NPS-HR-21-031

ACQUISITION RESEARCH PROGRAM
SPONSORED REPORT SERIES

Retention Analysis Modeling for the Acquisition Workforce II

February 10, 2021

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

To support the modern warfighters tasked with increasing demands in a constantly changing global environment, it is imperative that the defense acquisition system continue to evolve to maintain its capability and flexibility. In this effort, growing a talented, experienced, and well-qualified civilian workforce will be vital. As part of this broad effort, the Section 809 Panel has recommended change to the Department of Defense’s career management framework to grow and augment the workforce, and the *Acquisition Workforce Strategic Plan: FY 2016–FY 2021* (Department of Defense, 2015) has emphasized efforts since 2010 to restore and restructure the AWF after a period of 20 years of shrinkage.

This technical research report is the second in a proposed series of three linked studies to provide a cutting-edge modeling and simulation tool that leverages the increase in availability of acquisition workforce (AWF) data and the large increases in computing power in the last decades. Building on the proof-of-concept model created as part of the first-year effort, we continue our development of a dynamic retention model (DRM) designed from the ground up for the AWF.

Using a large personnel data set of the AWF as well as a representative data set of the civilian population from the Bureau of Labor Statistics, we estimate our DRM. DRM is a leading-edge technique that uses a powerful mathematical/econometric technique called dynamic programming. It takes a complex, multiperiod problem (such as the lifetime labor market decisions of an acquisition worker) and breaks it down into simpler, one-period subproblems in a recursive manner. Solving a single-period problem “nests” the future decisions that the worker will make, allowing the estimation and prediction of complex behavior in a surprisingly manageable framework.

With estimates from the model, we simulate how various modifications in personnel policies, such as changes in salary structure and bonuses, would have affected the labor market decisions of the workforce. In particular, our model takes into account civilian positions that the AWF may move into upon the decision to separate from the Department of Defense, allowing a more accurate prediction of the impact of monetary personnel policies, which must be evaluated in relation to what the worker could realistically earn in the civilian sector. In doing so, the model can help the AWF leadership in achieving the desired workforce size and structure.

We conclude this report by expanding on possible extensions to enrich the model to provide yet more accurate estimation and richer simulations, including evaluating the potential impact of the COVID-19 pandemic on the long-run career trajectory of the workforce.

NPS-HR-21-031

ACQUISITION RESEARCH PROGRAM
SPONSORED REPORT SERIES

Retention Analysis Modeling for the Acquisition Workforce II

February 10, 2021

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

Having a talented, well-qualified civilian workforce is vital to the continued success of the defense acquisition system. As domestic, international, economic, and political situations change rapidly, it is more important than ever to maintain a capable and flexible acquisition workforce to support the needs of the modern warfighter.

Central to growing an exceptional workforce is the need for accurate medium to long-term forecasts of the availability, size, and composition of the AWF as well as the private-sector workers. The Department of Defense (DoD) has addressed the long-run needs of the acquisition workforce (AWF) in two recent studies. The *Acquisition Workforce Strategic Plan: FY 2016–FY 2021* (Department of Defense [DoD], 2015) described the state of the AWF after 20 years of continuous and nondirected shrinkage. Figure 1, which is borrowed from the strategic plan, showed large losses in employees who were in the midcareer segment in Fiscal Year (FY) 2008, resulting in a “bathtub” depression and leading to an alarming lack of experienced career professionals who should have formed the bulk of the productive workforce. The strategic plan described the efforts to “fill in” the bathtub, largely restoring the AWF by FY2015. The Section 809 Panel has proposed further changes to the DoD’s career management framework to improve the capabilities of the AWF.

While efforts to reshape the AWF and improve the quality of the workforce are admirable and necessary, being able to foresee these issues and judiciously pruning the workforce while recruiting, retaining, and promoting highly qualified workers would have created fewer disruptions in the efficiency of the defense acquisition system.

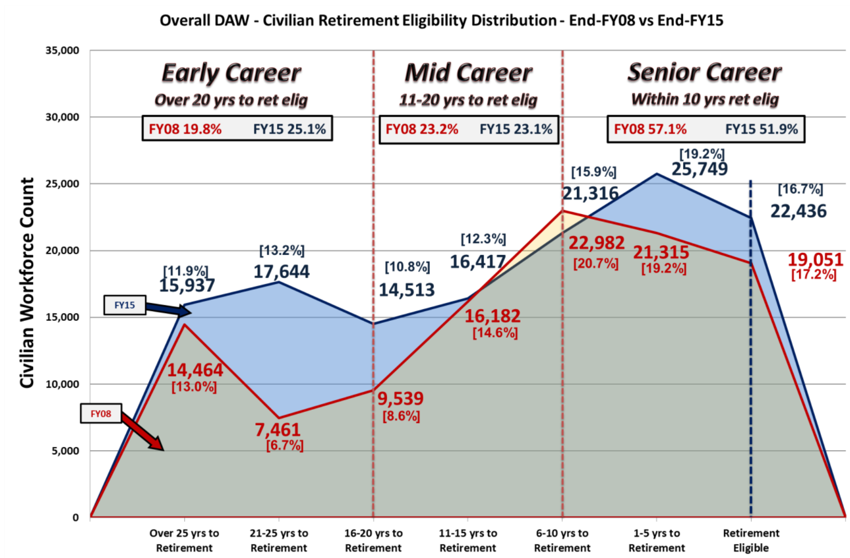

Figure 1. Civilian AWF Retirement Eligibility Distribution.
Adapted from DoD (2015).

Indeed, looking at the reshaped AWF in Figure 1, it is clear that there is much more left to be done. Approximately 50% of the AWF is within 10 years of retirement. As current midcareer employees (as of FY2015) comprise less than 25% of the AWF, the composition of the workforce is expected to change drastically within the next few years. While it is outside the scope of this study to examine whether having a “top-heavy” workforce structure is better than one where the majority of the workers are at midcareer, it is self-evident that such large changes in the number and composition of the workforce must be highly disruptive.

Unless drastic changes to the shape of the workforce is by design, such gyrations in the workforce can lead to both a degradation in performance as well as long-term instability in the personnel policies as the leadership undergoes sustained periods of large-scale hiring followed by mass layoffs to continually work at rebalancing the workforce.

In addition, a further goal of the AWF leadership is to improve the quality of the workforce, recruiting and retaining those with current (and future) in-demand skill sets, improving professionalism, and increasing diversity, all with the ultimate goal of maximizing effectiveness. To achieve all of this, a centralized, specific plan that outlines how many workers of what quality should be hired, retained, promoted, and separated is necessary. While the strategic plan and the Section 809 Panel elucidated the larger vision of how the workforce must change, they were aspirational, without the requisite specific details. To provide such specifics, a rigorous quantitative model approach is required.

In order to assist the DoD in meeting these objectives, this project continues our development of a dynamic retention model (DRM). The model is designed from the ground up for the civilian AWF. DRM is different from the simpler off-the-shelf regression analysis because it allows for complex and rigorous simulation exercises to be conducted to forecast the impact of policy changes.

DRM uses a power mathematical/econometric technique called dynamic programming, which takes a multiperiod decision problem (e.g., the lifetime labor market decisions of a representative AWF worker) and simplifies it into a single period problem, where the decisions made today impact the worker's environment tomorrow in a rigorous way, by "nesting" all future decisions the worker may make (assuming that the worker makes rational decisions) into an expression for "utility maximization" today.

To understand the relevant features of the AWF, we structure our report in the following manner:

1) We first quickly summarize our empirical findings from our Year 1 report. We had constructed a panel (longitudinal) data set to follow workers from the start of their careers to, in many cases, the end. We examined the impact of demographic and professional characteristics on career longevity, finding that highly educated workers in the AWF tended to stay longer in their careers.
2) We then review the intuitive and formal underpinnings of the DRM.
3) We then estimate the DRM using the AWF data and civilian-side data from the Bureau of Labor Statistics. The estimation itself is accomplished by using the Rust nested fixed-point algorithm.
4) With estimates from the model, we then experiment with changes in long-run personnel monetary policy changes. The estimated parameters allow us to predict how the representative worker would behave when a different pay policy is put in place. The behavior of the individual worker is then aggregated up to the workforce level to predict how the entire AWF would behave.
5) With these simulations, we provide guidance on how the shape of the workforce would evolve under passive and active hiring/firing policies.

The goal of this 3-year project is to create a DRM specifically tailored for the AWF to allow decision-makers to “look” into the future of the impact of large and potentially disruptive changes to personnel policies as well as predicted and unanticipated changes to the economic environment outside the DoD. This will aid in data-driven decision-making by AWF leaders. It is our goal that the easy-to-use analytics and visualization that can be created with output from the DRM will allow for a more robust management of risks associated with changes in hiring, promotion, and firing policies.

Beyond our original proposed output, we also propose the creation of a web-based app that will allow leadership to run a limited number of simulations directly from their desktop to examine the potential impact of policy changes.

The next section reviews the data. We then describe the simple summary, trend, and survival (regression) analysis that was done in the previous-year report. Then, after providing an intuitive description of DRM, we provide a technical review and estimate the empirical model. We then run several policy simulations to predict the impact of monetary policy changes and conclude with guidance on hiring/firing policies to actively shape the workforce.

Data

Data from Defense Manpower Data Center (DMDC)

We use a data extract (covering September 1987 to December 2018) from Defense Manpower Data Center (DMDC) of the 1,000-byte APF Civilian Personnel Master File that we obtained for our Year 1 report. The list of variables in the data set is in Table 1.

Table 1. Full List of Variables in the DMDC Extract. Source: Ahn & Menichini (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 restricted our sample to analysis of AWF workers who were ever in the Contracting, Industrial Property Management, or Purchasing fields (Occupation Codes 1102, 1103, and 1105). We also restricted our sample to workers who were born after January 1, 1950, and before December 31, 1980. Workers born prior to 1950 would have spent the majority of their careers in a labor market that may be less relevant for predicting AWF careers in the future. Workers born after 1980 would be mostly too young to provide information on long-run career outcomes.

Restricting the sample in this way, we track approximately 13,000 workers monthly, resulting in over 2 million observations. Table 2 presents some summary statistics for our sample.

