RAND CORPORATION

Research Report

MICHAEL G. MATTOCK, BETH J. ASCH

The Dynamic Retention Model

Theory, Estimates, Innovations, and Extensions

For more information on this publication, visit **www.rand.org/t/RRA2581-1**.

Published by the RAND Corporation, Santa Monica, Calif.

Cover: U.S. Army Reserve photo by Capt. Jerry Duong

About This Report

Military compensation is a pillar of the all-volunteer force. It is a fundamental policy tool for attracting and retaining personnel, and its structure—and the incentives implied by its structure—can affect U.S. service members' willingness to join, exert effort, demonstrate their leadership potential, remain in the military, and eventually exit the military at an appropriate time. Military compensation is a composite of existing pay and allowances, special and incentive pays, health benefits, disability benefits, retirement benefits, and other benefits.

Assessing the efficiency and flexibility of the existing and alternative compensation systems, especially alternatives that have yet to be tried, requires a model that allows analysis of how different military compensation changes affect the retention, cost, and productivity of individuals. The model should recognize the career decision processes of individual service members, the heterogeneity of their preferences, the uncertainty of the environment in which they make career decisions, the time path of these decisions, and the organizational structure and policy context in which they make these decisions. Furthermore, in the context of the military, given the greater operational role of the reserves and the importance of total force compensation and personnel policy, the model must consider both active and reserve career decisions. The purpose of this report is to synthesize previously published work on the RAND Corporation's dynamic retention model, a model that fulfills these criteria, into a single document. This past research includes both RAND analyses of the retention decisions of military personnel and of public sector employees, such as public school teachers, state employees, and federal civilians. We present theory and estimates, as well as innovations and extensions of the model to military and civilian populations. This report will mainly be of interest to policy analysis researchers and analysts and those who would like to understand the broad capability and applicability of the dynamic retention model. While some parts of this report are accessible to a general audience, others are quite technical.

The research reported here was completed in June 2023 and underwent security review with the sponsor and the Defense Office of Prepublication and Security Review before public release.

RAND National Security Reseach Division

This research was conducted within the Personnel, Readiness, and Health Program of the RAND National Security Research Division (NSRD), which operates the RAND National Defense Research Institute (NDRI), a federally funded research and development center (FFRDC) sponsored by the Office of the Secretary of Defense, the Joint Staff, the Unified Combatant Commands, the Navy, the Marine Corps, the defense agencies, and the defense

intelligence enterprise. This research was made possible by NDRI exploratory research funding that was provided through the FFRDC contract and approved by NDRI's primary sponsor.

For more information on the RAND Personnel, Readiness, and Health Program, see www.rand.org/nsrd/prh or contact the director (contact information is provided on the webpage).

Acknowledgments

We wish to thank our RAND colleagues Molly McIntosh and Daniel Ginsberg, the director and associate director, respectively, of the Personnel, Readiness, and Health Program within NDRI. And we would like to say to Jim Hosek, our colleague, mentor, and endless source of inspiration, our "7 o'clock Muse" who first turned (at least one of) us on to the dynamic retention model and who was critical in developing the theory to make it practical, thank you. It wouldn't have been nearly this much fun without you. Finally, we would like to thank David Knapp at the Center for Economic and Social Research at the University of Southern California for his excellent comments and review of an earlier draft.

Summary

The dynamic retention model (DRM) is a model that has been used to assess the retention effects of changes to compensation in both the military and civilian contexts. The DRM is part of a larger class of models known as *structural models of individual behavior* in which—in our context—people make retention decisions under uncertainty during their careers and have unique or heterogeneous tastes. That is, the DRM recognizes that individuals are forward-looking when they make their decisions (meaning they account for possible future events, such as future pay raises and how decisions might affect future events) and that people differ (meaning that people who face identical opportunities might make different choices). In the DRM, individuals make retention decisions in each year over their career and choose the alternative in each period—to stay or leave—that maximizes their utility, meaning they make the optimal choice each year given the information and opportunities available to them.

The advantage of this approach to modeling retention is that the approach has a solid foundation in the literature and in theories of how personnel make retention decisions over their careers. In particular, the DRM is formulated in terms of the parameters that underlie the retention decisionmaking process rather than on a specific compensation system and retention outcomes. When such a model is estimated with empirical data on individual retention decisions, it can be credibly used for counterfactual analyses, permitting assessments of alternative yet-to-be-implemented compensation systems. Furthermore, it includes a capacity to perform simulations of the retention and cost effects of relevant policy changes, including policies that have yet to be implemented.

The flexibility and logical consistency of this behavioral model has enabled applications to select populations, such as military personnel in both the active and reserve components; for subsets of military personnel, such as pilots; and for civilian personnel, such as public school teachers and state and federal employees. It has also enabled analyses of different compensation structures and systems; alternative retirement systems; compensation systems that require multi-year contracts; and different pay table structures, such as a time-in-service versus a time-in-grade pay structure. In addition, the model has enabled consideration of not only steady-state results but also the transition from one steady state to another.[1]

This report summarizes RAND Corporation research on the DRM from a technical standpoint, drawing from many previous documents that have used and further extended the

[1] *Steady state* refers to the period when the entire workforce has spent their entire careers under a given compensation system. For example, if the organization changes its retirement system and a full career spans 30 years, the new steady state after the change in the retirement system would occur after 30 years when entrants under the new retirement system have completed their entire careers. The transition period spans each of the 30 years from the baseline before the policy change to 30 years hence, when the new steady state has been achieved.

DRM. The purpose is to provide practitioners with the technical details of the DRM and recent extensions in one document rather than scattered across many. Our approach was to provide a review of RAND reports in which the DRM played a central role, focused on the model estimated in support of the 13th Quadrennial Review of Military Compensation (QRMC). While the focus in the main report is on the technical aspects of the model, we provide an informal introduction to the model in Chapter 2 and an annotated bibliography of the policy analyses conducted with the DRM and the published documents corresponding to each area of analysis in the appendix.

Conclusions

The DRM provides a practical capability for modeling the retention of military personnel, civil service employees, public school teachers, and state employees. It is based on a rigorous, logically consistent framework and has been successfully extended to cover multiple topics of interest, such as retirement reform and the structure of special and incentive pays.

Like any model, the DRMs used in our analyses have limitations. We will confine our remarks here to the military DRM, but similar remarks could apply to our other models. The DRM does not explicitly model other factors that can affect retention and retirement, including health status and health care benefits or household factors, such as family formation, spousal labor supply, or the presence of children at home. The analysis focuses on retention and does not model the decision to enlist (or, for officers, to access) into the military. Consequently, the model cannot address how changes to pension design might affect the types of people who become soldiers. Another limitation is that the model assumes risk neutrality. The utility function is assumed to be linear in compensation. While conceptually a more flexible functional form could be used, practically speaking, the computational challenges are formidable.

That said, the estimated models fit the observed data well. Our approach has several rich and realistic features that make it well suited for analyzing the retention effects of alternative compensation policies and pension reform. It is a life-cycle model in which retention decisions are made each year over an entire career, not just once. Those decisions are based on forward-looking behavior that depends on existing and future military and external compensation. The model allows for uncertainty in future periods and recognizes that people might change their mind in the future as they get more information about staying in the military and their external opportunities. Furthermore, the model is formulated in terms of the parameters that underlie the retention decision processes rather than on the average responses to historical changes in policy. Consequently, it is structured to enable assessments of alternative compensation reforms that have yet to be tried. Put differently, the DRM is particularly suited to assess major structural changes in the compensation system that do not have any historical antecedent.

Contents

About This Report iii
Summary v
Figures and Tables viii
Chapter 1. Introduction 1
- Organization of This Report 3

Chapter 2. Modeling Retention 4
- Modeling the Value of Flexibility—An Example 6
- A Retention Model 7
- Modeling Uncertainty—Taste 7
- Modeling Uncertainty—Shocks 10
- Mathematical Description of a Simple Version of the DRM 10
- Challenges Posed in Estimating the Model 14
- Summary 15

Chapter 3. The RAND Dynamic Retention Model 16
- Overview of RAND Dynamic Retention Model Research 16
- The 13th QRMC DRM Capability 19
- Simplifying Assumptions and Efficient Computing 30

Chapter 4. Model Estimates and Fit and Simulation 35
- Model Estimates 35
- Simulation and Assessment of Model Fits 41
- Simulation and Extension of the DRM to Model a Time-in-Grade Pay Table 49
- Summary 57

Chapter 5. Additional Innovations and Extensions to the DRM 58
- Adding Covariates to Shift the Taste Distribution 58
- Modeling the Effects of Bonuses and Special Pays That Involve a Multi-Year Contract Choice 61
- Modeling the Transition from One Steady State to a New Steady State 64
- Using Incumbents to Supplement Short Panels in Estimation 67
- Concluding Remarks 70

Chapter 6. Conclusion 71
- Looking Backward 71
- Looking Forward 72

Appendix. Thematic RAND DRM Bibliography 73
Abbreviations 81
References 82

Figures

Figures

Figure 2.1. Influences on Individual Retention Decisions 5
Figure 2.2. Simple Retention Model 7
Figure 2.3. Modeling Individual Taste and Uncertainty 8
Figure 2.4. Distribution of Taste for Air Force Pilot ROTC Graduates..9
Figure 2.5. Population Distribution of Taste Changes with Increase in Tenure 10
Figure 4.1. Model Fit Results: Active-Component Army Officers..44
Figure 4.2. Model Fit Results: Active-Component Navy Officers...44
Figure 4.3. Model Fit Results: Active-Component Air Force Officers..45
Figure 4.4. Model Fit Results: Active-Component Marine Corps Officers 45
Figure 4.5. Model Fit Results: Active-Component Army Enlisted..47
Figure 4.6. Model Fit Results: Active-Component Navy Enlisted...47
Figure 4.7. Model Fit Results: Active-Component Air Force Enlisted...48
Figure 4.8. Model Fit Results: Active-Component Marine Corps Enlisted 48
Figure 4.9. Years to Promotion by Ability Level, Army Enlisted Personnel...............................52
Figure 4.10. Years to Promotion by Ability Level, Army Officers..52
Figure 4.11. Calibrating the Parameters of the Ability Distribution, Army Enlisted Personnel...53
Figure 4.12. Calibrating the Parameters of the Effort Decision, Active-Component Army Enlisted Personnel 57

Tables

Table 4.1. Parameter Estimates and Standard Errors: Army and Navy Officers 36
Table 4.2. Parameter Estimates and Standard Errors: Air Force and Marine Officers 36
Table 4.3. Transformed Parameter Estimates: Officers 37
Table 4.4. Parameter Estimates and Standard Errors: Army and Navy Enlisted 39
Table 4.5. Parameter Estimates and Standard Errors: Air Force and Marine Enlisted 39
Table 4.6. Transformed Parameter Estimates: Enlisted 40
Table 5.1. How We Develop Retention Profiles by Time Elapsed: Cumulative Probabilities.....66
Table 5.2. How We Develop Retention Profiles by Time Elapsed: Value of the *s*.......................66

Chapter 1. Introduction

The dynamic retention model (DRM) is a model that has been used to assess the retention effects of changes to compensation in both the military and civilian contexts. The DRM is part of a larger class of models known as *structural models of individual behavior* in which—in our context—people make retention decisions under uncertainty during their careers and have unique or heterogeneous tastes. That is, the DRM recognizes that individuals are forward-looking when they make their decisions (meaning they account for possible future events, such as future pay raises and how decisions might affect future events) and that people differ (meaning that people who face identical opportunities might make different choices). In the DRM, individuals make retention decisions in each year over their career and choose the alternative in each period—to stay or leave—that maximizes their utility, meaning they make the optimal choice each year given the information and opportunities available to them.

The advantage of this approach to modeling retention is that the approach has a solid foundation in the literature and in theories of how personnel make retention decisions over their careers. In particular, the DRM is formulated in terms of the parameters that underlie the retention decisionmaking process rather than on a specific compensation system and retention outcomes. When such a model is estimated with empirical data on individual retention decisions, it can be credibly used for counterfactual analyses, permitting assessments of alternative yet-to-be-implemented compensation systems. Furthermore, it includes a capacity to perform simulations of the retention and cost effects of relevant policy changes, including policies that have yet to be implemented.

The flexibility and logical consistency of this behavioral model has enabled applications to select populations, such as military personnel in both the active and reserve components; for subsets of military personnel, such as pilots; and for civilian personnel, such as public school teachers and state and federal employees. It has also enabled analyses of different compensation structures and systems; alternative retirement systems; compensation systems that require multi-year contracts; and different pay table structures, such as a time-in-service versus a time-in-grade (TIG) pay structure. In addition, the model has enabled consideration of not only steady-state results but also the transition from one steady state to another.

While analysts agree that the stochastic dynamic programming approach is the state-of-the-art and ideal approach to modeling stay-leave decisions, such as retention decisions, it is only comparatively recently (within the last two decades) that researchers have been able to take this approach without making certain simplifications.[2] In part, the need for simplifications reflected

[2] In the context of retirement decisions, Blundell, French, and Tetlow (2016) discuss the pros and cons of a structural approach relative to reduced form approach. These arguments are also relevant to the retention decision.

the early challenges associated with the computational requirements of the model. The DRM requires computation of the value of staying and leaving under all possible career paths and the probabilities associated with that path, where these values and the probabilities are specified in terms of the underlying parameters of the model. Over a 30-year career, the number of paths can be quite large, especially when individuals are heterogeneous. In the early years of the all-volunteer force, researchers reduced the computational burden by using an approximation approach called the Annualized Cost of Leaving model (ACOL).[3] But the ACOL approach has been shown to be logically inconsistent, and it provides implausible estimates of the retention effects of compensation changes. Another approach is a method called *calibration* where the researchers choose rather than estimate the parameter values that fit observed data. The problem with this approach is that different sets of calibrated parameter values can fit the same data. That is, the parameters are poorly identified, and different parameters can lead to different results when policy analysis is conducted (Cooley, 1997). Furthermore, calibration does not provide standard errors of the estimates, so it is impossible to test for the statistical significance of the parameters.

A different approach to reducing computational burden is to simplify the problem by finding a way to avoid solving the stochastic dynamic program entirely, following an approach developed by Hotz and Miller (1993). Unlike the DRM approach that solves for the optimal decision rules over the entire career as we describe in later chapters, the Hotz-Miller approach estimates reduced form rather than structural parameters for individual decision probabilities. It then *inverts* these probabilities to estimate the value of staying, something that would otherwise require solving the dynamic program. This approach assumes *stationarity*, meaning that all entering cohorts are assumed to be identical so that parameters can be estimated with cross-sectional data or with limited time-series data. Early versions of the Hotz-Miller approach also assumed that observationally identical individuals facing identical opportunities would make identical decisions—no differences in unobserved preferences. Later versions of their approach allow for individual heterogeneity.

Since 2007, the RAND Corporation has successfully implemented the DRM, despite its computational burden, without resorting to simplifications, such as calibration or the Hotz-Miller approach. Instead, the DRM is estimated with individual longitudinal data on retention rather than cross-sectional data, and the process of estimation involves solving the dynamic program for each member for a trial set of parameters using a maximum likelihood algorithm to step toward a new set of trial parameters to seek a better fit to the data and repeating the process until *convergence* (no further appreciable improvement in model fit).

[3] The approximation was to exchange the expectation and maximization operators in the value function; that is, to compute the maximum of the expected values instead of the expected value of the maximum. This approximation does not capture the value associated with an individual being able to choose between alternatives once new information has been revealed.

This report summarizes RAND research on the DRM from a technical standpoint, drawing from many previous documents that have used and further extended the DRM. The purpose is to provide practitioners with the technical details of the DRM and recent extensions in one document rather than scattered across many. While the focus is on the technical aspects of the model, we provide an informal introduction to the model in Chapter 2 and an annotated bibliography of the policy analyses conducted with the DRM and the published documents corresponding to each analysis in the appendix.

Organization of This Report

In Chapter 2, we lay the groundwork for our presentation of RAND's DRM by walking through the Gotz-McCall model (Gotz and McCall, 1984). We also highlight more rigorously the challenges associated with estimating the model. In Chapter 3, we turn to RAND's implementation of the model. We provide a brief history of RAND DRM work since the early 1990s and then provide a technical presentation of a recent implementation of the DRM. Specifically, we present theory and the methodology for estimating the model, as well as descriptions of the data used. In Chapter 4, we show the estimates for the DRM we created in support of the 13th Quadrennial Review of Military Compensation (QRMC). We also present model fits and several extensions to the basic model, including the promotion process and how performance—as measured in terms of promotion speed relative to one's peers—might be affected by ability and effort supply. In Chapter 5, we consider four additional extensions to the model: (1) adding regression variables to change the mean and variance of the taste distribution, (2) modeling retention when individuals can choose multi-year contracts, (3) modeling the transition from one steady state to a new steady state, and (4) using incumbents to supplement short panels in estimation. In Chapter 6, we share some concluding remarks—reflecting on the past and looking toward the future. The appendix presents an annotated bibliography of the policy analyses conducted with the DRM and the published documents corresponding to each analysis.

Chapter 2. Modeling Retention

Simulation models using the DRM approach have been available since the late 1970s, but it has been difficult to use because it is computationally intensive. Advances in both computing software and hardware over the past two decades have eliminated this drawback. Using the DRM, we developed a method for statistically estimating model parameters using historical career data and simulating the effect of changes in personnel policies on retention. In this chapter, we first give an informal description of the basic concepts underlying the DRM and then give a mathematical description of a simplified version of the Gotz-McCall (1984) model, the original DRM. This chapter draws extensively from Mattock and Arkes (2007).

A key attribute of this approach is that it focuses on individual behavior. Figure 2.1 shows our concept of the process leading to a decision to stay or leave. In this figure, we use the Air Force as an example, but this figure could apply to any service. Two aspects of the figure merit comment. One is that individual retention decisions result from a complex interaction of many influences. Certainly, service compensation policies influence a service member's decision to stay or leave. However, the strength of that influence varies depending on the individual. A service member who really enjoys military service (has a "taste" for it) might elect to remain in the service for less compensation than would an individual with less of a taste for service. External influences are important as well. If the civilian job market is robust and the individual's skills are in demand, then the motivation to leave would be relatively greater than if the market is poor or the individual's skills are not prized. A second point worthy of comment (which might seem obvious) is that aggregate or group behavior is driven by individual decisions. Looking at how individuals make decisions gives us more insight into the retention process than does studying only the mythical average member.

It is important to focus on the behavior of specific individuals to arrive at parameter estimates that describe the preferences of service members regarding key aspects of their environment. Using the DRM, we model a service member's decision process and take into account the attempts that individuals make to optimize their futures. By modeling individual decisions, any given parameter estimate is less dependent on specific policies in effect during the period covered by the data than it would be using a traditional regression approach. If we construct a model of the internal decisionmaking process of a service member that takes into account regular military compensation, the military retirement system, and civilian career opportunities, then estimates of the remaining parameters depend less on these factors. This type of model can be used to predict the effect of a broader variety of compensation and personnel policy options. On the other hand, if we construct a retention model for service members that did not include the design of the military retirement system in the service member's decisionmaking calculus, then the estimates provided by that model could be used with confidence only if all aspects of the retirement system remained unchanged, because the estimated parameters implicitly depend on the specifics of the retirement system in effect during the period covered by the data used to estimate the parameters of the model. (By specifics, we mean, for example, the vesting rules, any defined benefit component, and any defined contribution component.) Many traditional regression models fall into this category.

Figure 2.1. Influences on Individual Retention Decisions

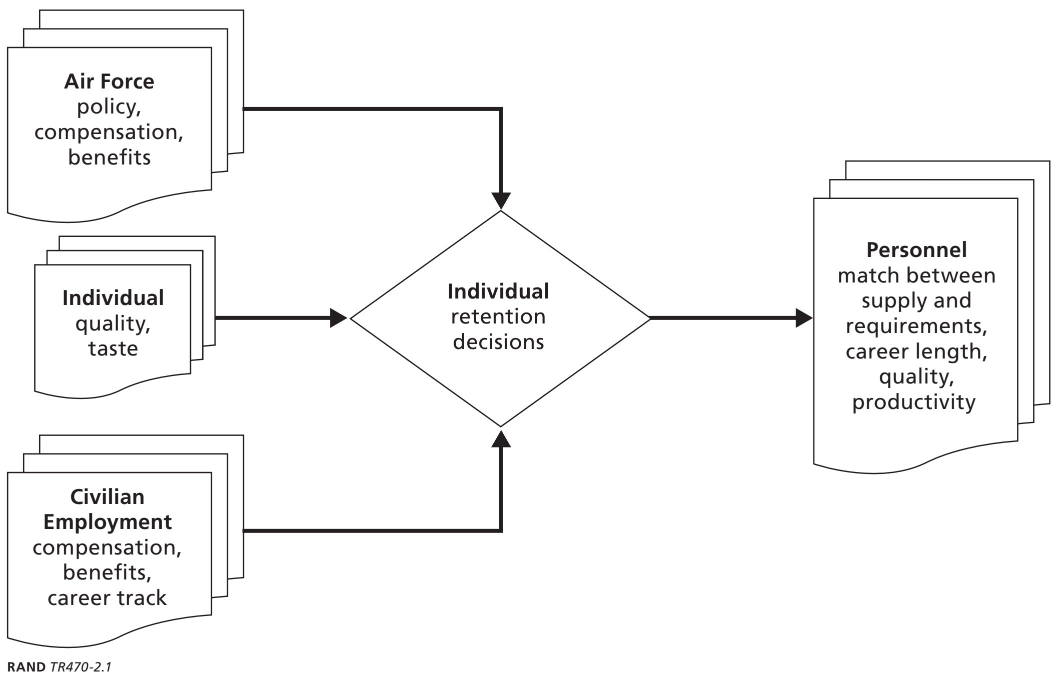

SOURCE: Reproduced from Mattock and Arkes (2007).

Explicitly modeling individual behavior also allows for individuals being different. People's behavior can differ because of both observable and unobservable characteristics. For example, a service member's decision to stay in the military can be affected both by their particular promotion history (an observable characteristic) or their taste for military service (an unobservable characteristic). The DRM allows for differences in both observed and unobserved characteristics, whereas some other models, such as ACOL, only allow for differences in observable characteristics.

One of the key features of the DRM is that it explicitly models a service member's decision calculus as taking into account future uncertainty. Other models of retention—such as ACOL,

ACOL 2,[4] or the Ausink and Wise (1993) "option value"[5] model—do not explicitly include future uncertainty in the service member's decision calculus. Including uncertainty enables us to model *flexibility*—the ability to make or change decisions when new information comes to light. This is sometimes referred to as an option value and is a common concept to those who trade in securities. The ability to buy or sell an option at a particular price has value: It enables a person to hedge risk. A concrete example might help illustrate this point. Readers who are more interested in the conceptual framework of the DRM might wish to skip to the next section.

Modeling the Value of Flexibility—An Example

Consider the case of betting on a coin flip. A head means that the individual wins $1, and a tail means $0. Thus, the expected value of the bet is

$$\frac{1}{2}\$1.00 + \frac{1}{2}\$0.00 = \$0.50.$$

Now consider a case in which (1) there are two coins that each have an equal chance of coming up heads or tails and (2) the bettor can choose either coin before it is flipped. The expected value from choosing a coin is the same as that in the previous example, $0.50. However, now consider the case in which the bettor can choose between the two coins after they are flipped. If both come up heads, the bettor can choose either one and receive $1. If only one comes up heads, the bettor can choose that one and still receive $1. If both come up tails, then the bettor receives nothing. The expected value of this bet is $0.75 because three times out of four the bettor can receive $1. The following formula describes this result:

$$\frac{1}{4}\$1.00 + \frac{1}{4}\$1.00 + \frac{1}{4}\$1.00 + \frac{1}{4}\$0.00 = \$0.75.$$

So, the ability to make an informed choice has an expected value of $0.25 ($0.75–$0.50 = $0.25). If all anyone cared about was the expected value of the return on the bet, then to get the bettor to give up the opportunity to choose after the coins had been flipped, they would have to be compensated by at least $0.25 because the value of the bet with no choice ($0.50) plus compensation for losing the opportunity to choose ($0.25) would just equal the value of a bet with choice ($0.75).

The exact magnitude of the option value will depend on the size of the random shocks (variation in, for example, civilian opportunities or health events) service members are subject to year to year. Service members can experience random shocks from both the civilian and military

[4] ACOL 2 is a version of the ACOL model that allows for unobserved heterogeneity.

[5] In the Ausink and Wise model, the member compares the utility of leaving now with the maximum value of expected future utilities associated with postponing leaving.

sides. On the military side, a service member might receive a good or bad assignment, might be passed over for promotion, and so on. On the civilian side, a service member might have the opportunity to take a high-paying civilian position, might see that civilian job opportunities have declined, might find that he or she needs to leave the service to care for an ailing parent, and so on. While we cannot directly observe the distribution of these shocks, we can statistically infer distribution of the difference between the military and civilian shocks in terms of dollars by using the DRM.

The approach taken by the ACOL, ACOL 2, and Ausink and Wise models is to calculate the maximum of the expected values rather than the expected value of the maximum. In the simple example we gave previously, the maximum of the expected values was \$0.50, while the expected value of the maximum was \$0.75. So, the maximum of the expected values can be a very poor approximation to the expected value of the maximum, because it does not reflect the value associated with being able to make an informed choice.

