NPS-HR-22-008

ACQUISITION RESEARCH PROGRAM
SPONSORED REPORT SERIES

Individualized and Optimal Talent Management of the AWF in Response to COVID-19: Dynamic Programming Approach

December 6, 2021

Dr. Sae Young (Tom) Ahn, Assistant Professor

Dr. Amilcar A. Menichini, Associate Professor

Graduate School of Defense Management

Naval Postgraduate School

Prepared for the Naval Postgraduate School, Monterey, CA 93943.

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

Executive Summary

This report is an extension of the originally proposed sequence of three studies that developed a cutting-edge modeling and simulation tool for the acquisition workforce (AWF). The initial objective of that sequence was to build a Dynamic Retention Model (DRM) from the ground up for the AWF to restore and maintain a capable and flexible AWF in support of the needs of the modern warfighter.

The current report uses the previous model to analyze the phenomenal and unprecedented impact of the COVID-19 pandemic on the U.S. civilian sector and its potential effects on the size and composition of the AWF in the coming years. After going steadily down for almost a decade and being at the historical low of 3.5% in February 2020, the U.S. unemployment rate spiked to almost 15% in April 2020. This event represented an unparalleled increase of more than 11% in just 2 months. As surprising as the initial increase was, the sharp fall in the U.S. unemployment rate that followed was just as remarkable. As of November 2021, just a year and a half after the peak, the unemployment rate is hovering around 4.6%, barely more than one percentage point above the previous historical low.

While the impact of COVID-19 so far has been much harsher on civilian-sector employment than on the government sector (and the AWF), it is unclear how the latter will evolve in the long-run after the fast, ongoing recovery of the private sector. We take advantage of the DRM developed in the previous studies and extend it to explore the potential consequences of economy-wide shocks (such as COVID-19) on the AWF as the economy shows signs of strong recovery.

We start analyzing the behavior of a representative AWF worker at the beginning of the pandemic, when the strength of the economic recovery was highly uncertain. We find that, under a number of different scenarios regarding the speed of recovery, it takes several years (in expectation) before the AWF employee returns to the pre-pandemic behavior. The main effect of the COVID-19 shock is to make the AWF job temporarily more attractive than a similar job in the private sector, inducing the AWF worker to stay much longer in the government.

A caveat of the previous analysis is that it assumes that the AWF employee is able to predict (in expectation) the recovery path of the economy. To address that unrealistic feature of the analysis, we extend the initial study by “forcing” the AWF worker to go through the strong economic recovery path observed after the outset of the pandemic. That is, we predict the agent behavior when the recovery paths are much more positive than originally forecasted. Not surprisingly, the initial higher valuation of the AWF job compared to the private sector quickly dissipates, and AWF attrition rates surge above pre-pandemic levels as employees who were planning to move to the private sector (and froze their plans due to the pandemic) resume their original courses of action. An important takeaway is that, while the COVID-19 shock may initially induce more employees to stay longer in the AWF, it is not a permanent solution to retain valuable workers. To this end, traditional personnel policy actions will be required by the AWF leadership.

We conclude the report by describing different possibilities to continue extending the model even further. These extensions augment the DRM to provide the AWF leadership more accurate and powerful predictions of future AWF worker behavior.

NPS-HR-22-008

ACQUISITION RESEARCH PROGRAM
SPONSORED REPORT SERIES

Individualized and Optimal Talent Management of the AWF in Response to COVID-19: Dynamic Programming Approach

December 6, 2022

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

In the current environment of fast evolution of the defense acquisition system, it is vital to have a talented, experienced, and qualified civilian workforce with the ability and flexibility to support the ever-changing needs and increasing demands of the modern warfighter. As part of this effort, the Section 809 Panel proposed changes to DoD's career management framework to improve and develop the workforce, and the *Acquisition Workforce Strategic Plan, FY2019–FY2024* has stressed efforts to restore and reshape the acquisition workforce (AWF) after 20 years of shrinkage. Both studies advocate for a reshaping of the AWF "by deliberate and targeted growth of specific career fields" (Defense Technical Information Center (2017) and Department of Defense (2015), p. 5) and stress that it is vital to ensure that the AWF is qualified (not merely certified) to accomplish the mission. The strategic plan defines goals to sustain the progress made in restoring the workforce, change force structure to retain required skill sets, improve employee quality and professionalism, and recruit and retain a more diverse workforce. These goals are laudable, yet policy cannot be driven solely on aspiration, and execution must be grounded on data. Inability to plan and respond proactively to recruitment, retention, training, and attrition challenges can lead to immediate and long-run deficits in qualified workers, which can take decades to rectify.

To assist leadership, we continue to develop a Dynamic Retention Model (DRM) designed for the AWF. The DRM is a leading-edge technique that takes a complex, multiperiod problem (e.g., lifetime labor market decisions of acquisition workers) and breaks it into simpler, one-period subproblems in a recursive manner. Solving a one-period problem "nests" future decisions the worker will make, allowing estimation and prediction of complex behavior in a surprisingly manageable framework. This setup allows us to simulate how modifications in incentive policy (salaries, retirement, bonuses), would have affected decisions of the workforce. Simpler versions of the DRM were used by the military to assess potential impacts of personnel policy changes on officers and enlisted Soldiers, most recently the Blended Retirement System (BRS). We seek to bring this capability to the AWF, to help the leadership achieve the desired workforce size, quality, and structure.

In particular, we extend the model to handle large and long-lasting shocks to the economy. The real-world example we have in mind is the impact of the COVID-19 pandemic on the U.S. civilian sector. The U.S. unemployment rate is at 4.6% (as of November 2021), which is roughly 1.2 percentage points higher than at the start of the pandemic (i.e., 3.5% in February 2020). However, the current state of the economy, viewed through the lens of where we were just 1 year ago, is remarkably strong. Unemployment spiked to 15% in April 2020. Most economies of the world contracted sharply in 2020. In early to mid-2020, governments and think tanks put forward sometimes wildly differing projections of how quickly (or if) the world economy would recover. However, by early 2021, there was some consensus of a robust recovery. The International Monetary Fund projected global economic growth to be about 5.5% in 2021. While a decline in the U.S. unemployment rate of over 10 percentage points in less than 1 year is a completely unprecedented event, the long-term trajectory of economic recovery is still somewhat unclear. As the recent surge in COVID cases due to the Delta variant has shown, return to pre-pandemic levels of economic activity will be dependent on a number of factors, such as the public health programs, the public's adherence to official guidance (e.g., mandated vaccination, mask-wearing, social distancing, etc.), the macroeconomic environment, as well as unpredictable mutations to the virus.

In the middle of the economic gyrations in the private sector, the AWF (and by extension the government sector) has been somewhat cushioned from the short-run impacts of COVID-19. While the civilian unemployment rate more than quadrupled, no similar increase was observed in the government sector. In this report, we examine the impact of large-scale, unpredictable, long-term changes to the private sector and how that may impact the AWF. We also study the retention implications of the observed, robust economic recovery.

Our prior dynamic programming model of worker attrition behavior is extended to include persistent shocks in the private sector explicitly. The model features a negative autoregressive AR(1) shock to the civilian sector. After calibrating the model parameters using a longitudinal data set of a large subset of the AWF, we simulate civilian-side labor market shocks that correspond to economic recoveries of varying speeds and forecast the retention behavior of the workforce.

In the next section, we describe the impact of COVID-19 on the civilian labor market and the long-run career trajectory of the representative AWF employee. We continue by reviewing our dynamic programming model. We then describe the data briefly and calibrate the model parameters to the data set. With the parameters calibrated, we proceed to run several simulations to demonstrate the impact of COVID-19 under varying recovery scenarios. These simulations project the attrition behavior of the workforce. Finally, we conclude and present future avenues for expanding the capabilities of the model.

THIS PAGE LEFT INTENTIONALLY BLANK

The Impact of COVID-19 on the Civilian Labor Force

The immediate impact of COVID-19 has been without historical precedent. The U.S. unemployment rate spiked to almost 15% from approximately 3.5% in just 2 months. The enormity of this change is demonstrated in Figure 1. Even during the Great Recession, the unemployment rate never reached 11%. After a contraction of around 3.5% of the U.S. economy in 2020, the growth prediction of the Congressional Budget Office (CBO) for 2021 was about 4.6%, which was a considerable upward revision of the initial CBO growth projection of 4% (Congressional Budget Office, 2021). It is worth noting that the actual growth rate will most likely be even higher. The growth rate in the second and third quarters of 2021 were over 6%. This remarkable recovery to date must be tempered by the fact that it remains unclear when the economy can fully return to "business as usual."

