynamic implications of a current choice affecting what happens in the future. The dynamic programming model allows the nesting of future periods in a compact manner, which allows for easier calculations that are logically consistent across the time period under evaluation. As we will describe below, one of the principal issues with previous attempts at estimating a dynamic model has been time-inconsistency. This means using the estimated parameters from these prior models to simulate employee behavior through time has them behaving in illogical ways (making choices that are counter to their best interests) when we look into the future.

Dynamic programming models are complex mathematic and econometric model of dynamic, optimal decision making across time. By “across time,” we mean that a decision made today has the potential to affect the agent’s labor market situation tomorrow, which may then affect his or her decision in the future period. The economics literature has produced several flavors of dynamic programming models over the past 50 years. The version most well-known to practitioners in the Department of Defense (DoD) is the Dynamic Retention Model (DRM), pioneered in the early 1980s by the RAND Corporation (Gotz & McCall, 1984). It remains one of the primary tools used by the DoD to examine the potential impacts of proposed personnel/talent management policy changes on service member retention. For example, the impact on exit behavior of new recruits due to the recent changes to the military retirement system (BRS) was examined with the DRM. Dynamic programming simplifies a complex, multi-period problem (for example, an officer’s lifetime labor market decisions) into a series of much simpler, single-period sub-problems using backward recursion. The single-period problem contains a value that captures future decisions that the officer will make, which allows the researcher to estimate and forecast complex, decades-long behavior in a manageable framework.

The strength of DRM then is its ability to map out a (labor market) lifetime behavior model of officers and enlisted men and women where they would make the best choices available to them at each point in time. Once estimation of the econometric model is finished, the model allows the researcher to simulate how policy alterations in salaries, retirement benefits, and bonuses, would affect the decisions of the average officer or enlisted soldier. The DRM and its many extensions have been the workhorse of manpower/retention analysis in the DoD for the past 30 plus years, yielding strong insights into the retention behavior of officers and enlisted personnel.

For its time, the DRM model was remarkable in its ability to accomplish this feat, given the limited computing power available. The important trade-off for the ability to compute these types of models was in the high degree of abstraction from the actual labor market. Ultimately, this forced parsimony in modeling has meant that DRM is attempting to describe the complex motivations and behaviors of officers and soldiers making life-altering labor market choices in a nuanced environment, with a small number of "regression parameters."

For example, assume that we wish to create a model in which we predict whether a soldier chooses to stay or leave. If we create a list of factors that may affect that decision, we may think about including: gender, age, specialty, education level, sensitivity to risk, health, income, benefits, marital status, number and age of dependents, location of workplace, proximity of station to home, income he or she could earn in the private sector market, etc. However, because of computational constraints, we are only allowed to select one or two pieces of information to make the prediction. As a result, we choose to attempt to predict labor market behavior based only on income and gender. These two elements may be very important in influencing the stay or leave decision of all soldiers, but we are now ignoring all of the other factors that may affect decision-making.

This, in effect, dramatically shrinks the state space (e.g., the set of information considered when making decisions) and drastically simplifies the model. The simple models allow for the prediction of retention behavior for officers and enlisted members by service, but not by specialty area, and does not adjust for the strength of the economy or service member quality. In addition, the model cannot handle non-monetary compensation, which is becoming increasingly important under current talent management initiatives.

The basic principles of dynamic programming can be demonstrated without reliance on sophisticated mathematics. For a more technical treatment, the reader is directed to the next subsection. In this simple example, borrowed from Ahn et al. (2019), a person has two choices, whether to select high (H) or low (L) in two periods. If choices are independent across time, the person selects whatever yields the greatest payoff at each period. So in Figure 9, for periods 1 and 2, the person would select (H,H) = $300 to maximize total payoff.

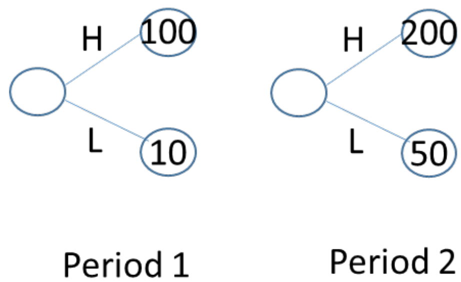

Figure 9. Simple Choice Across Independent Time Periods, Source: Ahn, et al. (2019).

Now, assume that choice in period 1 impacts possible choices in period 2. When there are a small number of periods and a limited number of choices, we can "brute force" solve for the solution by calculating the payoff for every path. As we see in Figure 10, since (H,H) = $300, (H,L) = $150, (L,H) = $60, and (L,L) = $1,010, it is optimal to select (L,L) to attain the maximum payout.

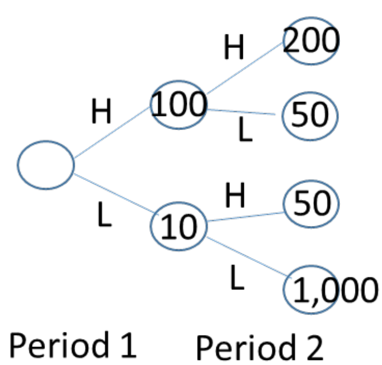

Figure 10. Choice Across Connected Time Periods Source: Ahn, et al. (2019).

While these calculations are relatively simple and quick, the scenario quickly changes once the time horizon increases or number of choices increase. The problem becomes much more complex. For example, keeping the number of choices at two (the simplest possible scenario), with one period, there are two possible outcomes, and with two periods, four possibilities, as we saw in Figures 9 and 10. With three or four periods, the number of choices (and thus calculations) increase to eight and sixteen, respectively. Over a 30-period span, there are 1,073,741,824 possible outcomes.[3] It would be very time-consuming and ultimately wasteful to calculate all 1+ billion outcomes, since most outcomes/scenarios would be undesirable and unlikely outcomes such that no rational person would make those choices. Researchers realized that it was possible to exploit a mathematical representation of this dynamic discrete choice problem by separating the payoff from one choice into the component received today plus a future term that is constructed by assuming that rational, optimal decisions will continue to be made by the individual into the final period. This is also called Bellman’s principle of optimality or Bellman’s equation.

The logic is as follows. If we are at the final period and choose between H and L, we can select the highest payoff. If we move back one period, we solve another easy problem. We already know what we would choose in the next period: the optimal one, so if we can describe this optimal decision as a number, we just have to do a single calculation. We continue this logic back to the start. This is called backward recursion. If, instead, we assume that we are myopic and attempt to make the optimal choice each period, without looking forward, we quickly run into situations where we make bad choices. Then, going back to our simple two-period example, we would choose (H,H) and attain $300 instead of the maximum possible $1,010.