A Retention Model

Figure 2.2 depicts a simple retention model. In this model, each service member makes a decision at the beginning of the period to either stay or leave. If the member stays, they collect the benefits associated with remaining in the military for a year, including the value of the option to stay or leave at the next decision point. If the member leaves, they get the value of a civilian career path starting in that period. In this simple model, behavior is deterministic. This model implicitly assumes that service members with identical observable characteristics would behave identically. It takes no account of the possibility that nominally identical service members might make different decisions about whether to stay or leave.

Figure 2.2. Simple Retention Model

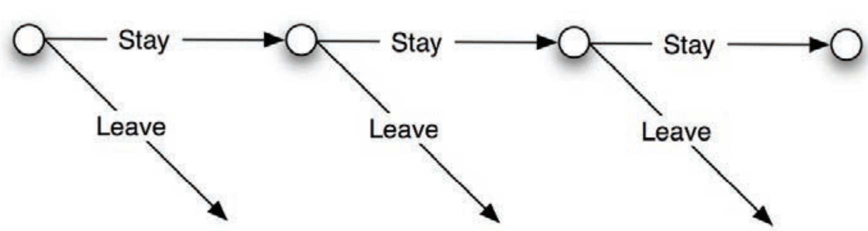

SOURCE: Reproduced from Mattock and Arkes (2007).

Modeling Uncertainty—Taste

This simple retention model is a start, but it is insufficient for our purposes because it does not allow for differences among observationally identical individuals. Allowing for differences in individual retention behavior requires the modeling of uncertainty. Figure 2.3 depicts a more sophisticated model that injects uncertainty and considers differences in individuals' characteristics or in the environment that an individual faces.

Figure 2.3. Modeling Individual Taste and Uncertainty

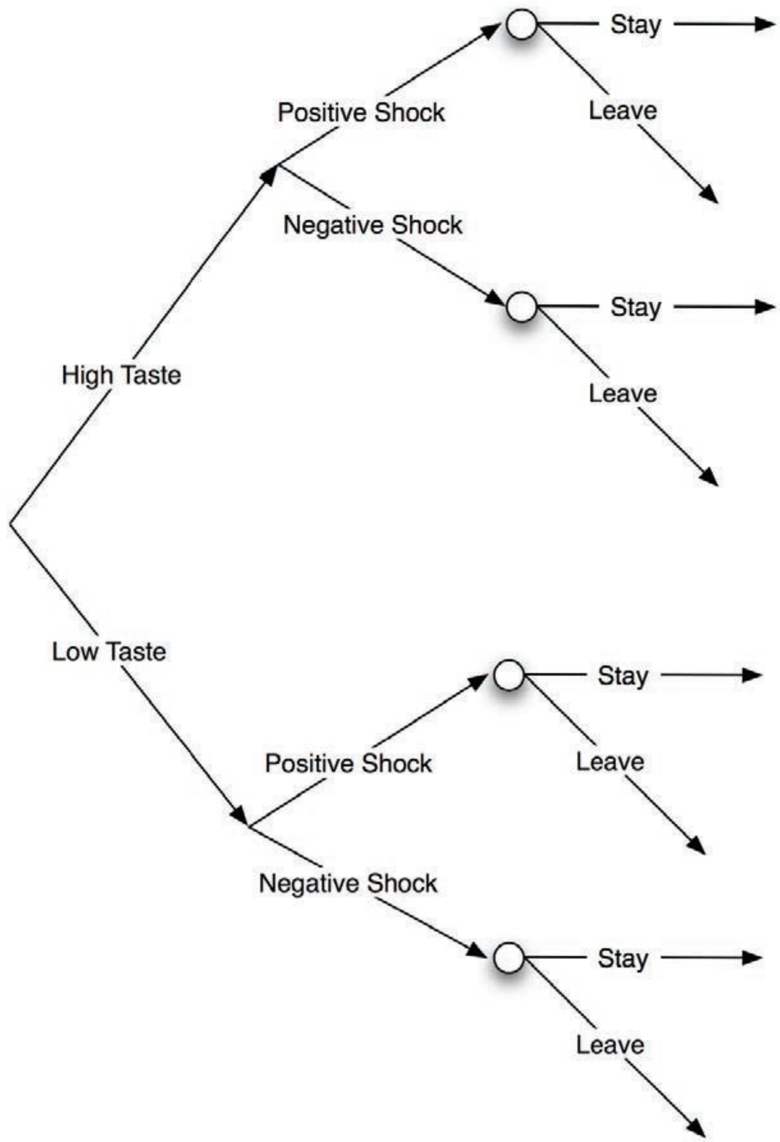

SOURCE: Reproduced from Mattock and Arkes (2007).

In this example, both the individual being modeled and the analyst face uncertainty. It begins with an individual who has a certain taste for military service. We assume that an individual is aware of their taste for military service and makes decisions accordingly, but this taste is not known to the analyst. The individual then experiences a positive or a negative shock. The value of the shock is unknown in advance to either the service member or the analyst. The shock affects the value that the individual places on staying in the military until the next decision. The shock can make an individual place either a higher value on staying (a positive shock) or a lower value on staying (a negative shock). Thus, the analyst faces uncertainty over both the value that a particular individual places on staying in the military and the value of the shock they might experience in any given period.

One analytical approach to this problem is to assume that taste is distributed across a population according to some parameterized distribution and then estimate the parameters of the taste distribution in a statistical model. Figure 2.4 presents one such estimate, in this case one developed for pilots who graduated from the Reserve Officer Training Corps (ROTC).

Figure 2.4. Distribution of Taste for Air Force Pilot ROTC Graduates

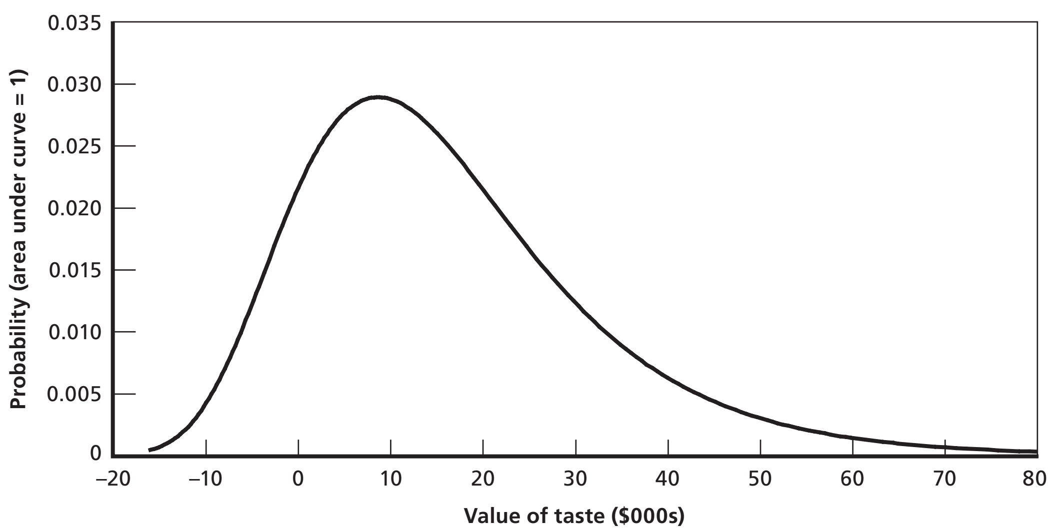

SOURCE: Reproduced from Mattock and Arkes (2007).

The figure shows the estimated distribution of the taste for military service held by the population of ROTC accession pilots when they reach their first stay-versus-leave decision points. The dollar values shown represent the monetary equivalent of the intrinsic value that an individual places on a year of military service (in addition to compensation and other benefits). A service member with a strong taste for military service would require relatively more money to be induced to leave than a member with a weak taste. This curve reflects the initial distribution of taste for the group. The shape of this curve will change over time as service members leave the service. That change is reflected in the curves displayed in Figure 2.5, which shows how the population distribution of taste changes over time.

Figure 2.5. Population Distribution of Taste Changes with Increase in Tenure

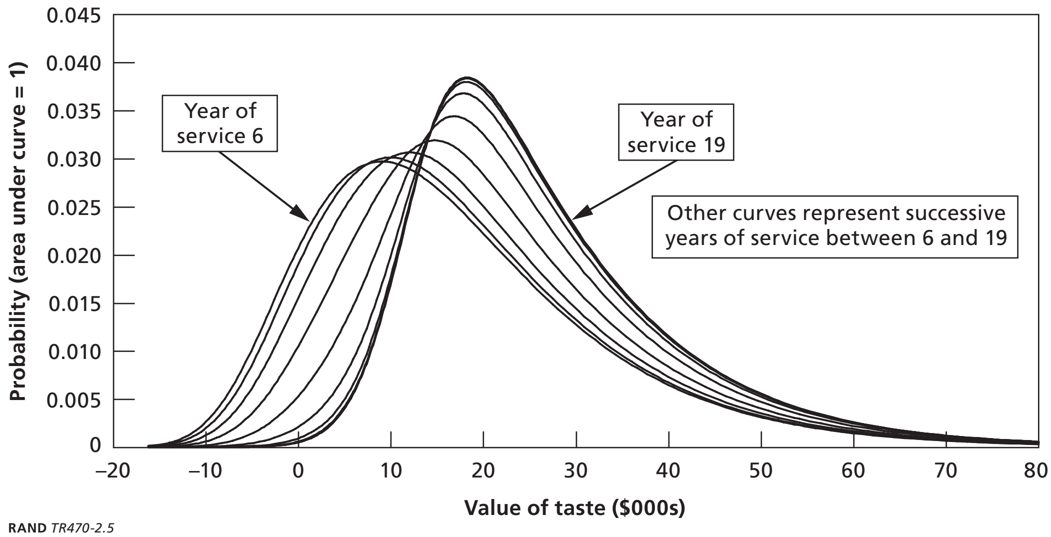

SOURCE: Reproduced from Mattock and Arkes (2007).

Figure 2.5 shows that the population distribution of taste for service increases with tenure for a notional service member population. This is relatively intuitive, because those who value active-duty service will tend to stay longer. The chart shows less taste for service among those with six years of service (the curve farthest to the left), the greatest taste among those with 19 years of service (the curve farthest to the right), and taste distributions gradually shifting from left to right with each successive year of service.

Modeling Uncertainty—Shocks

We assume that service members face shocks every year (which, as we noted earlier, are unanticipated events that will affect their desire to remain in active service). These shocks can be positive or negative. A positive shock is one that strengthens their preference for staying and a negative shock is one that has the opposite effect. Service members who choose to leave forgo the possibility of future positive shocks (e.g., a desirable assignment, an accelerated promotion, the opportunity for new and interesting training). The model assumes that the shocks are independently and identically distributed across the population.

Mathematical Description of a Simple Version of the DRM

We will begin explaining our approach to implementing the DRM by showing the equations for a simple stay-leave model. This model is very similar to the original Gotz-McCall (1984) model and the later Daula and Moffit model (1995). For simplicity, the model omits the Markov promotion process and the selection for reserve versus regular commission present in the original

Gotz-McCall model. We will then show how these equations can be used to derive an expression for the probability of staying and how the probability of staying expression can be used in conjunction with empirical career data to form a likelihood function that will allow us to estimate the structural parameters of our model.

The model consists of two equations, one giving the value of staying for an additional year and revisiting the stay-versus-leave decision and one giving the value of leaving. A service member decides to stay if the value of staying is greater than the value of leaving. To turn this into a statistical model, we add assumptions about the distribution of an independently and identically distributed random shock term and also about the distribution of the members' taste for military service.

The non-stochastic value of leaving is represented by

$$V_t^L = W_t^c + \sum_{\tau=t}^{T} \beta^{\tau-t} W_\tau^c + R_t^m,$$

and the stochastic value of leaving is represented by

$$V_t^L + \varepsilon_t^c$$

where

V_t^L is the value of leaving at time *t*,

W_t^c is civilian earnings at time *t*,

$$\sum_{\tau=t}^{T} \beta^{\tau-t} W_\tau^c$$

is the value of future civilian earnings (where β is the annual discount rate),

R_t^m is the retirement benefit accruing to the member if they leave at time *t*, and

ε_t^c is random civilian shock at time *t*.

The non-stochastic value of staying is represented by

$$V_t^S = \gamma^m + W_t^m + \beta \, \mathrm{E}_t[\mathrm{Max}(V_{t+1}^L, V_{t+1}^S)],$$

and the stochastic value of staying is represented by

$$V_t^S + \varepsilon_t^m$$

where

V_t^S is the value of staying at time *t*,

γ^m is individual taste for serving in the military,

W_t^m is military earnings in the current period t,

$\beta\,\mathrm{E}_t[\mathrm{Max}(V_{t+1}^L, V_{t+1}^S)]$ is the discounted expected value of being able to choose to stay or leave in the future, with β being the discount rate, and

ε_t^m is the random military shock at time t.

Retention probability in period t is represented by

$$\begin{aligned}\Pr[V_t^S + \varepsilon_t^m > V_t^L + \varepsilon_t^c] \\ &= \Pr[V_t^S - V_t^L > \varepsilon_t^c - \varepsilon_t^m] \\ &= \Pr[V_t^S - V_t^L > \varepsilon_t]\end{aligned}$$

where

$\varepsilon_t = \varepsilon_t^c - \varepsilon_t^m$ = the difference between the civilian and military shocks at time t.

The individual will decide to stay in the military if the value of staying is greater than the value of leaving. If a probability distribution is set for the difference between the random military shock and the random civilian shock, the probability that an individual service member with a particular taste will stay can be computed.

In this example we assume that ε_t is normally distributed with mean zero and variance σ. This leads to a closed-form solution for the expected value of the maximum of staying or leaving, specifically:

$$\begin{aligned}\mathrm{E}_t[\mathrm{Max}(V_{t+1}^L, V_{t+1}^S)] &= F((V_{t+1}^S - V_{t+1}^L)/\sigma)V_{t+1}^S \\ &+ (1 - F((V_{t+1}^S - V_{t+1}^L)/\sigma))V_{t+1}^L \\ &+ f((V_{t+1}^S - V_{t+1}^L)/\sigma)\sigma\end{aligned}$$

where $F(.)$ and $f(.)$ are the cumulative distribution function and probability density function (PDF) of the unit normal distribution.

In this equation, the first term is the probability of staying multiplied by the value of staying, the second term is the probability of leaving multiplied by the value of leaving, and the third term shows the expected value to the individual of being able to choose whether to stay or leave. This closed form solution is desirable; otherwise, the expected value of the maximum would have to be calculated via numerical integration, which is a relatively slow and imprecise process.

The probability of an individual with taste γ facing shock distribution $F(.)$ choosing to stay in period t can now be computed using:

$$\begin{aligned}\Pr[\mathrm{Stay}_t \mid \gamma, \sigma] &= \Pr[\mathrm{Stay}_t] \\ &= \Pr[V_t^S - V_t^L > \varepsilon_t] \\ &= F((V_t^S - V_t^L)/\sigma)\,.\end{aligned}$$

The probability of leaving is simply:

$$\Pr[\text{Leave}_t] = 1 - \Pr[\text{Stay}_t \mid \gamma, \sigma].$$

The following example applies this probability to assess the likelihood of a sequence of events. If we observe a particular individual whose service obligation completed at $t-1$ choosing to stay at $t, t+1$ and leaving at time $t+2$, we can compute the joint probability of this sequence of events like so:

$$\Pr[\text{Stay}_t \mid \gamma, \sigma] \Pr[\text{Stay}_{t+1} \mid \gamma, \sigma] \Pr[\text{Leave}_{t+2} \mid \gamma, \sigma].$$

Similarly, if we observe an individual leaving immediately at the conclusion of their service obligation, then the probability is $\Pr[\text{Leave}_t \mid \gamma, \sigma]$. If we observe an individual staying for two periods and then do not observe what they do in the third period, the probability is

$$\Pr[\text{Stay}_t \mid \gamma, \sigma] \Pr[\text{Stay}_{t+1} \mid \gamma, \sigma].$$

In general, the probability of observing someone staying for *s* periods will be

$$\prod_{\tau=t}^{t+s} \Pr[\text{Stay}_\tau \mid \gamma, \sigma],$$

and the probability of observing someone stay for s periods and then leaving will be

$$\Big(\prod_{\tau=t}^{t+s} \Pr[\text{Stay}_\tau \mid \gamma, \sigma]\Big) \Pr[\text{Leave}_{t+s+1} \mid \gamma, \sigma].$$

Of course, this probability is conditioned on the unobservable taste parameter γ. In this example we assume that γ is extreme-value distributed with mode α and scale parameter δ. We will use $g(.)$ to denote the cumulative distribution fnction of this distribution. We can use this distribution to construct an expected probability for a particular sequence of events.

If a person leaves immediately after the conclusion of their service obligation at *t*, the expected value of the probability is

$$L_i(\alpha, \delta, \sigma, \beta) = \int_{-\infty}^{\infty} \Pr[\text{Leave}_t \mid \gamma, \sigma]\, g(\gamma) d\gamma$$

If a person stays for *s* periods, the expected value will be

$$L_i(\alpha,\delta,\sigma,\beta) = \int_{-\infty}^{\infty} \prod_{\tau=t}^{t+s} \Pr[\text{Stay}_\tau \mid \gamma,\sigma]\, g(\gamma) d\gamma \,.$$

If a person stays for *s* periods and then leaves, the expected value will be

$$L_i(\alpha,\delta,\sigma,\beta) = \int_{-\infty}^{\infty} \prod_{\tau=t}^{t+s} \Pr[\text{Stay}_\tau \mid \gamma,\sigma] \Pr[\text{Leave}_{t+s+1} \mid \gamma,\sigma]\, g(\gamma) d\gamma.$$

Thus, the likelihood for the entire sample will be

$$L\ (\alpha,\delta,\sigma,\beta) = \prod_{i=1}^{n} L_i(\alpha,\delta,\sigma,\beta).$$

Challenges Posed in Estimating the Model

Dynamic programming models have a reputation for being difficult to estimate. For example, the pioneering work of Gotz and McCall (1984) did not provide standard errors for the model estimates because of the computing challenges. With the advent of faster computers and improvements in optimization algorithms, the task has become easier, but can still be challenging.

The main challenge is optimizing the likelihood function. Because of the approximation inherent in machine arithmetic—and, in particular, numerical integration—calculating numerical derivatives is not possible at all points of the likelihood surface. In addition, depending on the type of model, there can be many *local maxima*—many points on the surface where the gradient will be at or near zero. This means that most hill-climbing algorithms will generally fail.

An additional challenge is posed by the need to integrate out heterogeneity. As there is no closed-form solution for the integral in the likelihood function, it has to be estimated numerically. Unfortunately, many modern algorithms for integration, even those that recursively subdivide the interval of integration to a high depth, do not perform well when faced with functions that are near zero over most of the interval (Piessens et al., 1983).

The weak identification of the model parameters also poses a challenge, because it means that straight maximum likelihood estimates of the parameters might veer wildly from values that are credible. The method used in the original work of Gotz and McCall on Air Force officer retention to assure identification is not available to us because officers are no longer granted regular or reserve commissions when entering the officer ranks. Gotz and McCall posited that the individuals with higher gamma (higher taste) were more likely to be granted a regular commission. The Gotz and McCall likelihood function used a posterior distribution for gamma that was conditional on whether an officer was granted a regular or reserve commission, the relative proportion of the officers given a reserve commission, and a selectivity parameter that was estimated along with the other model parameters. (In a personal communication to one of the authors [Mattock] in 2002, Glenn Gotz stated that he could not get the original model to converge without adding this selectivity term.)

Later RAND projects addressed the weak identification challenge in different ways: In the Mattock and Arkes (2007) report on Air Force pilot retention, the authors found that model parameters could be identified given information on active-duty service obligation (ADSO) associated with pilot training and the structure of multi-year contracts. In Asch et al. (2008) and in later RAND reports on military personnel retention, authors found that including information on reserve component participation in addition to active retention helped to identify model parameters, in particular those associated with taste.

These challenges are difficult but not insurmountable. Judicious simplifications—such as using a wage trajectory rather than modeling the promotion process and using the pay table, as well as the other simplifications mentioned in Chapter 1—can help to make the dynamic program computationally feasible and improve computational efficiency by avoiding repeatedly solving identical subproblems by using caching speeds computation of dynamic programs. Furthermore, using Halton sequences to generate antithetic pairs of numbers when choosing support points for the taste distribution helps in calculating the weighted average over a population of interest (i.e., integrating out heterogeneity).

Summary

Modeling individual behavior requires modeling uncertainty explicitly. The individual components of uncertainty include elements that are known to the service member but unknown to the analyst, such as individual taste for military service relative to their next best alternative. The components of uncertainty also include elements that are unknown to both the service member and the analyst, such as future environmental shocks and future promotions. By modeling how these elements affect an individual's decisionmaking process, we can gain a better understanding of how an individual might value future career flexibility in the face of uncertainty. Furthermore, uncertainty can be explicitly captured in a mathematical model with structural parameters that can be estimated using empirical data—the DRM.

Chapter 3. The RAND Dynamic Retention Model

RAND researchers built on the Gotz-McCall (1984) DRM and addressed the challenges associated with estimation. This chapter provides technical details of RAND's approach, using the analysis conducted in support of the 13th QRMC in 2020 as an example. This model was developed and used to assess the retention, cost, and performance effects of reforming the military pay table. We will first give an informal description of the DRM used for the 13th QRMC then discuss the mathematical structure of the DRM and how we extended it to account for promotion. We will then discuss the data used for estimation and the estimation methodology. In the next chapter, we will discuss estimates, model fit, and the simulation capability that we developed to simulate pay table reform.

The example of the DRM developed for the 13th QRMC is a good one for showing the technical details of the model because it shows many of the innovations made to the DRM since RAND first began developing the DRM. However, we note that we have developed the DRM capability and applied it in other contexts for specific populations or policy questions, and, arguably, these other studies could be additional examples.[6] We therefore preface our description of the 13th QRMC DRM capability with an overview of these other studies. The appendix provides a list of the studies, organized by policy area. Furthermore, we provide an overview of four of these other examples in Chapter 5.

Overview of RAND Dynamic Retention Model Research

The purpose of this overview is to place the example DRM developed for the 13th QRMC in the context of the broader set of studies using the DRM. A more complete summary can be found in the appendix. Readers primarily interested in the 13th QRMC DRM capability might wish to proceed to the next section.

Except for work by Daula and Moffitt (1995), who estimated the DRM with data from two enlisted Army cohorts from the 1980s, no DRM analyses occurred following the seminal work by Gotz and McCall until the 1990s, owing to the computational complexity of the model given the computer technology available in the 1980s. However, that began to change in the 1990s. The Gotz-McCall model was extended in several ways in the 1990s. Asch and Warner (1994b)

---

[6] The DRM is not a model that can be taken off the shelf and applied by any analyst to any compensation-related question. Rather, it is a capability that can serve as a foundation for building models to address personnel policy questions that cannot already be addressed by existing DRMs. Straightforward policy explorations involving changing inputs (such as military or civilian earnings) may be done within any estimated DRM. More-complex policy explorations that introduce new choices by the service member (such as special pays that are conditional on the choice of an additional ADSO) might require an extension if this kind of choice is not already reflected in the DRM.

incorporated performance into the model, and Asch and Warner (1994a) calibrated a simulation model that enabled them to estimate the steady-state retention, performance incentives, and cost implications of alternative military compensation and retirement reform policy alternatives. Asch, Johnson, and Warner (1998) extended the simulation model further to estimate the retention, cost, and productivity effects of transitioning to a military retirement system that resembled the Federal Employees Retirement System. An updated version of the Asch-Warner simulation model (1994b) was used to assess retirement alternatives for the Defense Advisory Committee on Military Compensation (DACMC). Asch and Hosek (1999) also employed the simulation model to analyze the behavioral and cost implications of the TRIAD legislation included in the National Defense Authorization Act of 2000 that addressed concerns among military leadership in the late 1990s (when the services struggled to meet their military recruiting and retention targets) about adverse effects on retention and morale because of the reduced value of retirement benefits under the reform plan implemented in 1986 (often referred to as REDUX). Hosek et al. (2004) incorporated the enlistment decision into the DRM framework in a study of the recruitment and retention of information technology (IT) personnel. They also modeled skill accumulation—the learning of IT skills through training and experience provided by the military—with the assumption that the skills are transferable and thus increase the civilian opportunity wage. The authors included a switching cost that is imposed if the individual breaches their military contract by leaving before the end of the term. In calibrating their model, the authors estimated the distribution of the preference for military service in the youth population. They also analyzed the attractiveness of military IT occupations (compared with non-IT occupations), where IT occupations (by providing valuable, transferable training) provide a pathway to high-paying civilian jobs when the service member leaves the military.