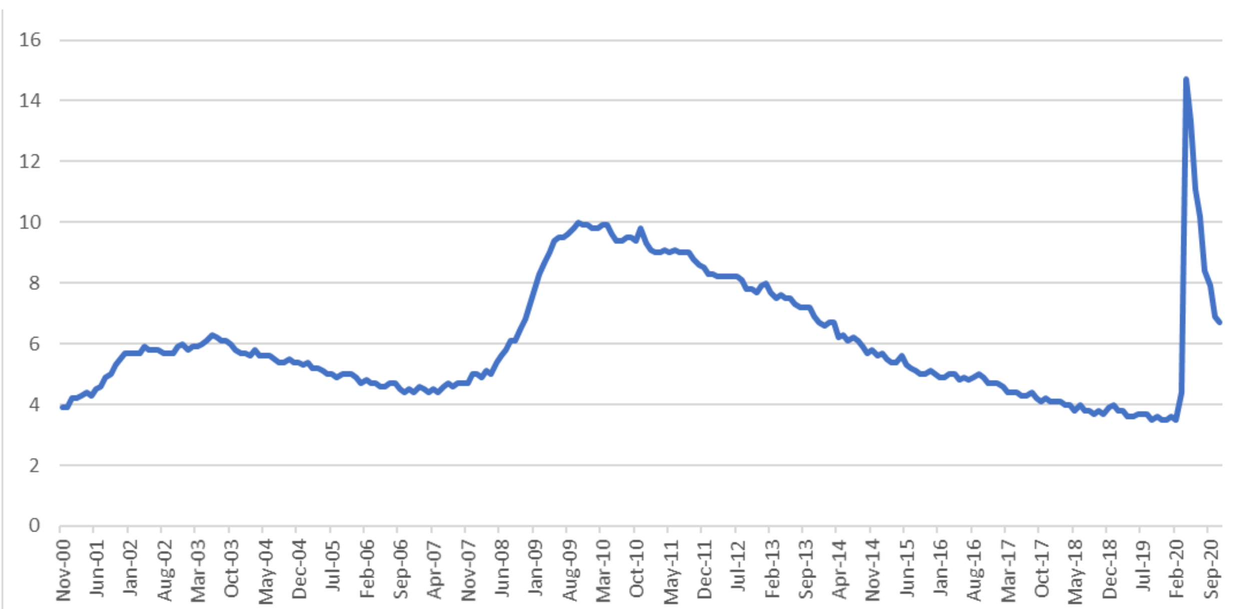

Figure 1. Civilian Unemployment Rate. Adapted from Bureau of Labor Statistics (2021).

While the initial shock of COVID-19 reverberated through almost every sector of the labor market, the AWF has been shielded from the worst of the impact. As seen in Figure 2, the government sector worker unemployment rate was at 4%, which is lower than education or health services workers. The lack of turmoil in the government sector implies that the long-run career trajectories of many AWF workers could remain largely unchanged. Figure 3, reproduced from Ahn and Menichini (2019), shows the attrition rate of AWF workers estimated from historical data prior to the COVID-19 pandemic, covering September 1987 to December 2018. Approximately 30% of workers have exited before 8 years of service. Approximately 75% of the original workforce has left by 25 years of experience. While these attrition rates are relatively low compared to the private sector, AWF leadership expresses a desire to hold on to highly skilled, senior civilian workers. For example, “Highly educated, skilled, and experienced government acquisition professionals are vital now and, in the future, to provide warfighters the products they need” Assistant Secretary of the Navy for Research, Development, and Acquisition James F. Geurts stated, “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” ((Department of the Navy [DoN], 2018, p. 3).).

While job stability has always been seen as a primary benefit of employment in the government sector, the increased uncertainty in the private sector may amplify this aspect of the job to induce longer careers by current AWF workers. This argument is parallel to what has been known in military recruiting: demand for military jobs is countercyclical to the state of the civilian economy. With the backdrop of this large, negative, persistent, and unpredictable shock, we project the long-run labor market decisions of AWF employees by extending our DRM.

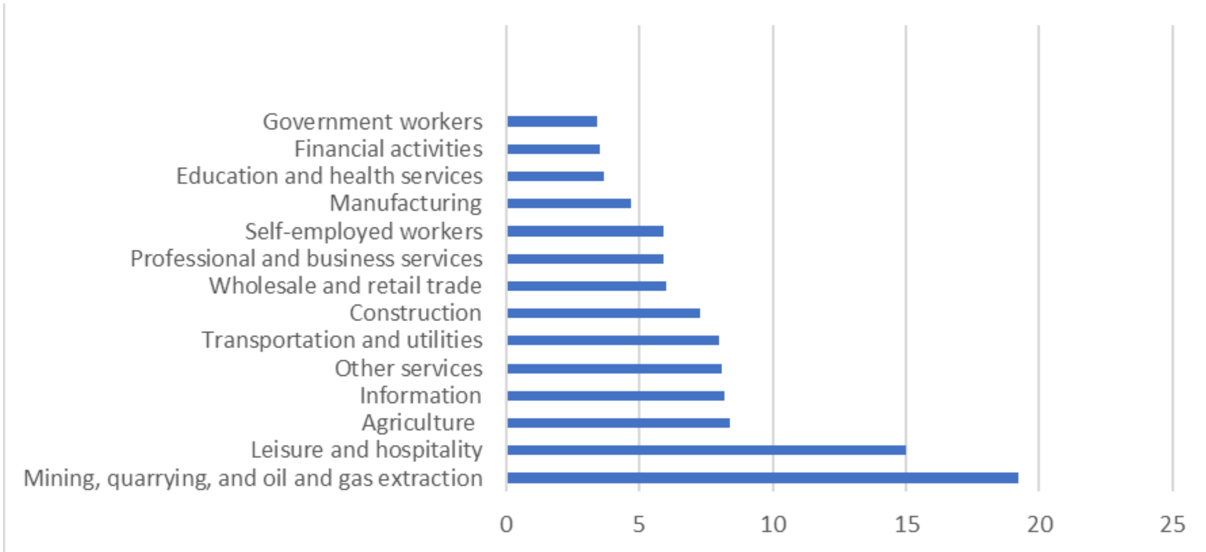

Figure 2. Unemployment Rate by Sector, November 2020. Adapted from Bureau of Labor Statistics (2020).

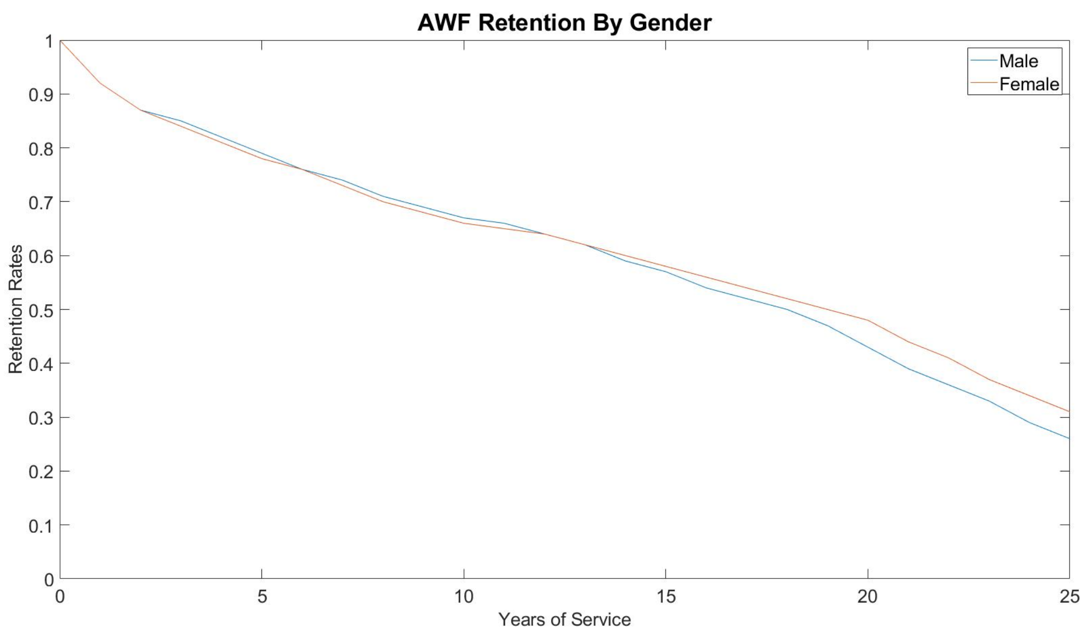

Figure 3. Career Trajectories of DoD AWF Employees. Source: Ahn & Menichini (2019).

THIS PAGE LEFT INTENTIONALLY BLANK

General Description of Dynamic Programming[1]

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 those previously estimated models to simulate worker behavior through time yields people 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, in turn, may then affect his 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 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 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

[1] 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).

make, which allows the researcher to estimate and forecast complex, decades-long behavior in a manageable framework.