Table 2. Summary Statistics for the Acquisition Workforce. Source: Ahn & Menichini (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 Postgraduate Degree	0.332
Gained Additional Education	0.441
Career Length in AWF (in months)	143.6 (103.8) [1 / 309]
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	95,144 (30,411) [27,397 / 189,600]
Observations	13,590

The representative AWF worker is white, female, and highly educated. (Over 50% have a college degree or higher at some point in their careers.) They begin their career at AWF after holding previous job(s)—the average starting age in AWF is 33—with a large fraction moving from active duty. Average tenure in the AWF is

12 years, which is almost double the average tenure observed in comparable civilian occupations.[1]

Most AWF workers hold professional or technical positions. This, along with the high education attainment and long tenure, implies that the AWF has a stable and highly capable workforce. However, it is somewhat worrying that over half the workforce has received a performance rating below fully satisfactory at some point in their careers.

Data from the Current Population Survey

We obtained a data extract of outgoing rotation group (ORG) of the Current Population Survey (CPS) from the National Bureau of Economic Research (NBER). The CPS is a Bureau of Labor Statistics (BLS) monthly survey of approximately 60,000 U.S. households. The CPS interviews a household for 4 successive months, followed by a lapse of 8 months, and then followed by another 4 months. We use data from the ORG—the households that are in their last month of interviews. This ensures that we do not count the same household multiple times in our data set. The CPS primarily focuses on employment data of persons in the household, asking about employment status and wage/income, as well as occupation characteristics, such as industry, full- or part-time status, and job tenure. Therefore, the data set serves as a representative snapshot of the U.S. civilian job market.

Table 3 shows that, compared to the AWF sample, the civilian labor force is more male, white, and less well educated. We use the civilian labor market data in our estimation to provide a reasonable estimate of the "outside option"—that is, what an AWF worker can reasonably expect to earn upon joining the civilian (private) workforce after separating.

Table 3. Summary Statistics for CPS Sample

Variables	Mean (Std. Dev) [Min / Max]
Female	0.442

[1] If interdepartmental transfers *within* the DoD are not treated as separation, AWF does not have issues holding on to workers.

Minority	0.105
Age	39.6 (11.6) [16 / 90]
Weekly Earnings	644.7 (392.2) [0 / 2884.6]
Has Postgraduate Degree	0.140
Work in Government Sector	0.143
Observations	125,828

Note: Sample covers ORG from 1987 to 2009.

THIS PAGE LEFT INTENTIONALLY BLANK

Review of Trend and Regression Analysis

In this section, we provide a brief summary of some of the relevant trends observed in the AWF data from our Year 1 report.[2] Figure 2 shows the gender differences in attrition rates. We confirm that the AWF is overall stable, with approximately 70% of the original workers still present after 8 to 9 years. After about 25 years of service, about 25% of workers remain. Attrition rate is similar across genders, although women tend to have slightly longer careers. This contrasts with the civilian labor force career trajectories. From the CPS, while average job tenure for men has declined from 8.3 years to 7.4 years from 1983 to 2012, average job tenure for women has increased from 5.8 years to 6.9 years. It should be noted that men still have longer careers in the civilian sector.

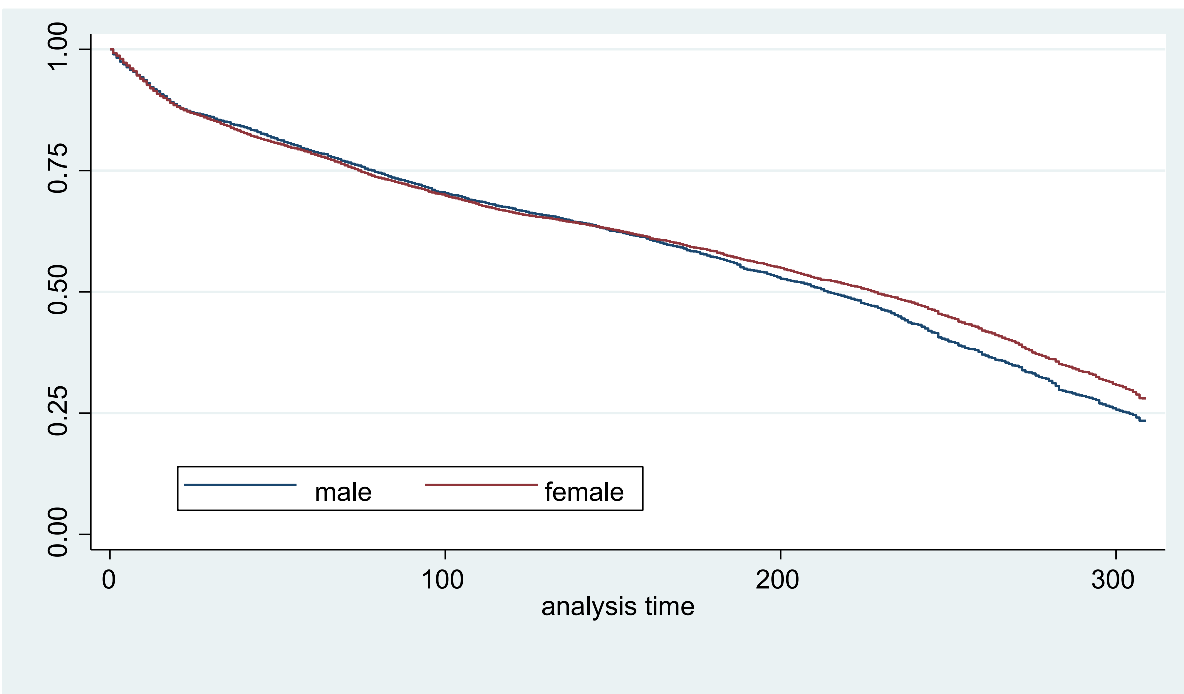

Figure 2. Career Trajectory of Workers by Gender. Adapted from Ahn & Menichini (2019).

Figure 3 shows differences in attrition by ethnicity. Minority (African American and Hispanic) and White (plus Asian) workers have similar career lengths. Minority

[2] Our summary of the trend analysis is necessarily compressed. Readers are encouraged to read Ahn & Menichini (2019) for a complete description of the AWF sample.

workers do have slightly shorter careers—especially comparing across workers who are at midcareer (between 7 and 15 years of service)—but they last longer in the AWF provided they get past midcareer. In addition, we also examined the impact of education on career longevity in Figure 4. Taking the worker's highest degree obtained (whether education was obtained prior to entry or acquired during), we find a strong positive relationship between education level and career length. While less than 50% of those with a high school degree or lower remain in the AWF past 15 years of service, those with post-undergraduate degrees are retained at over 50% well past 20 years of service. As we mentioned in the Year 1 report, while this is encouraging, it is possible that those with lower levels of education leave the AWF for schooling.

In addition to trend analysis, we estimated a more formal (reduced form) model of the AWF using duration (survival) analysis. Survival analysis has been extensively used in biostatistics, demography, economics, and actuarial sciences. We estimated four different specifications of the Cox proportional hazard model.[3]

[3] Readers are referred to Cox (1972) and Ahn & Menichini (2019) for a detailed description of the survival analysis set up.

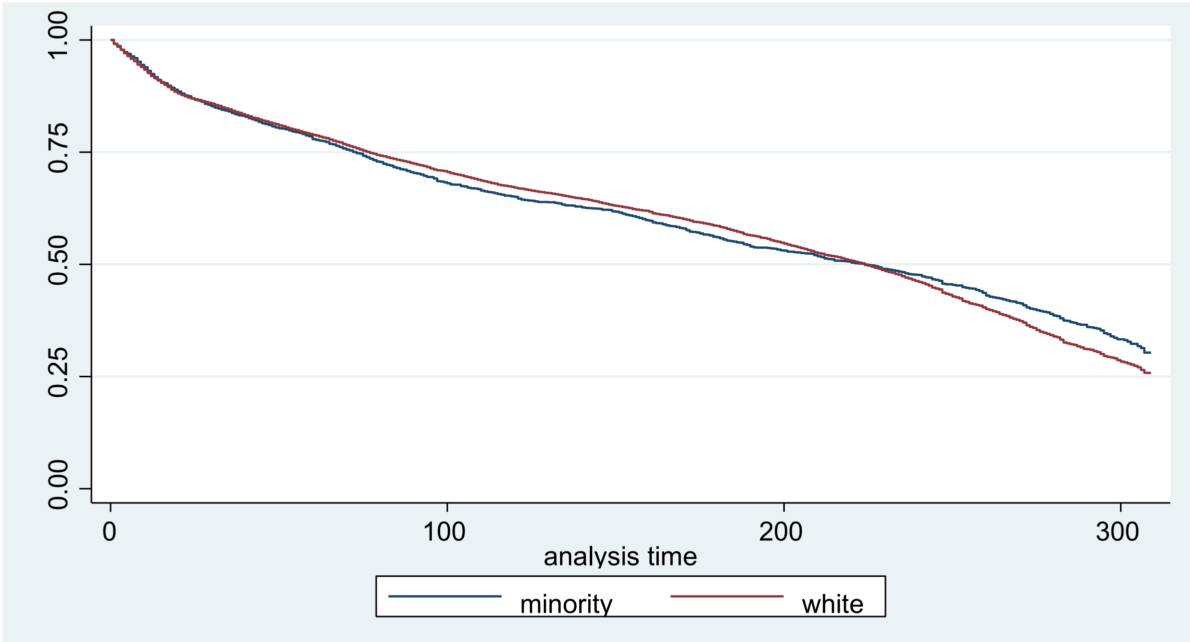

Figure 3. Career Trajectory of Workers by Minority Status. Adapted from Ahn & Menichini (2019).