Beginning in the 2000s, RAND estimated the DRM parameters rather than relying on a calibrated model. Mattock and Arkes (2007) adapted the Gotz-McCall approach to estimate a model that allowed analysis of incentive pay for Air Force officers, including a provision requiring a multiyear commitment to receive aviation bonus pay. Advances in computer hardware and software made estimation of the DRM more feasible. In Asch et al. (2008), RAND estimated a DRM of active and reserve retention using data provided by the Defense Manpower Data Center (DMDC) for the 10th QRMC. The model was estimated for each service branch (Army, Navy, Air Force, and Marines), and the parameter estimates were used to conduct policy simulations of alternative military retirement reform options using models for the enlisted force. Asch et al. (2008) provided a rigorous theoretical foundation for all the subsequent models we discuss in this report.

From 2008 to 2017, RAND conducted a series of studies that further developed the DRM and used the simulation capability to focus on retirement reform options, providing analyses that supported the Military Compensation and Retirement Modernization Commission and an Office of the Secretary of Defense working group that developed early versions of the Blended Retirement System (BRS). RAND's analysis of the retention and cost effects of the BRS are summarized in Asch, Mattock, and Hosek (2017) and included new model estimates for the U.S. Coast Guard, a service previously omitted in earlier DRM analyses. In addition, the DRM was used to consider alternative reserve retirement reform proposals (Mattock, Hosek, and Asch, 2012), and RAND extended the DRM for the Army reserve component to consider U.S. Army Reserve and Army National Guard retention separately (Asch, Mattock, and Hosek, 2019) and to simulate the retention effects of alternative retirement policies in the transition to the steady state for the reserve components (Mattock, Asch, and Hosek, 2014).

Other analyses focused on the retention and cost effects of special and incentive pay alternatives for selected populations. These include several studies on Air Force pilots and the effects of changes in aviator bonus and flight pay to improve retention in the face of increased major airline hiring (Mattock et al., 2016). These studies also includes analyses of reenlistment bonuses for career enlisted aviators in the Air Force (Tong, Mattock, and Asch, 2021), analyses of special pays for special operations (Asch et al., 2019), and for mental health care providers (Hosek et al., 2017). RAND also expanded the simulation capability to consider changes in the military pay table for those with more than 30 years of service (Asch et al., 2018; Asch, Hosek, Kavanagh, et al., 2016) and the role of separation pay in facilitating downsizing (Mattock, Hosek, and Asch, 2016).

RAND reestimated the DRM with four additional years of data (covering active-component retention and reserve-component participation for members who entered in 1990 and 1991 and were followed through 2016) and further expanded the DRM to include promotion in its analysis in support of the 13th QRMC (Asch, Mattock, and Tong, 2020). RAND researchers also incorporated metrics of performance in the simulation capability and the model was used to simulate the retention, cost, and performance implications of reforming the military pay table to a TIG rather than the existing time-in-service (TIS) format. To further describe the details of the DRM, we provide details of this model in the following section.

The DRM for Department of Defense Civil Service Personnel, Public School Teachers, and State Employees

We further built on the DRM capabilities that RAND developed for uniformed military personnel to focus on modeling the retention of civilian populations. We have now applied the DRM to Department of Defense (DoD) civil service personnel, public school teachers, and state employees.

Our first application of the DRM to DoD civil service personnel was in Asch, Mattock, and Hosek (2014), which modeled the retention behavior of a single cohort and used model estimates to assess the retention effects of pay freezes, changes in mandated retirement contributions, and other federal compensation changes. RAND further enhanced the model for DoD civilians by allowing taste to differ across multiple cohorts and by veteran status (Knapp, Asch, et al., 2016). This model was used to consider the downsizing effects of a voluntary separation incentive

(Asch, Hosek, Mattock, et al., 2016). Finally, Mattock et al. (2022) estimated a model of the retention of DoD cyber workers and considered using training as a retention lever.

Recently, many state governments have legislated reforms or are contemplating reforms to teachers' and state employees' retirement systems for new and future employees as a means of addressing large unfunded liabilities of their pension plans. Our first application of the DRM capability to state and local employees was to public school teachers for the Chicago Public Schools (Knapp, Brown, et al., 2016). The initial model was enhanced and broadened in its application to examine teachers and state employees in South Carolina (Knapp, Asch, and Mattock, 2021) and further applied to teachers across three states (Asch, Knapp, and Mattock, 2022). RAND also applied the model to examine ex ante prediction and ex post realization of a voluntary retirement incentive offer for teachers in Knapp et al. (2023). These studies led to a series of journal articles (Knapp et al., 2023; Hosek et al., 2023).

The 13th QRMC DRM Capability

Like the capability developed for the 10th QRMC, the DRM models estimated for the 13th QRMC allow service members to choose each year whether to continue in the active component or to leave; and once having left to either be a "pure" civilian or a civilian worker who also participates in the reserve component. Once having left, the individual can revisit the choice to participate in the reserve in each period. A key technical innovation in this variant of the DRM relative to earlier versions reviewed previously was to model the effect of promotion, which was key to being able to model a TIG pay table. In addition, we extended the model to consider how *performance*—as measured in terms of promotion speed relative to one's peers—might be affected by ability and effort supply. By *ability* we mean characteristics of individual service members that increase or decrease their promotion speed relative to their peers, including innate cognitive intelligence and other characteristics that lead to success (e.g., ability to work well in teams, work in a hierarchical organizational structure, resilience to changes such as frequent moves and new assignments). By *effort supply*, we refer to how hard and effectively members work in terms of achieving tasks that lead to faster promotion. We developed this capability so that we could run simulations and provide estimates of the effect of the TIG pay table on overall retention, retention of individuals with higher innate ability, and the average ability and level of effort exerted by individual service members. The results of our analysis are reported in Asch, Mattock, and Tong (2020). The text in this chapter draws heavily from this report and Asch et al. (2018).

The DRM Mathematical Structure

The DRM used for the 13th QRMC, similar to the DRM used for the 10th QRMC, is a model of the service member's decision—made each year—(1) to stay in or leave the active component, (2) for those who leave, to choose whether to participate in a reserve component, and

(3) if participating, whether to continue as a reservist. These decisions are structured as a dynamic program in which the individual seeks to choose the best career path, but the path is subject to uncertainty. The model was formulated in terms of parameters that are estimated with longitudinal data on retention in the active component and participation in the reserve component, and these data are then used to see how well the estimated model fits observed retention. We used the estimated parameters in policy simulations.

In the DRM, a set of parameters underlies the individual service member's retention decisions, and a goal of our analysis was to use individual-level data on active retention and reserve participation to estimate the parameters for both enlisted and officers for each service. We discuss the data we use in more detail later in this chapter, but, in short, we use the Defense Manpower Data Center's (DMDC's) Work Experience File (WEX) to track individual careers from 1990 to 2016.

Model Overview

In the behavioral model underlying the DRM, in each period, the individual can choose to continue on active duty, leave the military to hold a job as a civilian, or leave the military to join a reserve component and hold a job as a civilian. The individual bases their decision on which alternative has the maximum value. The model assumes that an individual begins their military career in an active component.

Individuals are assumed to differ in their preferences for serving in the military. Each individual is assumed to have given, unobserved preferences for active and reserve service, and these preferences do not change. The individual member, officer or enlisted, has knowledge of military pay and retirement benefits, as well as civilian compensation. In each period, there are random shocks associated with each of the alternatives, and the shocks affect the values of the alternatives. As shown next, the model explicitly accounts for individual preferences and military and civilian compensation, and, in this context, shocks represent current-period conditions that affect the value of being on active duty, being in the selected reserve while also being a civilian worker (or *reserve*, for short), or being a civilian worker and not in the reserve (*civilian*, for short). Examples of what might contribute to a shock are a good assignment; a dangerous mission; an excellent leader; inadequate training or equipment for the tasks at hand; a strong or weak civilian job market; an opportunity for on-the-job training or promotion; the choice of location; a change in marital status, dependency status, or health status; the prospect of deployment or deployment itself; or a change in school tuition rates. These factors might affect the relative payoff of being in an active component, being in a reserve component, or being a civilian. The individual is assumed to know the distributions that generate the shocks, as well as the shock realizations in the current period but not in future periods.

Depending on the alternative chosen, the individual receives the pay associated with serving in an active component, working as a civilian, or serving in a reserve component while also working as a civilian. In addition, the individual receives the intrinsic monetary equivalent of the preference for serving in an active component or serving in a reserve component. These values are assumed to be relative to that of working as a civilian, which is set at zero.

In considering each alternative, the individual considers their current state and type. *State* is defined by whether the member is active, reserve, or civilian, and by the individual's active year(s) of service (YOS), reserve YOS, total years since first joining the military, pay grade, and random shocks. *Type* refers to the level of the individual's preferences for active and reserve service. The individual recognizes that today's choice affects military and civilian compensation in future periods. Although the individual does not know when future military promotions will occur, he or she does know the promotion policy and can form an expectation of military pay in future periods. Furthermore, the individual does not know what the realizations of the random shocks will be in future periods. The expected value of the shock in each state is zero. Depending on the values of the shocks in a future period, any of the alternatives—active, reserve, or civilian—might be the best at the time. Once a future period has been reached and the shocks are realized, the individual can reoptimize (i.e., choose the alternative with the maximum value at that time). The possibility of reoptimizing is a key feature of dynamic programming models that distinguishes them from other dynamic models. In the current period with future realizations unknown, the best the individual can do is to estimate the expected value of the best choice in the next period; i.e., the expected value of the maximum. Logically, this will also be true in the next period, the one after it, and so forth—so the model is forward-looking and rationally handles future uncertainty. Moreover, the model presumes that the individual can reoptimize in each future period, depending on the state and shocks in that period. Thus, today's decision accounts for the possibility of future career changes and assumes that future decisions will also be optimizing.

Mathematical Formulation

We denote the value of staying in the active component at time *t* as

$$V^S(k_t) = V^A(k_t) + \epsilon_t^A,$$

where $\boldsymbol{k_t}$ is defined as

$$k_t = k_t(ay_t, ry_t, t, g_t)$$

or the vector of number of active years (ay_t) at time *t*, the number of reserve years (ry_t) total years since initial enlistment or accession, and grade (g_t). $V^A(k_t)$ is the non-stochastic value of the active alternative, and ϵ_t^A is a random shock.

The value of leaving at time *t* is

$$V^L(k_t) = \max[V^R(k_t) + \omega_t^R, V^C(k_t) + \omega_t^C] + \epsilon_t^L,$$

where the member can choose between reserve (*R*) and civilian (*C*). *Civilian* means working at a nonmilitary job, and *reserve* means participating in a reserve component and working at a nonmilitary job. The value of reserve is given by $V^R(k_t) + \omega_t^R$, where k_t is defined previously, while the value of civilian is given by $V^C(k_t) + \omega_t^C$. We model the reserve or civilian choice as a nest and assume that the stochastic terms follow an extreme value type I distribution, which leads to a nested logit specification in the estimation phase of this structural model.[7] The within-nest shocks to the reserve or civilian choice are given by ω_t^R and ω_t^C, while the nest-level shock is given by ϵ_t^L.

We allow a common shock for the reserve and civilian nest, ϵ_t^L (because an individual in the reserves also holds a civilian job) and shock terms specific to the reserve and civilian states, ω_t^R and ω_t^C. The individual is assumed to know the distributions that generate the shocks and the shock realizations in the current period but not in future periods. The distributions are assumed to be constant over time, and the shocks are uncorrelated within and between periods. Once a future year is reached, and the shocks are realized, the individual can reoptimize by choosing the alternative with the maximum value at that time. But in the current period, the future realizations are not known, so the individual assesses the future period by taking the expected value of the maximum (i.e., the expected value of civilian conditional on it being superior to that of reserve times the probability of that occurring, plus the expected value of reserve conditional on it being superior to civilian times the probability of that occurring). For instance, depending on the shocks and the compensation, there is some chance that $V^S(k_t)$ will be greater than $V^L(k_t)$, in which case $V^S(k_t)$ would be the maximum (and vice versa), and the individual makes an assessment of the expected value of the maximum, $Emax\left(V^S(k_t), V^L(k_t)\right)$.

The extreme value distribution, denoted *EV*, has location parameter *a* and scale parameter *b*; the mean is $a + b\phi$, and the variance is $\pi^2 b^2/6$ where ϕ is Euler's gamma (~0.577). As we derived in past studies (Asch et al., 2008; Mattock et al., 2016), this implies

$$\epsilon_t^{Leave} \sim EV\left[-\phi\sqrt{\lambda^2+\tau^2}, \sqrt{\lambda^2+\tau^2}\right]$$

$$\omega_t^R \sim EV[-\phi\lambda, \lambda]$$

$$\omega_t^C \sim EV[-\phi\lambda, \lambda]$$

$$\epsilon_t^L \sim EV[-\phi\tau, \tau]$$

[7] See Train, 2009, for a discussion of the logit and nested logit specifications.

where λ is the common scale parameter of the distributions of ω_t^R and ω_t^C, and τ is the scale parameter of the distribution of ϵ_t^L. In the nested structure of the model, leavers face a common shock for the "leave" nest, ϵ_t^L, as well as shocks for the reserve and civilian alternatives within the nest, ω_t^R and ω_t^C, which, all together, produce a leave shock distributed as extreme value type I, with location parameter

$$-\phi\sqrt{\lambda^2 + \tau^2} \text{ and scale parameter } \sqrt{\lambda^2 + \tau^2}.$$

The logit model requires that the scale parameters of the leave and stay shocks be equal, so we parameterize the model such that the stay scale parameter, which we denote k, has the same value as the leave scale parameter,

$$k = \sqrt{\lambda^2 + \tau^2}\,.$$

The values of the alternatives $V^A(k_t), V^R(k_t)$, and $V^C(k_t)$ depend on the current pay for serving in an active component or working as a civilian, $W^A(k_t)$ or $W^C(k_t)$. The service member's active pay is based on total YOS, ay_t, as well as their grade, g_t.

Our model includes promotion. The model assumes that the timing and probability of promotion at each grade is the same across all officers and is the same across all enlisted. Variation in the timing and probability of promotion for an individual service member is captured by the shock term. Promotion to a given grade occurs at a given number of YOS, but the probability of promotion differs by grade. Also, the probability of promotion is assumed to be invariant to policy change. Not being promoted decreases the value of continuing in the military and operates to decrease retention. Officers or enlisted service members that are promoted can look ahead to future promotion gates, and their value of staying is higher than that of service members that are not promoted.

The possibility of reoptimizing in future periods distinguishes dynamic programming models from other dynamic models. Reoptimization means that the individual can choose the best alternative in a period when its conditions have been realized (i.e., when the shocks are known). As mentioned, future realizations are unknown in the current period, and the best the individual can do is to estimate the expected value of the best choice in the next period (i.e., the expected value of the maximum). This will also be true in the following period, the one after it, and so forth—so the model is forward-looking and rationally handles future uncertainty. Thus, today's decision accounts for the possibility of future changes of state and assumes that future decisions will also be optimizing.

To be more specific, in developing a mathematical expression for the value of the value function $V^A(k_t)$, the DRM considers all possible future pathways, recognizing that each pathway depends on each probability of promotion to the next grade and the YOS when promotion can occur. Thus, the DRM views an officer or enlisted service member with a particular $\boldsymbol{k_t}$ as reasoning forward to identify the full set of possible future paths of staying or leaving. Then, the service member reasons backward starting from the final stay or leave decision year, called year T.

For each possible k_T, the model assumes that the service member considers whether to stay or leave. From the perspective of an earlier year *t*, the member's current year, there is no reason to commit to a decision at *T*, and, in fact, it would be short-sighted to do so because the member would not be able to base the decision on information that will be revealed when *T* arrives (i.e., when the shocks in *T* are realized). Instead, the service member at *t* develops a decision rule about whether to stay or leave at *T*, and that rule is to stay if the value of doing so is higher than the value of leaving, otherwise to leave. The service member can—in the context of the model—compute the expected value of making that optimal decision. Reasoning backward, this expression enters into the expression for the optimal stay-versus-leave decision at *t – 1* and so on back, year by year, to *t*.

At *t*, the value of continuing in the military for a member at grade g (now shown as a superscript) is

$$V^S(k_t) = V^A(k_t) + \epsilon_t^A = \gamma^A + W_t^{Ag} + \beta\ EMax(V^A(k_{t+1}) + \epsilon_{t+1}^A, V^L(k_{t+1}) + \epsilon_{t+1}^L) + \epsilon_t^A,$$

where $\boldsymbol{\gamma^A}$ is the individual's taste for active duty, $\mathrm{W}_{\mathrm{t}}^{\mathrm{Ag}}$ is active-duty pay, β is the personal discount factor, the ϵ terms are random shocks, and the operator $EMax$ finds the expected value of the maximum of the terms

$$V^A(k_{t+1}) + \epsilon_{t+1}^A \text{ and } V^L(k_{t+1}) + \epsilon_{t+1}^L).$$

Each of these terms has a nonrandom term and a random term.

Consider shocks that have an extreme value distribution with a mode of zero and a scale of kappa: $\epsilon \sim EV[0, \kappa]$. With an extreme value shock, the quantity $a + \varepsilon$ is distributed as $EV[a, \kappa]$. The mean of this distribution equals the scale factor times Euler's gamma plus the mode: $\phi\kappa + a$ where $\phi \approx 0.577$. If the mode is transformed by subtracting $\phi\kappa$, then $a - \phi\kappa + \varepsilon$ is distributed as $EV[a - \phi\kappa, \kappa]$ with a mean of a. (This transformation is equivalent to assuming that the shocks are distributed as $\varepsilon \sim EV[-\phi\kappa, \kappa]$; that is, that the shocks have mean zero and scale kappa.) Also, if two quantities $\mathrm{V^m}$ and V^n have the form $a + \varepsilon$ and we subtract $\phi\kappa$ from each, their maximum has an extreme value distribution, namely,

$$Max(V^m, V^n) \sim EV[\kappa \ln(e^{V^{\frac{m}{\kappa}}} + e^{V^{\frac{n}{\kappa}}}) - \phi\kappa, \kappa].$$

The mean of this distribution is $\kappa \ln(e^{V^m/\kappa} + e^{V^n/\kappa})$. The mean is the expected value of the maximum. This result implies that

$$EMax(V^A(k_{t+1}) + \epsilon^A_{t+1}, V^L(k_{t+1}) + \epsilon^L_{t+1}) = \kappa \ln\left(e^{V^A(k_{t+1})/\kappa} + e^{V^L(k_{t+1})/\kappa}\right).$$

To introduce promotion, we replace V^A with its expected value, where p is the probability of promotion:

$$V^A = p^{g+1}_{t+1} V^{A(g+1)} + (1\text{-}p^{g+1}_{t+1}) V^{Ag}.$$

In those YOS where no promotion occurs (that is, in those YOS when promotion is not possible), the probability of promotion is zero. In years where promotion might occur (i.e., in those YOS when promotion is possible), the probability of promotion is assigned a value relevant for the grade. In general, not all eligible individuals get promoted, particularly in the senior grades; as a result, the probability of promotion is typically strictly less than one. Note that in this model, the probability of promotion for an individual is solely a function of their YOS and does not depend on TIG or the inventory of service members in a grade; in simulations using estimated model parameters, we relaxed this assumption and allowed the timing of an individual's promotion to vary based on their (unobserved) ability or effort.

For simplicity, we assume that civilian pay only depends on YOS (or years since initial active enlistment or accession, if the individual has left active service). If the member is a reservist, he earns the civilian wage plus reserve pay, $W^C(k_t) + W^R(k_t)$. As with active pay, reserve pay depends on total years, including prior active years and reserve years.

The tastes for active and reserve duty, γ^A and γ^R, represent the individual's perceived net advantage of holding an active or reserve position, relative to the civilian state. Other things equal, more taste for active or reserve service increases retention. The tastes are assumed to be constant over time but vary across individuals. Also, tastes for active and reserve service are not observed but are assumed to follow a bivariate normal distribution among active component entrants.

The non-stochastic (in the current period) values of the reserve choice and civilian choice can be written as

$$V^R(k_t) = \gamma_r + W^C(k_t) + W^R(k_t) + \beta E\left[max[V^R(k_{t+1}) + \omega_r, V^C(k_{t+1}) + \omega_c]\right]$$

$$V^C(k_t) = W^C(k_t) + R(k_t) + \beta E\left[max[V^R(k_{t+1}) + \omega_r, V^C(k_{t+1}) + \omega_c]\right],$$

Where $R(k_t)$ in the civilian equation is the value of any active or reserve military retirement benefit for which the individual is eligible. NDAA 2016 created a new military retirement system, known as the BRS. Because our data cover retention decisions of personnel under the legacy retirement system, we use the formula for the legacy system for the purpose of our analysis given by

$$R(k_t) = 2.5\% \times ay_t \times W^A(k_t)$$

for the active retirement system where, in this formula, $W^A(k_t)$ is the highest three years of basic pay and is computed based on total active years, ay_t. For a member with 30 YOS, the multiplier $2.5\% \times ay_t$ is 75 percent, while it is 100 percent for a member with 40 YOS. (After 2007, the 75 percent cap on the multiplier was lifted, thereby permitting additional YOS beyond 30 to contribute to retired pay.)

The model has two switching costs, which enter the relevant value function as additive terms. *Switching cost* refers to a latent cost reflecting the presence of constraints or barriers affecting the movement from particular states and periods to other states, relative to the movement that would otherwise have been expected from the expressions shown previously for the values of staying and of leaving. Switching costs are not actually paid by the individual but, as estimated in the model, are a monetary representation of the constraints or barriers affecting the transition from one state to another at a given time. Furthermore, a switching cost can be either negative or positive. A negative value implies a loss to the individual when changing from the current status to an alternative status, while a positive value implies a gain or incentive for the change. The first switching cost is a cost of leaving the active component before an officer or an enlisted service member's ADSO is completed, or an enlisted service member's initial term of service is completed. This switching cost enters the value functions $V^R(k_t)$ and $V^C(k_t)$. The estimates, shown later, indicate that the switching cost has a negative value for all services, possibly reflecting the perceived cost of breaching the service contract. The second switching cost is a cost of switching into the reserve from the civilian state and enters the value function $V^R(k_t)$. This cost could represent difficulty in finding a reserve position in a desired geographic location or an adverse impact on one's civilian job, e.g., from not being available to work on certain weekends or for two weeks in the summer or being subject to reserve call-up. This impact is negative across all services.

Estimation Methodology

To estimate the DRM, we use the mathematical structure of the model together with assumptions on the distribution of tastes across service members and shock distributions. This allows us to derive expressions for the transition probabilities, given one's state, which are then used to compose an expression for the likelihood of each individual's years of active retention and reserve participation. Importantly, each transition probability is itself a function of the underlying parameters of the DRM. These are the parameters of the taste distribution, the shock distributions, the switching costs, and the discount factor. The estimation routine finds parameter values that maximize the likelihood.

The transition probability is the probability in a given period of choosing a particular alternative (i.e., active, reserve or civilian), given one's state. Because we assume that the model is first-order Markov,[8] that the shocks have extreme value distributions, and that the shocks are uncorrelated from year to year, we can derive closed-form expressions for each transition probability. For example, as Train (2009) shows, the probability of choosing to stay active at time *t*, given that the service member is already in the active component, is given by the logistic form

$$\Pr(V^S > V^L) = \frac{e^{\frac{V^A}{\kappa}}}{e^{\frac{V^A}{\kappa}} + \left(e^{\frac{V^R}{\lambda}} + e^{\frac{V^C}{\lambda}}\right)^{\frac{\lambda}{\kappa}}}.$$

We omit the state vector $\boldsymbol{k_t}$ in each expression for clarity. We can also obtain expressions for the probability of leaving the active component and, having left, the probabilities of entering (or staying in) the reserve component in each subsequent year.

The transition probabilities in different periods are independent and can be multiplied together to obtain the probability of any given individual's career profile of active, reserve, and civilian states that we observe in the data. Multiplying the career profile probabilities together gives an expression for the sample likelihood that we use to estimate the model parameters for using maximum likelihood methods.[9] Optimization is done using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm, a standard hill-climbing method. We compute standard errors of the estimates by using numerical differentiation of the likelihood function and taking the square root of the absolute value of the diagonal of the inverse of the Hessian matrix. To judge goodness of fit, we use parameter estimates to simulate retention profiles for synthetic individuals (characterized by tastes drawn from the taste distribution) who are subject to shocks (drawn from the shock distributions), then aggregate the individual profiles to obtain a force-level retention curve and compare it with the retention curve computed from actual data.

We estimate the following model parameters:

- the mean and standard deviation (SD) of tastes for active and reserve service relative to civilian opportunities, as well as their correlation (e.g., μ_a, μ_r, σ_a, σ_r, and ρ)
- a common scale parameter of the distributions of ω_t^R and ω_t^C, λ, and a scale parameter of the distribution of ϵ_t^L, or τ.
- a switching cost incurred if the individual leaves active duty before completing the ADSO or first term
- a switching cost incurred if the individual moves from civilian to reserve

[8] A first-order Markov assumption is that the probability of an event at time $t + 1$ only depends on the state at time t.

[9] This approach bears some resemblance to a (highly restricted) mixed logit model.