The strength of the 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.

In its beginnings, 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 the 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 area, 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 4, for Periods 1 and 2, the person would select (H,H) = $300 to maximize total payoff.

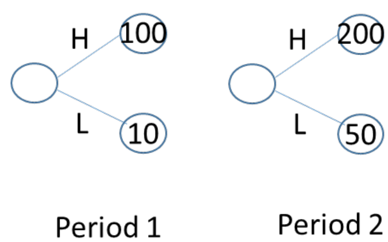

Figure 4. Simple Choice Across Independent Time Periods

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 5, 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 off.

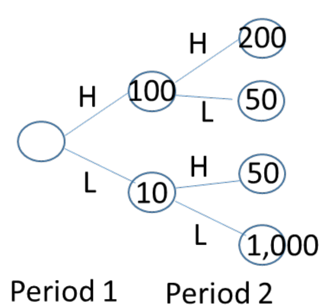

Figure 5. Choice Across Connected Time Periods

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 4 and 5. 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.[2] 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 until the final period. This is also called the 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 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 in (L, L).

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

[2] 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.

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

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

If they opt 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 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 information technology workers in the military. Asch et al. (2017) extended the 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 the DRM over other models of employee retention behavior, such as the traditional annualized cost of leaving (ACOL) model.[3]

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.[4]

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 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

---

[3] 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.

[4] 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 fixed-point algorithm, researchers 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).[5]

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).[6]

[5] 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.
[6] 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.

THIS PAGE LEFT INTENTIONALLY BLANK

Model

In this section, we describe the different parts of the dynamic programming model of employee attrition that will be used to produce policy simulations. Large portions of this section are repeated from Ahn and Menichini (2019) and Ahn and Menichini (2022). Interested readers are referred to those publications for additional details.

AWF workers are assumed to be rational decision-makers making career choices to maximize utility over their lifetimes. The worker evaluates all costs and benefits of each possible choice, both monetary and nonmonetary aspects, at each point in time. We describe the process in detail below. At the beginning of each period, the AWF worker chooses between exiting the AWF sector for the private sector and staying in the public sector for an additional period.[7]

We next describe how the worker evaluates the costs and benefits of labor market decisions made at each point in time. The primary monetary elements include

- compensation (including basic pay, health insurance, locality adjustment, bonuses, etc.) earned in the AWF, and
- comparable compensation earned in the private sector.

AWF workers are assumed to accrue pension benefits in the federal Civil Service Retirement System (CSRS).[8] Employees in civilian firms are assumed to be contributing to 401(k) plans where employers match up to 10% of gross pay.[9]

Nonmonetary elements of a career in the AWF include the worker's taste or preference for a public sector job. Such a preference can be due to preference for higher predictability and stability of public sector employment, even at the cost of a lower salary compared to the private sector. There may also be other aspects of the job, such as patriotism or satisfaction of working for the public benefit, which may not be easily

[7] We assume that exiting the AWF is an irreversible decision.
[8] The data set contains employees from the discontinued CSRS and the current Federal Employee Retirement System (FERS). We model the CSRS because there are more individuals belonging to that system than the FERS.
[9] The modal AWF employee has at least a bachelor's degree and earns close to $100,000 at their highest pay grade. Workers with these characteristics in the civilian sector most often are in jobs that offer matching 401(k) options.

convertible to a dollar figure. To account for these relative preferences, we use taste parameters reflecting monetary-equivalent preferences for careers in the private versus the public sectors.

To summarize, we use the following notation for the dynamic programming model:

- W_t^m is compensation in the AWF in period t
- W_t^c is compensation in the private sector in period t
- ω^m is the AWF taste parameter, capturing the preference for a career in the AWF
- ω^c is the civilian taste parameter, capturing preference for a private sector career
- T is the labor time horizon (number of working periods before retirement)
- $\beta = \frac{1}{1+r}$ is the discount factor, where r represents the subjective discount rate
- $E[.|\varepsilon_{t-1}]$ is the expectation operator, given the shock in the prior time period
- ε_t^m and ε_t^c are the random shocks affecting public and private jobs, respectively, in period t

The optimization problem for AWF workers is described by these equations:

$$V_t^L = W_t^c + \omega^c + \beta E[V_{t+1}^L | \varepsilon_t^c] + \varepsilon_t^c \tag{1}$$

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

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

Super-index S denotes the worker's choice to continue for one more period in the AWF (i.e., S = Stay). Super-index L indicates the worker's choice to leave the AWF and work in the private sector (i.e., L = Leave). Therefore, V_t^S is the present value of remaining in the AWF one more period, while V_t^L is the present value of leaving for the private sector. Equation 3 is the maximization problem that workers must "solve" every period: whether to stay or leave the AWF. The worker chooses to stay as long as $V_t^S > V_t^L$ and leaves as soon as the opposite is true.

We assume that the stochastic shocks ε_t^m and ε_t^c are mutually independent and mean reverting through time. The shocks are defined as follows:

$$\varepsilon_t^c = \mu_c + \rho_c \varepsilon_{t-1}^c + \tau_t^c, \qquad \tau_t^c \sim N(0, \sigma_c^2) \tag{4}$$

$$\varepsilon_t^m = \mu_m + \rho_m \varepsilon_{t-1}^m + \tau_t^m, \qquad \tau_t^m \sim N(0, \sigma_m^2) \tag{5}$$

$$\tau_t^c \text{ independent of } \tau_t^m \quad (6)$$

It is perhaps most intuitive to think about the shock as random fluctuations in the civilian and private sector salaries. In the private sector, fluctuations in the business cycle may drive wage changes from year to year. In the government sector, the political process can result in unpredictable increases in government sector wages. Ashenfelter and Card (1982) found that nominal wages are reasonably replicated as an AR(1) process.

These AR(1) shocks defined in Equations 4–6 are crucial for the model as they allow the impact of one-time shocks to carry over from period to period.[10] As we show in Table 3, we use the parameter ρ to define the rate at which the economy recovers from a shock (e.g., the COVID-19 outbreak). For the optimization problem of AWF workers described in Equations 1–3, random shocks ε_t^m and ε_t^c are state variables observed by the employee at the time the decision to stay or leave is made.

It is worth highlighting a key difference between our dynamic model and other popular models of military retention, such as the ACOL. The main advantage of our model is that it yields time-consistent individual preferences. That is, the original course of action that a worker defines continues to be optimal as times passes, even in the face of unanticipated shocks like COVID-19. This is contrary to what happens with ACOL, where the original plan of action might become suboptimal as uncertainty unfolds (i.e., the model is dynamically or time inconsistent).[11] However, it is also fair to mention that ACOL is usually simpler to implement, which partially explains its relatively higher popularity.

[10] This is different from white noise processes, where shocks do not persist over time (i.e., they return to the mean immediately), or random walk processes, where shocks do not return to the mean.
[11] In other words, the original plan of action does not satisfy Bellman's principle of optimality.

THIS PAGE LEFT INTENTIONALLY BLANK

Data Description and Model Calibration

In this section, we describe the AWF sample as well as the calibration of the parameter values necessary to implement the dynamic programming model described in the previous section. The data that we are using for this report is identical to our previous ARP reports. As such, the descriptions in the Data section are largely repeated from Ahn and Menichini (2019) and Ahn and Menichini (2022). In the next section, we show that those parameter values provide a good approximation of the long-run labor market outcomes for the representative worker in the AWF.

Data: The Acquisition Workforce

The AWF sample we draw from is comprised of approximately 150,000 employees, covering the period from September 1987 to December 2018. Civilians make up about 90% of the workforce, while active duty makes up the remaining 10%. The AWF's mission is the "timely and cost-effective development and delivery of warfighting capabilities to America's combat forces" (Department of Defense, 2015, p. 5). The AWF is responsible for overseeing the equipping and sustaining the military, spending over $1 trillion in Fiscal Year (FY) 2021. About 26% of the AWF belongs to the engineering career field, followed by contracting at 19%. Historically, the AWF was sharply reduced in size and capability during the 1990s. The DoD has been working to rebuild the AWF, starting in 2008, increasing the AWF by approximately 30,000 employees over 7 years.

We received our data from Defense Manpower Data Center (DMDC). The full list of variables in the extract follows in Table 1. For this technical report, we use the variables that are bolded.

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

The raw data was from the 1,000-byte Appropriated Funds (APF) Civilian Personnel Master File. The data files were transferred our secured workstations after encryption and anonymization by DMDC. The data were in a flat ASCII format and were 0.98 GB in size. We converted these data into Stata data file format (.dta) for analysis.