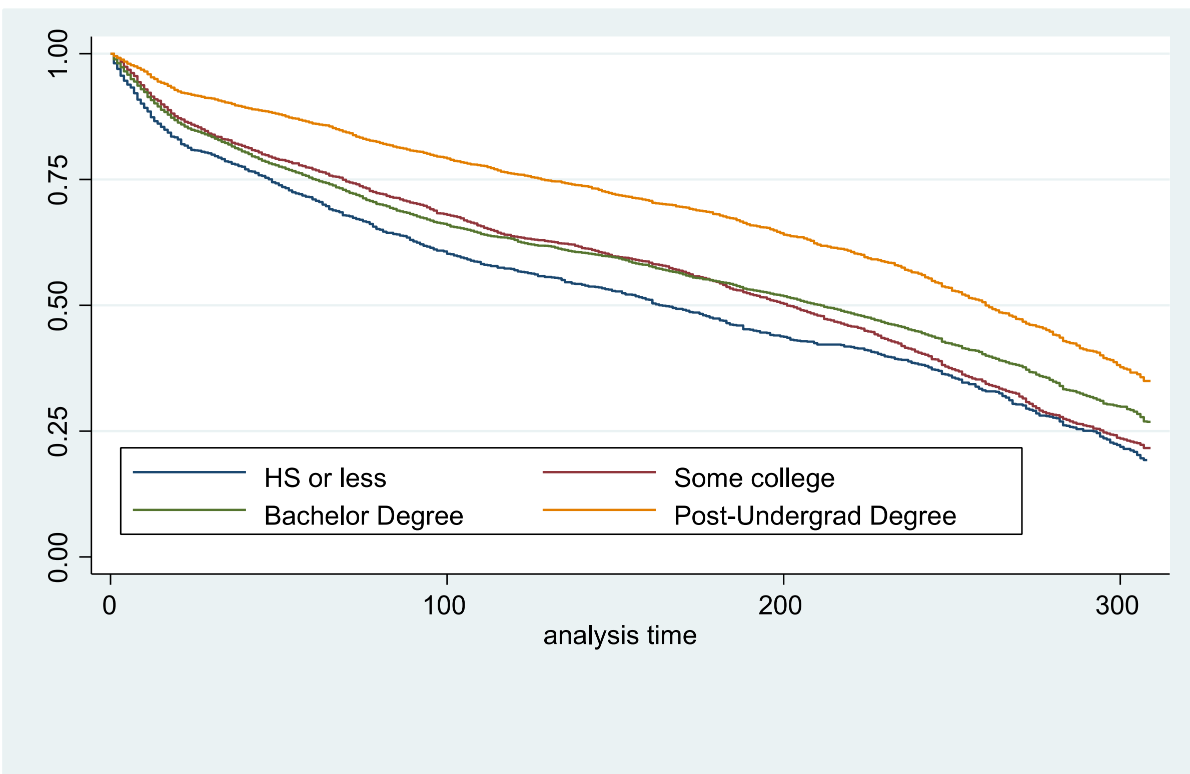

Figure 4. Career Trajectory of Workers by Education Level. Adapted from Ahn & Menichini (2019).

Coefficient estimates from the four different specifications, along with their hazard ratios are presented in Table 4. A positive coefficient estimate (with a hazard ratio greater than 1) indicates a higher probability of leaving the AWF. A negative estimate (with a hazard ratio less than 1) indicates the opposite.

Overall, we see that minority status does not impact attrition. Female and disability status yield moderately longer careers. By far the strongest predictor for a long AWF career is prior military experience. We speculate that this may be due to specific human capital match of those who were active duty in the past. There may be a high degree of overlap in culture and day-to-day tasks, allowing for these workers to succeed long-term in the AWF.

Table 4. Cox Proportional Hazard Model Parameter and Hazard Ratio Estimates. Adapted from Ahn & Menichini (2019).

	Model 1		Model 2		Model 3		Model 4	
	Coef.	Hazard	Coef.	Hazard	Coef.	Hazard	Coef.	Hazard
Female	−0.1866*	0.8298	-	0.7952	-	0.8505	−0.1126*	0.8935
	(0.0252)	(0.000)	0.2292*	(0.000)	0.1619*	(0.000)	(0.0261)	(0.000)
			(0.0259)		(0.0262)			
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	-	0.8771	-	0.8910	−0.0723§	0.9303
	(0.0327)	(0.000)	0.1312*	(0.000)	0.1154*	(0.000)	(0.0328)	(0.028)
			(0.0327)		(0.0327)			
Prior Military	−3.0036*	0.0496	-	0.0508	-	0.0516	−3.0574*	0.0470
	(0.0358)	(0.000)	2.9681*	(0.000)	2.9652*	(0.000)	(0.0384)	(0.000)
			(0.0361)		(0.0364)			
BA Degree	-	-	-	0.8986	−0.0050	0.9950	0.0319	1.0324
			0.1069*	(0.000)	(0.0275)	(0.841)	(0.0267)	(0.231)
			(0.0242)					
Post-BA	-	-	-	0.8523	−0.0051	0.9949	−0.0626§	0.9393
			0.1598*	(0.000)	(0.0297)	(0.863)	(0.0314)	(0.046)
			(0.0282)					
Add'n	-	-	-	-	-	0.6368	-0.3025*	0.7389
Degree					0.4513*	(0.000)	(0.0274)	(0.000)
					(0.0272)			
Professional	-	-	-	-	-	-	-1.2607*	0.2835
							(0.0295)	(0.000)
Technical	-	-	-	-	-	-	-1.0919*	0.3356
							(0.0359)	(0.000)
Deficient	-	-	-	-	-	-	-1.2102*	0.2981
Rank							(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 parenthesis. For hazard ratios, *p* values are in parenthesis.

We next describe how our dynamic model is constructed, starting with a nontechnical description.

THIS PAGE LEFT INTENTIONALLY BLANK

General Description of Dynamic Programming[4]

In this section, we provide a nontechnical 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 describe below, one of the principal issues with previous attempts at estimating a dynamic model has been time inconsistency. This means that using the estimated parameters from these prior models to simulate worker 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 models 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 their 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 DoD is the Dynamic Retention Model (DRM), pioneered in the early 1980s by the RAND Corporation. 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 Blended Retirement System (BRS) was examined with the DRM. Dynamic programming simplifies a complex, multiperiod problem (for example, an officer's lifetime labor market decisions) into a series of much simpler, single-period subproblems using backward recursion. The single-period problem contains a value that captures future decisions that the officer will make, which

[4] This section is a verbatim reproduction from our Year-1 report. We leave in the full description to aid the reader in understanding our main model, without having to refer back to Ahn & Menichini (2019).

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 they could earn in the civilian market, and so on. 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 do not adjust for the strength of the economy or service member quality. In addition, the model cannot handle nonmonetary 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 scenario, 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 5, for Periods 1 and 2, the person would select (H,H) = \$300 to maximize total payoff.

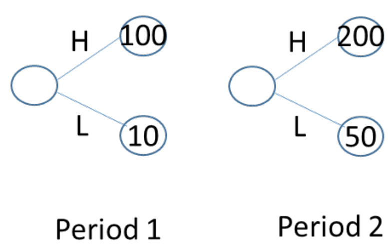

Figure 5. Simple Choice Across Independent Time Periods

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 6, 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 pay out.

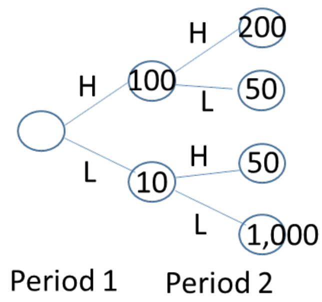

Figure 6. Choice Across Connected Time Periods

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. With two periods, there are four possibilities, as we saw in Figures 5 and 6. With three or four periods, the number of choices (and thus calculations) increases to eight and 16, respectively. Over a 30-period span, there are 1,073,741,824 possible outcomes.[5] It would be very time-consuming and ultimately wasteful to calculate all 1+ billion outcomes, since most outcomes/scenarios would be such undesirable and unlikely outcomes that no rational person would make such 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. As long as we can describe this optimal decision as a number, we

[5] 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. In Model 2.0 we plan to allow agents to make an additional third choice of attaining extra human capital while remaining in the AWF.

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 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 a worker is faced with the decision to retire or not, they will be comparing the monetary benefit of staying (quantifiable as $A) and the nonmonetary benefits (not necessarily quantifiable as B) against the monetary benefits ($C) and nonmonetary benefits (D) of leaving. If the worker stays in the AWF, then we know

$$\$\mathrm{A} + \mathrm{B} \geq \$\mathrm{C} + \mathrm{D}$$

If they opt to leave, we know

$$\$\mathrm{A} + \mathrm{B} < \$\mathrm{C} + \mathrm{D}$$

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 nonmonetary 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 also must distinguish how agents value money in relation to other nonmonetary 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 et al. (2001) and Asch and Warner (2001) analyzed 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 et al. (2002) extended the model to include the initial decision to enlist, looking specifically at IT workers in the military. Asch et al. (2017) extended 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 et al. (2017) examined the potential impact of changes to the BRS across the services. Gotz (1990) contained 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.[6]

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 increases to over 1 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 a number of required calculations that surpasses the number of atoms in the universe. This rapid growth in the "state space" that we have to track makes the computation burdensome (many times to the point of impossible) and is called the curse of dimensionality.[7]

---

[6] 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.

[7] 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.

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 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 (i.e., 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 include Slade (1998), Aguirregabiria (1999), Sanchez-Mangas (2002), and Rota (2004).[8]

[8] 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 et al. (2009), for example.

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 (Aguirregabiria & Mira [2007]; Arcidiacono & Ellikson [2011]; Arcidiacono & Miller [2011]; Kasahara & Simotsu [2007]).[9]

[9] Note that we do not make use of these empirical innovations in our Model 1.0. We may introduce these concepts in subsequent versions. Models become much more complicated and take longer to estimate once unobserved heterogeneity is introduced.

Technical Description of Dynamic Programming Models

In this section, we describe in detail the DRM that will be used in the estimation procedure as well as in the simulations. We begin describing the agent's decision problem, which is identical to the one in the Year 1 report.

We assume AWF employees are rational individuals who maximize their lifetime utility (or happiness) when they make career choices. This means that, in each decision, the agent considers all the associated costs and benefits, both monetary and nonmonetary (explained below). In the current DRM, the remaining life of an individual is divided in equal-length periods. At the beginning of each of those periods (e.g., monthly, quarterly, or yearly) the employee chooses to stay in the AWF for another period or to leave to work in the private sector. We keep the assumption that the decision to leave is irreversible—that is, returning to the AWF after departing is not possible. The possibility to return is a model feature that could be added in a future report and could give us information on actions that the AWF could take to recover employees.

As mentioned in the previous paragraph, the agent balances all pecuniary and nonpecuniary costs and benefits at the time of each decision. The monetary elements include (1) AWF compensation, such as basic pay, locality adjustment, health insurance, bonuses, and so on and (2) the analog compensation from the private sector. In particular, an expansion we make compared to the Year 1 report is the explicit modeling of the retirement systems. For the AWF worker, we assume they are covered by the Federal Employees Retirement System (FERS). For the private-sector employee, we assume the employer matches the employee's contributions to a 401K retirement account up to 10% of their gross pay.