In past DRM analyses, we also estimated a personal discount factor (see Asch, Hosek, and Mattock, 2014). We fixed the personal discount factor in this study because we found the model fits were better and parameter estimates were more reasonable relative to our expectations using past research.[10] The personal discount factor might not have been identified as strongly as in our prior models because this study made military compensation a function of probabilistic promotion as opposed to an assumed wage trajectory.[11] We set the personal discount factor for officers equal to 0.94 and for enlisted to 0.88, which are the values we have typically estimated for officers and enlisted in earlier work (e.g., Asch, Hosek, and Mattock, 2014; Mattock, Hosek, and Asch, 2012).

Once we have parameter estimates for a well-fitting model, we can use the logic of the model and the estimated parameters to simulate the active component cumulative probability of retention to each YOS in the steady state for a given policy environment, such as a change to the retired pay cap. By *steady state*, we mean when all members have spent their entire careers under the policy environment being considered. The simulation output includes a graph of the active component retention profile for officers and enlisted personnel by YOS. We can also produce graphs of reserve component participation and provide computations of costs, although we do not do so here. We show model fit by simulating the steady-state retention profile in the current policy environment and comparing it with the retention profile observed in the data.

Data

DMDC's WEX data contain person-specific longitudinal records of active and reserve service. WEX data began with service members in the active or reserve component on or after September 30, 1990. Our analysis files include active component entrants in 1990 and 1991, who are followed through 2016, providing up to 26 years of data for the 1990 cohort and up to 25 years of data for the 1991 cohort. In constructing the officer samples, we exclude medical personnel and members of the legal and chaplain corps because their career patterns differ markedly from those of the rest of the officer corps, suggesting that analysis of retention for these personnel needs to be conducted separately. We also excluded officers with prior enlisted service.

Another key source of data is information on civilian and military pay. For civilian pay opportunities for enlisted personnel, we used the 2007 median wage for full-time male workers

---

[10] The personal discount factor equals 1/(1+r) where *r* is the personal discount rate. For example, a personal discount factor of 0.88 corresponds to a discount rate *r* of 13.6 percent.

[11] Our inability to estimate a personal discount rate in this study could be because we examined different study period than previous studies or because we introduced promotion into the model. We were unable to explore why introducing promotion would affect identification of the discount rate, but future research should do so.

with associate's degrees.[12] For officers, we use the 2007 80th percentile of basic pay for full-time male workers with a master's degree in management occupations for civilian pay. The data are from the Census Bureau. *Civilian work experience* is defined as the sum of active years, reserve years, and civilian years since age 22, but here, pay does not vary by other factors, such as years since leaving active duty. We used 2007 military pay tables. Military pay increases are typically across the board; the structure of pay by grade and YOS remain the same.[13] Therefore, we did not expect our results to be sensitive to the choice of year. Annual military pay for active members is represented by regular military compensation (RMC) for fiscal year 2007, equal to the sum of basic pay, basic allowance for subsistence, basic allowance for housing, and the federal tax saved because the allowances are not taxed. Data on RMC and basic pay by grade and YOS were from the *Selected Military Compensation Tables*, also known as the Green Book (Office of the Under Secretary of Defense for Personnel and Readiness, Directorate of Compensation, 1980–2017). Reserve component members are paid differently from active component members, although the same pay tables are used. The method for computing reserve component annual pay is described in Asch, Mattock, and Hosek (2017). Military retirement benefits are related to the basic pay table, and we used the basic pay tables for 2007 for this computation.

We also required data on enlisted and officer promotion rates and promotion timing to each grade. Officer promotion rates were drawn from those used in Asch and Warner (1994a), while promotion rates for enlisted and promotion timing data for both officers and enlisted were based on computations of average time in service at promotion by grade and service for fiscal year 1993 to 2008 from DMDC. We chose these years because we sought promotion times that would be relevant to the 1990–1991 accession. In doing so, we assumed a static promotion schedule, and the individual expectations of the probability of promotion matched the realized probabilities of promotions (i.e., rational expectations). The advantage of this approach it that it is straightforward to implement and does not impose much computational burden. An alternative approach would have been to use the realized probabilities of promotion for each year observed. However, this would have required constructing a model whereby individuals would have priors over their probability of promotion and that would allow for individuals to update their estimated probability of promotion over time. This approach is desirable in that it would have higher fidelity to the observed changes in promotion probabilities but potentially imposes a significant

---

[12] Alternatively, we could have used time-dependent pay trajectories instead of using a single representative year. Using a single year has the advantage of reducing the computational complexity of the overall model and, with it, the run time associated with estimation. However, this approach is not without its weaknesses. For example, the continuous real decline in less-than-bachelor's degree salaries since 1900 could provide an increasing retention effect that would be missed by using a static wage value.

[13] An exception was the structural adjustment to the basic pay table in fiscal year 2000 that gave larger increases to midcareer personnel who had reached their pay grades relatively quickly (after fewer YOS). A second exception was the expansion of basic allowance for housing, which increased in real value from fiscal year 2000 to fiscal year 2005.

computational burden. Another, less computationally burdensome approach would be to have a dynamic promotion schedule and give members perfect foresight (i.e., rational expectations); this approach might hold the most promise for future research.

In the next chapter, we present the model estimates and model fits for the 13th QRMC DRM and discuss the technical details of the simulation capability.

Simplifying Assumptions and Efficient Computing

We made several simplifying assumptions in an effort to make the model practical to estimate. In addition, we took steps to compute the stochastic dynamic program in the most efficient manner.

Among the simplifying assumptions in this model is a simplified model of promotion, where the probability weight is placed on 1 YOS. That is, there is a single year in which an individual has a hazard of being promoted. A more realistic model of officers would, for example, include promotion below the zone, in the zone, and above the zone. As it is, the model we described previously models only in the zone promotion for estimation.

We use a single pay table for estimating the model, the 2007 pay table. Therefore, we would miss any phenomena related to structural differences from the 2007 pay table (we have relaxed this in some of our work related to Air Force pilot retention).

Individual tastes are assumed to be constant. A more realistic model might account for job characteristics and allow taste to vary over a career. For example, Air Force pilots might be modeled as having a taste for flying, and thus their taste in a given YOS would be a function of their flying hours.

Similarly, in this model the taste distribution is held constant within the individual populations for which we estimate the model. Alternatively, we could model both the mean and SDs of the taste distribution as being functions of, for example, demographic characteristics that might vary over a population.

We do not allow for serial correlation in this model. All environmental shocks are assumed to be independently and identically distributed.

We also have a linear functional form for utility. Thus, we are implicitly assuming risk neutrality on the part of the individual. Alternatively, we could use nonlinear utility functions, such as the constant absolute risk aversion or constant relative risk aversion, but only at considerably more computational complexity.

We assume that every individual in a population has the same discount rate. Alternatively, we could assume that an individual's discount rate is drawn from some distribution. Similarly, we assume that an individual's discount rate does not change as they age.

We do not model an individual's savings decisions. This is relevant to modeling the choice of the level an individual contributes to a Thrift Savings Plan or to other retirement-related decisions.

These simplifications were made to make the model computable. As software and computing hardware improve over time, we anticipate relaxing some of these assumptions.

Computing the Stochastic Dynamic Program Efficiently

To make computation more efficient, the RAND model makes extensive use of caching (that is, storing the results of functions that are expensive to evaluate) to avoid repeatedly solving identical subproblems. The identical subproblems are an artifact of using recursion to compute the dynamic program; such programs can often end up needing to solve identical subproblems that appear at different points of the branching tree of possibilities. So, for efficiency, we cache all functions that are computationally expensive to evaluate.

In addition, we exploit the fact that any one iteration of the likelihood function consists of many subproblems that can be computed in parallel. That is, for a given value of the parameters and each point in the taste distribution, each individual's stochastic dynamic program corresponding to their career history can be solved in parallel. In practice, however, there appears to be little benefit to expanding beyond the eight to 12 processors on existing computing architectures. This is a case where a quantum computer could help. With a quantum computer of sufficient capacity, we could fully exploit the parallelism of our problem.

Finally, we use Halton sequences that generate antithetic pairs (pairs symmetric around a mean) when choosing support points for the taste distribution, which aids in efficiently averaging over the population of interest (i.e., integrating out heterogeneity) when calculating probabilities. To put this in more concrete terms, we use a Halton sequence to generate two uncorrelated standard normal distributions and use the current value of the mean and SD parameters for active component and reserve component tastes, as well as the taste correlation, to generate a bivariate normal distribution with the desired properties. The way we do this is documented in Asch et al. (2008). Briefly, following Train (2009), we take two independent draws from a standard normal distribution and use a Cholesky decomposition to transform them into random variables that are jointly normally distributed. This results in a set of sample points from a bivariate normal distribution that smoothly varies with the changing values for the mean, SD, and correlation parameters during successive evaluations of the likelihood function while trying to find the parameter values that maximize the likelihood.

Optimization Algorithms

Typically, researchers use a hill-climbing algorithm to find the values of the parameters of a maximum likelihood model.[14] Sophisticated hill-climbing algorithms, such as BFGS, generally perform well in estimating the DRM. However, there are some variants of the DRM in which hill-climbing algorithms do not perform well, such as when there are thresholds governing

---

[14] This discussion of optimization algorithms draws from Mattock and Arkes (2007).

member behavior. However, we have found some approaches to this problem that worked well, as we describe in the following paragraphs. Readers who are not interested in these details might want to skip the remainder of this subsection and continue with the next.

Numerical integration routines sample functions over a finite number of points and estimate the value of the integral using this sample. If the interval of interest is small, as can sometimes be the case with the DRM, then the numerical integration routine might miss the interval. This means that small parameter changes might result in large changes in the calculated number for the log likelihood. These large changes are because of the sample points of the numerical integration routine hitting or missing the possibly small interval. These large changes spell trouble for any hill-climbing algorithm that relies on numerical estimates of the gradient of the likelihood function. Even hill-climbing algorithms that do not rely on estimating the gradient, such as the Nelder-Mead multidimensional simplex algorithm, can run into trouble, becoming trapped in a local maximum.

This problem can be addressed in two ways. One is to attempt to increase the precision of the numerical integration routine. This is not always an effective strategy for the DRM because simply increasing the number of sample points does not obviate the interaction of the numerical integration routine with any thresholds that might govern member behavior in this model. The second way to address this problem is to use a search algorithm that is well suited to problems in which there are many local maxima. Simulated annealing (also known as the Metropolis-Hastings algorithm) is one such approach. The simulated annealing algorithm randomly jumps to a point and compares the value of the objective function at the new point with the old best value. With some probability (governed by a Boltzman distribution, for example), it accepts the new point, even if the new point has a lower value for the objective function than the current best point; in this way it avoids getting stuck at a local maximum. As the algorithm progresses, it chooses random points closer and closer to the current best point; it does this according to a "cooling schedule" governed by an optimizer parameter called "temperature." This algorithm bears some resemblance to models of materials undergoing an annealing process, hence the name *simulated annealing*.

When the simulated-annealing algorithm finds a maxima of the likelihood function, a hill-climbing algorithm (either Nelder-Mead [also known as the *downhill simplex method*, which is a derivative-free method] or BFGS [which uses numerical estimates of the gradient]) can be used to further refine the parameter estimates.

Another approach is to use an expectation-maximization (EM) algorithm. This approach avoids the problems caused by the approximate nature of numerical integration noted previously. The EM algorithm consists of two steps (Dempster, Laird, and Rubin, 1977). In the first step, a likelihood function that depends on the unobservable data (in this case, the individual values of gamma) and assumed values for the rest of the parameters is used to estimate the unobservable data. The second step consists of using the estimated values of the unobservable data in a likelihood function for the complete data model, which is then used to generate estimates of the model parameters. The model parameters generated in the second step are then used as the assumed values in the first step, and the two steps are repeated until the algorithm converges on a set of parameter values.

More formally, we can write the likelihood function for each individual so that it depends on the value of gamma, the individual data x_i, and the model parameters $\theta = (\alpha, \delta, \sigma, \beta)$. For example, if an officer stays for s periods and then leaves, the likelihood would be

$$L_i(\gamma_i |\, x_i, \theta) = \prod_{\tau=t}^{t+s} \Pr\,[Stay_\tau \mid \gamma_i, x_i, \theta] \Pr[Leave_{t+s+1} |\ \gamma_i, x_i, \theta] g(\gamma|\theta).$$

Step 1. Find the value of γ_i that maximizes the individual likelihood, assuming that the parameter vector is $\hat{\theta}^0$:

$$\hat{\gamma}_i^1 = \underset{\gamma_i}{arg\,max}\, L_i(\gamma_i \mid x_i, \hat{\theta}^0).$$

Step 2. Find the value of θ that maximizes the individual likelihood given the values for the individual gammas estimated in step 1 and the data:

$$\hat{\theta}^1 = \underset{\theta}{arg\,max}\, L(\hat{\gamma}^1|x, \hat{\theta}^1).$$

Then iterate over the two steps, substituting the n^{th} estimate of θ into the $(n+1)^{\text{th}}$ iteration of step 1 to estimate

$$\hat{\gamma}_i^{\text{n}+1} = \underset{\gamma_i}{arg\,max}\, L_i(\gamma_i \mid x_i, \hat{\theta}^{\text{n}})$$

and then use the new estimates of γ_i to generate new estimates of

$$\hat{\theta}^{\text{n}+1} = \underset{\theta}{arg\,max}\, L(\hat{\gamma}^{\text{n}+1}|x, \hat{\theta}^{\text{n}})$$

until the algorithm converges on a value of $\hat{\theta}$.

This algorithm generally converges to the maximum likelihood solution. However, convergence is not guaranteed.

This algorithm has been generalized to the case where in each step values of the unobservable data and the parameters of the full data model are chosen that merely improve the likelihood function (Neal and Hinton, 1998). If each step assures some improvement, then the algorithm can be shown to have similar convergence properties to an algorithm using the optimal values of the unobserved data and the model parameters. This is the Generalized Expectation Maximization (GEM) algorithm.

The GEM algorithm is very useful for estimating models with unobserved individual heterogeneity. As can be seen from the aforementioned description, the algorithm avoids the need to integrate out heterogeneity. This leads to a significant reduction in the computer time needed to estimate models with unobserved individual heterogeneity and facilitates exploration of additional sources of heterogeneity (e.g., in the individual discount rate for future earnings).

Finally, we can also use a grid search to help narrow down parameter values of interest. However, selecting the appropriate ranges and intervals within each range so that the number of sample points is computationally feasible can be challenging. A grid search can be used to help find starting values for some of the optimization algorithms already mentioned, or the values computed for the grid can be used to construct an interpolating function that is easier to compute than the likelihood function itself. In this case, the interpolating function is used with another optimization algorithm to find parameter estimates.

In general, we find that using Nelder-Mead to get initial DRM parameter estimates and then refining the parameter estimates using the BFGS algorithm works well. For those DRMs where the Nelder-Mead/BFGS combination performs poorly, we can use alternative methods such as grid search, simulated annealing, and the EM or GEM algorithms.

Chapter 4. Model Estimates and Fit and Simulation

This chapter presents the model estimates and illustrates how the model predictions fit the observed data, a measure of goodness of fit.[15] We then illustrate the technical details of the simulation capability, focusing on the TIG pay table reform. In this application, the simulation capability required that we extend the capability to include metrics of performance. We discuss these extensions.

Model Estimates

Tables 4.1 and 4.2 show the estimated parameters and standard errors for the retention model of officers. To make the numerical optimization easier, we did not estimate most of the parameters directly but instead estimated the logarithm of the absolute value of each parameter, except for the taste correlation, for which we estimated the inverse hyperbolic tangent of the parameter.[16] All of the parameters are statistically significant in the Navy and Air Force models, and all but the between-nest scale parameter τ are significant in the Army and Marine Corps models. To recover the parameter estimates, we transformed the estimates. Table 4.3 shows the transformed parameter estimates for each service. The estimates are denominated in thousands of 2007 dollars, except for the assumed discount rate and the taste correlation.

---

[15] In this example, we assess goodness-of-fit by comparing predicted retention with observed retention. However, we discuss and present other approaches to testing external validity of the DRM capability in other examples; specifically, our analysis of retention of Chicago public school teachers (Knapp et al., 2018) and the retention of state public school teachers in South Carolina, Tennessee, and Pennsylvania (Asch, Knapp, and Mattock, 2022).

[16] Using the logarithm helps to keep all the parameters on roughly the same scale, which can be an important consideration with numerical optimization algorithms. In addition, using logarithms means that an optimization approach that relies on numerical derivatives, such as BFGS, uses deltas that are the same proportion across all parameters, which can help to speed convergence.

Table 4.1. Parameter Estimates and Standard Errors: Army and Navy Officers

	Army		Navy	
Parameter	Estimate	Standard Error	Estimate	Standard Error
Log(Scale Parameter, Nest = τ)	–1.36	33.83	5.20	0.04
Log(Scale Parameter, Alternatives within Nest = λ)	4.69	0.03	3.40	0.06
Log(–1*Mean Active Taste = μ_a)	3.19	0.04	3.00	0.05
Log(–1*Mean Reserve Taste = μ_r)	5.63	0.05	4.01	0.05
Log(SD Active Taste = σ_a)	3.76	0.04	3.87	0.05
Log(SD Reserve Taste = σ_r)	5.26	0.05	3.88	0.06
Atanh(Taste Correlation = ρ)	0.67	0.02	0.94	0.01
Log(–1*Switch Cost: Leave Active < ADSO)	4.81	0.03	5.20	0.04
Log(–1*Switch Cost: Switch from Civilian to Reserve)	6.05	0.03	4.90	0.05
Personal Discount Factor β (Assumed)	0.94	N/A	0.94	N/A
–1*Log Likelihood	24,141		32,139	
N	5,318		6,445	

SOURCE: Author calculations using DMDC WEX data on cohorts of personnel entering active duty as officers in 1990–1991.

NOTE: N/A = not applicable. The scale parameter κ governs the shocks to the value functions for staying and for the reserve-versus-civilian nest and equals $\sqrt{\lambda^2 + \tau^2}$. The means and SDs of tastes for active and reserve service relative to civilian opportunities are estimated, as are the costs associated with leaving active duty before completing ADSO and switching from civilian status to participating in the reserves. The personal discount factor was assumed to be 0.94 in these models.

Table 4.2. Parameter Estimates and Standard Errors: Air Force and Marine Officers

	Air Force		Marine Corps	
Parameter	Estimate	Standard Error	Estimate	Standard Error
Log(Scale Parameter, Nest = τ)	4.79	0.09	1.02	3.49
Log(Scale Parameter, Alternatives within Nest = λ)	3.96	0.35	4.37	0.05
Log(–1*Mean Active Taste = μ_a)	2.92	0.07	2.65	0.07
Log(–1*Mean Reserve Taste = μ_r)	6.20	0.53	4.93	0.08
Log(SD Active Taste = σ_a)	3.24	0.09	3.16	0.07
Log(SD Reserve Taste = σ_r)	5.78	0.55	4.51	0.08
Atanh(Taste Correlation = ρ)	0.45	0.01	0.56	0.04
Log(–1*Switch Cost: Leave Active < ADSO)	4.73	0.06	4.89	0.05
Log(–1*Switch Cost: Switch from Civilian to Reserve)	5.52	0.34	5.63	0.05
Personal Discount Factor β (Assumed)	0.94	N/A	0.94	N/A
–1*Log Likelihood	8,871		9,086	
N	2,339		1,757	

SOURCE: Author calculations using DMDC WEX data on cohorts of personnel entering active duty as officers in 1990–1991.

NOTE: The scale parameter κ governs the shocks to the value functions for staying and for the reserve-versus-civilian nest and equals $\sqrt{\lambda^2 + \tau^2}$. The means and SDs of tastes for active and reserve service relative to civilian opportunities are estimated, as are the costs associated with leaving active duty before completing ADSO and switching from civilian status to participating in the reserves. The personal discount factor was assumed to be 0.94 in these models.

Table 4.3. Transformed Parameter Estimates: Officers

Parameter	Army	Navy	Air Force	Marine Corps
Scale Parameter, Nest = τ	0.26	181.83	120.73	2.78
Scale Parameter, Alternatives within Nest = λ	109.15	29.96	52.67	78.68
Mean Active Taste = μ_a	–24.30	–20.06	-18.51	–14.14
Mean Reserve Taste = μ_r	–279.98	–55.37	–490.71	–138.94
SD Active Taste = σ_a	42.89	47.77	25.50	23.53
SD Reserve Taste = σ_r	191.57	48.66	324.13	90.75
Taste Correlation = ρ	0.58	0.74	0.42	0.51
Switch Cost: Leave Active < ADSO	–122.34	–180.42	–113.49	–133.39
Switch Cost: Switch from Civilian to Reserve	–425.02	–133.41	–248.92	–277.81
Personal Discount Factor β (Assumed)	0.94	0.94	0.94	0.94

NOTE: Transformed parameters are denominated in thousands of 2007 dollars, with the exception of the taste correlation and personal discount factor. Definitions of variables are provided in the note for Table 2.1.

We found that mean active taste is negative for the Army and equal to –\$24,300. A negative value is consistent with past studies estimating the mean active taste among military officers and suggests that the military must offer relatively high pay to compensate for the requirements of service on active duty relative to not being in the military. For the Navy, the point of estimate of mean active taste is negative but smaller in absolute value than for the Army, equal to –\$20,060. The mean active taste is also smaller in absolute value for both the Air Force and Marine Corps, at –\$18,510 and –\$14,140 respectively. All estimates of mean active taste are statistically different from zero.

Mean taste for reserve duty is negative: –\$279,980 for Army officers, –\$55,370 for Navy officers, –\$490,710 for Air Force officers, and –\$138,940 for Marine Corps officers. As for the variance in tastes, we found that the SD of active-duty taste is larger for the Army and the Navy, at \$42,890 for Army officers and \$47,770 for Navy officers, while the SD of active-duty taste is smaller for Air Force and Marine Corps officers, at \$25,500 and \$23,530, respectively. The SD of reserve taste is largest for the Air Force at \$324,130, followed by the Army at \$191,570, the Marine Corps at \$90,750, and the Navy at \$48,660.

The estimated scale parameter for the between-nest shock in the Navy model is much larger than the means and SDs of tastes, while the within-nest shock is of the same order of magnitude. These scale parameters provide information on the SD of the common random shock for the reserve or civilian nest, as well as the within civilian or reserve nest shocks. The model nests the reserve and civilian alternatives because most reservists also hold a civilian job; hence, a shock to civilian is also likely to be felt by reserve. The scale parameter for the active and reserve-civilian shock is $\sqrt{\lambda^2 + \tau^2}$, while the within-civilian or reserve nest shock is λ. We estimate λ to be \$29,960 and τ to be \$181,830 for the Navy. These estimates imply that the scale parameter for the total shock κ is \$184,278. The relative magnitudes of the scale parameters suggest that movement between the active nest and the reserve or civilian nest is largely driven by random shocks rather than by diverse tastes among Navy members (i.e., taste heterogeneity), while the movement between civilian and reserve statuses is equally driven by diverse tastes and random shocks.

For the Air Force, we found that the between-nest shock τ is larger than the mean and SD of active taste, but smaller in absolute value than the mean and SD of reserve taste. We estimated a τ of \$120,730, about six times the absolute value of the active mean taste of –\$18,510 and about five times the SD of the active taste of \$25,500. However, the estimated value of τ is about one-fourth of the absolute value of the reserve mean taste at –\$490,710 and about one-third of the SD of reserve taste, \$324,130. The within-nest shock λ is estimated to be \$52,670, which, like the estimate for τ, places it between the absolute values of the estimates for the mean and SD of active taste and the mean and SD of reserve taste. The relative sizes of these parameters suggest that movement between the active nest and the reserve-civilian nest are driven by a combination of both service members' individual tastes and random shocks.

For the Army, we found that τ is small and not statistically significant from zero, so that the scale parameter for the active and reserve-civilian shock is essentially reduced to λ. The interpretation of τ being close to zero is that the reserve-civilian shocks are uncorrelated, which might be the case when the year-to-year shocks experienced by civilians who participate in the reserve component are markedly different from pure civilians, such as when members of the reserve component are called on to support members of the active component. We estimated a λ of \$109,150, approximately four times the estimated mean active taste, –\$24,300, and about half the value of the (absolute value of the) estimated mean reserve taste, –\$191,570, implying that tastes and shocks play roles in explaining shifts into and out of active, reserve, and civilian statuses for the Army.

Similarly, for the Marine Corps, we found that τ is small and not statistically significant from zero. As a result, the scale parameter for the active and reserve-civilian shock is essentially reduced to λ. The estimated value of λ is \$78,680, significantly larger than the mean and SD of active taste at –\$14,140 and \$23,530 respectively, and smaller than the mean and SD of reserve taste at –\$138,940 and \$90,750 respectively.

The switching costs for leaving active duty early, before completing ADSO, are –\$122,340 for Army officers, –\$180,420 for Navy officers, –\$113,490 for Air Force officers, and –\$133,390 for Marine Corps officers. The cost of switching to a reserve component after being a civilian is –\$425,020 for Army officers, –\$248,920 for Navy officers, –\$113,490 for Air Force officers, and –\$277,810 for Marine Corps officers. These high costs might reflect the difficulty of finding an available reserve position or an implicit cost to one's civilian career and lifestyle.