An additional difficulty arises in evaluating the behavior of economic agents. Whether we are examining the decisions of officers or civilian employees in the AWF to stay or retire, we must be cognizant of the fact that we are not simply evaluating monetary

[3] It should be noted that a stay-or-leave model, where leaving implies permanent exit, is much simpler in terms of the potential number of outcomes, as long as “staying” leads deterministically to one and only one state. Currently, our model 1.0 assumes this type of decision-making. A natural extension of the model would allow agents to make an additional third choice of attaining extra human capital while remaining in the AWF.

payoff as in the simple example above. While there are undoubtedly monetary considerations, the retirement decision is inextricably tied to family, health, geographic, and professional reasons that are very difficult to monetize.

In a simple one-period framework, if an employee is faced with the decision to retire or not, he or she will be comparing the monetary benefit of staying (quantifiable as $A) and the non-monetary benefits (not necessarily quantifiable as B) against the monetary benefits ($C) and non-monetary benefits of leaving (D). If the employee stays in the AWF, then we know:

$$\$A + B \geq \$C + D$$

If he or she opts to leave, we know:

$$\$A + B < \$C + D$$

So while we would be able to tell that the sum of benefits from one option is more attractive, it is difficult to know by how much: we need an “exchange rate” between the non-monetary characteristics and salary. We need to rely on the econometric technique to translate B or D into dollars in order to make policy recommendations. So then, a DRM must not only solve the backward recursion problem, but it must also distinguish how agents value money in relation to other non-monetary characteristics of the job.

The first DRM in the military economics literature was developed by Gotz and McCall (1984) working at the RAND Corporation. They analyzed the stay/leave decisions of Air Force officers facing diverse compensation incentives at different moments in their careers. The DRM has been extended in various ways to tackle a myriad of other topics in military talent management policy. Asch, Johnson, and Warner (1998) and Asch and Warner (2001) analyze how changes to the retirement benefit system and basic pay would impact retention. The latter paper also adds individual ability and effort to the model. Hosek, Asch, Fair, Martin, and Mattock (2002) extend the model to include the initial decision to enlist, looking specifically at IT employees in the military. Asch, Mattock, and Hosek (2013) extend DRM to calculate retention cohort size as new policies are introduced and follow them through time, estimating the transition path until the new stable equilibrium. Asch, Mattock, and Hosek (2017) examine the potential impact of changes to the Blended Retirement System across the services. Gotz (1990) contains a detailed discussion of the advantages of DRM over other models of employee retention behavior, such as the traditional Annualized-Cost-of-Leaving (ACOL) model.[4]

In estimating a dynamic programming model, we deal with two computational problems. First, note that our simple example only contains two potential "states" each period. The agent can choose H to get to one state, or L to get to the other. Even in such a simple problem, across 30 periods, the number of states increase to more than one billion. Since choices in the previous periods matter, a person's sequence of selecting H or L each period each creates a new state. If there is a third choice available, there will be 205,891,132,094,649 states at the 30th period. With small increases in the number of states/periods (say, by including race/gender), we easily approach such a number of required calculations that approach and surpass the number of atoms in the universe. This rapid growth in the "state space" that we have to keep track of makes the computation burdensome (many times to the point of impossible) and is called the curse of dimensionality.[5]

Second, even the substantial simplification by the use of Bellman's equation requires us to calculate the future value of the subsequent choices to be made each period. This future term is traditionally derived through a nested-fixed-point-algorithm. This relies on a mathematical concept called contraction mapping which starts with a random guess at the value and loops through the problem continuously, at each iteration getting a better estimate of the future value until the difference in future value across iterations shrinks to some very small number. The computational burden to solve a modest model would traditionally require weeks of computing time at a supercomputer. Any alteration of the model would require calculations to be redone. Together, this has meant that any dynamic discrete choice model would have to walk a fine-line between computational tractability and fidelity of the model to the real world.

The literature in the recent past has attempted to overcome the computational burdens of dynamic programming by abandoning *exact* value function calculations and

---

[4] This is not an exhaustive list of extensions and applications of the original Gotz-McCall model, but it does represent a good cross-section of the ways in which the model has been pushed forward.

[5] The retention problem is usually cast as an "optimum stopping problem," where the decision to separate is an absorbing state. Once that decision is made, the individual receives the outside option and the problem is terminated. This reduces potential state space significantly, but not enough to allow "brute forcing" the solution.

focusing on approximate solutions that can reduce the computational burden. Among "full solution" methods, which still require the explicit calculation of the value function using the nested-fixed-point-algorithm, authors have successfully reduced the time to estimate the model through discretization, approximation and interpolation of the "Emax" function, and randomization.

Recently in the literature, estimation methods that do not require solving the full dynamic programming problem have been applied across a range of labor economics problems. The most promising is the Conditional Choice Probability (CCP) method, created by Hotz and Miller (1993). The model uses nonparametric estimations of the choice and transition probabilities (how likely are individuals to make certain career choices and how likely is the state space to change?) to circumvent the need to calculate the value functions. Some recent examples that have used the CCP method includes Slade (1998), Aguirregabiria (1999), Sanchez-Mangas (2002), and Rota (2004).[6]

An important limitation of CCP was its inability to accommodate permanent unobserved heterogeneity. If the individuals differed in an important way leading them to make different choices given identical pay structure, but we lacked the ability to observe how these individuals were different, the model would be unable to account for these behaviors. Advances in estimation have enabled the incorporation of finite mixture models to extend models to accommodate permanent unobserved heterogeneity. See Aguirregabiria and Mira (2007), Arcidiacono and Miller (2011), Kasahara and Simotsu (2007), and Arcidiacono and Ellickson (2011).[7]

[6] There have also been advances in using Bayesian statistical techniques to lessen computational burden. These techniques are newer and have not been as robustly applied. See Imai, Jain, and Ching (2009), for example.

[7] Note that we do not make use of these empirical innovations in our current model. We may introduce these concepts in subsequent versions. Models become much more complicated and take longer to estimate once unobserved heterogeneity is introduced.

THIS PAGE LEFT INTENTIONALLY BLANK

Technical Description of Dynamic Programming Models

In this section, we provide a more detailed description of the proof-of-concept DRM that could be used to analyze the AWF. The basic assumption about employees is that they are fully rational agents that, in every moment, make decisions regarding their careers that maximize their lifetime utility (i.e., lifetime “happiness”). This assumption implies that, at the beginning of each period (for instance, a month or a year), the individual selects to stay for another period or leave the workforce after trading off all the benefits and disadvantages involved in the decision, including monetary (e.g., wages and bonuses) and nonmonetary components (e.g., preference for working in the public sector). At this stage of the model development, we will assume that the decision to leave the workforce is irreversible, that is, once an employee chooses to leave, he or she cannot be rehired later. The initial model is also not considering potential periods when voluntary separation is restricted (e.g., after receiving a bonus with a promise to stay for two years.) These features will be added to the model in the future, increasing its precision to predict retention behavior.