The nonmonetary components include the preference or taste of the decision-maker for an AWF versus a private-sector job. For instance, there might be agents who prefer a job in the public sector due to its higher predictability and stability, even when they know they could augment their income should they switch to the private sector. On the other hand, there might be individuals who prefer private-sector jobs due to their usually higher compensation schemes, even if they are aware of the lower stability of those jobs. These relative nonpecuniary elements are introduced into the DRM as "taste parameters" that reflect monetary-equivalent preferences for public- versus private-sector jobs.

Basic notation of the DRM is as follows:

- W_t^m denotes the AWF compensation (including all monetary elements) in period t
- W_t^c indicates the compensation that the agent obtains in the private sector in period t
- ω^m is the taste parameter capturing the monetary equivalent preference for an AWF job
- ω^c is the taste parameter capturing the monetary equivalent preference for a job in the private sector
- T indicates the time horizon (number of periods before retirement)
- $\beta = \frac{1}{1+r}$ is the discount factor, where r represents the subjective discount rate
- $E[.]$ denotes the expectation operator
- ε_t^m and ε_t^c are stochastic shocks affecting the utility (and the decision) of the individual in every period.

The next equations describe the maximization problem of the AWF employee:

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

$$V_t^S = W_t^m + \omega^m + \beta E[V_{t+1}] + \varepsilon_t^m, \quad (2)$$

$$V_t = Max[V_t^L, V_t^S] \quad (3)$$

where super-index L refers to the decision to leave the AWF to work in the private sector, while super-index S indicates the decision to continue working in the AWF one more period. Accordingly, we let V_t^L denote the (present) value for the agent of leaving the AWF and V_t^S indicate the (present) value for the individual of staying in the AWF for another period. The decision-maker will keep choosing to stay in the AWF while the value of staying, V_t^S, exceeds the value of leaving, V_t^L. In the period in which the opposite happens (i.e., $V_t^L > V_t^S$), the agent leaves the AWF. To introduce the model in a simple way, we have omitted subindex *i* in Equations 1–3, which refers to individual *i*. We will introduce it in later sections.

In contrast to popular, ad hoc models of military retention (e.g., adjusted cost of leaving [ACOL]), the DRM has the advantage of generating time-consistent courses of action. This means that the original plan of action remains always optimal as time passes. This optimality feature is not always present in other retention models, such as ACOL, as they yield courses of action that may become suboptimal as time passes (i.e., the model is dynamically or time inconsistent).[10] However, the DRM is usually more computationally burdensome compared to the simpler models.

Jointly with the other retention models (e.g., ACOL), the DRM makes a number of important assumptions. For instance, the agent knows the time horizon T is able to foresee both income streams (i.e., W_t^m and W_t^c) and has constant taste parameters (i.e., ω^c and ω^m) over the agent's life. In addition, the AWF employee is assumed to have linear utility functions and to know the parametric distribution of the random shocks (i.e., ε_t^c and ε_t^m). These assumptions are made mainly to keep the computational tractability of the model.

Private-Sector Wage Calculations

An empirical innovation worth discussing is our treatment of the private-sector wage (W_t^c). Most of the elements we need are raw data from DMDC or parameters derived from the model. The exception is the private-sector wage, which serves as the "outside option" for AWF workers. Each period, the worker must gauge the stream of income they expect to earn in the AWF as well as outside in the private sector. While we have usable data of AWF salaries from DMDC, once a worker decides to separate, their salary in the private sector is unobserved.

Being able to accurately assess a worker's outside option is critical to evaluating their labor market decisions. Imagine a scenario where it is assumed that there is no difference between earnings in the government and private sector. Then,

---

[10] Technically speaking, the course of action does not satisfy Bellman's principle of optimality.

any decision on the part of the worker to leave or stay must be purely dependent on preference for government work compared to private-sector work. While personal preferences no doubt play a large role in decisions of where and/or whether to work, it is somewhat disingenuous to examine potential impacts of personnel policy changes, most of which will involve changes to compensation structure, without accurately accounting for the worker's potential earnings outside the AWF.

In addition, the geographic location of the worker can also impact the relative attractiveness of the outside option. While the government sector does offer locality adjustments to account for large cost-of-living differences across different states and cities, it is unclear whether this adjustment accounts for differences in consumer goods prices accurately across the entire United States. In addition, it is unclear how much these adjustments track private-sector wage differences across regions. For example, the average private-sector salary in California (which has the highest cost-of-living in the continental United States) is 59.1% higher than Mississippi (lowest cost of living). The locality adjustment in San Francisco currently stands at 41.44%. This implies that if government and private-sector workers in Mississippi are comparably compensated, then government workers are being underpaid compared to private-sector workers in California.

We rely on the data from the CPS to construct the gap between AWF and private-sector wages. In the CPS, the survey respondent identifies which state they reside/work in, whether they are working in the government or private sector, and time spent in the labor market. Using this information, it is possible to identify for each state and experience level the gap in pay between government and private-sector workers. We match this information to the AWF worker in the data set, which then allows us to accurately determine, at each point in time, the outside monetary option facing the worker.

Estimation and Econometric Results

In this section, we describe the path-breaking estimation procedure proposed by Rust (1987). It is a nested algorithm that Rust called NFXP-ML, which stands for *nested fixed point–maximum likelihood*. Broadly speaking, it is a maximum likelihood estimation procedure that consists in two nested algorithms: an "inner" algorithm and an "outer" algorithm. In the inner step, the DRM described in the previous section is solved, assuming certain values for the parameters. In the outer step, a search is done for the parameter values that maximize the likelihood function. We next explain the procedure and its implementation.

The first step is to make some additional assumptions regarding the DRM. The key assumption in the Rust model is the parametric distribution of the stochastic innovations (i.e., ε_t^c and ε_t^m), which are supposed to follow an independently and identically distributed (*iid*) bivariate Type-I Extreme Value process. The importance of this assumption is that it ultimately enables us to use the discrete choice model of McFadden (1974). In particular, Rust assumed that ε_t^c and ε_t^m are *iid* Type-I Extreme Value random variables with location parameter -γ (where γ≈0.5772157) and scale parameter 1. This yields a normalized mean and variance of 0 and $\pi^2/6$, respectively, for both shocks. It also means that ε_t^c and ε_t^m are *iid* over individuals (*i* dimension, to be added later) and over time (*t* dimension), which rules out unobserved heterogeneity across individuals and persistent shocks over time. In terms of the maximization problem cast in Equations 1–3, the "error terms" ε_t^c and ε_t^m are state variables that are observed by the AWF employee at the time of making the decision but unobserved by the econometrician.

With the previous assumptions about the random innovations, we can continue developing the model as

$$V_t^L \sim Extreme\ Value(W_t^c + \omega^c + \beta E[V_{t+1}^L] - \gamma, 1), \quad (4)$$

$$V_t^S \sim Extreme\ Value(W_t^m + \omega^m + \beta E[V_{t+1}] - \gamma, 1), \quad (5)$$

$$E[V_t] = ln\left\{exp\left[E[V_t^L]\right] + exp\left[E[V_t^S]\right]\right\} \quad (6)$$

The previous three equations mean that the value functions can be calculated in closed form, substantially reducing the computational requirements of the maximization problem.

The next step is to acknowledge that Equations 1–6 are also individual specific and, thus, we add subindex *i* to refer to a particular agent. Then, Equation 3 becomes

$$V_{it} = Max\left[V_{it}^{L}, V_{it}^{S}\right] \quad (7)$$

$$V_{it} = Max[W_{it}^{c} + \omega_i^{c} + \beta E[V_{it+1}^{L}] + \varepsilon_{it}^{c}, W_{it}^{m} + \omega_i^{m} + \beta E[V_{it+1}] + \varepsilon_{it}^{m}] \quad (8)$$

We let $d_{it} \in \{0,1\}$ be an indicator function of the event "individual *i* leaves the AWF in period *t*." That is, $d_{it} = 1$ when agent *i* leaves the force in period *t*, and $d_{it} = 0$ if they decide to stay one more period. In terms of utility, we have

$$d_{it} = 1 \; if \; W_{it}^{c} + \omega_i^{c} + \beta E[V_{it+1}^{L}] + \varepsilon_{it}^{c} \; \geq \; W_{it}^{m} + \omega_i^{m} + \beta E[V_{it+1}] + \varepsilon_{it}^{m} \quad (9)$$

$$d_{it} = 0 \; otherwise \quad (10)$$

Furthermore, we let

$$\tilde{\varepsilon}_{it} = \varepsilon_{it}^{m} - \varepsilon_{it}^{c} \quad (11)$$

$$\widetilde{U}_{it} = \left[W_{it}^{c} + \omega_i^{c} + \beta E[V_{it+1}^{L}]\right] - \left[W_{it}^{m} + \omega_i^{m} + \beta E[V_{it+1}]\right] \quad (12)$$

The assumptions made about the stochastic terms lead us to

$$\tilde{\varepsilon}_{it} \sim Logistic(0,1) \quad (13)$$

and the conditional choice probability of employee *i* leaving the AWF in period *t* becomes

$$P(d_{it} = 1) = F_{\tilde{\varepsilon}}\left(\tilde{\varepsilon}_{it} \leq \widetilde{U}_{it}\right) \quad (14)$$

$$P(d_{it} = 1) = \frac{1}{1 + exp(-\widetilde{U}_{it})} \quad (15)$$

Finally, the Logit model to be estimated is

$$P(d_{it} = 1|\theta) = \frac{1}{1 + exp(\theta' Z_{it})} \quad (16)$$

where θ' is the vector containing the parameters to be estimated and Z_{it} is a vector containing the data described above as well as the value functions derived in Equations 1–8. The individual characteristics in vector Z_{it} help us capture the effects of factors affecting the difference in taste parameters (i.e., $\omega_i^m - \omega_i^c$). That is, positive values of the parameters imply a lower probability of leaving (i.e., larger probability of staying) or, equivalently, higher preference for the AWF versus the civilian sector.

In summary, the model proposed by Rust (1987) involves two nested algorithms:

1. An “outer” optimization algorithm that searches for the values of the unknown parameters in θ to maximize the likelihood function implied by Equation 16.
2. An “inner” optimization algorithm that calculates the value functions described in Equations 1–8, given the data and the parameter values set in the outer algorithm.