Model Estimates for Enlisted Personnel

Tables 4.4 and 4.5 show the estimated parameters and standard errors for the enlisted DRM for the Army, Navy, Air Force, and Marine Corps, respectively. As with the officer models, to make the numerical optimization easier, we did not estimate most of the parameters directly but instead estimated the logarithm of the absolute value of each parameter, except for the taste correlation, for which we estimated the inverse hyperbolic tangent of the parameter. All but the between-nest scale parameters τ are statistically significant in the models. To recover the parameter estimates, we transformed the estimates. Table 4.6 shows the transformed parameter estimates for each service. The estimates are denominated in thousands of 2007 dollars, except for the assumed discount rate and the taste correlation.

Table 4.4. Parameter Estimates and Standard Errors: Army and Navy Enlisted

	Army		Navy	
Parameter	**Estimate**	**Standard Error**	**Estimate**	**Standard Error**
Log(Scale Parameter, Nest = τ)	–1.18	16.62	0.73	2.27
Log(Scale Parameter, Alternatives within Nest = λ)	3.25	0.04	2.99	0.05
Log(–1*Mean Active Taste = μ_a)	2.78	0.03	2.96	0.03
Log(–1*Mean Reserve Taste = μ_r)	3.93	0.05	4.59	0.07
Log(SD Active Taste = σ_a)	1.79	0.11	2.06	0.09
Log(SD Reserve Taste = σ_r)	3.45	0.05	4.04	0.08
Atanh(Taste Correlation = ρ)	0.68	0.03	0.70	0.05
Log(–1*Switch Cost: Leave Active < ADSO)	3.00	0.06	2.87	0.06
Log(–1*Switch Cost: Switch from Civilian to Reserve)	4.71	0.04	4.39	0.05
Personal Discount Factor β (Assumed)	0.88	N/A	0.88	N/A
–1*Log Likelihood	24,656		15,691	
N	5,540		4,863	

SOURCE: Parameter estimates from cohorts of enlisted personnel entering active duty in 1990–1991.
NOTE: The scale parameter κ governs the shocks to the value functions for staying and for the reserve versus-civilian nest and equals $\sqrt{\lambda^2 + \tau^2}$. The means and SDs of tastes for active and reserve service relative to civilian opportunities are estimated, as are the costs associated with leaving active duty before completing ADSO and switching from civilian status to participating in the reserves. The personal discount factor was assumed to be 0.88 in these models. Army and Navy models were estimated using a 5-percent random sample of the data.

Table 4.5. Parameter Estimates and Standard Errors: Air Force and Marine Enlisted

	Air Force		Marine Corps	
Parameter	**Estimate**	**Standard Error**	**Estimate**	**Standard Error**
Log(Scale Parameter, Nest = τ)	1.09	1.38	–3.27	177.29
Log(Scale Parameter, Alternatives within Nest = λ)	3.25	0.05	3.01	0.05
Log(–1*Mean Active Taste = μ_a)	2.67	0.03	3.79	0.04
Log(–1*Mean Reserve Taste = μ_r)	5.16	0.15	7.34	0.25
Log(SD Active Taste = σ_a)	2.13	0.09	3.28	0.06
Log(SD Reserve Taste = σ_r)	4.75	0.15	6.94	0.26
Atanh(Taste Correlation = ρ)	0.49	0.01	0.43	0.00
Log(–1*Switch Cost: Leave Active < ADSO)	3.06	0.06	4.09	0.05
Log(–1*Switch Cost: Switch from Civilian to Reserve)	4.86	0.05	4.30	0.06
Personal Discount Factor β (Assumed)	0.88	N/A	0.88	N/A
–1*Log Likelihood	10,312		11,217	
N	2,576		4,442	

SOURCE: Parameter estimates from cohorts of enlisted personnel entering active duty in 1990–1991.
NOTE: The scale parameter κ governs the shocks to the value functions for staying and for the reserve-versus-civilian nest and equals $\sqrt{\lambda^2 + \tau^2}$. The means and SDs of tastes for active and reserve service relative to civilian opportunities are estimated, as are the costs associated with leaving active duty before completing ADSO and switching from civilian status to participating in the reserves. The personal discount factor was assumed to be 0.88 in these models. Air Force and Marine Corps models were estimated using a 5-percent and 10-percent random sample of the data, respectively.

Table 4.6. Transformed Parameter Estimates: Enlisted

Parameter	Army	Navy	Air Force	Marine Corps
Scale Parameter, Nest = τ	0.31	2.06	2.99	0.04
Scale Parameter, Alternatives within Nest = λ	25.66	19.83	25.83	20.19
Mean Active Taste = μ_a	–16.16	–19.32	–14.37	–44.30
Mean Reserve Taste = μ_r	–50.84	–98.45	–174.10	–1,545.44
SD Active Taste = σ_a	5.99	7.84	8.42	26.47
SD Reserve Taste = σ_r	31.44	56.93	115.73	1,030.80
Taste Correlation = ρ	0.59	0.61	0.46	0.41
Switch Cost: Leave Active < ADSO	–20.08	–17.65	–21.33	–59.50
Switch Cost: Switch from Civilian to Reserve	–110.95	–80.42	–129.35	–73.97
Personal Discount Factor β (Assumed)	0.88	0.88	0.88	0.88

NOTE: Transformed parameters are denominated in thousands of 2007 dollars, with the exception of the taste correlation and personal discount factor. Definitions of variables are provided in the note for Table 2.4.

We found that mean active tastes are negative and equal to –\$16,600, –\$19,320, –\$14,370, and –\$44,300 for the Army, Navy, Air Force, and Marine Corps, respectively. The negative values are consistent with past studies and suggest that the military must pay a relatively high wage to compensate for the rigors of military life and retain enlisted members. All estimates of mean active taste are statistically different from zero.

The mean reserve tastes are also negative and are equal to –\$50,840, –\$98,450, –\$174,100, and –\$1,545,440 for the Army, Navy, Air Force, and Marine Corps respectively. As for the variance in tastes, we found that the SD of active-duty taste is largest for the Marine Corps at \$26,470, while the SD of active-duty taste is smaller for Army, Navy, and Air Force enlisted members at \$5,990, \$7,840, and \$8,420, respectively. Similarly, the SD of reserve taste is largest for the Marine Corps at \$1,030,800, followed by the Air Force at \$115,730, the Navy at \$56,930, and the Army at \$31,440.

The mean taste for reserve duty in the Marine Corps is large and negative, and this might be because it reflects not only individual preference for serving in the Marine Corps, but also the demand of the Marine Corps for members with prior active service to participate in the reserve component. In short, in addition to individual preference, the mean reserve taste also reflects the cost or difficulty associated with finding a position in the reserve component that works for the individual member, or simply the lack of available positions.

The estimated scale parameter for the between-nest shock in all the models is much smaller than the means and SDs of tastes, while the within-nest shock is of the same order of magnitude of the active tastes and uniformly smaller than the means and SDs of the reserve tastes. The estimated scale parameters for the between-nest shock in the Army, Navy, Air Force, and Marine Corps are \$310, \$2,060, \$2,990, and \$40, respectively, none of which are significantly different from zero. The within-nest shock parameters are \$25,660, \$19,830, \$25,830, and \$20,190 respectively and are all significantly different from zero. The relative magnitudes of the scale parameters suggest that movement between the active nest and the reserve-civilian nest is equally driven by random shocks and diverse tastes among enlisted members in all four services, while the movement between civilian and reserve statuses tends to be more driven by taste than by random shocks.

The switching costs for leaving active-duty early, before completing the first term, are –$20,080 for Army enlisted members, –$17,650 for Navy enlisted members, –$21,330 for Air Force enlisted members, and –$59,500 for Marine Corps enlisted members. The cost of switching to a reserve component after being a civilian is –$110,950 for Army enlisted members, –$80,420 for Navy enlisted members, –$129,350 for Air Force enlisted members, and –$73,970 for Marine Corps enlisted members. These high costs might reflect the difficulty of finding an available reserve position within traveling distance of where the former active member has settled down.

Simulation and Assessment of Model Fits

Once the model has been estimated, it can be used to assess model fits and for policy simulations.[17] The first step in conducting a simulation is to create a population of synthetic individuals. Within the context of the model, an individual is an entity with specific preferences for active and reserve service and a specific set of random shocks for each alternative in each period. Therefore, the simulation creates the individual by a random draw of active and reserve preferences from the preference distributions and a set of random draws from the shock distributions. Consistent with the model, the individual is assumed to know their preferences, the values of the shocks in the current decision period, and the distributions of the shocks (i.e., the scale parameters which are used in the individual's computation of the expected value of the maximum in the next period). That is, even though the analyst knows the shocks for each period in the individual's work life, in any period, the individual does not know the values of shocks in future periods.

The second step in a simulation is to specify the compensation structure. Our policy analysis generally focuses on comparisons between the current compensation structure and alternative structures. The current structure was coded into the model when it was estimated; the observed active and reserve retention behavior was conditional on the current compensation structure. New coding is required for each alternative structure. The policy alternatives under consideration involve changing the structure of the military pay table, as we discuss in the following section.

The third step is to put the synthetic population into the model; compute each person's value functions recursively, as described previously; and let the person choose an alternative at each decision point. The result is a career path that is optimal for the individual, given the compensation structure and particular shocks he or she faced in each period.

[17] These prefatory comments on the use of simulation draw from Asch et al. (2008). The subsequent model fit results draw from Asch, Mattock, and Tong (2020).

The fourth and final step is postprocessing. The career information for our synthetic individuals includes period-to-period information about their state (active years, reserve years, total years, and pay grade). We combine this to create information about the synthetic population (its active-duty retention, participation in the selected reserve, highest grade attained, expected years of active and reserve service, and compensation cost). We can manipulate this information to make cost comparisons subject to holding active-duty personnel strength constant or, alternatively, make strength comparisons holding cost constant. When discounting is required, we use the discount factor appropriate for the calculation (i.e., the individual's discount factor for calculations from the perspective of the individual or the organizational discount factor for calculations from the organization's perspective).

The cost concept that we use is current cost rather than life-cycle cost. Life-cycle cost is often used in weapon system procurement costing, and it could be used in manpower costing if "procuring" a cohort of new entrants is considered purchasing a new asset. However, policymakers are accustomed to viewing manpower costs as current outlays, so current costing seems more appropriate. With respect to our simulation, we simulate the career behavior of a population entering active duty, and to convert our results into a current setting, we adopt the assumption that the personnel force is in a steady state. By implication, the active and reserve retention behavior that we simulate can be interpreted as the force structure that one would see in the cross-section (i.e., in the current period). Under this assumption, we compute two costs: current compensation and deferred compensation. *Current compensation* includes regular military compensation plus any gate pays (i.e., continuation pays conditional on reaching YOS milestones) and separation pays. *Deferred compensation* includes outlays required to fund defined benefit, and, for the BRS, the defined contribution portion of retirement for vested personnel upon their departure from service. For example, if a service member left active duty after completing 24 YOS, we register a cost equal to the present discounted value of the stream of retirement benefits expected to be paid to the individual in all future periods, under the terms of the retirement-benefit system we were simulating and allowing for survival probabilities.

Assessing Model Fit Through Simulation

To assess model fit, we used the parameter estimates to simulate the behavior of 10,000 synthetic service members represented by tastes drawn from the active and reserve taste distribution and subject to shocks drawn from a shock distribution with a scale parameter equal to the estimated value. Given active and reserve tastes, current-period shock values, knowledge of the expected pay and promotion environment in the military and the civilian world, and knowledge of the shock scale parameter, each synthetic individual, behaving as a dynamic-program decisionmaker, makes a stay-leave decision in each YOS in the active component. This generates a career length of service in the active component. After leaving active service, the individual becomes a civilian and makes a yearly decision regarding reserve participation. If the individual is not in the reserves, the decision is whether to participate; if the individual is in the reserves, the decision is whether to continue to participate. These decisions generate information about reserve participation by year for the years after active component service. We obtained the predicted active component retention profile by adding together these simulated active component retention profiles across many simulated individuals, and we similarly combined individual reserve participation profiles to obtain the predicted reserve participation profile for the population of simulated individuals. The predicted profiles are plotted against the actual profiles to assess goodness of fit.

Model Fit for Officers

Figures 4.1 through 4.4 show the model fit graphs for the active component for each of the four services using the simulation approach described previously. The red lines are simulated cumulative retention, and the black lines are retention observed in the data. The figures show the Kaplan-Meier survival curves, and the dotted lines show the 95 percent confidence intervals for the Kaplan-Meier estimates for the observed data. The horizontal axis counts years since the individual was observed beginning active service. The vertical axis shows the cumulative probability of retention on active duty until that year. For example, at entry, YOS is zero, the fraction of personnel retained is 1, and the fraction of the force retained falls over an active career as officers leave active duty. The solid black line shows the actual retention of individuals in our cohorts, and the red line shows the predicted retention. The numbers beneath the x-axis correspond to the model parameters shown in Tables 4.1 or 4.2 and help to ensure that a given figure matches a particular set of estimates. We assess goodness of model fit by visual inspection, that is, in terms of how well the black and red lines coincide.

Visual inspection reveals that model fit for the active component is good for the Army, Air Force, and Marines, and that the model captures the general sweep of Navy retention. In all cases, the simulated retention line lies close to the observed retention line and reflects the pattern of retention seen in the data with attrition first being high, then slowing after mid-career as vesting in the defined-benefit retirement plan approaches, and then falling quickly once the vesting point is reached.

Figure 4.1. Model Fit Results: Active-Component Army Officers

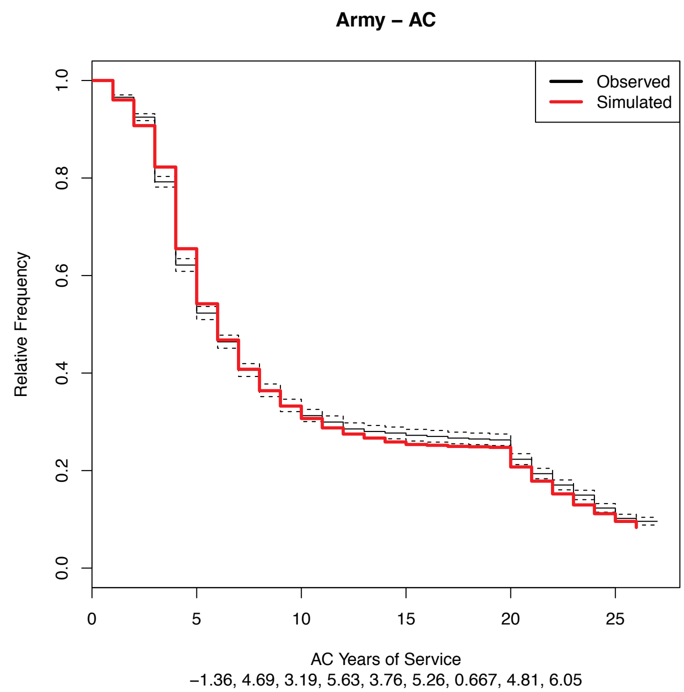

SOURCE: Authors’ computations using DMDC WEX files.

Figure 4.2. Model Fit Results: Active-Component Navy Officers

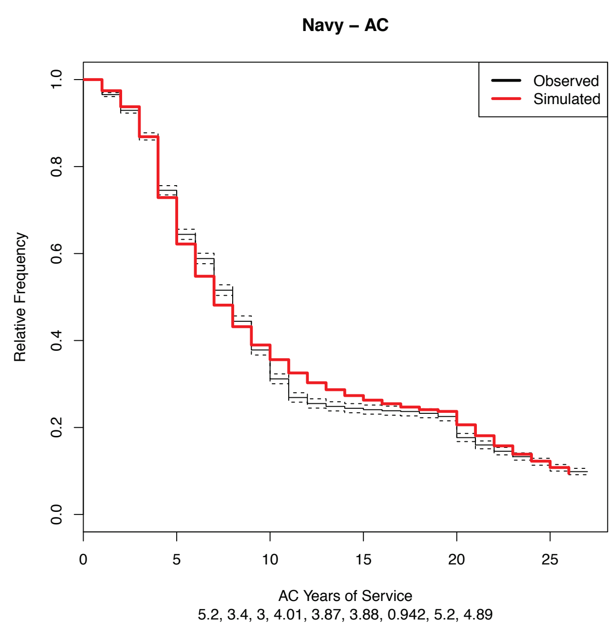

SOURCE: Authors’ computations using DMDC WEX files.

Figure 4.3. Model Fit Results: Active-Component Air Force Officers

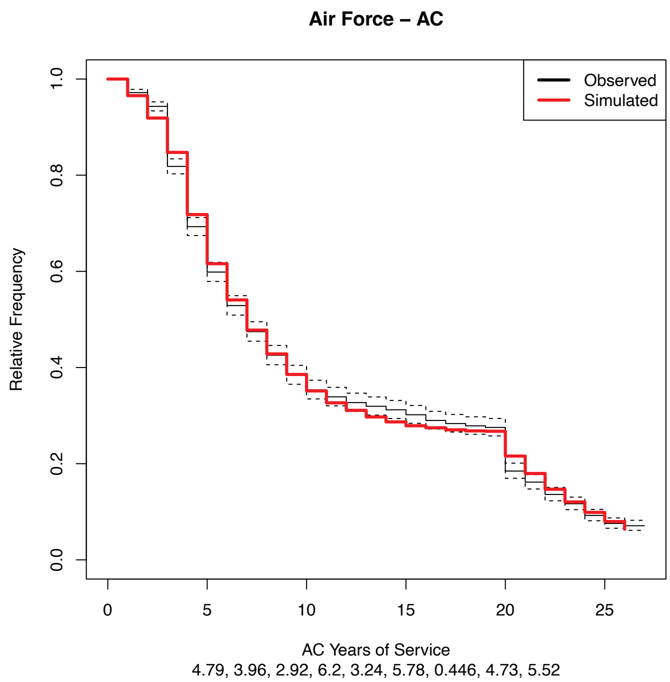

SOURCE: Authors' computations using DMDC WEX files.

Figure 4.4. Model Fit Results: Active-Component Marine Corps Officers

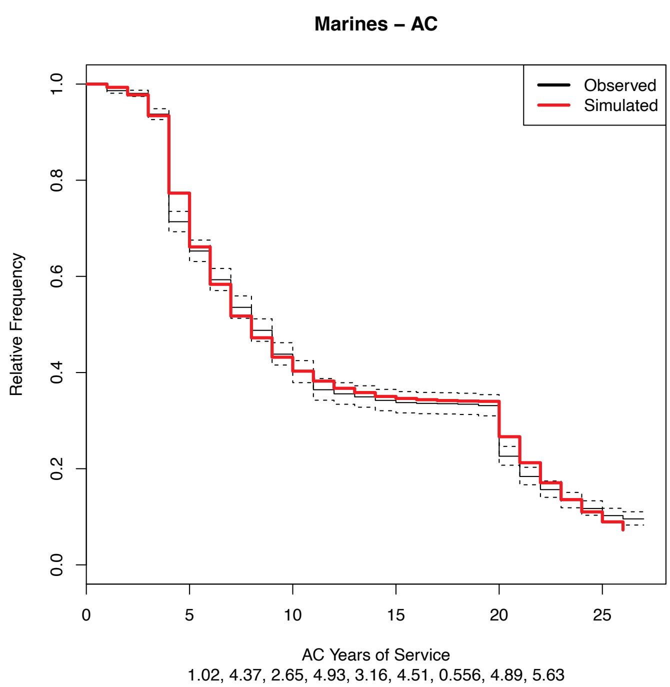

SOURCE: Authors' computations using DMDC WEX files.

Model Fit for Enlisted

Similar to the models of officer retention behavior, to assess model fit, we used the parameter estimates to simulate the behavior of synthetic personnel represented by tastes drawn from the active/reserve taste distribution and subject to shocks drawn from a shock distribution with a scale parameter equal to the estimated value. Figures 4.5 through 4.8 show the model fit graphs for the active component for each of the four services. The red lines are simulated cumulative retention, and the black lines are retention observed in the data. The figures show the Kaplan-Meier survival curves, and the dotted lines show the 95 percent confidence intervals for the Kaplan-Meier estimates for the observed data.

The horizontal axis counts years since the individual began active service. The vertical axis shows the cumulative probability of retention on active duty until that year. The solid black line shows the actual retention of individuals in our cohorts, and the red line shows the predicted retention.

Visual inspection shows that the model fit for the active component is good for the Army, Navy, and Air Force, and that the model under-predicts retention for the Marine Corps for YOS 3–7 and slightly over-predicts retention beyond YOS 10. In all cases, the simulated retention line lies close to the observed retention line and reflects the pattern of retention seen in the data, with attrition first being high, then slowing after mid-career as vesting in the defined-benefit retirement plan approaches, and then falling quickly once the vesting point is reached.

Figure 4.5. Model Fit Results: Active-Component Army Enlisted

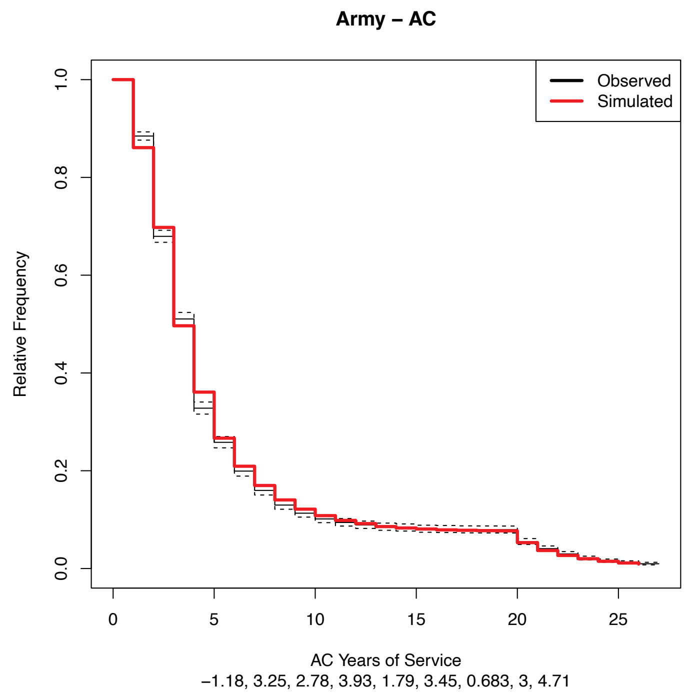

SOURCE: Authors' computations using DMDC WEX files.

Figure 4.6. Model Fit Results: Active-Component Navy Enlisted

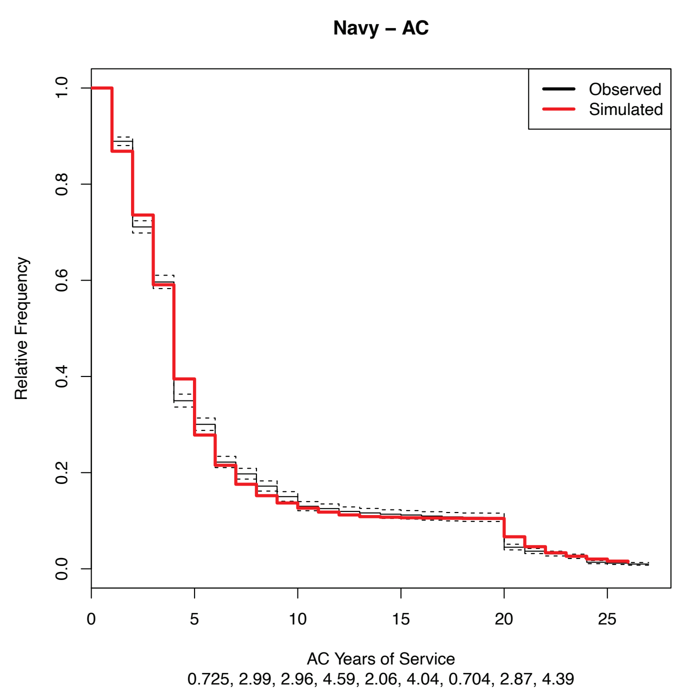

SOURCE: Authors' computations using DMDC WEX files.

Figure 4.7. Model Fit Results: Active-Component Air Force Enlisted

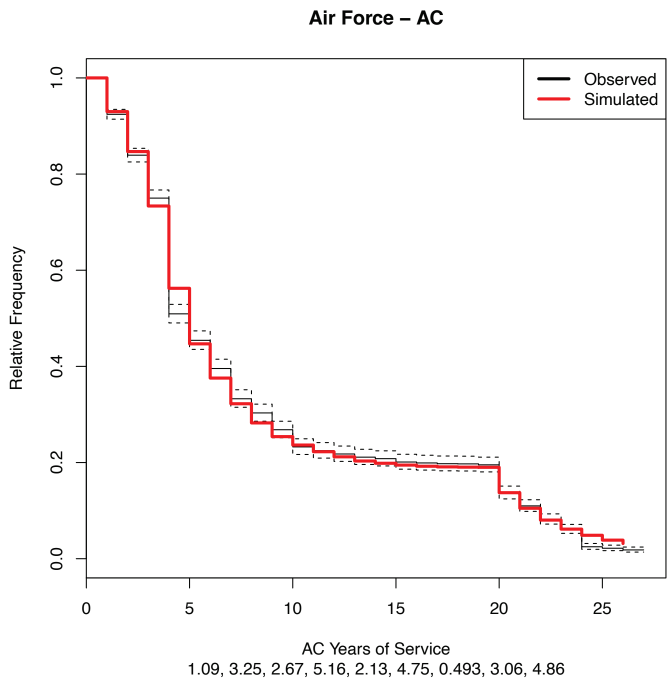

SOURCE: Authors' computations using DMDC WEX files.