Note that the assumption of the fully rational agent is a very high standard. Employees must, in theory, be able to correctly predict professional outcomes (in and out of the acquisition field) at every point in their career, as well as changes to personal status, and make the optimal decision at every point in time. Considering the fact that we will require the use of a cluster server (the Hamming supercomputer at the Naval Postgraduate School) to solve this model, our model may appear to be unrealistic. However, there are four reasons to choose dynamic programming as the modeling framework.

The first reason is that the fully rational assumption can be supported in a more modest sense when we consider that people do try to think through important choices, gather advice from friends, family, and senior employees, and the DoD itself has outreach programs to try to educate personnel on making smart retirement decisions. Even if a civilian employee making retirement decisions is not a supercomputer, as long as he or she is making rational decisions following best practices and good advice from experts, dynamic programming models are expected to have good predictive powers. The second reason is that dynamic programming models have been shown in the academia to have good performance in actually modeling and predicting dynamic systems in a number of different labor markets, as well as industrial organization problems. The third reason is that with the estimates from the model, we are able to generate powerful simulation exercises to explore a myriad of policy changes. The final reason is that prior models used to generate predictions or recommendations are fundamentally flawed. We discuss shortcomings of these other models later in the report.

In this initial proof-of-concept model, the decision-maker considers the following monetary elements: 1) regular compensation and bonuses and 2) compensation that the individual could obtain in the private sector. We also assume that the nonmonetary components are given by the employee's preference or "taste" for the government (versus working in the private sector.)

The basic model variables are:

- W_t^m indicates the regular compensation that the individual can obtain in period t (including bonuses)
- W_t^c denotes the compensation that the employee can obtain in the private sector in period t
- T represents the time horizon of the decision problem (e.g., the expected number of periods until final retirement)[8]
- $\beta = \frac{1}{1+r}$ indicates the discount factor, and r is the subjective discount rate of the employee
- ω^m denotes the taste parameter that captures the monetary equivalent of the preference for the AWF and military work
- ω^c denotes the taste parameter that captures the monetary equivalent of the preference for the civilian life
- $E_t[.]$ is the expectation operator given the information in period t
- ε_t^m and ε_t^c are *iid* random variables

[8] In this version of the model, we assume that retirement income is the same in the private and public sector, regardless of the individual decisions. This simplification allows us to disregard income after the time horizon *T*. This feature will be added in the next year's report.

Super-index L represents the choice to leave the AWF voluntarily while super-index S denotes the decision to stay one more period. Then, the decision-maker's maximization problem can be cast as:

$$V_t^L = W_t^c + \omega^c + \beta E_t[V_{t+1}^L] + \varepsilon_t^c = \sum_{\tau=t}^{T} \beta^{\tau-t}(W_\tau^c + \omega^c) + \varepsilon_t^c, \tag{1}$$

$$V_t^S = W_t^m + \omega^m + \beta E_t[V_{t+1}] + \varepsilon_t^m, \tag{2}$$

$$V_t = Max[V_t^L, V_t^S] \tag{3}$$

where V_t^L denotes the present value of leaving the acquisition field to enjoy a career in the private sector while V_t^S indicates the present value of remaining in the AWF one more period. In this context, the employee will choose to remain when the present value of staying, V_t^S, is larger than the present value of departing, V_t^L. The individual will face this decision at the beginning of each period until he or she leaves.

To start simulating the proof-of-concept model for a representative employee, we initially chose the parameter values showed in Table 6.

Table 6. Initial Parameter Values

Parameter	*Value*
W_t^m	1
W_t^c	1
T	30
β	0.95
ω^m	0.1
ω^c	0.1
$\mu_{\varepsilon,m}$	0
$\mu_{\varepsilon,c}$	0
$\sigma_{\varepsilon,m}$	0.1
$\sigma_{\varepsilon,c}$	0.1

Table 6 shows that the parameters are constant over time (i.e., during the work life of the employee). We initially assume that military/AWF and private sector compensation (W_t^m and W_t^c, respectively) are equal at the value of 1. The time horizon of the decision problem (T) is 30 and implies that the individual works for 30 years. The discount factor (β) is set at 0.95, which implies a subjective discount rate (r) of around 5.26% per year. The military/AWF and private-sector taste parameters (ω^m and ω^c, respectively) are equal at the value of 0.1. Finally, the *iid* random shocks (ε_t^m and ε_t^c) have zero mean and standard deviations equal to 0.1.

This initial parameterization describes a representative employee who could earn roughly the same salary by working for the AWF or the private sector. It also assumes that the individual has similar tastes for both types of jobs. We start simulating the model using these parameter values. While varying those parameter values would naturally change the simulation figures, the main conclusions would remain roughly the same. In fact, we later analyze the effects of changing the parameter capturing military compensation (W_t^m) on employee behavior. Specifically, we study the retention impact of a permanent increase in military pay of 2% after 15 years of service and of paying a one-time bonus of 15% at year-of-service 15.

While our proof-of-concept model simplifies reality considerably, we show in the figures below that it already captures some key features observed in the data. We start by replicating the bathtub problem described by the DoD AWF Strategic Plan—FY 2016–FY 2021 in Figure 11. It refers to the lack of employees with five to 15 years of experience in the AWF as a consequence of several years of personnel reduction. The vertical axis shows the number of employees, while the horizontal axis depicts the number of years until retirement.

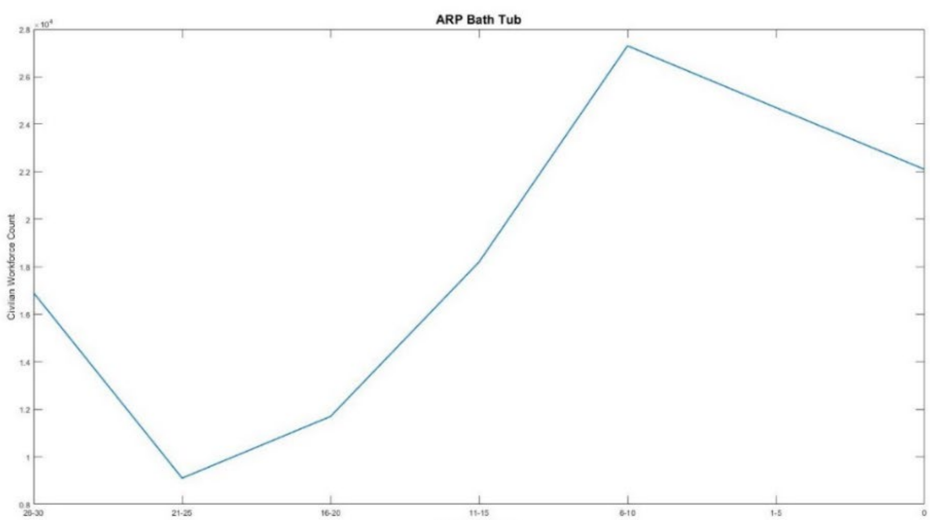

Figure 11. AWF Bath Tub Problem

We simulate the model with the parameterization outlined in Table 6 and show its prediction about retention behavior for a representative employee in Figure 12. As expected, the figure shows that attrition is high in the initial years and diminishes continuously as time passes.