Restricting the AWF sample substantively decreased the size of the population to be analyzed. In addition, we dropped workers born before January 1, 1950 and those born after December 31, 1980. Workers born prior to 1950 would have retired or be nearing retirement, which may have resulted in sample selection bias. Additionally, it would have been difficult to model decisions for this population based on unobservable factors such as health or family circumstances. Further, these workers' primary labor market experiences, in the 1970s and 1980s, may have been less relevant for predicting the behavior of current or future workers. Employees born in 1981 or later may have been too young to provide relevant information on long-term career decisions.

After restricting the sample, we obtained more than 2 million worker-month variables. More than 13,000 AWF employees were tracked in the sample. Table 2 shows some summary statistics for our sample.

Table 2. Summary Statistics

Variables		Mean	Std. Dev.	Min	Max
Female		0.632			
Minority		0.278			
Disability		0.202			
Prior Military Service		0.619			
Has Bachelor's Degree		0.547			
Has Postgraduate Degree		0.332			
Gained Additional Education in AWF		0.441			
Career Length in AWF (in years)		12.0	(8.6)	0.1	25.8
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 Rated Not Fully Satisfactory		0.575			
Highest Salary		95,143.67	(30,410.74)	27,397	189,600
Observations		13,590			

To rigorously assess the impact of the civilian sector on the attractiveness of the AWF position, every worker in the data set must be linked to a private sector wage that the worker can reasonably expect to earn. We estimate a hedonic regression using the outgoing rotation group (ORG) of the current population survey (CPS). As this data set contains a representative sample of workers in the United States, including, most importantly, those who are in the government sector, it is possible to make an apples-to-apples comparison with workers in the private sector. See Ahn and Menichini (2020) for a detailed description.

We run a hedonic regression using the individual sociodemographic characteristics, professional and educational experience, and locality indicators from the

ORG of the CPS, which broadly match the AWF variables summarized in Table 1, to obtain predicted civilian and government sector wages. The difference in the wages across private and public sectors, conditioned on individual characteristics, defines the government sector "wage penalty."

Calibration Results

The simulation of the model in Equations 1–3 requires choosing values for the parameters. We show those values in Table 2. It is worth noting that all parameters but compensation do not change over time.

The results from the hedonic regressions suggest that the ratio of income in the private sector (i.e., W_t^c) to income in the AWF (i.e., W_t^m) is roughly 1.1761. After normalizing W_t^m=1, we let W_t^c=1.1761. The next step is to add the retirement income, which increases as the individual stays more in the AWF, making compensation vary over time.

We assume a discount rate of around 5.3% per year, which implies a personal discount factor of 0.95. Finally, we assume T=25, as that is the longest career observed in the data. [12]

We then calibrate the taste parameter ω^m to make the model-predicted survival curve match as much as possible the empirical survival curve. More precisely, we perform a grid search to find the value of ω^m that minimizes the squared distance between the vector of points of the model-predicted survival curve and the vector of points corresponding to the empirical survival curve. After normalizing $\omega^c = 1$, the calibration procedure yields $\omega^{\mathrm{m}} = 1.2782$.[13] The greater value of ω^m suggests that AWF employees prefer, on average, to work in the AWF instead of the private sector.

[12] This interest rate is similar to the average 30-Year T-Bond Constant Maturity Rate reported by the Federal Reserve Bank of St. Louis for the period covered by the data set.

[13] Ahn and Menichini (2020) estimated a similar dynamic model where economic shocks to the civilian and public sectors are i.i.d. (independent and identically distributed) with mean zero. They found the difference between military and private sector taste parameters (i.e., $\omega^m - \omega^c$) to be around 0.2, which, reassuringly, is not far from the result of the calibration exercise.

The final parameters in Table 3 describe the random process followed by the error terms ε_t^m and ε_t^c. We follow Ashenfelter and Card (1982) to parameterize those parameters. That is, parameters μ_m and μ_c are assumed to be zero, the standard deviation of the error terms, σ_m and σ_c, are supposed to be 0.005, and the mean-reversion coefficients, ρ_m and ρ_c, are defined to be 0.9. These parameter values are commonly found in the macroeconomics literature and depict the historical behavior observed for the error terms. It is worth noting that the high value of the mean-reverting coefficient means that wages are highly persistent over time, which, in turn, implies that it takes a long time for shocks to dissipate.

Table 3. Parameter Values

Parameter	*Value*
W_t^m	1
W_t^c	1.1761
T	25
β	0.95
ω^m	1.2782
ω^c	1
μ_m	0
μ_c	0
ρ_m	0.90
ρ_c	0.90
σ_m	0.005
σ_c	0.005

Model Solution and Policy Simulations for a One-Time Shock

In this section, we describe our *initial* policy simulations to forecast the evolution of the behavior of the representative AWF worker under a number of scenarios with differing speed rates of economic recovery from a large, abrupt, and unanticipated negative impact (i.e., COVID-19) to the private sector. This is a major systematic event that adversely affects all sectors of the economy, except the government sector.[14] The latter is consistent with the assumption of independent errors in Equation 6. This section largely repeats our findings in Ahn and Menichini (2022).

In a nutshell, we "shock" the model with a large negative civilian error draw at a specified point in time. Then, we allow the system to recover and converge back to the steady state. In an "optimistic" scenario, we assume that the economy bounces back quickly and robustly. In a "pessimistic" scenario, the recovery is slow and weak. We add an "expected" scenario where the recovery falls between the two extremes. We "control" the speed of recovery of the economy by setting the autoregressive term, ρ, which controls the velocity at which shocks gradually disappear over time. We describe the simulation procedure in more detail next.

It is important to note that due to the structure of the model, this initial policy simulation assumes that AWF workers are only surprised in the initial period. The initial "shock" is unanticipated in the sense that the mean of the shock term is zero. However, once this shock occurs, the gradual "recovery" of the economy, governed by ρ and σ, are known to the workers. That is, AWF workers are able to forecast the future (in expectation) after the initial shock and make optimal decisions from that point forward. This is a strong assumption about the information set of AWF workers. In subsequent sections, we explore the impact of incorrect predictions about the future state of the economy in more than one period.

[14] While our negative shock is the COVID-19 pandemic, any future unanticipated national shock to the economy and/or public health that is concentrated in the private sector can be expected to operate in a similar manner.

We start solving the model described in Equations 1–3 by backward induction. The solution describes the retention behavior of a representative AWF employee in all possible states of the economy. We then stochastically simulate the model forward (i.e., over time) 10,000 times, which produces the stay/leave decisions of 10,000 employees in all possible different situations over the labor period. These simulations summarize the retention behavior of the representative employee, which we show in Figure 6. The figure shows the survival curve of the representative individual and displays the cumulative probability of the worker staying in the force after a certain period of time. For example, the figure suggests that the likelihood that the employee is still part of the AWF after 10 years is about 40%.

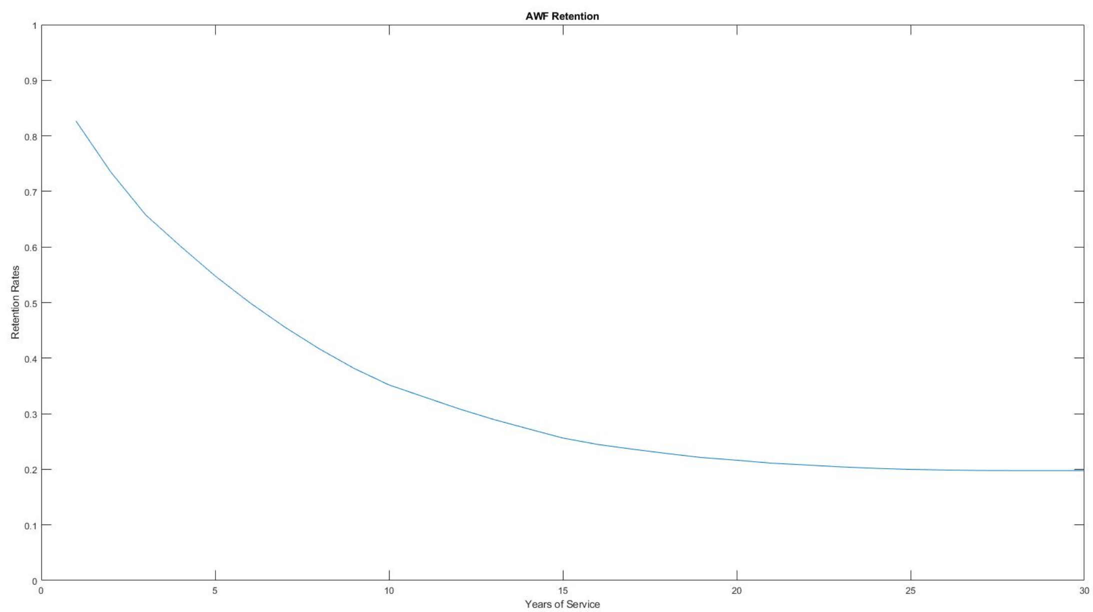

Figure 6. Retention Behavior

We then proceed to shock the model with a large negative error on the civilian side at Year 10. The shock is equivalent to 3 standard deviations below the mean and is intended to capture the large effect of the sudden appearance of COVID-19. The fact that the error terms (ε_t^m and ε_t^c) are mean reverting over time implies that the impact of the negative civilian shock disappears gradually over time, as opposed to immediately.