Table 7 shows the parameter values estimated by the procedure outlined above for three different models. As we mentioned, a positive (and larger) parameter value indicates a higher preference for the AWF as opposed to the private sector and, thus, lower probability of leaving the force. The opposite is true for negative (and lower) values of the parameters. The *p* values are in parentheses.

Each of the three models contains a different set of variables. Model 1 includes only exogenous variables—that is, individual characteristics that are set before the employee starts working in the AWF. As we showed in the Year 1 report, the coefficient on female is significantly positive, suggesting longer careers in the AWF. We find analogue results for prior military experience. Regarding race, we find that Hispanics and Native Americans also tend to have longer tenure in the force, while the opposite is true for African Americans and Asians.

Model 2 also includes the education level attained by the individual. An important difference with the variables in Model 1 is that education is a decision of the employee and, thus, is not exogenous. In fact, we observe in the data set that a considerable number of employees acquire more education while working in the

AWF. This lack of exogeneity implies that the resulting associations with longevity need to be interpreted more cautiously. The table shows that individuals with higher education, such as post–high school studies, baccalaureate degrees, and post-baccalaureate degrees, are significantly positively associated with longer careers.

Finally, Model 3 adds some other individual characteristics, such as being a job of a professional, technical, white-collar, and/or blue-collar nature, as well as whether the employee was ever ranked as deficient in their job. In agreement with the outcomes in the Year 1 report, the parameter estimates related to job nature turn out to be significantly positive, which suggests that those types of jobs are associated with individuals working more time in the AWF. Regarding ever being ranked deficient, the parameter estimate is also significantly positive. This result might suggest that those employees who are never ranked deficient (i.e., good workers) tend to leave the AWF earlier. Overall, the results from the estimation of the DRM are consistent with the outcomes obtained in the Year 1 report.

Table 7. Model Parameter Estimated Values

Variable	*Model 1*	*Model 2*	*Model 3*
Intercept	-0.02***	-0.07***	-0.36***
	(0.00)	(0.00)	(0.00)
Female	0.03***	0.01**	0
	(0.00)	(0.02)	(0.43)
African-Am.	-0.01**	0	0.02***
	(0.01)	(0.38)	(0.00)
Asian	-0.04***	0.09***	0.06***
	(0.00)	(0.00)	(0.00)
Native Am.	0.02	13.76***	3.67***
	(0.21)	(0.00)	(0.00)
Hispanic	0.01	0.08***	-0.01
	(0.16)	(0.00)	(0.31)
Disability	0	0.02***	0.01*
	(0.40)	(0.00)	(0.05)
Prior Military	0.19***	0.2***	0.16***
	(0.00)	(0.00)	(0.00)
Some Post HS		0.09***	0.03***
		(0.00)	(0.00)
BA Degree		0.03***	0
		(0.00)	(0.41)
Post BA		0.02***	0.03***
		(0.00)	(0.00)
Professional			0.15***
			(0.00)
Technical			0.08***
			(0.00)
White Collar			0.08***
			(0.00)
Blue Collar			7.35***
			(0.00)
Deficient Rank			0.3***
			(0.00)

Note: *, **, *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

THIS PAGE LEFT INTENTIONALLY BLANK

Simulation Results

In this section, we show the results from simulating the DRM using the parameter values estimated in the previous section. For the simulations, we depart from the assumption in the last section that the stochastic innovations (i.e., ε_t^c and ε_t^m) follow an *iid* bivariate Type-I Extreme Value process and instead assume they follow an *iid* bivariate normal process. The objective of this replacement is twofold. First, the following simulations become more comparable with the simulations of the Year 1 report, for which we also assumed normal random variables. Second, simulations of multivariate processes are computationally easier to handle under the normality assumption. From a parametric perspective, that change has minimal impact on outcomes, as both distributions share a somewhat similar shape determined by a location and scale parameters.

To simulate the model described in Equations 1–3, we first need to select its parameter values. We use the parameter values estimated in the previous section as well as some features observed in the data to simulate the model for a representative AWF worker. The initial values are shown in Table 6.

Table 6. Parameter Values

Parameter	*Value*
W_t^m	1
W_t^c	1.15
T	30
β	0.95
ω^m	1.2
ω^c	1
$\mu_{\varepsilon,m}$	0
$\mu_{\varepsilon,c}$	0
$\sigma_{\varepsilon,m}$	0.1
$\sigma_{\varepsilon,c}$	0.1

We can see in Table 6 that the parameter values are constant over the career of the representative employee. Regarding compensation, we assume income from the private sector (i.e., W_t^c) is around 15% higher than in the AWF (i.e., W_t^m). This relationship is roughly observed in the available data for representative workers. For that reason, we let $W_t^c = 1.15$ and $W_t^m = 1$. We also assume that individuals have a time horizon (*T*) of 30 years, implying that they expect to work for that number of years before final retirement. We let the discount factor (β) be 0.95, which is equivalent to a subjective yearly discount rate of roughly 5.26%.

In the previous section, we used the Rust procedure to estimate three different models that help us capture the difference in taste parameters (i.e., $\omega_i = \omega_i^m - \omega_i^c$). Employing Model 1, which includes only exogenous variables, that difference is estimated at around 0.2, which suggests that the representative AWF worker has a higher preference for the AWF as opposed to the private sector. Accordingly, Table 6 displays a value of $\omega_i^m = 1.2$ and $\omega_i^c = 1$. Finally, we let the stochastic innovations (i.e., ε_t^c and ε_t^m) be *iid* normal random variables with zero mean and standard deviations equal to 0.1. This final assumption matches that in the Year 1 report.

The following step is to simulate the model with the parameterization just described. The most basic model prediction is the retention behavior of the representative AWF employee. Figure 7 shows the likelihood that the worker leaves the force in each particular year of their career. The probability of leaving is relatively low in all years with a clear downward slope as time passes, implying that the likelihood of leaving diminishes as time passes and the employee settles in their job. For instance, the likelihood that the individual leaves the AWF in Year 5 is close to 10%.

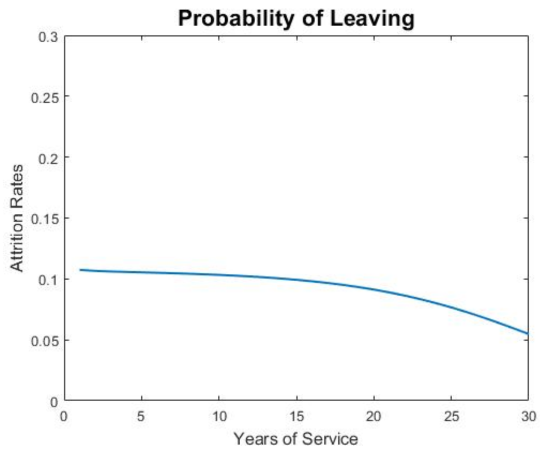

Figure 7. Probability of Leaving for a Representative AWF Employee

Associated with the previous probability of leaving the AWF is the survival curve of the individual, which we display in Figure 8. The survival curve shows the cumulative probability of the employee remaining in the force after a certain period of time. The figure shows that, for instance, the likelihood that the individual is still working in the AWF after 5 years is approximately 60%.

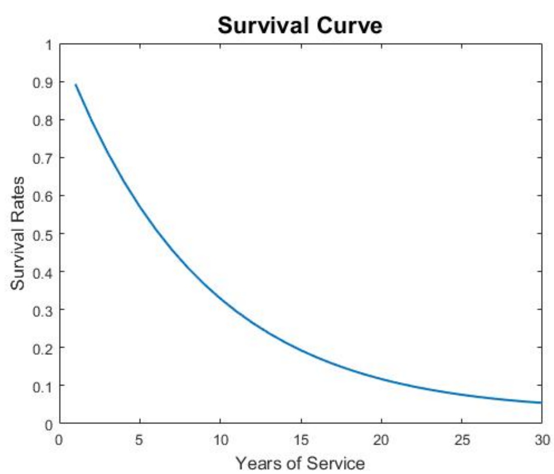

Figure 8. Survival Curve for a Representative AWF Employee

The probability of staying in the AWF and the associated survival curve can be influenced in several ways. In this report, we consider two policy changes that can be readily implemented by the AWF leadership—that is, bonus payments (i.e., lump sums) and salary increments (i.e., permanent sums). We start analyzing the impact of bonuses by assuming the AWF decides to pay a 20% bonus at Year of Service 15. This means that the salary of the employee is $W_t^m = 1$ in every period, except in their Career Year 15, when their income is $W_t^m = 1.20$. We also assume the employee knows about the future bonus at moment 0. Figure 9 displays the effect of the bonus on the probability of leaving the AWF. It is clear that the main effect of the bonus is the relatively lower attrition rate before the payment and the relatively increased probability of leaving the AWF in the periods right after it is paid.

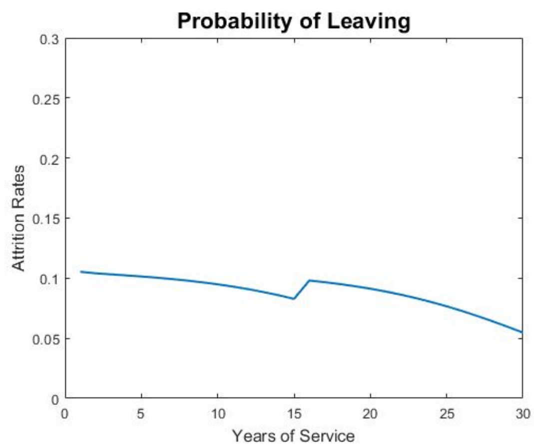

Figure 9. Probability of Leaving With a 20% Bonus at 15 Years of Service

The impact of the bonus on the survival curve is shown in Figure 10. Compared to the baseline situation in Figure 8, the bonus creates an upward kink in the survival curve around the time it is paid (i.e., Year of Service 15), reflecting the relatively lower probability of leaving the force. This behavior is very similar to that of service members (i.e., enlisted members and officers) before and after the 20-year-of-service mark.