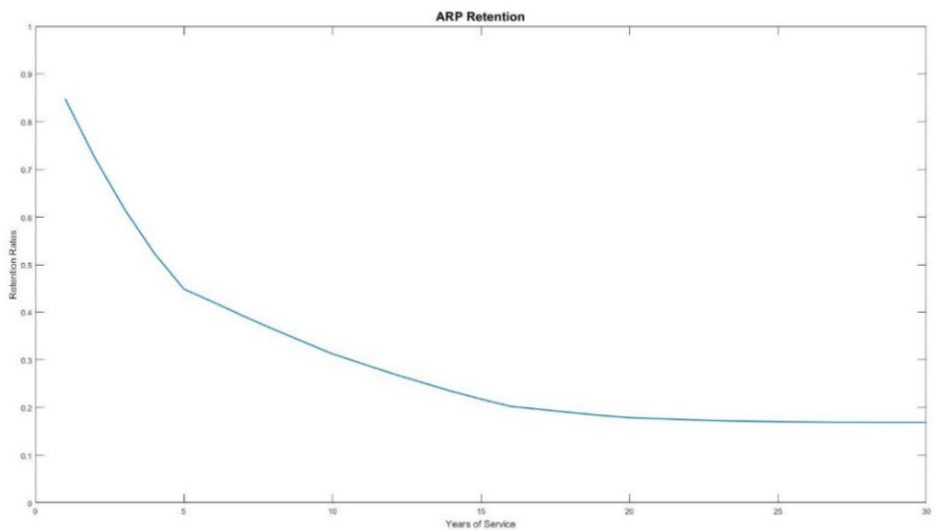

Figure 12. AWF Retention Rate

With the predicted behavior of a representative employee over his or her career, we simulate the model to forecast the evolution of the workforce distribution over time in the context of a “myopic” personnel policy. This policy assumes the AWF hires employees only at the beginning of their careers (i.e., in Year 1) and has no long-run planning of the structure of the force. The result is displayed in Figure 13, which suggests that the bathtub problem would move to the right over time, progressively reducing the number of experienced leaders as time passes.

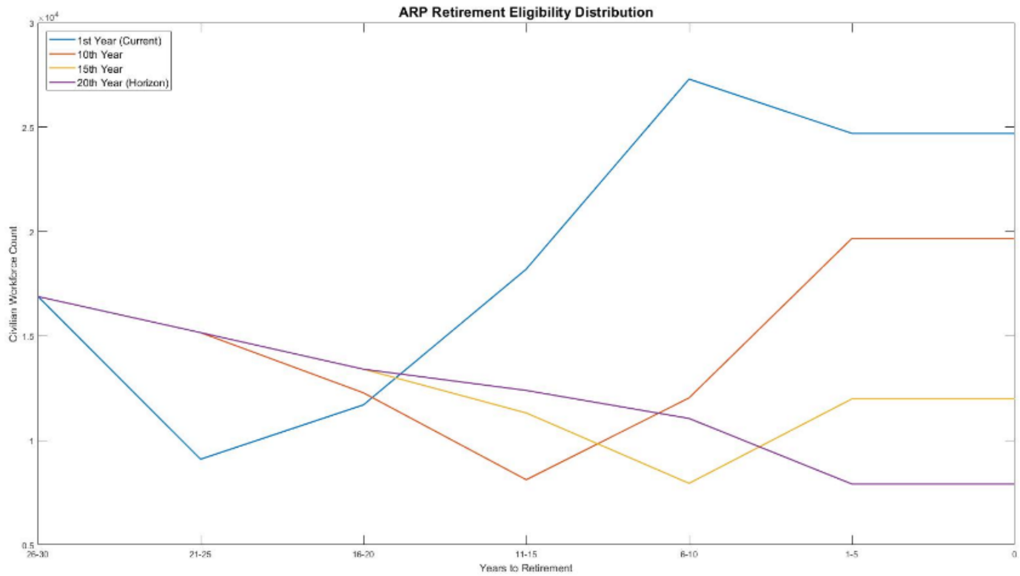

Figure 13. Evolution of AWF Without Active & Forward-Looking Intervention

The Strategic Plan clearly describes the bathtub problem around employees with five to 15 years of experience and suggests that, within five years, roughly 50% of the workforce would be retired or eligible for retirement. The leadership has addressed and ameliorated this problem by implementing substantial hiring and retention efforts. While it is easy to conclude that higher salaries will induce employees to stay longer, it is quite a bit more complicated to predict how individuals will react to different types of incentives at dissimilar moments of their career. That is one of the main aspects where the DRM we have developed in the first year can start helping the leadership when analyzing changes in compensation (e.g., raising wages or paying bonuses).

Figure 14.1 shows the employee retention behavior with a 2% pay increase after 15 years of service. That is, the salary of the individual is W_t^m =1 in Years 1 through 15 and 1.02 in Year 16 through 30. The model assumes the individual knows at moment 0 he or she will receive the 2% raise starting in Year 16. As compared to Figure 12, Figure 14.1 shows that the wage rise induces employees to stay more in the early and mid-career years. Another interesting feature of Figure 14.1 is its general similarity with the actual career trajectories shown in Figure 2. Both figures suggest that employee attrition is relatively high in the initial career years, diminishing during the mid-career years, and increasing again close to the end of the employee's tenure.

Figure 14.2 shows how the bathtub problem would evolve over time in the context of a salary rise. The main difference with respect to Figure 13 is increased retention during the early and mid-career years, though the central bathtub problem persists.