As we mentioned before, the speed of return to the pre-shock state will depend on the mean-reversion coefficient, ρ. In Figure 7 we show, given an initial negative shock, how the shocks are expected to evolve over time for three different values of the coefficient of mean-reversion. The blue bars represent the case of a relatively fast return to the pre-COVID economy, which corresponds to the optimistic recovery scenario with ρ = 0.3. On the other extreme are the yellow bars, reflecting a slow recovery to normality with ρ = 0.7, which depict the pessimistic scenario. In between are the red bars showing the expected recovery with ρ = 0.5. Even in the optimistic recovery scenario, it is clear that the effect of the large negative shock remains in effect for some years.[15]

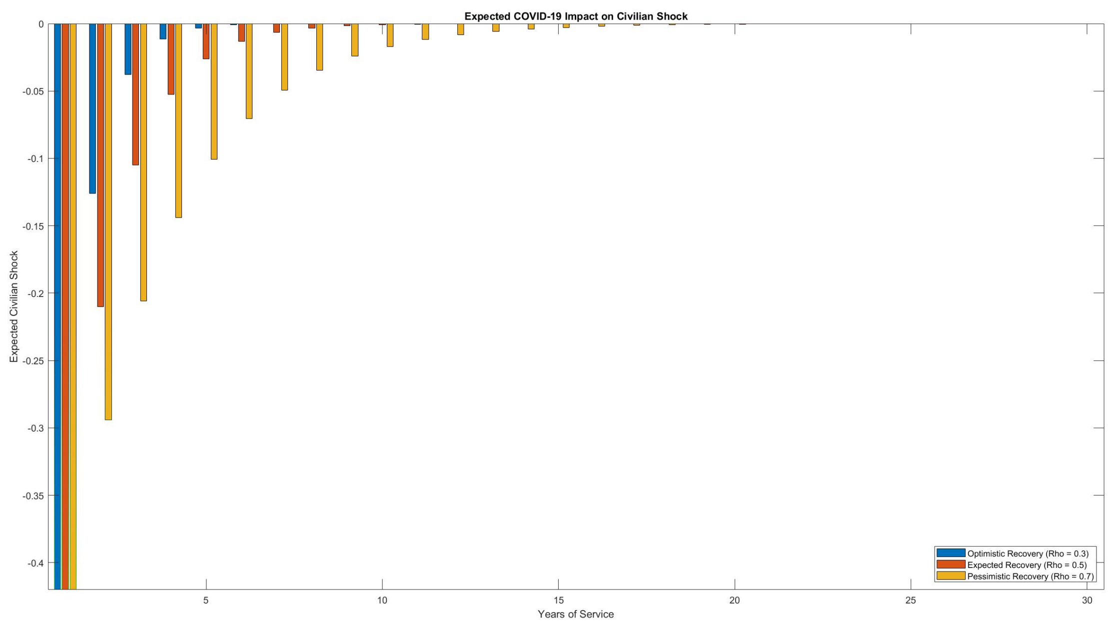

Figure 7. Expected Impact of COVID-19 on Civilian Shock

The effect on retention behavior of the representative AWF worker can be observed in Figure 8. The figure shows that, during the initial 10 years, the retention behavior is equivalent to that in Figure 6. This represents the baseline retention behavior that uses ρ = 0.5. At Year 10, the COVID-19 shock happens, and the retention

[15] The magnitude and persistence of shocks are speculative, although they are informed by very recent (and on-going) research. Many scholars are attempting to forecast the long-run impact of COVID-19 on the economy. See Petrosky-Nadeau and Valetta (2020), for example.

behavior changes considerably. As we mentioned before, we study the attrition behavior in three different contexts. The optimistic scenario, represented by the yellow line, assumes $\rho = 0.3$ and implies a fast economic recovery. The pessimistic scenario, depicted by the purple line, shows a slow economic recovery assuming $\rho = 0.7$. The orange line in between reflects the expected recovery scenario where ρ stays at 0.5. In all cases there is a kink and sudden flattening of the curve, suggesting that the individuals will stay longer in the AWF, avoiding the sharp effect of the virus on the civilian labor market. Depending on the speed of recovery, it might take more than a decade for the employee to return to the pre-shock retention behavior. For instance, in the pessimistic scenario (purple line), the AWF worker needs roughly 15 years to return to the original behavior. Accordingly, the red line, corresponding to the intermediate scenario, suggests it takes around 10 years for the individual to return to the pre-COVID retention behavior. This long-lasting effects on retention behavior have important implications for the hiring policies of the public sector.

It is worth noting that the 10- to 15-year lag in return to “original” behavior specified above does *not* mean that all workers will choose to leave the AWF by a decade or more due to the impact of COVID-19. Instead, all workers will process the negative shock in the civilian economy as making the AWF job more attractive. Until the shock fully dissipates, the DoD position will be more attractive than had there not been the global pandemic. However, given the substantial wage premium in the civilian sector, the pandemic shock does not need to completely disappear before workers who were planning to move to the civilian sector resume their plans.

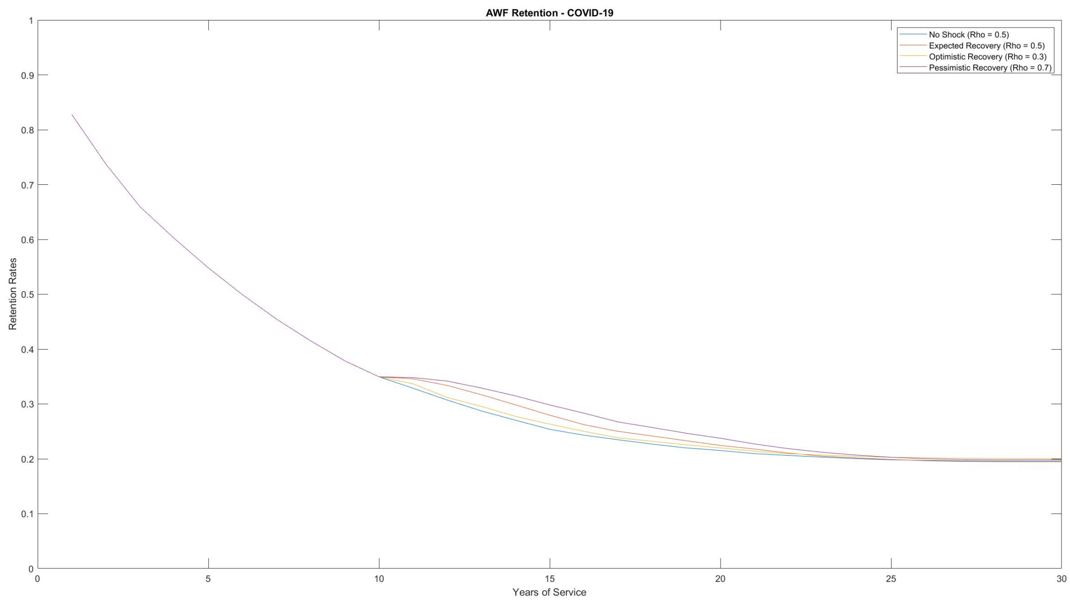

Figure 8. Retention Impact of COVID-19

In addition, we perform robustness checks for parameters ρ and σ to ascertain their impact on the main results of the previous section. In technical terms, these parameters affect the variance of the unconditional distribution of the errors, $\varepsilon_t{}^m$ and $\varepsilon_t{}^c$. An increase in σ, the conditional standard deviation of the mean-reverting process, naturally increases the unconditional variance of the errors. An increment in ρ makes large shocks persist more, on average, which also augments the unconditional variance of the errors. Accordingly, a larger unconditional variance should diminish retention because it increases the chance of getting large (positive) shocks in the private sector, which may induce the individual to leave the AWF more frequently.

Next, we analyze how different values of ρ and σ change employee attrition in the context of the COVID-19 shock, as compared to the results presented in the previous section. We start evaluating changes in parameter ρ. We used a value of ρ = 0.5 to represent the baseline situation. We then change that value to 0.4, that is, higher speed of mean-reversion of the shocks. The pessimistic recovery scenario is assumed to happen with a ρ = 0.6, while we assume ρ = 0.2 for the optimistic recovery scenario. Finally, the expected scenario, like the baseline scenario, assumes ρ = 0.4. Figure 9 shows the results of such changes. It can be easily observed that the main effect is an overall increase in retention, which is consistent with the above explanation about the unconditional variance. Regarding the effects of the COVID-19 shock on retention behavior, the conclusions remain unchanged. That is, employees stay more in the AWF for several years, as compared to a no-shock situation.