Figure 4.8. Model Fit Results: Active-Component Marine Corps Enlisted

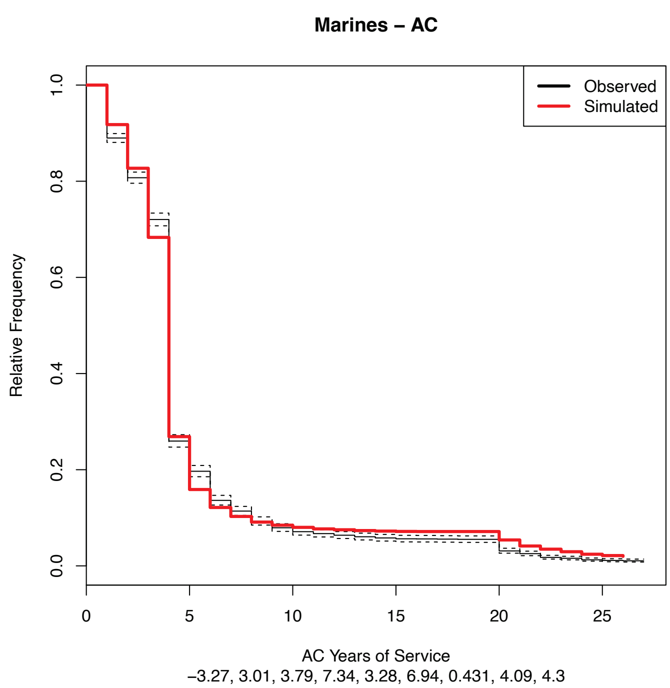

SOURCE: Authors' computations using DMDC WEX files.

Simulation and Extension of the DRM to Model a Time-in-Grade Pay Table

We also developed a simulation capability to specifically consider the policy scenarios under consideration by the 13th QRMC related to the TIG pay table. To simulate the effect on retention of changing to a TIG pay table, we needed to extend the DRM in two ways: (1) adapt the model to track TIG, i.e., the number of YOS since a member was last promoted and (2) ensure that military pay in the model uses TIG rather than TIS.

The DRM was estimated using data on the behavior of officer and enlisted members under a TIS pay table, where the compensation that an individual received was a function of their grade and YOS, which could conceptually be written as

$$W_t^{Ag} = W(ay_t,\ g_t).$$

Under a TIG pay table, the compensation a member receives is a function of their grade and the number of YOS since they were promoted to that grade. If we let py_t be the number of YOS since a member was last promoted, then we can write their wage as

$$W_t^{Ag} = W(py_t,\ g_t).$$

If we change the definition of k_t by adding py_t as follows,

$$k_t = k_t(ay_t, ry_t, t, g_t, py_t),$$

then the rest of the mathematical expressions developed earlier in this chapter still follow through. As a result, we could use the parameters estimated with the historical career data and TIS pay table to simulate the retention effects of replacing the TIS pay table with the TIG pay table. We could also simulate the effects on performance and cost. We discuss how we incorporate performance in the next subsection.

Incorporating Performance into the DRM Simulation Capability

A major impetus for considering a TIG pay table is that it increases the incentives for performance, as discussed in Asch, Mattock, and Tong (2020). We incorporated performance into our analysis by focusing on two aspects of individual service members that can affect their performance in the military: innate ability and how hard they work. This focus on the inputs of performance on the part of the service member is consistent with two of the key objectives of the military compensation system related to individual performance: (1) to motivate personnel to work hard and effectively and (2) to induce higher-ability personnel to stay and seek advancement to more-senior grades, where it is likely that ability has a bigger impact than in the lower ranks.[18]

Asch and Warner were the first to incorporate ability and effort supply into a DRM, and they used the model to assess the retention, performance, cost effects of alternative retirement reform proposals, and policies to restructure the military pay table (Asch and Warner, 1994a, 1994b, 2001). In their model, higher-ability personnel and those who exert more effort are promoted faster and have higher promotion probabilities, but higher-ability personnel also have better external opportunities, while expending effort involves a cost or disutility to the member (under the assumption that individuals would prefer to exert less effort for the same amount of financial benefit or return to effort). Compensation policy can affect the financial returns to exerting more effort and the financial benefits to staying for higher-ability personnel. Asch and Warner used their DRM to provide simulations of how compensation reforms affected overall retention, the retention of higher-ability personnel, ability-sorting into higher grades, average effort supply, and personnel cost.

The Asch-Warner simulations used a calibrated model in which key parameters, such as the mean and SD of taste for service, were assumed to replicate the observed retention profile. In contrast, the parameters of the DRM shown here and in Asch, Mattock, and Tong (2020) were estimated, not calibrated. We built on the Asch and Warner modeling of ability and effort supply and incorporated their approach into our DRM simulation capability to evaluate a TIS versus TIG pay table. Ideally, we would consider both effort supply and ability simultaneously as factors affecting promotion probabilities, an approach taken by Asch and Warner (1994b). But we found we were better able to incorporate ability and effort supply by considering them separately, as we will discuss in more detail in the following section. In the rest of this subsection, we first discuss how we incorporated ability and then effort supply.

Ability

We can use the structure of the DRM along with the estimated parameters and assumptions about how innate ability affects the speed of promotion to examine how selective the TIG and TIS pay tables are on ability. To incorporate ability into the DRM, we made three assumptions:

1. the extent to which ability differs among military entrants[19]

---

[18] The objectives of military compensation are listed in Under Secretary of Defense for Personnel and Readiness (2018) and have been articulated by past QRMCs as well as the DACMC.

[19] We assumed that the distribution of ability at entry is fixed and the same under TIS and TIG pay tables. Because we did not consider the effects of a TIG pay table on recruiting in this study, we did not consider the possibility that a TIG pay table might be more attractive to higher-ability recruits, thereby shifting the mean of the ability distribution. The implication is that a TIG pay table could have a greater effect on ability of the force than what we found in Asch, Mattock, and Tong (2020).

2. the extent to which ability affects promotion speed[20]
3. the effect of ability on external civilian opportunities.

We discuss each of these in turn.

First, we assumed that any given individual has a fixed level of ability at entry, drawn from a normal distribution and rounded to the nearest integer. The SD of the distribution indicates the extent to which ability differs among military entrants. Regarding rounding, individuals with ability drawn from a normal distribution with mean zero and SD 0.5 (and then rounded) would typically have values of ability of –1, 0, or 1. We assumed a different mean and SD for each service and for enlisted personnel and officers within that service. The values of the mean and SD for each distribution we used in our simulations were calibrated to replicate the steady state retention profiles of enlisted and officers under the baseline TIS pay table, given the other two assumptions we make.

Second, we assumed that higher-ability personnel are promoted faster. We implemented this concept by subtracting the (rounded) draw from the normal distribution for a given individual from the TIS between promotions. This increase in promotion speed was modeled to start happening between E-5 and E-6 for enlisted members and between O-3 and O-4 for officers. Thus, an enlisted member with an innate ability of 1 would reach E-6 one year faster, reach E-7 two years faster, and so on. An officer with an innate ability of 1 would reach O-4 one year faster, reach O-5 two years faster, and so on. Consequently, the effect of ability on promotion speed to more-senior grades is larger than for more-junior grades because the effects on promotion timing are cumulative. Figure 4.9 shows how years to promotion from E-6 to E-9 vary with ability for Army enlisted personnel, while Figure 4.10 shows how years to promotion from O-4 to O-7 vary with ability for Army officers. Results will differ for the other services insofar as the assumed parameters of the ability distribution differ. As mentioned in the previous paragraph, the assumed parameters are calibrated to best fit the retention profile for that service and grade category.

[20] The model only considers individual attributes in promotion timing and probability, so it does not allow for the possibility of the ability distribution skewing higher under TIG, which would result in slowing down the promotion of individuals who might have been promoted early under TIS.

Figure 4.9. Years to Promotion by Ability Level, Army Enlisted Personnel

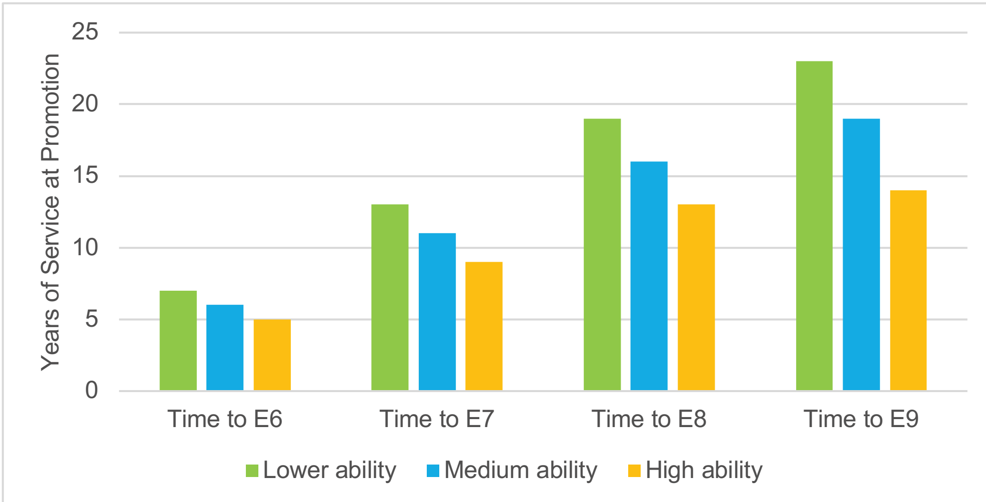

SOURCE: Authors' computations using DMDC WEX files.

Figure 4.10. Years to Promotion by Ability Level, Army Officers

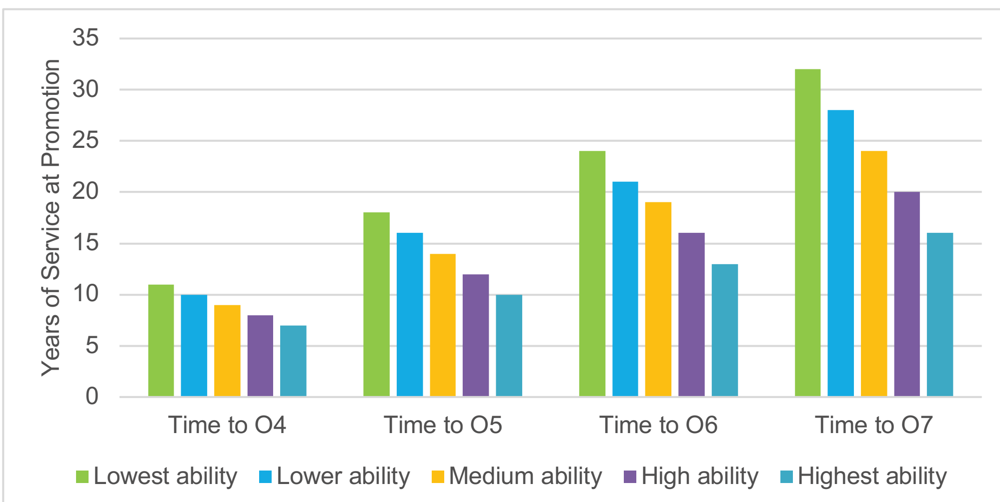

SOURCE: Authors' computations using DMDC WEX files.

Third, we assumed that higher-ability members also have better external opportunities. We modeled this by multiplying the civilian opportunity wage by one plus 0.1 times the ability distribution SD times the individual's ability draw, or $(1 + 0.1\ \times\ \sigma_a)$ where σ_a is the SD of the draw. This had the effect of increasing the civilian opportunity wage for high-ability individuals and decreasing the civilian opportunity wage for low-ability individuals. For example, an individual drawn from a normal distribution with mean zero and SD 0.5 with innate ability of 1 would have an opportunity wage that is 5 percent greater than that of the average individual, while an individual with innate ability –1 would face a civilian opportunity wage that is 5 percent less.

We illustrate how we calibrated the mean and SD of the normal distribution to fit the observed retention profile in Figure 4.11 for Army enlisted personnel. In the process of calibration, we systematically varied the mean and SD within the TIS DRM and chose the mean and SD that most closely replicated the historically observed retention, as indicated by the Kaplan-Meier curve. The right panel shows the observed retention profile versus the simulated retention profile when we mis-calibrate the mean and SD to equal 0 and 1.5, respectively. The simulated retention profile is too high relative to the observed profile. We chose a SD of 0.5 instead resulting in a good fit, as shown in the left panel.

Figure 4.11. Calibrating the Parameters of the Ability Distribution, Army Enlisted Personnel

Calibrated Mean = 0, SD = 0.5	**Mis-calibrated** Mean = 0, SD = 1.5

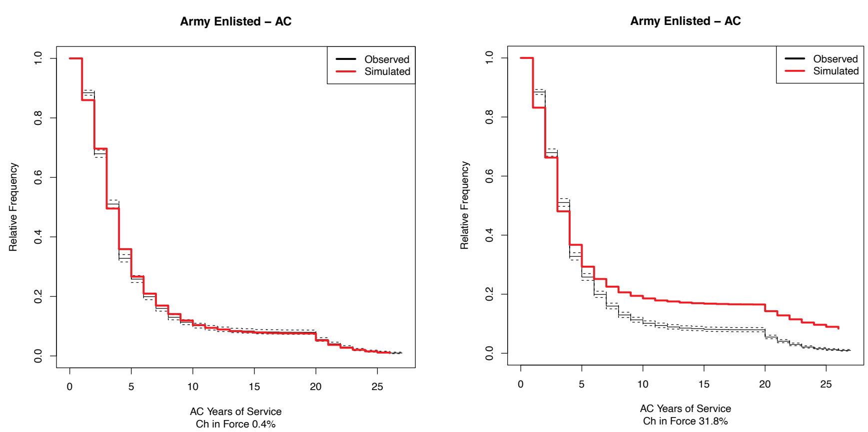

SOURCE: Authors' computations using DMDC WEX files.

The three assumptions we make regarding how ability enters the model could affect our simulation results and, in particular, the effects of the TIG pay table on retention, ability-sorting, and cost. Consequently, in our presentation of results in Chapter 4 of Asch, Mattock, and Tong (2020), we included sensitivity analyses where we varied these three underlying assumptions regarding ability.

Modeling Effort

In addition to native ability, a service member's promotion performance can depend on the amount of effort they exert. The main idea is that with other things held constant, the more effort a member exerts, the more likely it is that they will be promoted. The structure of the model allows us to derive the optimal amount of effort an individual would exert given assumptions about how effort affects the probability of an individual being promoted and assumptions about the disutility of effort.

Following Asch and Warner (1994b), we added disutility of effort to the value function in the DRM presented previously. The individual's problem is to choose the level of effort to exert in the current period to maximize their utility:

$$\max_{e_t} V^A(k_t) - Z(e_t).$$

To simplify notation, we define $\bar{V}^A(k_t)$ to be the value of staying in the active component net the disutility of effort, like so:

$$\bar{V}^A(k_t) \equiv V^A(k_t) - Z(e_t).$$

The first order condition for the optimal level of effort is

$$\frac{\partial \bar{V}^A(k_t)}{\partial e_t} = \beta \Pr\left(\bar{V}^S(k_{t+1}) > V^L(k_{t+1})\right)\left(\bar{V}^{A(g+1)}(k_{t+1}) - \bar{V}^{Ag}(k_{t+1})\right)\frac{\partial p_{t+1}^{g+1}}{\partial e_t} - Z'(e_t) \equiv 0$$

or

$$\Pr\left(\bar{V}^S(k_{t+1}) > V^L(k_{t+1})\right)\beta\left(\bar{V}^{A(g+1)}(k_{t+1}) - \bar{V}^{Ag}(k_{t+1})\right)\frac{\partial p_{t+1}^{g+1}}{\partial e_t} \equiv Z'(e_t).$$

The interpretation of this expression is that the product of the probability of staying in the next period, the discounted difference of the value of being active and promoted and the value of being active and not promoted, and the marginal effect of effort on the probability of promotion equals the marginal disutility of effort. Or, to put it more simply, the expected marginal return to effort equals the marginal disutility of effort.

If we make some assumptions regarding the functional form of the disutility of effort function and the probability of promotion as a function of effort, we can solve for optimal effort at time t. Similar to Asch and Warner, we let the disutility of effort be

$$Z(e_t) = \frac{\eta_0}{2} e_t^2$$

and let the probability of promotion be

$$p_{t+1}^{g+1} = \mu^{g+1} \bar{p}_{t+1}^{g+1} e_t,$$

where μ^{g+1} is a parameter that captures the relationship between effort and the probability of promotion for a given individual and $\bar{p}_{t+1}^{g+1}$ is the average promotion probability to grade g+1 at time t+1. We can rewrite the first order condition as:[21]

$$\beta \Pr\left(\bar{V}^{S}(k_{t+1}) > V^{L}(k_{t+1})\right)\left(\bar{V}^{A(g+1)}(k_{t+1}) - V^{Ag}(k_{t+1})\right)\mu^{g+1}\bar{p}_{t+1}^{g+1} - \eta_0 e_t \equiv 0$$

and solve for e_t as

$$e_t = \frac{\beta \Pr\left(\bar{V}^{S}(k_{t+1}) > V^{L}(k_{t+1})\right)\left(\bar{V}^{A(g+1)}(k_{t+1}) - \bar{V}^{Ag}(k_{t+1})\right)\mu^{g+1}\bar{p}_{t+1}^{g+1}}{\eta_0}.$$

Given assumptions for the values of the parameters η_0, μ^{g+1}, and $\bar{p}_{t+1}^{g+1}$, along with our DRM parameter estimates, we can solve for e_t and then simulate how the average level of effort among service members differs under the TIS pay table versus the TIG pay table.

Modeling the Effect of Effort in Multiple Periods to Promote to the Next Grade

In the previous formulation, the individual has some probability of being promoted in each period *t* and the probability of promotion was dependent on effort in the immediately preceding period. In our model, as we described earlier in the chapter, we assume that the probability of promotion to a given grade occurs at a given number of YOS, but the probability of promotion differs by grade. That is, in our model, promotion occurs at a given point in time for a particular grade. An implication of this approach to modeling promotion is that individual's promotion chances might depend on effort over multiple periods. We accommodate this feature by changing the assumed form of the probability of promotion function. Instead of the probability being dependent on effort in a single period as follows,

$$p_{t+1}^{g+1} = \mu^{g+1}\bar{p}_{t+1}^{g+1}e_t$$

it can depend on effort in multiple periods, as in this example:

$$p_{t+1}^{g+1} = \mu^{g+1}\bar{p}_{t+1}^{g+1}\sum_{i=t-k}^{t} e_i.$$

The expressions for e_{t-1}, e_{t-2}, and so on take on a similar form to the expression for e_t. For example, the expression for e_{t-1} is:

[21] The derivation of this expression requires several steps shown in the appendix.

$$e_{t-1} = \frac{\beta^2 \Pr\left(\bar{V}^A(k_t) > V^L(k_t)\right) \Pr\left(\bar{V}^S(k_{t+1}) > V^L(k_{t+1})\right)\left(\bar{V}^{A(g+1)}(k_{t+1}) - \bar{V}^{Ag}(k_{t+1})\right)\mu^{g+1}\bar{p}_{t+1}^{g+1}}{\eta_0}.$$

Note that the values of $V^A(k_t)$ and $V^A(k_{t+1})$ depend on the value of e_t, e_{t+1}, e_{t+2}, and so on, so we cannot compute the value of e_{t-1} without knowing all the future levels of effort, as well as any past levels of effort associated with the same promotion point with which e_{t-1} is associated. In general, if a promotion point probability depends on multiple years of effort, we need to solve for all levels of effort associated with a promotion point simultaneously. In our simulations, we use an iterative procedure to solve for a set of levels of effort that are stationary; that is, we start off with a guess of the optimal level of effort in each period and then (1) solve for the optimal level of effort in each period, given that all others are fixed, (2) update the levels of effort, and (3) iterate until the computed levels of effort cease to change. We solve for the levels of effort associated with the senior-most promotion point first, then the levels of effort associated with the next-most-senior promotion point, and so on, until we work our way backward to the initial promotion point.

Solving for the optimal effort supply decision in each year of service for each service member in our simulations is a nontrivial task. In the model, these decisions depend on only two parameters: the disutility of effort parameter and the relationship between promotion and effort. As with the ability parameters, we calibrated the effort-related parameters to replicate the cumulative retention profile. Figure 4.12 shows the fit for the Army enlisted model after calibrating the effort-related parameters where we ignore ability in the model. The simulated profile broadly tracks the observed profile, but the fit is not as good as the one where we calibrate only the ability parameter, as shown in Figure 4.11. Consequently, in our presentation of results related to the effects of the TIG pay table on effort in Chapter 4 of Asch, Mattock, and Tong (2020), we only showed results for Army enlisted personnel and considered our results as exploratory.

Figure 4.12. Calibrating the Parameters of the Effort Decision, Active-Component Army Enlisted Personnel

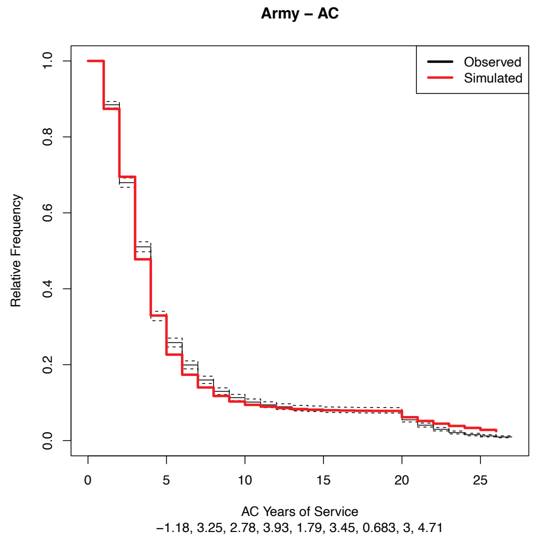

SOURCE: Authors' computations using DMDC WEX files.

Summary

The DRM is a model with a relatively simple structure, but it can support a rich variety of analyses. In this chapter, we extended the DRM to model the promotion process, presented parameter estimates, and showed the model fits for enlisted personnel and officers for each service. We also discussed how we had extended the simulation capability to permit analysis of the TIG pay table, incorporated ability, and the effort supply decision.

Chapter 5. Additional Innovations and Extensions to the DRM

Analyses of different populations, both military and civilian, have required different innovations and extensions to the basic DRM. In the previous chapters, we provided an overview of RAND DRM analyses, extensions to model promotion, ability, and effort, as well as the capability to model a TIG pay table that were done in support of the 13th QRMC. In this chapter, we discuss some additional innovations and extensions that we have implemented to meet the needs for selected military communities, such as Air Force pilots, military mental health care providers, and special operations forces (Mattock et al., 2016; Mattock and Arkes, 2007; Hosek et al., 2017; Asch et al., 2019); as well as civilian communities, such as DoD civil servants (Knapp, Asch, et al., 2016), state employees (Knapp, Asch, and Mattock, 2021) and public school teachers (Knapp et al., 2023). To be specific, we consider four different extensions to the model: (1) adding covariates to shift the mean and variance of the taste distribution, (2) modeling retention when individuals are offered bonuses and special pays that involve multi-year contracts, (3) modeling the transition from one steady-state to a new steady-state, and (4) using incumbents to supplement short panels in estimation.

Adding Covariates to Shift the Taste Distribution

In the basic DRM, there is a single taste distribution at entry, estimated for an entire population. One straightforward extension is to have indicator variables for different subpopulations and estimate coefficients on these indicator variables that shift the mean and SD. A simple example of this type of extension is in the model of Air Force pilot retention in Mattock and Arkes (2007), in which an indicator variable for ROTC graduates was used to shift the mode and scale of the taste distribution (which was assumed to be extreme-value distributed). The omitted group was Air Force Academy graduates, and the parameter estimates showed that, while the taste distribution mode was significantly lower for ROTC graduates, there was no significant difference between the two groups in the scale parameter.

A more sophisticated example of this approach appears in a model of DoD civil service retention by Knapp, Asch, et al. (2016), where indicator variables are used to distinguish between different cohorts and by veteran status. We will discuss this in more detail in the following section. Our exposition draws heavily from Knapp, Asch, et al. (2016).

The DRM in Asch, Mattock, and Hosek (2014) used data on the 1988 entry cohort of DoD general schedule civilian employees. Knapp, Asch, et al. (2016) extended the DRM by using data on entry cohorts for 1988 to 2000 and considering the possibility that veterans have a different distribution of tastes for DoD civil service from that of nonveterans. That is, we estimated the

DRM with additional entry cohorts and allowed for group differences in tastes within an entry cohort and across entry cohorts.

Taste Distribution Differences Across Entry Cohorts and Across Groups

In the DRM used in Asch, Mattock, and Hosek (2014), we estimated the mean and SD of the taste distribution of the 1988 entry cohort. Together with the other parameter estimates, we showed that the model fit is extremely good, and, in fact, we conducted out-of-sample predictions and found that the model performs very well. Because of these estimates, we simulated the retention effects of the 2011–2013 pay freeze and furloughs.