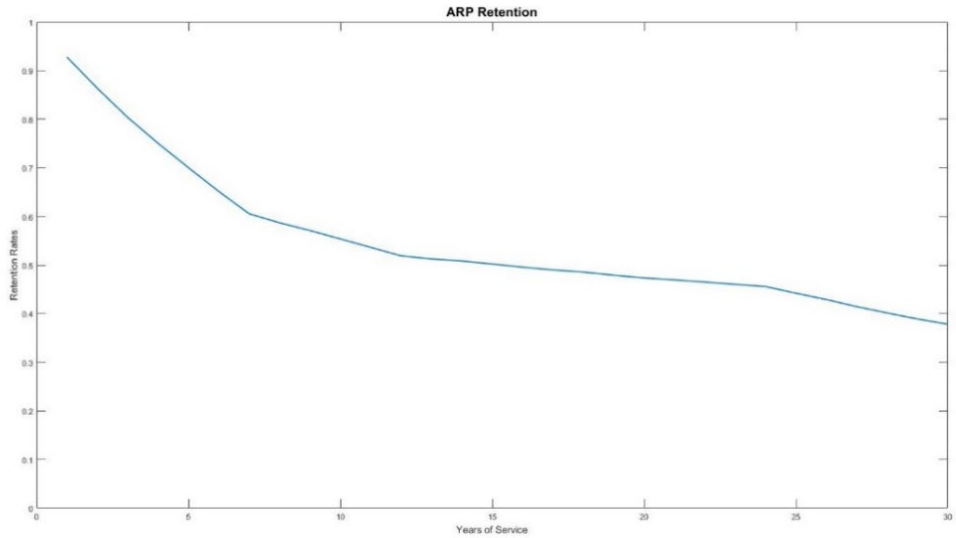

Figure 14.1. AWF Retention Rate with 2 Percent Pay Increase

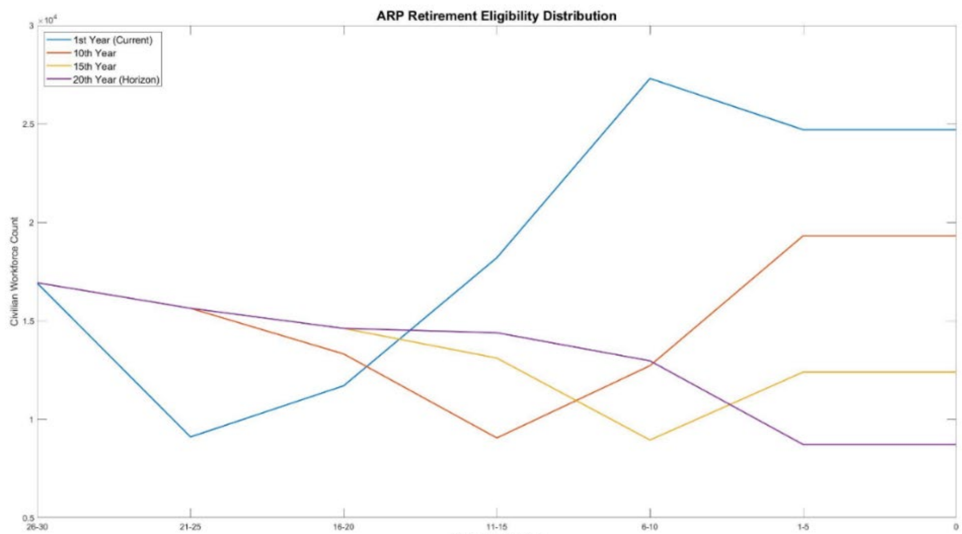

Figure 14.2. AWF Evolution Without Active and Forward-Looking Intervention with 2% Pay Increase

Figure 15.1 exhibits the retention behavior of an employee that receives a one-time bonus at Year 15 equivalent to 15% of yearly pay. That is, the employee earns a wage W_t^m =1 in all years, except in Year 15, when he or she earns W_t^m =1.15. As before, the model assumes the individual is aware at moment 0 of the bonus he or she will receive in Year 15. The main difference with respect to Figure 12 is the increased retention in the years right before year-of-service 15, as well as increased attrition in the years right after the 15-year mark. This particular behavior creates a clear kink at year-of-service 15, similar to that observed with military members around year-of-service 20.

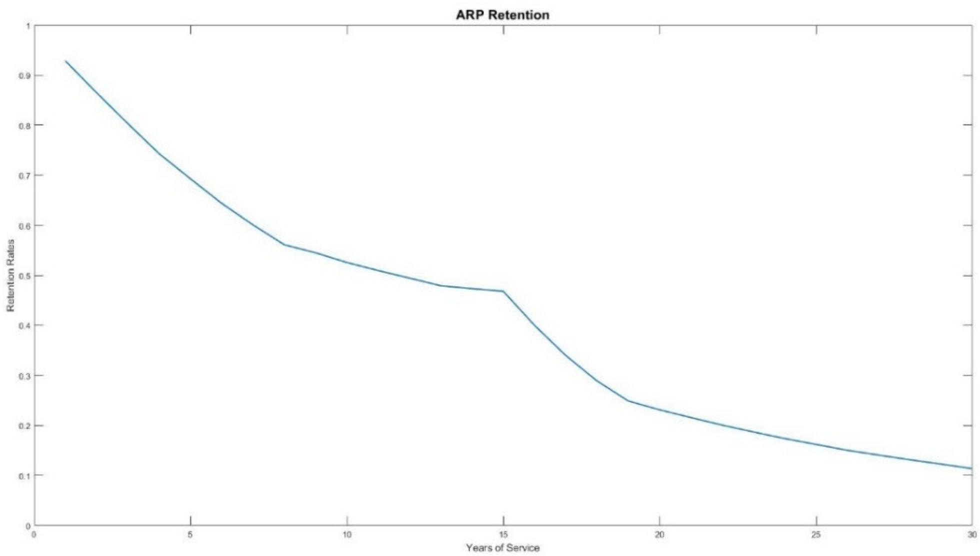

Figure 15.1. AWF Retention Rate With 15% Bonus at Year 15

The time evolution of the bathtub problem in the context of the 15-year bonus is displayed in Figure 15.2. The main departure from Figure 13 is higher retention in the initial and mid-career years but, as expected, the undesired bathtub problem remains over time.

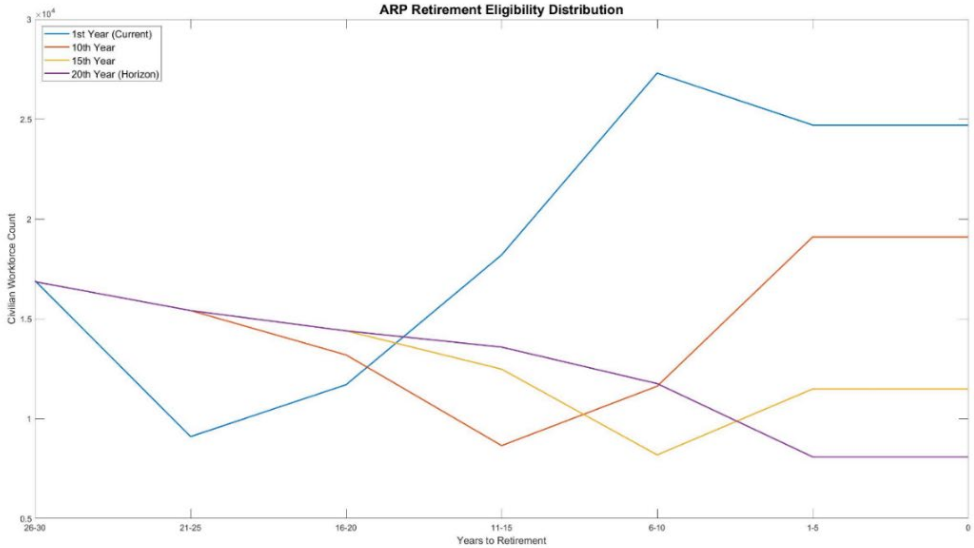

Figure 15.2. AWF Evolution Without Active and Forward-Looking Intervention with 15% Bonus at Year 15

The previous simulations of the model suggest that it can be used to forecast the evolution of the entire workforce distribution, allowing decision-makers to have a better sense of the long-term trajectory of the workforce.