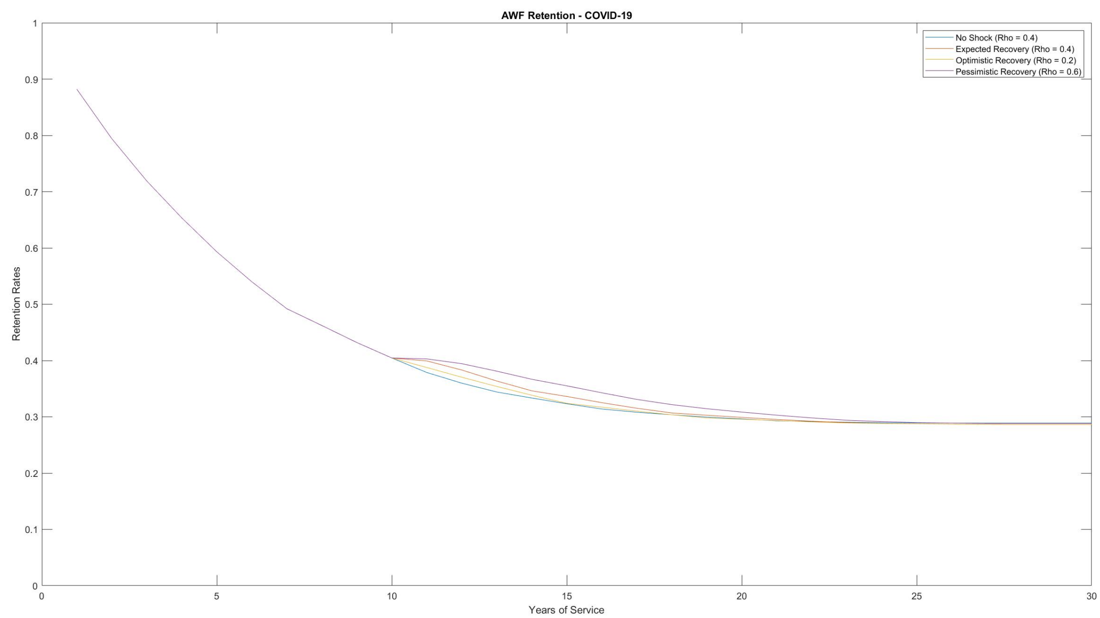

Figure 9. Retention Impact of COVID-19 With Faster Economic Recovery

Then we evaluate increasing the value of the mean-reverting coefficient in the baseline situation from $\rho = 0.5$ to $\rho = 0.6$. The pessimistic, expected, and optimistic scenarios are simulated at $\rho = 0.8$, $\rho = 0.6$, and $\rho = 0.4$, respectively. Overall, this situation (i.e., $\rho = 0.6$) represents a lower speed of mean-reversion of the shocks compared to the results described in the previous section. Accordingly, as Figure 10 shows, overall retention falls. Nevertheless, the effects of the COVID-19 shock on retention behavior are long-lasting as well. That is, employee retention increases considerably and for a long period after the shock.

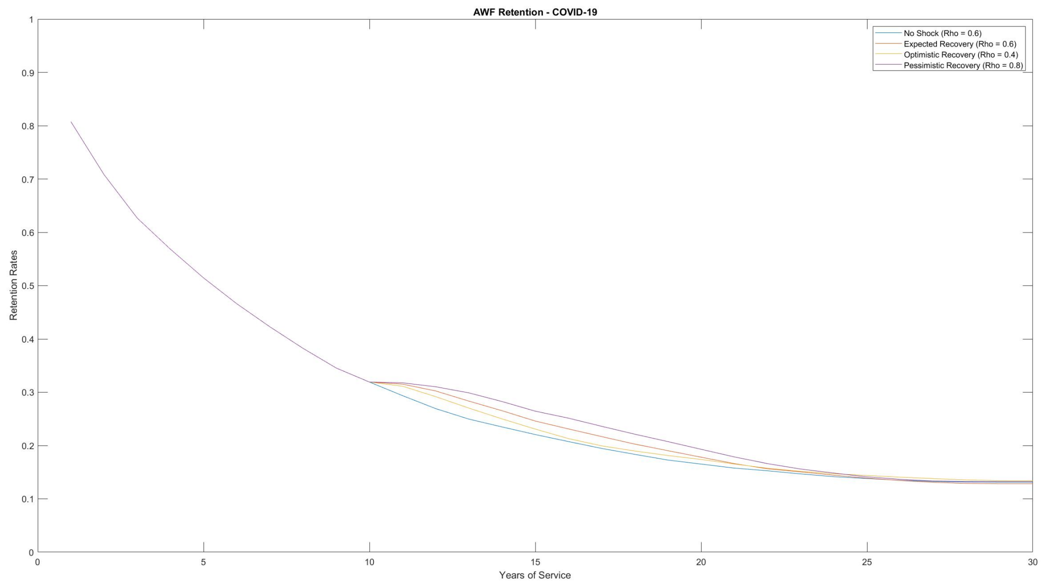

Figure 10. Retention Impact of COVID-19 With Slower Economic Recovery

Finally, we analyze the impact of changing the conditional variance, *σ*, of the errors. As we mentioned above, the overall effect of increasing this parameter is decreased retention, as larger shocks in the private and public sectors will appear during the working period, increasing the chances that the individual leaves the AWF. Figures 11 and 12 show that effect for values of *σ* of 1.5 and 2, respectively. Regarding the negative shock generated by COVID-19, the results are very similar to those in the previous section. The virus shock has an impact on retention that may last for several years. That is, retention increases for several years before returning to the pre-virus levels.

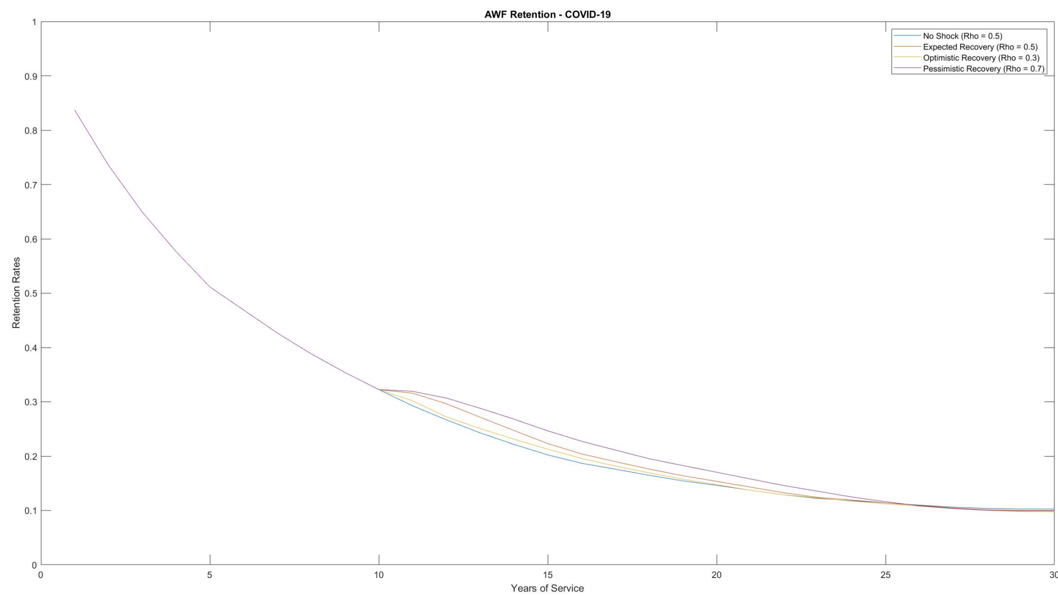

Figure 11. Retention Impact of COVID-19 With Greater Variability of Shocks (σ = 1.5)