A natural question is whether the experience and retention behaviors of the 1988 entry cohort are representative and whether we would continue to find similar estimates and retention responses to policy changes using data for other cohorts. At the heart of this question is whether the taste and shock distribution parameters differ across entry cohorts. We focus here on the taste distribution.

There are two reasons that the mean and SD of the taste distribution might change across entry cohorts. The first is that individual tastes can change over time. For example, it is common for the popular press to point out differences in the career aspirations and cultural attitudes of generation X (born between 1965 and 1980, according to the Pew Research Center), millennials (born between 1980 and 2000), baby boomers (born between 1945 and 1965), and Depression-era and wartime (born between 1925 and 1945). The common argument is that younger generations evaluate their labor market opportunities and establish priorities that differ from those of older generations. However, Stafford and Griffis (2008) reviewed available data and literature on how the millennial generation differs from earlier generations and argued that some of the characteristics attributed to millennials—or, for that matter, any specific generation at a point in time—might be because of life-stage effects that are found in all generations as they age.[22] That is, different generations might respond in similar fashions when they are at the same age. An alternative view is that calendar year, not birth year, is the most relevant for decisionmaking. For example, one could argue that changes in societal attitudes and culture over time affect all individuals in a similar fashion, regardless of birth year. Thus, the tastes of all individuals entering public service in 2000 might differ from those entering in 1988, regardless of age.

Ultimately, whether the taste distribution has changed over time across entry cohorts is an empirical issue. We investigated this issue by estimating a combined model using data for all entry cohorts that include a separate mean taste parameter for each entry calendar year. We estimated the combined model for a specific demographic subset of the data to limit the

---

[22] None of the entrants in our data were millennials because the youngest person was age 22 in 2000, born in 1978. The sample nonveterans included baby boomers and generation X-ers, while the veterans included mostly baby boomers.

possibility that the observed differences are because of changes in the demographic composition of the entry cohorts over time. In particular, we estimated a combined model for nonveterans, ages 30 or younger, with bachelor's degrees and no further education. We then tested the statistical significance of the mean taste parameters across years for evidence of a shift in tastes across entry cohorts, holding group composition constant.

We found that point estimates of mean tastes differ across entry cohorts but are not always statistically significant when we hold group composition constant as mentioned. The responsiveness to a pay change also differs across entry cohorts, but the differences are not large.

A second reason that the estimated taste distribution parameters might differ across entry cohort is that entry cohort composition changes over time. Knapp, Asch, et al. (2016), presented tabulations of the demographic characteristics of the entry cohorts and how they changed over time. The most dramatic change was in the percentage of entrants who were veterans, increasing from 5.5 percent in 1991 to 23.5 percent in 1994. This increase accounted for much of the increase in mean age over the same period. The changes in veteran representation also affected changes in the grade distribution, education distribution, and gender distribution, although these changes also occurred among nonveteran entrants.

Although demographic differences do not automatically translate into differences in tastes for DoD civil service, the distribution of tastes for DoD service could differ between veterans and nonveterans, and especially between veterans who are military retirees and nonveterans. Military retirees have already spent at least 20 years working in support of the military mission and have thereby revealed a high taste for DoD employment. Indeed, our DRM estimates for military personnel show that mean taste increases with years of service because of selective retention: Lower-taste personnel separate and higher-taste personnel stay, while the SD of tastes decreases because personnel who stay have more-homogeneous tastes (Mattock, Asch, and Hosek, 2014). Consequently, we tested whether military veterans have a different taste distribution from that of nonveterans and, specifically, have a higher mean taste and lower SD of taste. We estimated the mean and SD of the taste distribution for nonveterans and included parameters that allow the mean and SD of the taste distribution for veterans to differ linearly. We also allowed the mean and SD of tastes to shift across entry cohorts separately for veterans and nonveterans in our combined models and entry cohort–specific models.

We found that the taste distributions for civil service differ between veterans and nonveterans. We estimated higher mean taste among veterans but more-homogeneous tastes (lower SD), and these estimated differences are statistically significant from 0. In addition, we found evidence that mean taste has changed across entry cohorts because some differences across entry cohorts in mean tastes are statistically significant.

Ideally, we would want to add parameters that allow the taste distribution to shift for different demographic groups. This would allow us to test whether the distribution differs across groups (e.g., baby boomers and generation X-ers). However, adding parameters also increases the computational burden of estimating the model. For this reason, we focused on differences in the taste distributions between veterans and nonveterans.

Modeling the Effects of Bonuses and Special Pays That Involve a Multi-Year Contract Choice

The services often offer bonuses linked to a multi-year obligation to service members who are in occupations with a high civilian opportunity wage, who are costly to train, or both. Modeling the behavior of members choosing multi-year contracts requires a model of the option value associated with being free to choose whether to stay or leave when new information is revealed; retention models where members are modeled as knowing with certainty when they will leave in the future (such as the ACOL) are unable to model the choice over contract length with fidelity (Mattock and Arkes, 2007). The computation of the expected value of the maximum in the DRM yields the option value, and, thus, the DRM is well-suited to modeling multi-year contracts. An ability to model multi-year contracts is useful in modeling the retention of Air Force pilots (Mattock et al., 2016), mental health care workers (Hosek et al., 2017) and special operations forces (Asch et al., 2019). We will discuss how we extended the DRM for Air Force pilots in some detail here; readers interested in the details of this extension for mental health care workers or special operations forces can consult the previously cited documents. Our discussion draws heavily on Mattock et al. (2016).

In Mattock et al. (2016), we extended the DRM in several ways, including by incorporating a new method to model the pilot's choice of a multi-year contract under the aviation bonus (AvB) program. Pilots who choose a longer contract receive AvB for more years, but they are also locked into their contract and forego the opportunity to take advantage of better opportunities that might present themselves during the contract period. The new method involved recognizing that the multi-year contract length choice is a nested choice made under uncertainty. The uncertainty arises from not knowing the specific future conditions (e.g., assignments, flying time, deployments) that accompany these choices. This extension requires estimation of an additional parameter in the model related to the variance of the shock associated with the multi-year contract choice. In this study, we found the parameter estimate to be statistically significant, indicating that this portrayal of the multi-year contract choice is an improved approach to modeling the AvB versus contract choice over that in Mattock and Arkes (2007).

Extending the DRM to Include the Aviator Bonus

Over the period covered by our data, Air Force pilots were eligible for multi-year contracts where they would be paid a retention AvB that typically increased with the length of the service commitment the individual elected. The availability of and rules governing eligibility for these multi-year contracts varied over time. Consequently, in our model, we incorporated AvB choice into both the estimation computer code and the simulation code.

Following Mattock and Arkes (2007), we extended the DRM to include the AvB choice by adding equations that express the value of the AvB program for different obligation lengths. The DRM previously described involves two equations: The first is the value of staying active while the second is the value of leaving, which is a nest of the reserve-civilian choice. Because our focus here is on the multi-year choice, while a member is on active duty, we will ignore the reserve-civilian nest aspect of the model and describe the value of leaving at time *t* simply as V_t^L.

The equation V_t^S gives the value of staying active for one additional year at time *t*. Thus, we can write the value of staying active for one more year as

$$V_t^{\frac{S}{1}} = V_t^{\frac{A}{1}} + \epsilon_t^A = \gamma_a + W_t^a + \beta Emax[V_{t+1}^L, V_{t+1}^S] + \epsilon_t^{\frac{A}{1}}$$

where W_t^a includes aviation incentive pay.

We can write the value of staying active and taking the AvB with a three-year obligation as

$$V_t^{\frac{S}{3}} = V_t^{\frac{A}{3}} + \epsilon_t^{\frac{A}{3}} = \sum_{n=0}^{2} \beta^n \left[\gamma_a + W_t^{\frac{a}{3}}\right] + \beta^3 Emax[V_{t+3}^L, V_{t+3}^S] + \epsilon_t^{\frac{A}{3}},$$

where $W_t^{\frac{a}{3}}$ includes AvB for the 3-year contract and AP.

Similarly, we can write the value staying active and taking AvB with a *k*-year obligation as

$$V_t^{\frac{S}{k}} = V_t^{\frac{A}{k}} + \epsilon_t^{\frac{A}{k}} = \sum_{n=0}^{k-1} \beta^n \left[\gamma_a + W_t^{\frac{a}{k}}\right] + \beta^k Emax[V_{t+k}^L, V_{t+k}^S] + \epsilon_t^{\frac{A}{k}}.$$

An eligible pilot compares the value of leaving V_t^L with the maximum of the value of staying for

one year, $V_t^{\frac{S}{1}}$,

three years, $V_t^{\frac{S}{3}}$,

five years, $V_t^{\frac{S}{5}}$,

or *k* years where *k* could be ten or more years in the case of an until-20 YOS option. If the 3-year, 5-year, and until-20 YOS options are offered, the probability that an initially offered pilot stays active is:

$$\Pr\left(max\left[V_t^{\frac{S}{1}}, V_t^{\frac{S}{3}}, V_t^{\frac{S}{5}}, V_t^{\frac{S}{20\text{YAS}}}\right] > V^L\right).$$

Similar to the reserve-civilian choice, the contract length choice can be handled as a nested choice. If we assume the random shocks of the contract length choice follow an extreme value distribution, then we can write

$$\epsilon_t^{\frac{A}{k}} \sim EV[-\phi\lambda_2, \lambda_2],$$

where λ_2 is the shape parameter and is subscripted with a 2 to distinguish it from the shape parameter associated with the within reserve-civilian nest shock, defined previously, which we will now denote as λ_1 (e.g., $\omega_t^R \sim EV[-\phi\lambda_1, \lambda_1]$). Thus, the AvB choice adds an additional parameter to be estimated. In the model without the contract length choice nest, the scale of the error in the value function for leaving was $\kappa = \sqrt{\lambda^2 + \tau^2}$, which we now relabel as

$$\kappa = \sqrt{\lambda_1^2 + \tau_1^2}.$$

By similar logic, the scale in the value function for staying with the contract length choice nest can be written

$$\sqrt{\lambda_2^2 + \tau_2^2}.$$

Imposing the requirement that the scales be equal, we have $\kappa = \sqrt{\lambda_2^2 + \tau_2^2}$. When estimating the model, we estimate κ, λ_1, and λ_2 and treat τ_1 and τ_2 as slack variables (that is, variables that are set to levels implied by the relationship between τ_1 and τ_2 and the other variables).

Similar to the reserve-civilian choice, service members might have the option to make multiple contract choices over their career. For example, they might choose a 1-year contract at first, then choose a 5-year contract, and then follow that with a 3-year contract before leaving. Because our data did not indicate which contract choice pilots made or the sequence of contract choices, we instead calculated the probability of observing a pilot staying a particular number of years and then leaving or being censored by summing up all possible sequences of contract decisions for the purposes of constructing the likelihood function. The method we used to compute the probability of all possible paths follows the logic in Mattock and Arkes (2007). As discussed in that paper, most paths have a near-zero probability. We exploited this fact in our calculations by noting that if one term of a product of probabilities is zero, the entire expression is zero. This saved us from having to explicitly calculate the other terms in the cumulative probability expression.

The multi-year contract extension enabled us to provide recommendations for AvB policy in Mattock et al. (2016), and also gave us the capability to assess the relative cost-effectiveness of retaining versus accessing Air Force pilots in Mattock et al. (2019). More generally, it also enabled us to consider other cases where bonus pays entail a multi-year contract, such as selective retention bonuses (SRB)s, and multi-year special pays for psychiatrists and other mental health care providers, as well as certain special pays for special operations forces.

Modeling the Transition from One Steady State to a New Steady State

In many cases, the DRM has been used to assess the effects of policy changes in the steady state. In the case of the military, where the typical military career is 30 years, it would take 30 years to reach the new steady state as a result of a policy change. However, policymakers are often concerned about the effects of a policy change in the transition to the steady state (i.e., during the 30-year period before the new steady state is reached) and how different implementation strategies can affect the 30-year time path. In this section, we discuss our extension to the DRM to enable modeling the transition to a new steady state. The following exposition draws heavily from Asch, Mattock, and Hosek (2013).

A common implementation strategy is to grandfather in existing members so only new entrants are covered by any policy change. Grandfathering is often desirable because policymakers do not want to break the implicit contract with existing service members. The problem with this approach is that it can take a long time before the effects of a policy are realized. Policymakers must wait until existing service members flow through and separate and new members get enough experience to be affected by the policy. One solution to this problem, as we describe in the following section, is to grandfather existing members but also give them the choice to switch to the new system. By offering a choice, the shift to a new policy allows service members under the existing policy to continue with it or, if they prefer, to opt for the new policy. More people will be under the new system more quickly if substantial numbers choose to switch, so it allows policymakers to move toward the steady state faster. Furthermore, faith has not been broken, and those who decide to change would do so only if they expect to be better off under the new policy.

Existing methodologies typically used to assess the transition phase and the effects of transition strategies are either severely limited or logically inconsistent. For example, personnel inventory projection models cannot be used to analyze the effects of allowing grandfathered members to switch to a new system because it does not include a model of decisionmaking that would logically allow members to change their behavior during the transition. Similarly, the so-called ACOL approach, in which estimates of retention responsiveness to pay are used to simulate the retention effects of pay changes over some time period, has been shown to be inconsistent with rational optimizing behavior and assumes away the possibility that individuals might change their minds when new information is revealed to them.

The DRM has neither of these disadvantages; it is logically consistent and can permit analysis of behavioral changes among incumbent members during the transition period. To do this, we extended the mathematical model that defines the DRM to incorporate a fourth time clock. This is in addition to the three time clocks that track years of service in the active component, years of service in the reserve component, and total years since initial accession or enlistment. Specifically, we add a clock that accounts for the service member's state when the policy occurs. We call this state the service member's cohort, defined by the service member's years of service when the policy change occurs. (Note that this cohort numbering convention differs from the usual, in that lower numbers are associated with newer cohorts and higher numbers with older cohorts.) We then use recent DRM parameter estimates to develop computer code that implements the extended model and permits us to simulate retention behavior for each cohort. Importantly, the extended model allows us to simulate both the retention behavior of each cohort over time and the retention behavior of all cohorts in the aggregate (i.e., a cross-section across all cohorts) for each time period since the policy change occurred. Thus, the total force can be observed in each period as a force planner or programmer might want to see it. We can simulate retention behavior in the 30-year transition.

Simulating how the retention curve evolves over time in the transition to a new steady state when a compensation policy change occurs requires that we do a year-by-year simulation of the retention curve starting from the year in which the policy changed to when the new steady state is reached. This simulation uses as input the cohort-specific retention survival curves described in the previous subsection.

Representing time properly requires careful attention. Consider the first year after the policy change is enacted. If s is time elapsed since the policy occurred, then $s = 0$ is the year when the policy occurred and $s = 1$ is the first year after the policy change. At time $s = 1$, those in YOS 2 were those who were in YOS 1 at $s = 0$. Similarly, those in YOS 3 at $s = 1$ were those in YOS 2 at $s = 0$. An alternative way to express this relationship is to say that at time $s = 1$, those in YOS 2 were in cohort = 1 at $s = 0$, and those in YOS 3 at $s = 1$ were in cohort = 2 at $s = 0$. Thus, the retention survival probability to YOS = t_a at $s = 1$ is the survival probability of the t_a-1 cohort, aged one year. More generally, the retention probability to YOS = t_a at s is the survival probability for cohort $c = t_a - s$, aged s years. Thus, elapsed time since the policy occurred is given by $s = t_a - c$.

We can illustrate how we select the relevant probabilities from each cohort retention survival curve as time elapses by considering Tables 5.1 and 5.2. The columns in Table 5.1 indicate YOS and are numbered from $t_a = 1$ to 30, where 30 is assumed to be the maximum length of an active component career. The rows indicate the cohort, defined by the YOS that the member is in when the policy occurs, also spanning from 1 to 30. Each cell in the table represents the cumulative probability of reaching YOS = t_a for cohort c. For example, the element in the first row, fifth column is the cumulative probability that an entrant in cohort 1 stays in active component service through YOS $t_a = 5$. Similarly, the elements in the last column are the cumulative probability that an entrant in given cohort stays for a full 30-year career.

To illustrate how we compile the retention survival curves over time as s increases in Table 5.1, we show only the time clock s associated with the cumulative probabilities in Table 5.1. The baseline steady state is given by $s = 0$. We compile the baseline steady state retention curve by selecting the diagonal elements in Table 5.1 and labeled *s* = 0 in bold font in Table 5.2. We compile the retention curve at *s* = 1 (one year elapsed after the policy occurs) as the first off-diagonal elements in Table 5.1—the cells labeled *s* = 1 in Table 5.2. The cells that are in light font in Table 5.2—the lower diagonal section of the table—are probabilities that are for career segments that occur before the policy was enacted. For example, years of service 1 through 3 for cohort = 4 (e.g., those with 4 YOS when the policy was enacted) occur prior to the policy. Therefore, we assign the baseline steady state retention survival probabilities to these cells. These are also denoted as *s* = 0 in Table 5.2 (in light font).

Table 5.1. How We Develop Retention Profiles by Time Elapsed: Cumulative Probabilities

	t_a = 1	t_a = 2	t_a = 3	t_a = 4	t_a = 5	…	t_a = 29	t_a = 30
c = 1	π(1,0,1,1)	π(2,0,2,1)	π(3,0,3,1)	π(4,0,4,1)	π(5,0,5,1)		π(29,0,29,1)	π(30,0,30,1)
c = 2	π(1,0,1,1)	π(2,0,2,2)	π(3,0,3,2)	π(4,0,4,2)	π(5,0,5,2)		π(29,0,29,2)	π(30,0,30,2)
c = 3	π(1,0,1,1)	π(2,0,2,2)	π(3,0,3,3)	π(4,0,4,3)	π(5,0,5,3)		π(29,0,29,3)	π(30,0,30,3)
c = 4	π(1,0,1,1)	π(2,0,2,2)	π(3,0,3,3)	π(4,0,4,4)	π(5,0,5,4)		π(29,0,29,4)	π(30,0,30,4)
c = 5	π(1,0,1,1)	π(2,0,2,2)	π(3,0,3,3)	π(4,0,4,4)	π(5,0,5,5)		π(29,0,29,5)	π(30,0,30,5)
…								
c = 29	π(1,0,1,1)	π(2,0,2,2)	π(3,0,3,3)	π(4,0,4,4)	π(5,0,5,5)		π(29,0,29,29)	π(30,0,30,29)
c = 30	π(1,0,1,1)	π(2,0,2,2)	π(3,0,3,3)	π(4,0,4,4)	π(5,0,5,5)		π(29,0,29,29)	π(30,0,30,30)

NOTE: The cells in the table are cumulative probabilities where each probability is given by π(ta, tr, tt, c). The columns are year of active service ta and the rows indicate cohort c.

Table 5.2. How We Develop Retention Profiles by Time Elapsed: Value of the *s* (Time-Elapsed)

	t_a = 1	t_a = 2	t_a = 3	t_a = 4	t_a = 5	t_a = 6	t_a = 7	…	t_a = 29	t_a = 30
c = 1	**s = 0**	s = 1	s = 2	s = 3	s = 4	s = 5	s = 6		s = 28	s = 29
c = 2	s = 0	**s = 0**	s = 1	s = 2	s = 3	s = 4	s = 5		s = 27	s = 28
c = 3	s = 0	s = 0	**s = 0**	s = 1	s = 2	s = 3	s = 4		s = 26	s = 27
c = 4	s = 0	s = 0	s = 0	**s = 0**	s = 1	s = 2	s = 3		s = 25	s = 26
c = 5	s = 0	s = 0	s = 0	s = 0	**s = 0**	s = 1	s = 2		s = 24	s = 25
c = 6	s = 0	s = 0	s = 0	s = 0	s = 0	**s = 0**	s = 1		s = 23	s = 24
c = 7	s = 0	s = 0	s = 0	s = 0	s = 0	s = 0	**s = 0**		s = 22	s = 23
…										
c = 29	s = 0	s = 0	s = 0	s = 0	s = 0	s = 0	s = 0	s = 0	**s = 0**	s = 1
c = 30	s = 0	s = 0	s = 0	s = 0	s = 0	s = 0	s = 0	s = 0	s = 0	**s = 0**

NOTE: The table cells only show s, the time-elapsed clock for the cumulative probability in Table 2.2 where s is defined as t_a-*c*. The columns are year of active service t_a and the rows indicate cohort c.

We have used the capability to analyze the transition to a new retirement system (Asch, Mattock, and Hosek, 2013), policies for drawing down the force (Mattock, Hosek, and Asch, 2016), the effect over time of making the reserve component retirement system more similar to the active component system (Mattock, Asch, and Hosek, 2014) and the dynamics of a pay freeze on DoD civil servants (Asch, Mattock, and Hosek, 2014), as well as in many other reports. The general approach we take to extending the DRM to the transition period between steady states should be readily adaptable to alternative structural models of retention.

Using Incumbents to Supplement Short Panels in Estimation

The model we described in the previous chapters was estimated using a population that we observed from the beginning of their careers. Fortunately, we had a panel for estimation of up to 26 years. But, in some applications, the longitudinal datasets are shorter. For example, data might not be available, as we found for certain state and local pension systems. Or, in the case of the military, a group is relatively new, such as a new occupation or force (e.g., the Space Force). The lack of a long time series can be challenging when the analysis is focusing on retirement behavior because there will be no observations of years during which, for example, a population might be retirement-eligible. (This need not be a completely debilitating problem; some initial empirical estimates of the military DRM were done using datasets where no one had yet reached retirement vesting at 20 YOS.)

We first extended the DRM to the use incumbents to supplement our longitudinal data when we were modeling teacher retention (Knapp, Brown, et al., 2016). The teacher retention model is a stay-versus-leave model, and, thus, the taste distribution is univariate normal (as opposed to the bivariate normal distribution in the model described in the previous chapter). The main concern in using data from incumbent teachers is that the teachers who are present are the result of a process that had resulted in a posterior population taste density that might be quite different from the initial taste density. Put another way, we cannot simply assume that the taste distribution of incumbent teachers is normal with the same mean and SD as initial entrants. What we need to do is find a way of computing the posterior taste distribution of incumbent teachers when we first observe them. To do so, we need to assume the taste distribution for entering cohorts is stationary over time, and that the retention process is also stationary (that is, the probability that a person is retained in a particular period given their taste is the same over time). Given these assumptions, we can compute the posterior taste distribution given an incumbent has been present for a set number of years. The discussion in the following section that describes how we go about this draws heavily on Asch, Knapp, and Mattock (2022).

Before discussing how we can use data from incumbent teachers, it will be helpful to discuss how we form the likelihood function for those teachers we observe for their entire careers. Given independent shocks in each period, the cumulative probability that teacher *i* will stay through service year $t-1$ may be written:[23]

$$cumulativePr(Stay)_{i,t} = \prod_{s=0}^{t-1} Pr_{i,s+1}(Stay).$$

The cumulative probability that teacher *i* stays for $t-1$ years and leaves at *t* is

$$cumulativePr(Leave)_{i,t} = \prod_{s=0}^{t-2} Pr_{i,s+1}(Stay)(1 - Pr_{i,t}(Stay)).$$

These probabilities are conditioned on the unobserved taste parameter γ. We assume the taste parameter has a normal distribution $g(\gamma)$ with mean μ and SD σ. We use this information to formulate the expected cumulative probability of a given career path, or the likelihood of that path. Thus, for teacher *i* in our data who stays through $t-1$ and leaves at *t*, the likelihood of that career path is

$$\mathcal{L}_i(\mu, \sigma, \lambda, \beta) = \int_{-\infty}^{\infty} \prod_{s=0}^{t-2} Pr_{i,s+1}(Stay)(1 - Pr_{i,t}(Stay))\, g(\gamma) d\gamma. \quad (1)$$

Similarly, if the individual stays through *t* and is then censored, the likelihood is

$$\mathcal{L}_i(\mu, \sigma, \lambda, \beta) = \int_{-\infty}^{\infty} \prod_{s=0}^{t-1} Pr_{i,s+1}(Stay)\, g(\gamma) d\gamma.$$

Thus, the likelihood for the entire data sample, *N*, is given by

$$\mathcal{L}(\mu, \sigma, \lambda, \beta) = \prod_{i=1}^{N} \mathcal{L}_i(\mu, \sigma, \lambda, \beta).$$

To consider incumbent teachers who entered prior to the period of observation we need to extend our model. For concreteness, let us consider South Carolina public school teachers. Our data included new entrants from 2008 to 2015 who were followed to 2020; however, such a sample provides no observations of years where individuals are retirement eligible. To augment the sample, we extended the DRM to allow inclusion of teachers who were incumbent in 2008, on whom we had longitudinal data from 2008 forward to 2020. The extension assumed their taste distribution at entry was the same as the taste distribution of the 2008–2015 new entrants. Under this assumption, we expressed their conditional taste distribution as of 2008 in terms of the new entrant taste distribution and the cumulative probability that individuals of a given taste who entered in years before 2008 and stayed until 2008. Like 2008–2015 new entrants, they were then followed forward to 2020 and in each year could choose to stay or leave.