The DoD has been using the original DRM for more than 30 years to predict the effect of changes in compensation on service-member retention behavior, period during which the computing power has increased dramatically. These gains in computer capabilities allow us to start introducing into the model relevant demographics of AWF employees. Thinking about extending on this proof-of-concept model, we can use the data set used in this report to estimate the model to allow us to characterize employees by groups of interest. For instance, we may estimate model parameters for employees in the Contracting field (Occupation Code 1102), which will allow the AWF leadership to understand more precisely the attrition behavior of this specific group of employees. The same can be done with other groups of interest, such as employees in the Industrial Property Management and Purchasing fields (Occupation Codes 1103 and 1105, respectively), as well as by gender, race, prior military service, disability, and level of education (or any desired combination of available demographics).

In addition, the model can be extended in different important directions. For instance, we may incorporate the effect of the FERS pension and Thrift Savings Plan

(TSP) on retention behavior. We can also extend the model to help leaders make hiring/firing decisions at various points of the employee experience distribution to achieve a target shape of the force in a certain number of years (e.g., inverted-U shape, flat, descending, or virtually any desired shape of employee distribution.) Figures 16.1, 16.2, and 16.3 show some potential desired shapes for the workforce employee distribution. The model will also be able to describe the evolution of AWF over time in special contexts, such as periods of hiring freeze or of restricted layoffs.