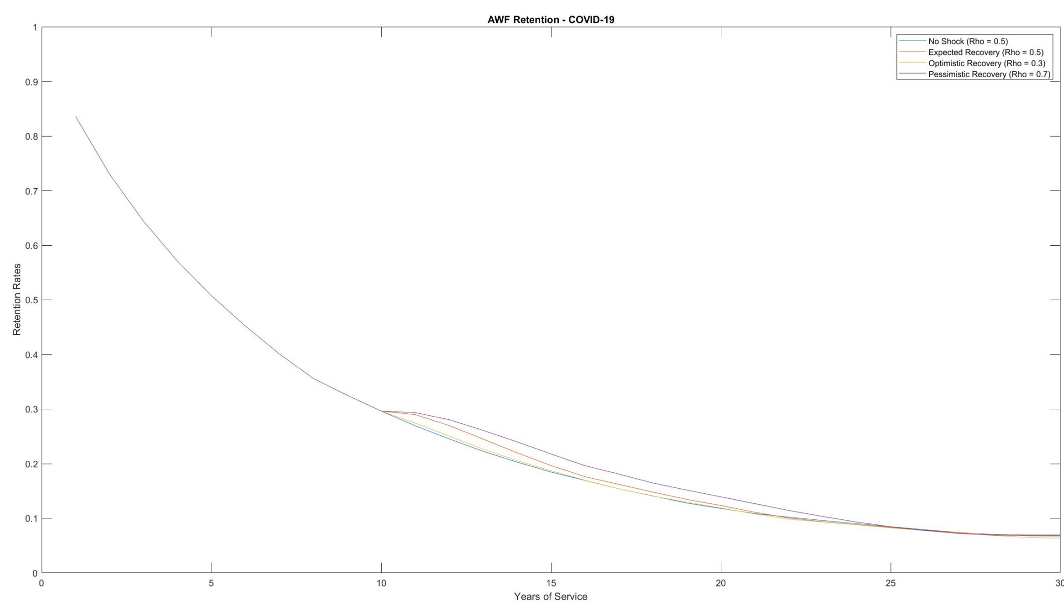

Figure 12. Retention Impact of COVID-19 With Greater Variability of Shocks (σ = 2)

To summarize, while the values assumed for parameters ρ and σ do affect employee retention in general, they do not change the main effects of the COVID-19 shock on retention.

It is worth emphasizing once again that the model examined above makes one glaring simplification. Once the shock is introduced in Year 10, the expected recovery of the economy is assumed known by the agents. That is, although the initial shock is unexpected, ρ and σ are assumed known by AWF workers from that point forward.

THIS PAGE LEFT INTENTIONALLY BLANK

Policy Simulations for a Multiperiod Shock

In this section, we change the nature of the shock to emphasize the unpredictable nature of the recovery process itself. In particular, after the initial unexpected shock, we introduce additional shocks in subsequent periods that do not align with the AWF workers' expectation of the evolution of the shock process.

This policy simulation is especially relevant for COVID-19. As explained above, the unprecedented negative impact on the economy as well as the largely robust recovery were largely unanticipated by governments, industries, and academics. It is then perhaps unrealistic to assume, as we did in the previous section, that AWF employees would be able to predict the recovery path with any degree of accuracy.

The goal of this section is to replicate (i.e., to "force" the AWF agent to go through) the economic path observed since the outbreak of the pandemic:

1. 2020: COVID-19 shock implying a large negative economic impact (i.e., -3σ)
2. 2021: significant economic recovery, but still a slightly negative economic shock (i.e., $-3\alpha\sigma$)
3. 2022: expected (as of today) continued economic recovery, with slightly positive economic shock (i.e., $3\alpha\sigma$)

The shocks to the economy in this section follow the following pattern. In Year 10, as done in the previous section, the economy is impacted by a −3 standard deviation shock (i.e., -3σ), simulating the initial COVID-19 impact on the civilian economy.

In the subsequent time period, in Year 11, a shock to the economy is *artificially* introduced that corresponds to a fraction, α, of the original shock (i.e., $-3\alpha\sigma$), such that

$$\alpha = \frac{UR_{aug21} - UR_{feb20}}{UR_{apr20} - UR_{feb20}} = 0.168$$

where *UR* is the U.S. domestic unemployment rate at the month and year indicated by the subscript. The unemployment rates for August 2021, February 2020, and April 2020 were, respectively, 5.4%, 3.5%, and 14.8%. The August 2021 unemployment rate is our current unemployment rate. The February and April 2020 unemployment rates are the minimum and maximum values of unemployment rate observed over the last several years. The denominator represents the (total) required change in unemployment rate for the labor market to return to its pre-pandemic levels, starting from the worst of the pandemic. The numerator is the remaining amount of recovery since August 2021 to return to the pre-pandemic levels. The ratio between the two is the fraction of recovery left to get back to the pre-pandemic labor market.

The year after, in Year 12, another artificial shock that is −1 times the shock from Year 11 is introduced (i.e., $3\alpha\sigma$). Thus a modest economic boom is assumed to occur 2 to 3 years after the pandemic. See Figure 13.

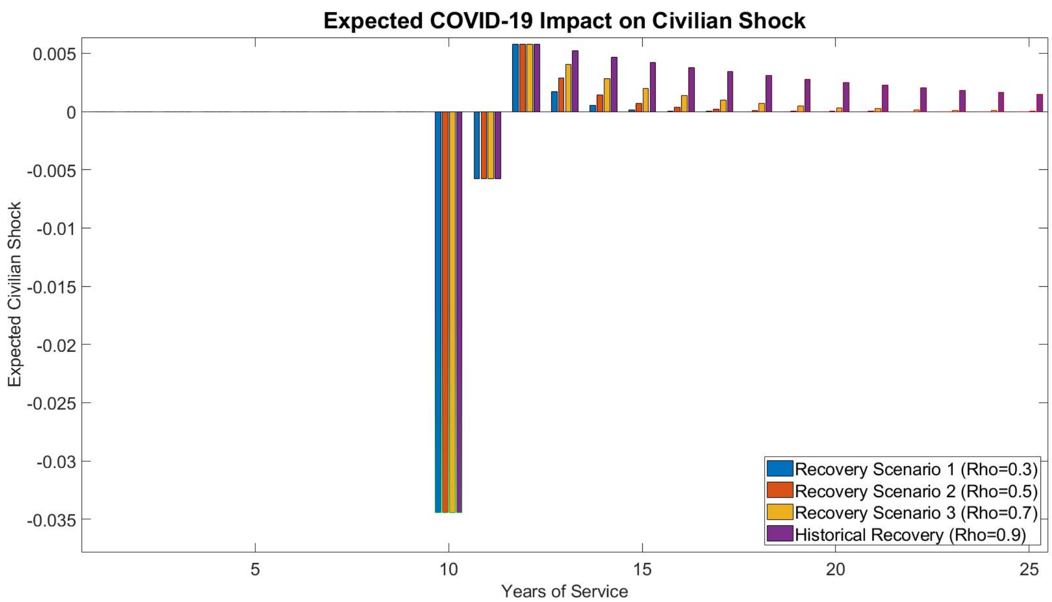

Figure 13. Expected Impact of COVID-19 on Multiperiod Shock

After the unanticipated shock in Year 12, the labor market is allowed to gradually settle down, with the mean of the shock reverting ever-closer to zero according to ρ.

It is worth re-emphasizing the reason for these artificial shocks. Without manual insertion of these shocks, all workers "know" the parameters of the shock term (ρ and σ). Then, beyond the initial surprise at Year 10, all workers are able to accurately predict the evolution of the labor market accurately and take optimal actions through time. As this may imply too much sophistication and knowledge on the part of the average AWF employee, the shocks in Years 11 and 12 insert the extreme amount of uncertainty workers faced (during COVID-19) or will face (in future economic gyrations).

Figure 14 shows the retention rates of workers. Worker retention rates are identical to Figure 8. At Year 10, when the initial shock occurs, the worker response to this shock is also identical to what is observed in Figure 8: workers, anticipating long-term negative impacts on the civilian labor markets are reluctant to exit the AWF sector.

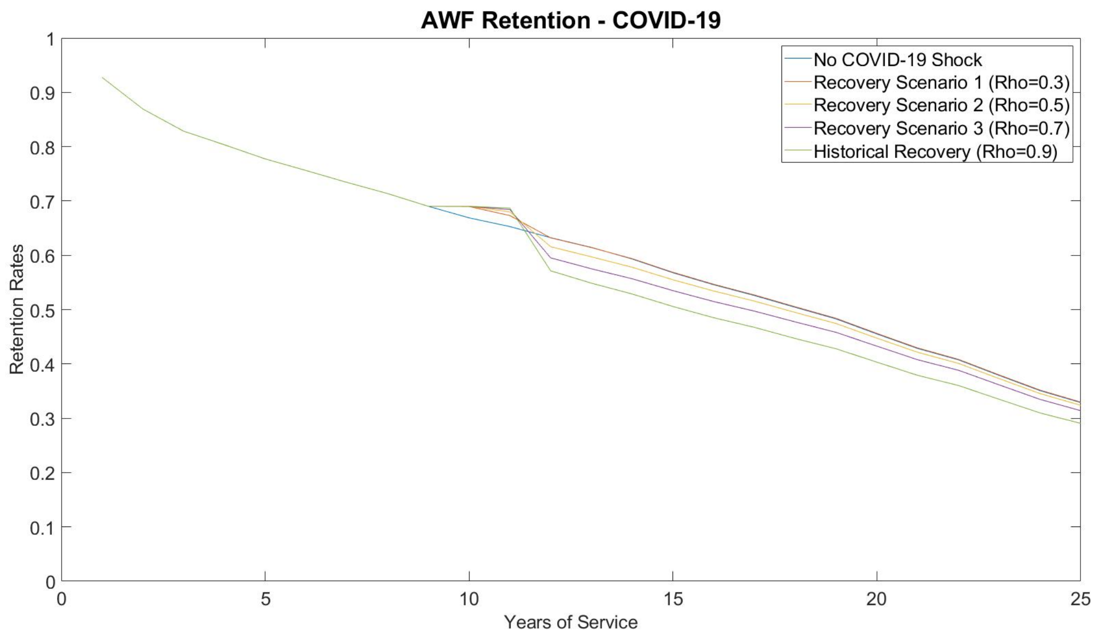

Figure 14. Retention Impact of COVID-19 with Multiperiod Shock

However, in Years 11 and 12, workers are surprised by the second and third shocks that are much more positive than originally anticipated. In response, workers who held back in Year 10 on leaving the AWF for the private sector rush to exit in Years 12 and onward. The sharp downturn in retention rates and the curve dipping even lower than the blue line (no COVID-19 shock scenario) show that ultimately, the AWF may lose workers faster to the civilian sector than anticipated, due to the speed of recovery.