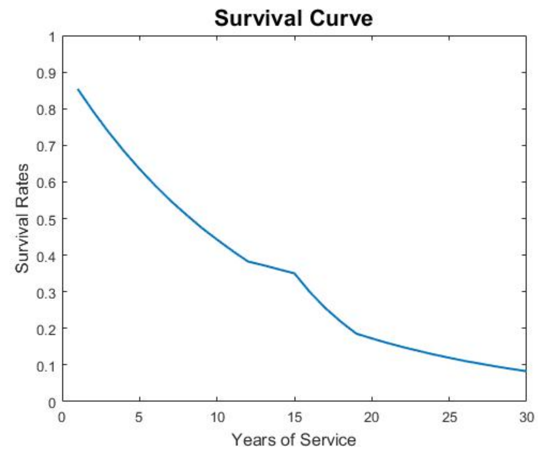

Figure 10. Survival Curve With a 20% Bonus at 15 Years of Service

As benchmark, we investigate the effects of paying two bonuses instead of only one. The first bonus is 10% of the employee salary, paid in Year of Service 10, and the second bonus is also 10% of the individual income, paid in Year of Service 20. As before, the agent knows they will receive those bonuses in the corresponding years. The impact of those lump sum payments on the probability of leaving the AWF is shown in Figure 11. Similar to Figure 9, the main effect of the bonuses is decreased attrition in the years prior to the bonus payments and increased attrition in the periods right after those payments.

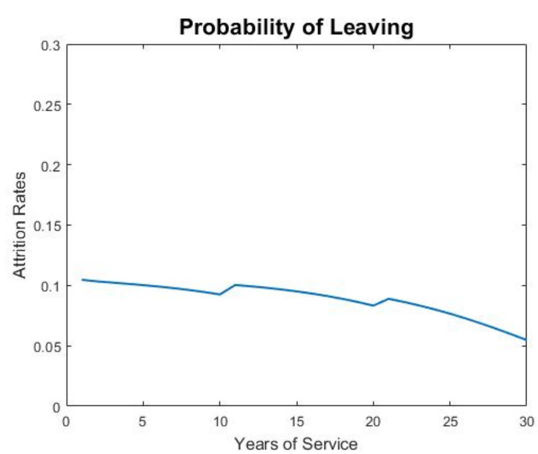

Figure 11. Probability of Leaving With a 10% Bonus at 10 Years of Service and a 10% Bonus at 20 Years of Service

Compared to the base case in Figure 8, as expected, the payment of two bonuses affects the survival curve by creating two kinks at the payment dates. This effect is shown in Figure 12, where the upward kinks can be observed at Years of Service 10 and 20.

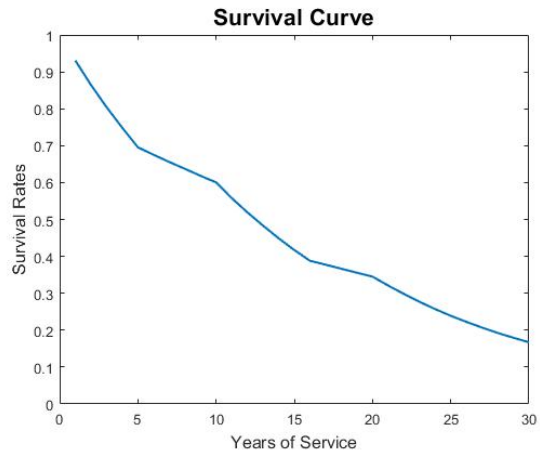

Figure 12. Survival Curve With a 10% Bonus at 10 Years of Service and a 10% Bonus at 20 Years of Service

The second policy change we investigate is the influence of a permanent salary increase. We start assuming a 5% pay increment starting at Year of Service 15. This implies that the employee income is $W_t^m = 1$ in the first 15 years and $W_t^m = 1.05$ in the final 15 years. We further assume that, at moment 0, the individual is aware of this future salary change. The effect of the wage increase is shown in Figure 13. The figure suggests that the likelihood of leaving the AWF is highest in the early career years, falling to a relative minimum during the midcareer years—more specifically, the relative minimum coincides with the start of the salary increase—and finally slightly falling in the final years.

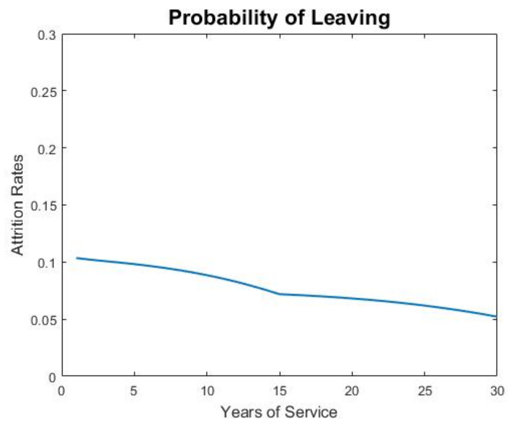

Figure 13. Probability of Leaving With a 5% Salary Increase from Year of Service 15

The survival curve associated with the previous probability of leaving is shown in Figure 14. Compared to the baseline situation in Figure 8, it is clear that the pay increase considerably diminishes employee attrition in the midcareer years, making the survival curve more horizontal in final career years.

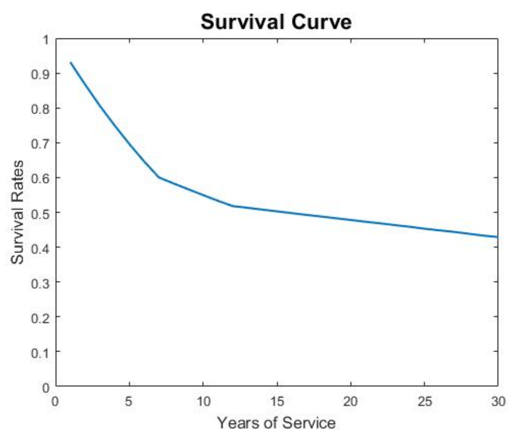

Figure 14. Survival Curve With a 5% Salary Increase from Year of Service 15

As a final exercise, we consider the impact of two permanent salary increases instead of one. The first one is a 3% pay increment starting in Year of Service 10, while the second one is another 3% salary increase since Year of Service 20. Figures 15 and 16 display the results of such policy change. Consistent with Figures 13 and 14, the key impact of the income increments is diminished attrition during the midcareer years of the employee and the leveling of the survival curve in the late career years.

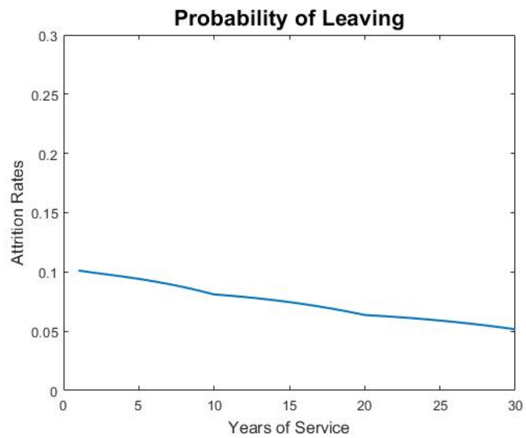

Figure 15. Probability of Leaving With a 3% Salary Increase from Year of Service 10 and Another 3% Salary Increase from Year of Service 20

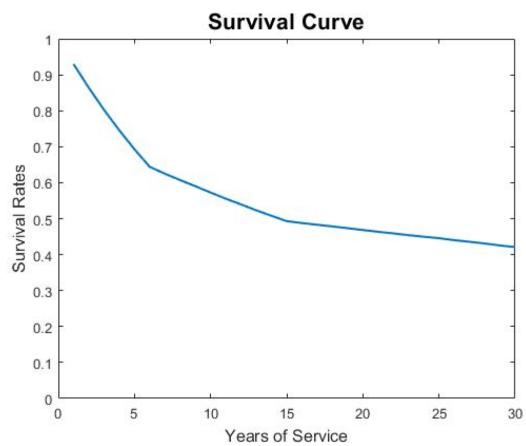

Figure 16. Survival Curve With a 3% Salary Increase from Year of Service 10 and Another 3% Salary Increase from Year of Service 20

The DoD's (2015) *Acquisition Workforce Strategic Plan: FY 2016–FY 2021* describes the "bathtub problem" as the lack of employees between 5 and 15 years of experience in the AWF, after a long sequence of years of personnel reduction. We reproduce that situation in Figure 17, where the horizontal axis shows the time (i.e., number of years) before retirement, while the vertical axis depicts the total number of employees. The figure considers a total number of 130,000 AWF employees, which is in line with the information in the strategic plan at the time of the bathtub problem.

Figures 7 and 8 show the retention behavior predicted by the model for a representative AWF worker, highlighting that attrition starts at a relatively high level and diminishes monotonically as the employee progresses in their career. Considering this basic behavior, we simulate the evolution of the bathtub problem assuming a myopic personnel policy. That is, we assume the AWF only hires new employees at the start of their careers (i.e., in the first year) with the sole objective to maintain the existing number of employees in Year 1. This assumption is equivalent to the AWF leadership having no long-run objective for the overall structure of the AWF force. While probably quite unrealistic, this exercise allows us to understand the evolution of the force structure (i.e., the bathtub problem) over time in a context of no active intervention by the AWF. Figure 18 shows the results of applying this policy for a period of 20 years. It is apparent that the bathtub problem would be moving to the right of the figure as time passes, suggesting a gradual diminishing in the proportion of experienced employees. In the long run, the overall structure of the AWF would look similar to the retention behavior shown in Figure 8.

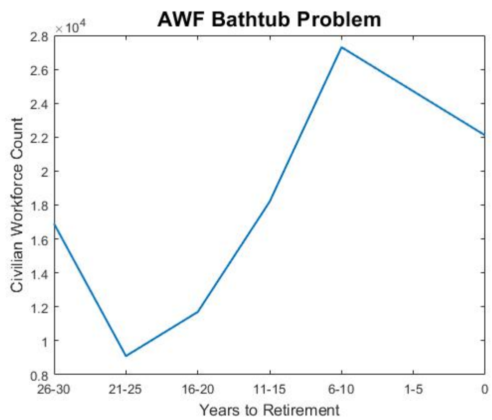

Figure 17. AWF Bathtub Problem

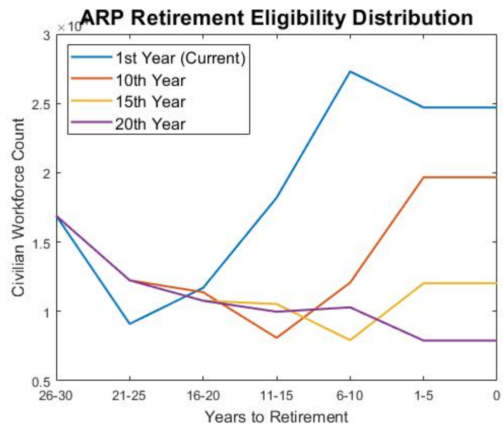

Figure 18. Evolution of AWF Structure Without Active Intervention

While the previous figure shows the evolution over time of the AWF structure with no active intervention, it is reasonable to expect that the AWF will use all the available tools to change that path. Among those tools, we can mention direct hiring/laying off of employees as well as retention efforts such as salary raises and bonuses. In fact, according to the strategic plan, the AWF leadership has indeed been addressing the bathtub problem by executing a considerable amount of hiring and retention policies in recent years. In this context, one of the main advantages of the DRM is that it can assist the AWF in the implementation of those policies by explaining expected employee behavior as well as the projected time evolution of the entire force structure. We start evaluating retention policies (i.e., bonuses and pay raises) and then we analyze hiring/laying off of personnel.