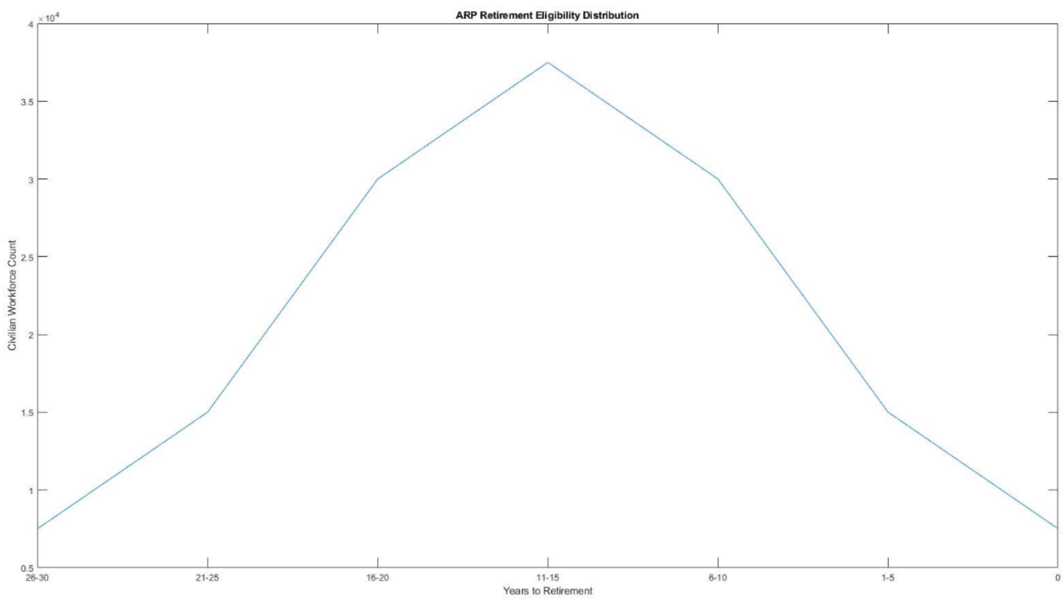

Figure 16.1 Inverted-U shape AWF Target Distribution

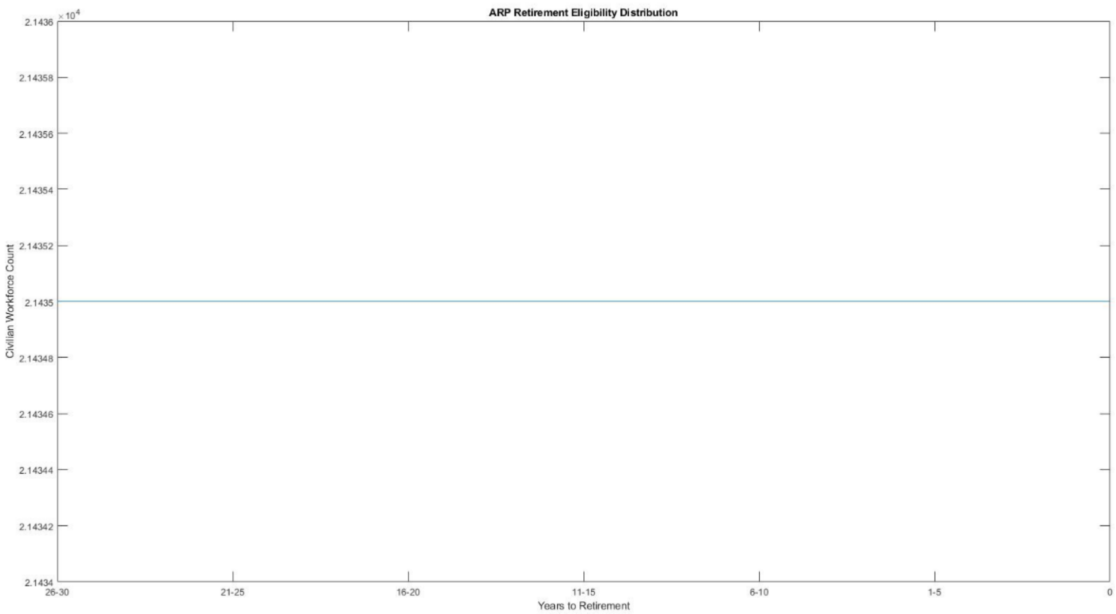

Figure 16.2. Flat Shape AWF Target Distribution

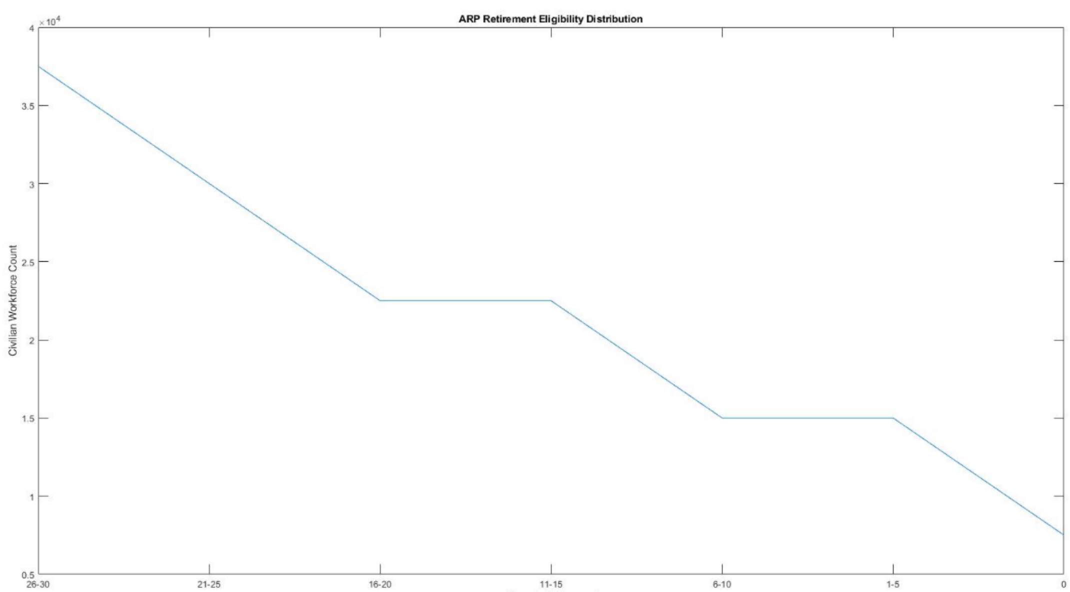

Figure 16.3 Descending Shape AWF Target Distribution

Prior Alternative Retention Models

The proof-of-concept model described in the previous section is based on the Bellman's Principle of Optimality, which guarantees that the optimal solution found at time zero for the entire maximization problem (i.e., the solution for the *T* periods) remains optimal as time passes. In other words, the individual correctly knows at time zero the optimal choices in all future years, regardless of what the future unfolds. This is known as time consistency. The caveat associated with this nice feature is the computer power required to solve the maximization problem. In an attempt to address the computing time issue, different ad hoc simplifications have been proposed over time. Inevitably, the cost of the simplifications has always been to abandon the Bellman's Principle of Optimality and accept solutions that are time inconsistent. This situation implies that the "optimal" plan of action proposed by the alternative models at time zero (most likely) becomes suboptimal over time, leading the individual to choose differently from the original plan of action.

The most well-known departure from the DRM was the Annualized Cost of Leaving (ACOL) model developed by Enns, Nelson, and Warner (1984) and Warner and Goldberg (1984). Part of the popularity of the latter model was due to its speed to solve the maximization problem and, thus, estimate model parameters, in spite of the known time inconsistency problem. However, the continuous advances in computing speeds have made working with the DRM more and more feasible over time.[9]

[9] Arkes, Ahn, Menichini, and Gates (2019) contains a more thorough review of the literature and prior models.

Conclusion

It is difficult to envision the size and shape of an efficient and effective workforce without a quantitative model that will allow decision-makers to see how many and what kind of employees must be recruited, maintained, and let-go each year. The model should inform what the workforce looks like today, what the transition path is, and what it will take to get to the desired goal a set number of years into the future.

The overarching goal of our project is to create a custom dynamic model of the AWF from the ground up that will allow the leadership to manage manpower proactively by expanding the time horizon for which labor demand, supply, and transition will be predictable. While we would not claim that the model will allow the leadership to avoid any and all disturbances or inefficiencies in shaping the workforce, it is our hope that the model will allow a more strategic, long-term planning of talent management as well as quicker and more effective responses to unforeseen political and economic gyrations that affect the AWF's mission and/or resources.

In this report, we accomplished two foundational tasks that will support the effort to build, estimate, and simulate using the final model. First, we collected and examined a subset of the workforce data, paying close attention to the career trajectory of these employees. We found that while there is relatively little separation in behavior of these employees across gender or ethnicity, several factors, such as prior military experience and education level have large impacts on career longevity. Identification of factors that affect (and perhaps just as important, those factors that do *not* affect) career decisions will be central in building and estimating the dynamic programming model.

Our second task was to build a proof-of-concept model to demonstrate some of its capabilities. After describing the model in some detail, we showed that the model is able to create an economic agent (employee) who makes rational, time-consistent decisions about whether to stay in the workforce or leave at each point in time. When we are able to create a useful model of behavior for one agent, it is then possible to create as many thousands of these virtual employees as is necessary to simulate the entire workforce. We then showed that it is possible to forecast what such a workforce would look like as it matured through time. In addition, we demonstrated that the model is capable of simulating changes in monetary compensation policies, such as a 2% annual raise or a one-time 15% bonus at a certain point in the career. When these changes take effect, the employee re-evaluates his or her decision-making and changes career trajectory logically in response. Once again, if we can correctly model this change for one employee, we can aggregate up to the entire workforce. Thus, we demonstrated the impact on the entire AWF of such a policy change as well.

The salient conclusions from our empirical analysis is that higher education level is positively correlated with career longevity. This is good news, as it seems to indicate that the AWF has little trouble retaining high quality employees. In addition, we found that employees who acquired additional education while working are even more likely to stay longer. Subject to some caveats, this implies that one way the leadership can upgrade productivity *and* retain its best people is to encourage and perhaps even subsidize continuing education for the workforce.

In subsequent analysis, we should be able to use the data to estimate the parameters of the proof-of-concept model. That is, we can generate Figures 12 to 16 such that, the behavior of the individual employee and the AWF will reflect actual labor market behavior extracted from the data. In addition, our analysis of the proof-of-concept model can lend addition insights into how active intervention from the leadership in terms of hiring and firing of employees can help to shape the AWF into its most productive and stable form.

Ultimately, it is difficult to plan for the future when we do not understand how decisions made today alter decisions and environments through time. Our research project will ensure that the AWF leadership can approach talent management with better information and more confidence to grow a diverse, agile, highly qualified, and motivated workforce of acquisition professionals.

References

Aguirregabiria, V. (1999). The dynamics of markups and inventories in retailing firms. *The Review of Economic Studies, 66*, 275–308.

Aguirregabiria, V., & Mira, P. (2002). Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models. *Econometrica, 70*(4), 1519–1543.

Aguirregabiria, V., & Mira, P. (2007). Sequential estimation of dynamic discrete games. *Econometrica, 75*(1), 1–54.

Aguirregabiria, V., & Mira, P. (2010). Dynamic discrete choice structural models: A survey. *Journal of Econometrics, 156*(1), 38–67.

Ahn, T., Menichini, A. & Tick, S. (2019) *Retention analysis model (RAM) for Navy manpower and personnel analysis* (Technical Report, GSBPP-19-006). Monterey, CA: Naval Postgraduate School.

Arcidiacono, P., & Ellickson, P. (2011). Practical methods for estimation of dynamic discrete choice models. *Annual Review of Economics, 3*, 363–394.

Arcidiacono, P., & Miller, R. A. (2011). Conditional choice probability estimation of dynamic discrete choice model with unobserved heterogeneity. *Econometrica, 79*(6), 1823–1867.