Figure 15 shows the model-predicted yearly probabilities of leaving the AWF. Again, the green line shows the attrition behavior in the historical recovery scenario. The red, yellow, and purple lines each reflect a slower mean-reversion rate. In all four scenarios, the likelihood of leaving the AWF goes roughly to zero in the year of the shock and springs upward quite suddenly as the unanticipated positive shocks to the economy hit. As the labor market slowly reverts back to normalcy, annual attrition rates also converge back to the blue line (no COVID-19 shock).

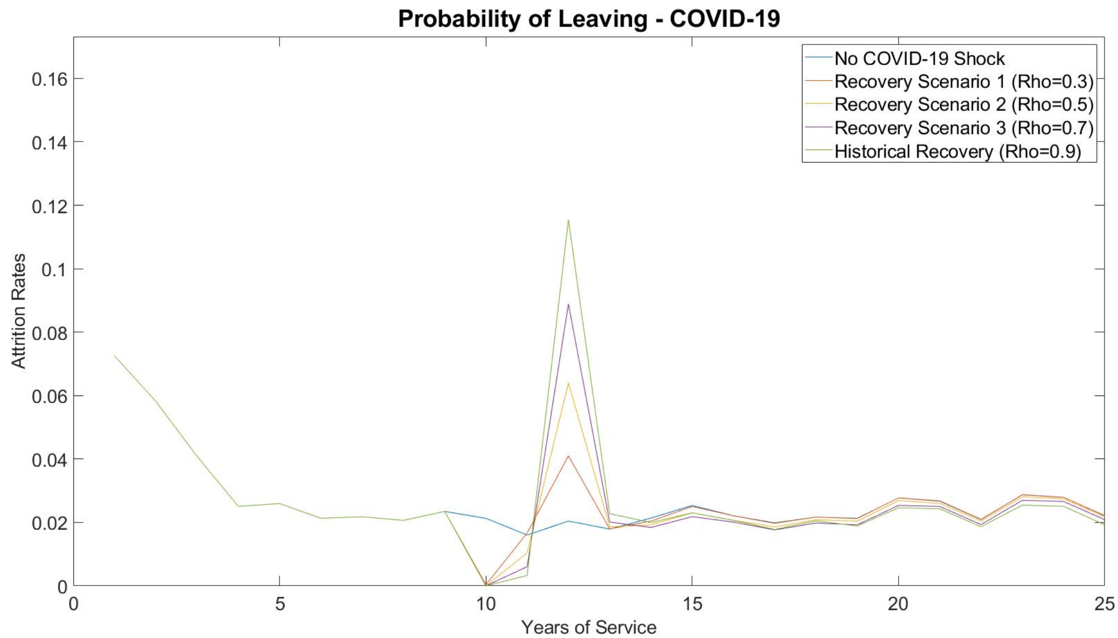

Figure 15. Annual Attrition Rates of Workers in Response to Multiperiod Shocks

Conclusion

In this report, we built upon the foundational tasks completed in the Year 3 report 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 when unanticipated economic gyrations such as the COVID-19 pandemic hits the civilian sector labor market.

As of early 2021, the overall unemployment rate in the United States stands at 6.7%, an 8 percentage point decrease from the worst unemployment rate in almost 90 years arising from the COVID-19 global pandemic, in just 8 months. While the recovery was as dramatic as the decline, the future remains very much in doubt. For example, in December 2020, payrolls shrank by 140,000.

In this environment, this report analyzed the potential impact of the economic recovery on the labor market trajectory of the AWF. The contrast in stability of jobs in the government versus the private sector should increase the attractiveness of DoD jobs, especially if the recovery proves to be slow. We built and estimated a dynamic programming model with a negative persistent shock to the civilian sector and simulated different recovery paths.

Our results show that government positions become more attractive the larger the magnitude of the negative shock to the civilian economy, and the slower the economic recovery, such that workers may value government positions more highly compared to the pre-pandemic period for a decade or more. However, the unpredictability of recovery in scenarios like a global pandemic means that additional challenges await the AWF leadership. Specifically, if the economic downturn is not a one-time unanticipated shock, but a series of negative and positive gyrations, worker behavior will fluctuate as the degree of uncertainty increases.

While a generally depressed private sector economy can reduce attrition of the average worker from the AWF, leadership should understand that eventually, recovery of the civilian sector will push down the relative desirability of government jobs. In addition, as the economy recovers, there may be fundamental structural changes to the labor market that remain, changing the valuation of both government and private sector jobs in unpredictable ways. Forward-looking leaders should regard these simulation results not as predictions of the future, but as guides to help set personnel policies that are flexible enough to adjust to and even take advantage of gyrations in the civilian economy.

The model we have developed these past 4 years will be useful for the AWF leadership as one of a set of predictive tools to manage manpower proactively by extending the time horizon over which the AWF size and shape can be forecast and controlled.

To conclude, we preview potential next steps to push the capability of the model further:

- **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.
- **Perform further analyses for additional 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 we have introduced this year, 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 the AWF (change in General Schedule (GS) scale); change in Federal Employee Retirement System pension annuity computation formula (akin to the BRS in active duty); economic expansions and recessions (modeled as random, unforeseen macroeconomic shocks); and other scenarios as requested by AWF leadership.

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* (Report No. NPS-HR-20-001). Acquisition Research Program. https://dair.nps.edu/handle/123456789/2775

Ahn T., & Menichini, A. (2020). *Retention analysis modeling for the acquisition workforce II* (Report No. NPS-HR-21-031). Acquisition Research Program. https://dair.nps.edu/handle/123456789/4317

Ahn T., & Menichini, A. (2022). Optimal talent-management of the acquisition workforce in response to COVID-19: Dynamic programming approach. *Defense Acquisition Research Journal*, 29(1), 50–77.

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* (Report No. 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* (Report No. 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.

Ashenfelter, O., & Card, D. (1982). Time series representations of economic variables and alternative models of the labour market. *The Review of Economic Studies*, *49*(5), 761–781.

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.

Bureau of Labor Statistics. (2021) Civilian unemployment rate. https://www.bls.gov/charts/employment-situation/civilian-unemployment-rate.htm

Congressional Budget Office. (2021). *An overview of the economic outlook: 2021 to 2031*. https://www.cbo.gov/publication/56965

Cox, D. R. (1972). Regression models and life tables (with discussion). *Journal of the Royal Statistical Society*, *34*(2), 187–220. https://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.

Defense Manpower Data Center. (2019). *APF civilian master file*.

Defense Technical Information Center. (2017) *Advisory panel on streamlining and codifying acquisition regulations: Section 809 panel - Interim report*

Department of Defense. (2015). *Acquisition workforce strategic plan, FY 2016–FY 2021*. https://asc.army.mil/web/wp-content/uploads/2019/03/DoD_Acq_Workforce_HSCP_FY16-21.pdf

Department of the Navy. (2018). *DoN acquisition workforce FY 19-24 strategic plan*.

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* (Report No. 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* (Report No. MR-1565-OSD). RAND. 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

International Monetary Fund. (2021). *World economic outlook update, January 2021*. https://www.imf.org/en/Publications/WEO/Issues/2021/01/26/2021-world-economic-outlook-update

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 and 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. www.its.caltech.edu/~mshum/gradio/papers/pesendorferschmidt.pdf

Petrosky-Nadeau, N., & Valetta, R. (2020). Unemployment paths in a pandemic economy [Working paper]. Federal Reserve Bank of San Francisco.

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

www.acquisitionresearch.net