Figures 7 through 16 described the effect of bonuses and salary raises on employee retention behavior. We next study the impact of these retention efforts on the evolution of the bathtub problem over time. Figure 19 displays the results of the AWF paying a 20% bonus at Year of Service 15 to all employees arriving at that mark. Compared to Figure 18, Figure 19 shows higher retention in the initial and midcareer years, though in the long run the force exhibits a decreasing pattern similar to the survival curve in Figure 8.

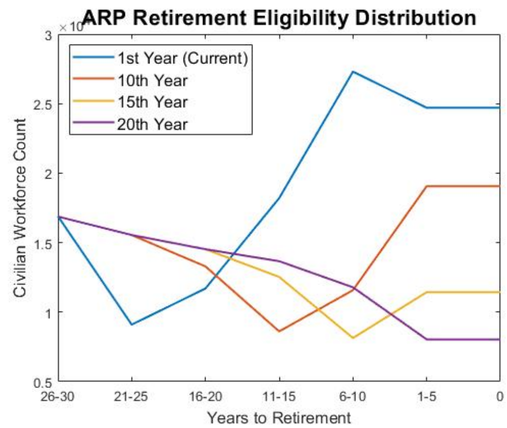

Figure 19. Evolution of AWF Structure Without Active Intervention With a 20% Bonus at 15 Years of Service

The second retention effort we study is a permanent salary raise. In particular, we assume a 5% pay increment at the beginning of Year of Service 15. Figure 20 shows the evolution of the bathtub problem over time in the context of that wage increase. Similar to the bonus, the main departure from Figure 18 is higher retention in the early and midcareer years, with the overall AWF structure displaying a decreasing pattern over time.

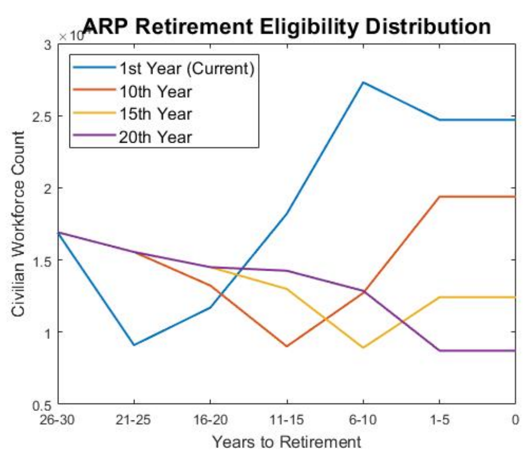

Figure 20. Evolution of AWF Structure Without Active Intervention With a 5% Wage Increase at 15 Years of Service

The main takeaway from the previous analyses is that targeted bonuses and salary raises have limited usefulness if the AWF leadership desires a long-run shape for the force different from the decreasing one. That is, the AWF needs to actively make strategic hiring/laying-off decisions at different points of the distribution of employee tenure if it wants to achieve a particular objective force shape. It is probably in this situation when the benefits of the DRM are the highest point. In this line, we next explain how to use the DRM to help the AWF leadership make such decisions with the objective to achieve a certain target force structure in a determined number of years. The DRM is flexible enough to accommodate virtually any shape of employee distribution as well as any time horizon to achieve it. The AWF leaders just need to define the desired shape and horizon, and the DRM will produce the optimal policies that will achieve those objectives. For instance, Figures 21, 22, and 23 show some possible shapes of employee distributions that the AWF might desire (e.g., inverted-U shape, flat, or descending, respectively). The three figures are created assuming that, starting from 130,000 employees, the AWF also aims to have a total number of 160,000 employees, which is in line with the situation described by the strategic plan.

To illustrate more clearly how the DRM could help with this feat, we next assume that—starting from the bathtub problem shown in Figure 17, with an initial number of 130,000 employees—the objective of the AWF is to achieve a flat force structure (like the one in Figure 22) in a 10-year horizon period, with a total of 160,000 employees. The simulation of the DRM creates the optimal hiring/laying-off decisions at each point of the worker experience distribution in each year until the final horizon. These decisions also incorporate the effect of the expected employee retention behavior depicted in Figure 8. The results from this experiment are shown in Figure 24. The figure displays how the overall shape of the AWF distribution evolves over time under the implementation of the hiring policy suggested by the DRM. It starts with the blue line as the current situation (i.e., the bathtub problem) and finishes with the purple line in 10 years. The latter represents the shape of the AWF distribution at the horizon, which is flat as desired. The red and yellow lines represent intermediate personnel distributions (at the 4th and 7th years, respectively) before the final objective is achieved.

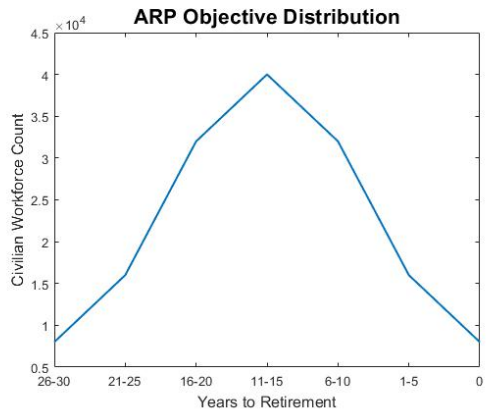

Figure 21. Inverted-U Shape AWF Target Distribution

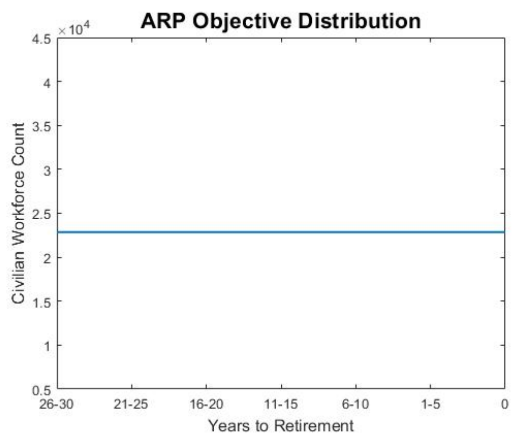

Figure 22. Flat Shape AWF Target Distribution

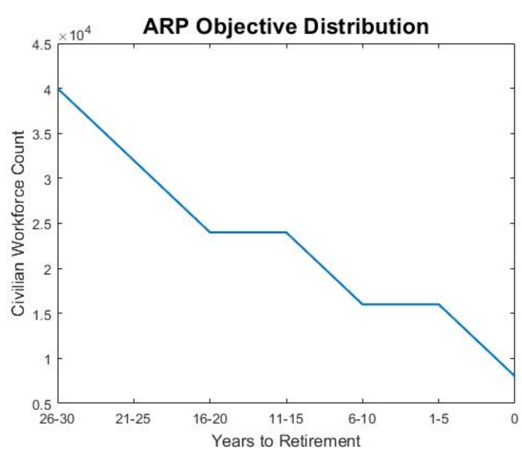

Figure 23. Descending Shape AWF Target Distribution

The previous AWF personnel distribution at the horizon date (i.e., 10 years) is the result of the implementation of the hiring/laying-off decisions proposed by the DRM, starting from the bathtub problem in Year 1. The specific hiring/laying-off decisions over time are displayed in Figure 25. The latter shows the number of people to hire or lay off in each of the 10 years. It is clear from the picture that achieving a flat shape with a total of 160,000 employees (starting from the bathtub with 130,000 employees) implies heavy hiring in the initial career years, with a slow decrease in hiring in the middle and final career years. The blue line displays the number of employees to hire/lay off during the first year at each point of the worker career, while the green line does the same for the final 10th year, when the overall flat shape for the AWF distribution is achieved (the purple line in Figure 24).

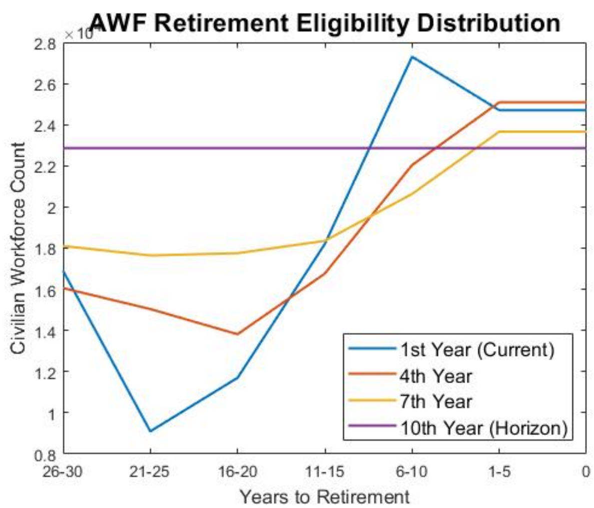

Figure 24. Flat Shape AWF Target Distribution Over a 10-Year Horizon

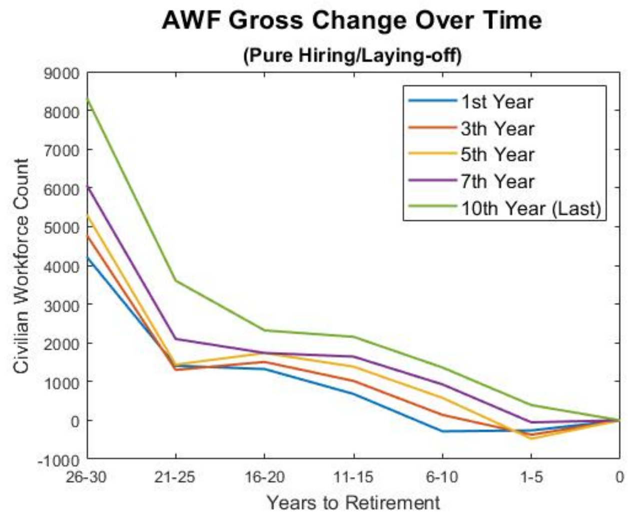

Figure 25. AWF Gross Change Over a 10-Year Horizon

Figure 25 shows the exact number of AWF employees to hire/lay off (i.e., the gross change), but not the actual or final change in the number of AWF employees (i.e., the net change). This difference happens because the figure does not account for the expected natural worker attrition. When the latter is incorporated into the analysis, we obtain the net change in AWF personnel, which shows how the AWF total number of employees at each point of the worker career will be changing over time. We show the net change in Figure 26.