Arkes, J., Ahn, T., Menichini, A., & Gates, B. (2019) *Retention Analysis Model (RAM) for Navy Manpower Analysis* (Technical Report, GSBPP-19-003). Monterey, CA: Naval Postgraduate School.

Asch, B. J., Johnson, R. W., & Warner, J. T. (1998). *Reforming the military retirement system.* Santa Monica, CA: RAND Corporation.

Asch, B. J., Mattock, M. G., & Hosek, J. (2013). *A new tool for assessing workforce management policies over time: Extending the dynamic retention model*. Santa Monica, CA: RAND Corporation.

Asch, B. J., Mattock, M. G., & Hosek, J. (2017). *The blended retirement system: Retention effects and continuation pay cost estimates for the armed services.* Santa Monica, CA: RAND Corporation.

Asch, B. J., & Warner, J. T. (2001). A theory of compensation and personnel policy in hierarchical organizations with application to the U.S. military. *Journal of Labor Economics, 19*, 523–562.

Bajari, P., Benkard, L., & Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica, 75*(5), 1331–1370.

Black, M., Moffitt, R., & Warner, J. (1990). The dynamics of job separation: The case of federal employees. *Journal of Applied Econometrics, 5*(3), 245–262.

Burke, R. P. (2018). NavAdmin message 263/18. Retrieved: www.public.navy.mil/bupers-npc/reference/messages/Documents/NAVADMINS/NAV2018/NAV18263.txt

Cox, D. R. (1972). Regression models and life tables (with discussion). *Journal of the Royal Statistical Society, B*(24), 406–424.

Daula, T., & Moffit, R. (1995). Estimating dynamic models of quit behavior: The case of military reenlistment. *Journal of Labor Economics, 13*, 449–523.

Defense Manpower Data Center (DMDC). (2019). *APF Civilian Master File*.

Department of Defense (DoD). (2015). *DoD Acquisition Workforce Strategic Plan—FY 2016–FY 2021*. Washington, DC.

Department of Navy (DoN). (2018). *DoN Acquisition Workforce FY 19-24 Strategic Plan.* Washington, DC.

Dominitz, J., & Manski, C. (1996). Eliciting student expectations of the returns to schooling. *Journal of Human Resources, 31*, 1–26.

Dominitz, J., & Manski, C. (1997). Using expectations data to study subjective income expectations. *Journal of American Statistical Association, 92*(439), 855–867.

Doraszelski, U., & Judd, K. (2010). Avoiding the curse of dimensionality in dynamic stochastic games. *Qualitative Economics, 3*(1), 53–93.

Eckstein, Z., & Wolpin, K. I. (1990). Estimating a market equilibrium search model from panel data on individuals. *Econometrica, 58*(4), 783–808.

Enns, J. H., Nelson, G. R., & Warner, J. T. (1984). Retention and retirement: The case of the U.S. military. *Policy Sciences, 17*, 101–121.

Flinn, C., & Heckman, J. (1982). New methods for analyzing structural models of labor force dynamics. *Journal of Econometrics, 18*(1), 115–168.

Gotz, G., & McCall, J. (1984). *A dynamic retention model for Air Force officers: Theory & estimates*. Santa Monica, CA: RAND Corporation.

Gotz, G. A. (1990). Comment on "The dynamics of job separation: The case of federal employees." *Journal of Applied Econometrics, 5*(3) 263–268.

Hosek, J., Asch, B. J., Fair, C. C., Martin, C., & Mattock, M. G. (2002). *Married to the military: The employment and earnings of military wives compared with those of civilian wives*. Santa Monica, CA: RAND Corporation.

Hotz, V. J., & Miller, R. A. (1993). Conditional choice probabilities and estimation of dynamic models. *The Review of Economic Studies, 60*, 497–529.

Imai, S., Jain, N., & Ching, A., (2009). Bayesian estimation of discrete choice models. *Econometrica, 77*(6), 1865–1899.

Kasahara, H., & Shimotsu, K. (2007). *Nonparametric identification of finite mixture models of dynamic discrete choices* (Manuscript). Ontario, Canada: The University of Western Ontario, Department of Economics.

Keane, M. P., Todd, P. E., & Wolpin, K. I. (2011). The structural estimation of behavioral models: Discrete choice dynamic programming methods and applications. *Handbook of Labor Economics*. San Diego, CA: North Holland.

Keane, M. P., & Wolpin, K. I., (1994). The solution and estimation of discrete choice dynamic programming models by simulation and interpolation: Monte Carlo evidence. *The Review of Economics and Statistics, 76*, 648–672.

Keane, M. P., & Wolpin, K. I. (1997). The career decisions of young men. *Journal of Political Economy, 105*, 473–522.

Lancaster, T. (1990). *The econometric analysis of transitional data.* Cambridge, England: Cambridge University Press.

Manski, C. F. (2004). Measuring expectations. *Econometrica, 72*(5), 1329–1376.

Miller, R., (1984). Job matching and occupational choice. *Journal of Political Economy 92*, 1086–1120.

Pakes, A. (1986). Patents as options: Some estimates of the value of holding European patent stocks. *Econometrica, 54*(4), 755–784.

Pakes, A., Ostrovsky, M., & Berry, S. (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *The RAND Journal of Economics, 38*(2), 373–399.

Pesendorfer, M., & Schmidt-Dengler, P. (2008). Asymptotic least square estimators for dynamic games. *The Review of Economic Studies, 75*, 901–908.

Postel-Vinay, F., & Robin, J. (2002). Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica, 70*(6), 2295–2350.

Rota, P. (2004). Estimating labor demand with fixed costs. *International Economic Review, 45*, 25–48.

Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica, 55*(5), 999–1033.

Rust, J. (1997). Using randomization to break the curse of dimensionality. *Econometrica, 65*(3), 487–516.

Sanchez-Mangas, R. (2002). *Pseudo maximum likelihood estimation of a dynamic structural investment model* (Working Paper 02-62, Statistics and Econometrics Series). Madrid, Spain: Universidad Carlos III de Madrid.

Sargent, T., & Ljungqvist, L. (2012). *Recursive macroeconomic theory*. Cambridge, MA: MIT Press.

Slade, M. (1998). Optimal pricing with costly adjustment: Evidence from retail grocery stores. *The Review of Economic Studies, 65*(1), 87–107.

Van der Klaauw, W. (2000). On the use of expectations data in estimating structural dynamic models: An analysis of career choices. *Mimeo*. University of North Carolina.

Warner, J. T., & Goldberg, M. S. (1984). The influence of non-pecuniary factors on labor supply: The case of Navy enlisted personnel. *Review of Economics and Statistics, 66*, 26–35.

Wolpin, K. (1984). An estimable dynamic stochastic model of fertility and child mortality. *Journal of Political Economy, 92*, 852–874.

Acquisition Research Program
Graduate School of Defense Management
Naval Postgraduate School
555 Dyer Road, Ingersoll Hall
Monterey, CA 93943

www.acquisitionresearch.net