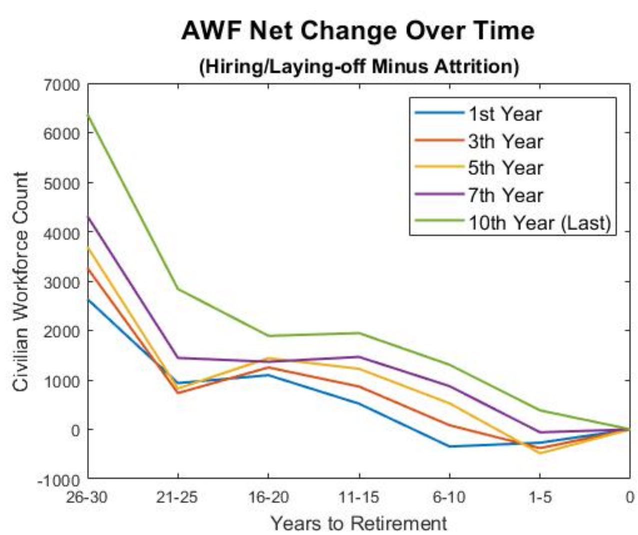

Figure 26. AWF Net Change Over a 10-Year Horizon

The previous exercise shows an instance of how to use the DRM to help the AWF leadership manage its personnel, taking into account the optimal employee behavior. It is also straightforward to extend the previous analysis to use the DRM jointly with one or more bonuses and/or salary raises, as we described. This flexibility of the DRM creates a very large number of potential policies that the AWF could explore in order to find the one that achieves its short- and long-run personnel objectives and, at the same time, satisfies its budgetary restrictions.

Conclusion

In this report, we built upon the foundational tasks completed in the Year 1 report (data acquisition and proof-of-concept model) to estimate the DRM using AWF and CPS data, generate coefficient estimates, and run policy simulations to predict behavior of individual workers as well as the evolution of the shape of the workforce through time in response to these policy actions.

The estimation of the model and simulations of the DRM with changes to pay structures reveal that workers will respond to one-time pay bumps in the form of bonuses by temporarily delaying separation (just long enough to attain the bonus) and then quickly exit. Permanent salary increases induce workers to stay longer, but this is significantly more expensive.

The final set of simulations showed exactly how many workers at particular experience levels should be hired or laid off in order to achieve different shapes of the AWF. Attrition behavior of workers at different years of service is summed up to calculate the workforce without active intervention. Then, assuming a final desired shape, the DRM accurately prescribes hiring behavior within a defined time frame.

The qualitative conclusions are not surprising, nor should they be. If our model is designed and estimated correctly, the predicted behavior of individual workers should match our intuition and prior observations of their labor market outcomes. The estimated model coefficients are reasonable, and the simulation of individual attrition probabilities broadly lines up with our reduced form results from our Year 1 report.

The value of the analysis, then, lies in our ability to take the rational agents (workers) we have created for the DRM estimation, run them through complex policy changes, and generate precise predictions on how many and what kind of workers will remain in the AWF at particular points in time.

Of course, we do not claim to be able to perfectly predict the future. The model predictions arise from the concept of a stable equilibrium, where once the policy changes are implemented, worker preferences, AWF personnel policies, and economic circumstances are static: they do not actively change. In reality, all of the above necessarily react to each other and change, leading to worker behavior that will diverge from the predictions of the model.

Therefore, this model does not allow AWF leadership to make a once-and-for-all policy change that will unalterably set the course of evolution of the workforce into its desired shape and size. Instead, we view this model as providing the best guess at the evolution of the AWF given current circumstances, which allows the leadership to make data-driven decisions about how many workers to hire and retain to chart the correct general course. As the economic environment, demographic makeup of the workforce at large, and personnel policies change, the model should be re-estimated and new simulations run to provide more up-to-date guidance.

DRM will be most useful as one of the predictive tools to allow the leadership to manage manpower proactively by extending the time horizon over which workforce size and shape will be predictable. The model should allow for a more strategic plan for talent management to more efficiently carry out the AWF's mission.

Preview of Year 3 Report

In next year's report, we will extend the model to examine the following:

- **Assess the impact of employee quality.** When the leadership identifies AWF deficits and alters policies to impact retention behavior, care must be taken to ensure that high ability workers are retained while those with the lowest level of skills, training, or education are encouraged to attrite.
- **Incorporate the effects of the state of the economy.** As retention will be impacted by the state of the economy as well as the national and international environments, the model must take into account macroeconomic changes. For example, the COVID-19 pandemic is expected to drastically alter the state of the labor market for the foreseeable future. In particular, private demand for acquisition workers may be significantly reduced, and workers everywhere may be induced to work from home over a longer time frame.

- **Perform the previous analyses for specific career fields.** The civilian AWF is composed of the fourth estate, Defense Acquisition University, Defense Contract Management Agency, Defense Logistics Agency, business, contracting, engineering groups, information technology, and so on—all having their unique workforces with different goals for recruitment and retention.
- **With the extensions to the model above, we will conduct additional policy simulations,** including one or more permanent pay increases at specific career years; one or more bonuses paid at specific career years; increased rate of pay increase for AWF (change in GS scale); change in FERS pension annuity computation formula (akin to the BRS in active duty); economic expansions; recessions (modeled as random, unforeseen macroeconomic shocks); and other scenarios as requested by AWF leadership.

THIS PAGE LEFT INTENTIONALLY BLANK

References

Aguirregabiria, V. (1999). The dynamics of markups and inventories in retailing firms. *Review of Economic* Studies, *66*(2), 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. (2019). *Retention analysis modeling for the acquisition workforce*. Naval Postgraduate School. https://dair.nps.edu/bitstream/123456789/2775/1/NPS-HR-20-001.pdf

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.

Asch, B. J., Johnson, R. W., & Warner, J. T. (2001). *Reforming the military retirement system* (MR-748-OSD). RAND. https://www.rand.org/pubs/monograph_reports/MR748.html

Asch, B. J., Mattock, M. G., & Hosek, J. (2017). *The blended retirement system: Retention effects and continuation pay cost estimates for the armed services* (RR-1887-OSD/USCG). RAND. https://www.rand.org/pubs/research_reports/RR1887.html

Asch, B. J., & Warner, J. T. (2001). A theory of compensation and personnel policy in hierarchical organizations with application to the US military. *Journal of Labor Economics*, *19*(3), 523–562.

Bajari, P., Benkard, L., & Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, *75*(5), 1331–1370. https://web.stanford.edu/~lanierb/research/Estimating_Dynamic_Models_EMA.pdf

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.

Cox, D. R. (1972). Regression models and life tables (with discussion). *Journal of the Royal Statistical Society*, *34*(2), 187–220. http://www.biecek.pl/statystykaMedyczna/cox.pdf

Daula, T., & Moffit, R. (1995). Estimating dynamic models of quit behavior: The case of military reenlistment. *Journal of Labor Economics*, *13*(3), 499–523.

Department of Defense. (2015). *Acquisition workforce strategic plan: FY 2016–FY2021*. https://asc.army.mil/web/wp-content/uploads/2019/03/DoD_Acq_Workforce_HSCP_FY16-21.pdf

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

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

Doraszelski, U., & Judd, K. (2012). Avoiding the curse of dimensionality in dynamic stochastic games. *Quantitative Economics*, *3*(1), 53–93.

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.

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

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* (R-3028-AF). RAND. https://www.rand.org/pubs/reports/R3028.html

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* (MR-1565-OSD). RAND. https://www.rand.org/content/dam/rand/pubs/monograph_reports/2009/MR1565.pdf

Hotz, V. J., & Miller, R. A. (1993). Conditional choice probabilities and estimation of dynamic models. *Review of Economic Studies*, *60*(3), 497–529.

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

https://openaccess.city.ac.uk/id/eprint/5366/1/Bayesian%20jain_econometrica.pdf

Kasahara, H., & Shimotsu, K. (2007). *Nonparametric identification of finite mixture models of dynamic discrete choices* [Unpublished manuscript].

Keane, M. P., Todd, P. E., & Wolpin, K. I. (2011). The structural estimation of behavioral models: Discrete choice dynamic programming methods and applications. In O. Ashenfelter & D. Card (Eds.), *Handbook of Labor Economics* (vol. 4B, pp. 331–461). 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. *Review of Economics and Statistics*, *76*(4), 648–672.

Keane, M. P., & Wolpin, K. I. (1997). The career decisions of young men. *Journal of Political Economy*, *105*(3), 473–522.

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

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

McFadden, D. (1984) “Econometric analysis of qualitative response models.” In Handbook of Econometrics, (vol 2, pp. 1395-1457). Elsevier.

Miller, R. (1984). Job matching and occupational choice. *Journal of Political Economy*, *92*(6), 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). *RAND Journal of Economics*, *38*(2), 373–399. https://web.stanford.edu/~ost/papers/dynamic_games.pdf

Pesendorfer, M., & Schmidt-Dengler, P. (2008). Asymptotic least square estimators for dynamic games. *Review of Economic Studies*, *75*, 901–928. http://www.its.caltech.edu/~mshum/gradio/papers/pesendorferschmidt.pdf

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*(1), 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]. Statistics and Econometrics Series. Universidad Carlos III de Madrid.

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

Slade, M. (1998). Optimal pricing with costly adjustment: Evidence from retail grocery stores. *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* [Unpublished manuscript]. 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*(1), 26–35.

Wolpin, K. (1984). An estimable dynamic stochastic model of fertility and child mortality. *Journal of Political Economy*, *92*(5), 852–874.

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

www.acquisitionresearch.net