[23] At entry, each teacher is assumed to decide to stay for the first period. In other words, when a teacher enters, it is assumed that the teacher has in effect decided to stay for the first period: $Pr_{i,1}(Stay) = 1$. Hence, the first stay-versus-leave decision occurs at the beginning of the second period.

The density of taste γ at the start of year of service t conditional on staying continuously from entry is

$$p(\gamma|s_0, s_1, \dots, s_{t-1}) = p(\gamma, s_0, s_1, \dots, s_{t-1})/p(s_0, s_1, \dots, s_{t-1}) =$$

$$\frac{p(s_0, s_1, \dots, s_{t-1}|\gamma)g(\gamma)}{p(s_0, s_1, \dots, s_{t-1})}. \quad (2)$$

Here, $p(s_0, s_1, \dots, s_{t-1}|\gamma)$ is the probability that a teacher stays continuously to complete $t-1$ years of service (i.e., stays to the beginning of period *t*) given a particular value of taste drawn at entry into teaching. As before, the density of taste for new entrants is $g(\gamma)$. The denominator, $p(s_0, s_1, \dots, s_{t-1})$, is the probability of staying continuously to complete $t-1$ YOS averaged over all values of taste (that is, taste is integrated out).

The DRM is a first-order Markov process, so the probability of staying in $t-1$ given that one has stayed continuously from entry through $t-2$ is just the probability of staying in $t-1$ given staying in $t-2$, and so forth. The expression in the numerator of (2) can then be written

$$p(s_0, s_1, \dots, s_{t-1}|\gamma) = p(s_{t-1}|\gamma)p(s_{t-2}|\gamma) \dots p(s_0|\gamma).$$

Also, the denominator in (2) is this probability averaged over taste:

$$p(s_0, s_1, \dots, s_{t-1}) = \int_{-\infty}^{\infty} p(s_{t-1}|\gamma)p(s_{t-2}|\gamma) \dots p(s_0|\gamma)g(\gamma)d\gamma.$$

These results imply that (2) can be written as

$$p(\gamma|s_0, s_1, \dots, s_{t-1}) = \frac{p(s_{t-1}|\gamma)p(s_{t-2}|\gamma)\dots p(s_0|\gamma)g(\gamma)}{\int_{-\infty}^{\infty} p(s_{t-1}|\gamma)p(s_{t-2}|\gamma)\dots p(s_0|\gamma)g(\gamma)d\gamma}.$$

The usefulness of this expression for the conditional probability of taste given some period of staying (left-hand side) comes from breaking it into a product of per-period stay probabilities of known form multiplied by the a priori taste distribution, also of known form (assumed to be normal), divided by an average value that can be computed from the same expressions.

Returning to our South Carolina example and using the conditional density of taste for an incumbent teacher's YOS as of 2008, we can construct probability expressions for the incumbent's retention decisions in years from 2008 forward in the same fashion as for new entrants, where the unconditional density of taste was used. For example, consider teachers who served continuously from entry and were making a stay-versus-leave decision at the beginning of YOS 20 in 2008. These teachers began in 1989 and had already completed 19 YOS. The conditional taste distribution for these teachers is

$$\frac{p(s_{19}|\gamma)p(s_{18}|\gamma)\dots p(s_0|\gamma)g(\gamma)}{\int_{-\infty}^{\infty} p(s_{19}|\gamma)p(s_{18}|\gamma)\dots p(s_0|\gamma)g(\gamma)d\gamma}.$$

In developing the likelihood for these teachers, this taste distribution was used in place of the original taste distribution $g(\gamma)$ in (1), and their retention decisions were tracked from 2008 through 2020, the last period observed in the data set.

So, if we are willing to assume stationarity in the taste distribution for initial entrants, we can derive expressions for the posterior distribution of taste that allow us to use observations on the retention of incumbents. This is advantageous when the available longitudinal data are short relative to the policies of interest (such as the structure of a retirement system).

Concluding Remarks

In this chapter, we showed four additional innovations of and extensions to the basic DRM model, from a relatively simple innovation (allowing regression variables to shift the mean and SD of the taste distribution) to the more complex (multi-year contracts, transition to a new steady state, and using information from incumbents). Throughout, we see that the rigor and richness of the DRM can support substantive innovations to address new challenges in a logically consistent and coherent manner.

Chapter 6. Conclusion

The DRM has proven to be a practical capability for modeling the retention of military personnel, civil service employees, public school teachers, and state employees. It uses a rigorous, logically consistent framework and has been successfully extended to assess multiple policies of interest, such as retirement reform, the structure of the pay table, and the structure of special and incentive pays. Here we will discuss what we have learned in the past, including the limitations of the DRM, and what we hope for the future from the DRM and succeeding models.

Looking Backward

The DRMs used in our analyses have several limitations. We will confine our remarks here to the military DRM, but similar remarks could apply to our other models. The DRM does not explicitly model other factors that can affect retention and retirement (including health status and health care benefits) or household factors (such as spousal labor supply or the presence of children at home). The analysis focuses on retention and does not model the decision to enlist or access into the military. Consequently, the model cannot address how changes to pension design might affect the types of people who become service members.

Another limitation is that the model assumes risk neutrality. The utility function is assumed to be linear in compensation. While, conceptually, a more flexible functional form could be used, practically speaking, the computational challenges are formidable.

Another limitation of the model is that it assumes that the individuals making decisions fully understand the implications of their decisions. That is, the model makes strong assumptions about the rationality of the individuals who are making retention decisions.

That said, the estimated models fit the observed data well. Our approach has several rich and realistic features that make it well suited for analyzing the retention effects of alternative compensation policies and pension reform. It is a life-cycle model where retention decisions are made each year over an entire career and not just once. Those decisions are based on forward-looking behavior that depends on existing and future military and external compensation. The model allows for uncertainty in future periods and recognizes that people might change their minds in the future as they get more information about staying in the military and their external opportunities. Furthermore, the model is formulated in terms of the parameters that underlie the retention decision processes rather than on the average responses to historical changes in policy. Consequently, it is structured to enable assessments of alternative compensation reforms that have yet to be tried. Put differently, the DRM is particularly suited to assess major structural changes in the compensation system that do not have any historical antecedent.

Looking Forward

The main challenge to expanding the DRM, either by taking into account additional demographic variables or by using a more flexible functional form to account for risk aversion, is computational. As the state space gets larger, the “curse of dimensionality” rears its ugly head, and models rapidly become unwieldy. Part of the art of modeling is being able to identify the key elements that need to be modeled and designing a parsimonious and computable model.

In addition, there are challenges associated with modeling the evolution of expectations of individuals when, for example, promotion rates change over time or special and incentive pays change. As of this writing, our models assume rational expectations, i.e., that individuals have perfect foresight of future conditions, but a more realistic mode of expectations formation is desirable.

Furthermore, all our models have so far been partial-equilibrium models of individual labor supply. Extending the model to consider demand-side factors, such as the impact of increased retention on promotion rates, time-to-promotion, and the future trajectory of military earnings is an area that should be explored.

However, there is reason for optimism. Estimating the DRM became feasible over the past four decades because of improvements in both software and hardware. While technically the DRM is *embarrassingly parallel* because the stochastic dynamic programming problems for each of the individual support points in the taste distribution can be solved in parallel, the individual stochastic dynamic programs are irreducible. However, new technologies in the form of quantum computing and artificial intelligence/machine learning–aided optimization might yet come to the rescue. Alternative estimation strategies, such as those used by Hotz and Miller (1993) and surveyed in Aguirregabiria and Mira (2010) might yet prove helpful in identifying fruitful starting values for estimating a DRM by more-conventional methods. Improvements in speed will aid in creating models that model expectations in a more realistic manner and the possibility of moving beyond purely partial-equilibrium models.

Appendix. Thematic RAND DRM Bibliography

This appendix provides a list of RAND studies organized by policy area. These studies were published on or after 2007 and all feature analyses based on DRMs in which the parameters were estimated using empirical data rather than calibrated.

Military Compensation

Studies for the Office of the Secretary of Defense (OSD)

OSD—Military Retirement Reform

This thread of research explored military retirement reform (Asch et al., 2008) and documented support provided to decisionmakers within DoD (Asch, Hosek, and Mattock, 2014) and the Military Compensation and Retirement Modernization Commission (Asch, Mattock, and Hosek, 2015) that culminated with work in support of OSD implementation of the BRS (Asch, Mattock, and Hosek, 2017). Of note: the active, civilian, and reserve DRM for enlisted members of the Army, Navy, Air Force and Marine Corps was first documented in Asch et al., 2008, later extended to officers in Mattock, Hosek, and Asch, 2012 (see under heading "OSD—Reserve Compensation and Retirement Reform"), and further extended to the U.S. Coast Guard in Asch, Mattock, and Hosek, 2017.

Asch, Beth J., James Hosek, and Michael G. Mattock, *Toward Meaningful Compensation Reform: Research in Support of DoD's Review of Military Compensation*, RAND Corporation, RR-501-OSD, 2014.

Asch, Beth J., James R. Hosek, Michael G. Mattock, and Christina Panis, *Assessing Compensation Reform: Research in Support of the 10th Quadrennial Review of Military Compensation*, RAND Corporation, MG-764-OSD, 2008.

Asch, Beth J., Michael G. Mattock, and James Hosek, *Reforming Military Retirement: Analysis in Support of the Military Compensation and Retirement Modernization Commission*, RAND Corporation, RR-1022-MCRMC, 2015.

Asch, Beth J., Michael G. Mattock, and James Hosek, *The Blended Retirement System: Retention Effects and Continuation Pay Cost Estimates for the Armed Services*, RAND Corporation, RR-1887-OSD/USCG, 2017.

OSD—Pay Table and Related Analyses

This thread of research used the DRM to the study the impact of alternative pay table policies on retention. Asch et al. (2016) used the DRM to show the effects of the 2007 policy changes, effects of the 2014 policy change, effects of reverting to the 30-year table, and cost and cost savings under a 40-year versus 30-year pay table. Asch et al. (2018) showed the effect of capping retired pay for senior (over 30 YOS) field-grade officers under the 40-year pay table, particularly for those officers with prior enlisted service. Asch, Mattock, and Tong (2020) simulated the effect of switching to a TIG pay table on retention, cost, retention of higher-ability members, and individual member effort supply.

Asch, Beth J., James Hosek, Jennifer Kavanagh, and Michael G. Mattock, *Retention, Incentives, and DoD Experience Under the 40-Year Military Pay Table*, RAND Corporation, RR-1209-OSD, 2016.

Asch, Beth J., Michael G. Mattock, James Hosek, and Patricia K. Tong, *Capping Retired Pay for Senior Field Grade Officers: Force Management, Retention, and Cost Effects*, RAND Corporation, RR-2251-OSD, 2018.

Asch, Beth J., Michael G. Mattock, and Patricia K. Tong, *Analysis of a Time-in-Grade Pay Table for Military Personnel and Policy Alternatives*, RAND Corporation, RR-A369-1, 2020.

OSD—Special and Incentive Pays

This thread of research focused on special and incentive (S&I) pays. The methodology developed to study multi-year contracts for Air Force pilots (Mattock, Hosek, and Asch, 2016; see under heading "Air Force") was applied to military mental health care providers (Hosek et al., 2017), as well as Army and Navy special operations forces (Asch et al., 2019). The Air Force pilot model was used as a means to compare retention incentives tied to a service obligation with a straight wage differential with no associated service obligation in Hosek, Mattock, and Asch (2019). Finally, the Air Force pilot model was used to assess the potential impact of paying the full rate of S&I pays to members of the reserve components in Marrone et al. (2022).

Asch, Beth J., Michael G. Mattock, James Hosek, and Shanthi Nataraj, *Assessing Retention and Special and Incentive Pays for Army and Navy Commissioned Officers in the Special Operations Forces*, RAND Corporation, RR-1796-OSD, 2019.

Hosek, James, Michael G. Mattock, and Beth J. Asch, *A Wage Differential Approach to Managing Special and Incentive Pay*, RAND Corporation, RR-2101-OSD, 2019.

Hosek, James, Shanthi Nataraj, Michael G. Mattock, and Beth J. Asch, *The Role of Special and Incentive Pays in Retaining Military Mental Health Care Providers*, RAND Corporation, RR-1425-OSD, 2017.

Marrone, James V., Michael G. Mattock, Beth J. Asch, and Hannah Acheson-Field, *Payment of the Full Rate of Special and Incentive Pays to Members of the Reserve Components*, RAND Corporation, RR-A669-1, 2022.

Mattock, Michael G., James Hosek, Beth J. Asch, and Rita T. Karam, *Retaining U.S. Air Force Pilots When the Civilian Demand for Pilots Is Growing*, RAND Corporation, RR-1455-AF, 2016.

OSD—Reserve Component Compensation and Retirement Reform

Two reports for OSD focused on reserve component compensation and retirement reform: Asch, Hosek, and Mattock (2013) and Mattock, Hosek, and Asch (2012). Mattock, Hosek, and Asch (2012) featured DRM estimates for both officers and enlisted members of the Army, Navy, Air Force, and Marine Corps; while the study was about reserve compensation, the estimated model produced results for both the active and reserve component. Asch, Hosek, and Mattock (2013) applied the (then) recently estimated model to reserve retirement reform.

Asch, Beth J., James Hosek, and Michael G. Mattock, *A Policy Analysis of Reserve Retirement Reform*, RAND Corporation, MG-378-OSD, 2013.

Mattock, Michael G., James Hosek, Beth J. Asch, *Reserve Participation and Cost Under a New Approach to Reserve Compensation: Research in Support of the 11th Quadrennial Review of Military Compensation*, RAND Corporation, MG-1153-OSD, 2012.

OSD—DRM Extensions

These two reports were devoted to efforts to extend the capability of the DRM. The first and most significant report (Asch, Mattock, and Hosek, 2013) is an extension to the DRM that permits the analysis of the transition from one steady state to a new steady state, recognizing that policy interventions often have differing effects on different cohorts and that it might take up to three to four decades until all members have spent their entire careers under the new policy. The second report (Mattock et al., 2014) uses a simple stay-versus-leave DRM of officer retention to demonstrate the effects of different compensation policies and also includes the complete source code for a Microsoft Excel implementation of the DRM that can be used to replicate the results in the report.

Asch, Beth J., Michael G. Mattock, and James Hosek, *A New Tool for Assessing Workforce Management Policies Over Time: Extending the Dynamic Retention Model*, RAND Corporation, RR-113-OSD, 2013.

Mattock, Michael G., Beth J. Asch, James Hosek, Christopher Whaley, and Christina Panis, *Toward Improved Management of Officer Retention: A New Capability for Assessing Policy Options*, RAND Corporation, RR-764, 2014.

OSD—Other

These two papers used previously estimated DRMs to explore policy options. The first, Mattock, Hosek, and Asch (2016), used the transition model first introduced in Asch, Mattock, and Hosek (2013) (see heading "OSD—DRM Extensions") to (1) examine alternative voluntary separation incentives for efficiently drawing down the force, (2) calculate optimal incentive pays, and (3) allow for anticipation that the incentive pays would be offered. The second, Rennane et al. (2022) evaluated using the value of a lost military career as calculated by taking the difference of the DRM value functions (the value of staying minus the value of leaving) to compensate disabled veterans with a military career–ending disability. This use of the value functions to calculate the value of a lost military career was first done in support of DoD's review of military compensation and is documented in Asch, Hosek, and Mattock (2014) (see under heading "OSD—Military Retirement Reform").

Mattock, Michael G., James Hosek, and Beth J. Asch, *Policies for Managing Reductions in Military End Strength: Using Incentive Pays to Draw Down the Force*, RAND Corporation, RR-545-OSD, 2016.

Rennane, Stephanie, Beth J. Asch, Michael G. Mattock, Heather Krull, Douglas C. Ligor, Michael Dworsky, and Jonas Kempf, *U.S. Department of Defense Disability Compensation Under a Fitness-for-Duty Evaluation Approach*, RAND Corporation, RR-A1154-1, 2022.

Studies for the Department of the Army

The thread of research for the Army started with an examination of the effect of making the reserve component retirement system more like the active component retirement system (Mattock, Asch, and Hosek, 2014) and then branched out to examine possible alternative (that is, service-specific) retirement accrual charges that would take into account the differing retention patterns between officers and enlisted and among the services in Hosek, Asch, and Mattock (2017). Asch, Mattock, and Hosek (2019) extended the Army DRM to separately model the U.S. Army Reserve and the Army National Guard; previously, the reserve component had been treated as a single composite. The model was used to evaluate the retention effect of the BRS on the Army Reserve. Asch et al. (2021) used DRMs estimated at the individual military occupation specialty level to evaluate the effectiveness of alternative SRB policies in retaining those members with higher innate ability. Finally, Calkins et al. (2023) modeled the retention of Army aviators (both commissioned officers and warrant officers) and used the model to evaluate alternative incentive pays based on reaching career milestones. Similar to the earlier work on SRBs, Calkins et al. (2023) examined the effect of alternative incentive pays in retaining individuals with higher ability.

Asch, Beth J., Michael G. Mattock, and James Hosek, *Effects of the Blended Retirement System on United States Army Reserve Participation and Cost*, RAND Corporation, RR-2591-A, 2019.

Asch, Beth J., Michael G. Mattock, Patricia K. Tong, and Jason M. Ward, *Increasing Efficiency and Incentives for Performance in the Army's Selective Reenlistment Bonus (SRB) Program*, RAND Corporation, RR-A803-1, 2021.

Calkins, Avery, Michael G. Mattock, Beth J. Asch, Ryan Schwankhart, and Tara L. Terry, *Army Aviation Special and Incentive Pay Policies to Promote Performance, Manage Talent, and Sustain Retention*, RAND Corporation, RR-A2234-1, 2023.

Hosek, James, Beth J. Asch, and Michael G. Mattock, *Toward Efficient Military Retirement Accrual Charges*, RAND Corporation, RR-1373-A, 2017.

Mattock, Michael G., Beth J. Asch, and James Hosek, *Making the Reserve Retirement System Similar to the Active System: Retention and Cost Estimates*, RAND Corporation, RR-530-A, 2014.

Studies for the Department of the Air Force

The need to study Air Force pilot compensation and retention drove the initial RAND work that used the DRM to model multi-year contracts. Mattock and Arkes (2007) used a Gotz and McCall (1984) style stay-versus-leave model to examine multi-year contracts for Air Force pilots. Mattock et al. (2016) marked a substantial improvement over the initial model, including the reserve component in the analysis and allowing for the choice among contracts to be subject to uncertainty. Robbert et al. (2018) used the DRM to model the retention for two supplemental career tracks, either a warrant officer track or an aviation technical track for commissioned officers. Mattock et al. (2019) looked at whether retaining more Air Force pilots was more cost-effective than accessing and training additional pilots. Mattock and Asch (2019) used the DRM to assess the efficiency of a novel pay proposal by the Air Force relative to the existing aviation bonus program. Tong, Mattock, and Asch (2021) and Tong et al. (2020) estimated a set of DRMs for selected career enlisted aviator occupations and assessed the trade-off between accessing or retaining more career enlisted aviators via S&I pays. Most recently Robbert, Tong, and Hardison (2022) used the DRM to examine how the retention behavior of enlisted maintenance, logistics, and munitions personnel might be influenced by the BRS when they reach their retention decision points.

Mattock, Michael G., and Jeremy Arkes, *The Dynamic Retention Model for Air Force Officers: New Estimates and Policy Simulations of the Aviator Continuation Pay Program*, RAND Corporation, TR-470, 2007.

Mattock, Michael G., and Beth J. Asch, *An Initial Look at the U.S. Air Force Aviation Professional Pay Proposal*, RAND Corporation, PE-309-AF, 2019.

Mattock, Michael G., Beth J. Asch, James Hosek, and Michael Boito, *The Relative Cost-Effectiveness of Retaining Versus Accessing Air Force Pilots*, RAND Corporation, RR-2415-AF, 2019.

Mattock, Michael G., James Hosek, Beth J. Asch, and Rita Karam, *Retaining U.S. Air Force Pilots When the Civilian Demand for Pilots Is Growing*, RAND Corporation, RR-1455-AF, 2016.

Robbert, Albert A., Michael G. Mattock, Beth J. Asch, John S. Crown, James Hosek, and Tara L. Terry, *Supplemental Career Paths for Air Force Pilots: A Warrant Officer Component or an Aviation Technical Track?* RAND Corporation, RR-2617-AF, 2018.

Robbert, Albert A., Patricia K. Tong, and Chaitra M. Hardison, *Retention of Enlisted Maintenance, Logistics, and Munitions Personnel: Analysis and Results,* RAND Corporation, RR-A546-1, 2022.

Tong, Patricia K., Michael G. Mattock, and Beth J. Asch, *Cost-Benefit Analysis of Special and Incentive Pays for Career Enlisted Aviators*, RAND Corporation, RR-A189-1, 2021.

Tong, Patricia K., Michael G. Mattock, Beth J. Asch, James Hosek, and Felix Knutson, *Modeling Career Enlisted Aviator Retention in the U.S. Air Force*, RAND Corporation, RR-3134-AF, 2020.

DoD Civilians—Studies for OSD

RAND work applying the DRM to DoD civil service personnel started with Asch, Mattock, and Hosek (2014), which features a stay-versus-leave model estimated on a single cohort. The authors applied the model to examine the impact of pay freezes and unpaid furloughs. Knapp, Asch, et al. (2016) expanded the model to multiple cohorts and to consider the behavior of veterans compared with nonveterans, and Asch et al. (2016) used the estimated model to examine the use of voluntary separation incentive programs versus involuntary separations. Most recently, Mattock et al. (2022) used a DRM estimated on DoD civilian cyber workers to examine retention responsiveness to enhanced training opportunities.

Asch, Beth J., James Hosek, Michael G. Mattock, David Knapp, and Jennifer Kavanagh, *Workforce Downsizing and Restructuring in the Department of Defense: The Voluntary Separation Incentive Payment Program Versus Involuntary Separation*, RAND Corporation, RR-1540-OSD, 2016.

Asch, Beth J., Michael G. Mattock, and James Hosek, *The Federal Civil Service Workforce: Assessing the Effects on Retention of Pay Freezes, Unpaid Furloughs, and Other Federal-Employee Compensation Changes in the Department of Defense*, RAND Corporation, RR-514-OSD, 2014.

Knapp, David, Beth J. Asch, Michael G. Mattock, and James Hosek, *An Enhanced Capability to Model How Compensation Policy Affects U.S. Department of Defense Civil Service Retention and Cost*, RAND Corporation, RR-1503-OSD, 2016.

Mattock, Michael G., Beth J. Asch, Avery Calkins, and Daniel Schwam, *Civilian Cyber Workers in the U.S. Department of Defense: Demographics, Retention, and Responsiveness to Training Opportunities*, RAND Corporation, RR-A730-3, 2022.

Public School Teachers and State Employees

RAND's initial foray into studying public school teacher retention was with the Chicago Public Schools in Knapp, Brown, et al. (2016). Knapp et al. (2018) explores the potential effect of a voluntary retirement incentive on Chicago Public School teacher retention. Knapp, Asch, and Mattock (2021) extends the DRM to state employees in addition to teachers in South Carolina. Asch, Knapp, and Mattock (2022) estimates the same DRM on teachers from three states with substantially different retirement plans and finds that the model works well in modeling each state, even though the retention profiles differ from state to state. Knapp et al. (2022) and Hosek et al. (2023) are two recent journal publications documenting RAND research on the Chicago Public School voluntary retirement incentive plan.

Asch, Beth J., David Knapp, and Michael G. Mattock, *The Effects of Public Sector Retirement Plan Reform on Workforce Retention: Evidence from Teachers in Three States*, RAND Corporation, WR-A816-2, 2022.

Hosek, James, David Knapp, Michael G. Mattock, and Beth J. Asch, "Incentivizing Retirement: An Analysis of Cash Retirement Incentives for Chicago Teachers," *Educational Researcher*, Vol. 52, No. 2, March 2023.

Knapp, David, Beth J. Asch, and Michael G. Mattock, *Public Employee Retention Responses to Alternative Retirement Plan Design: South Carolina Teachers and State Public Employees*, RAND Corporation, WR-A816-1, 2021.

Knapp, David, Kristine Brown, James Hosek, Michael G. Mattock, and Beth J. Asch, *Retirement Benefits and Teacher Retention: A Structural Modeling Approach*, RAND Corporation, RR-1448-RC, 2016.

Knapp, David, James Hosek, Michael G. Mattock, and Beth J. Asch, “Predicting Teacher Retention Behavior: Ex Ante Prediction and Ex Post Realization of a Voluntary Retirement Incentive Offer,” *Economics of Education Review*, Vol. 93, April 2023.

Abbreviations

ACOL	Annualized Cost of Leaving (model)
ADSO	active-duty service obligation
AvB	aviation bonus
BFGS	Broyden-Fletcher-Goldfarb-Shanno (algorithm)
BRS	Blended Retirement System
DACMC	Defense Advisory Committee on Military Compensation
DMDC	Defense Manpower Data Center
DoD	Department of Defense
DRM	dynamic retention model
EM	expectation-maximization
GEM	Generalized Expectation Maximization
IT	information technology
QRMC	Quadrennial Review of Military Compensation
RMC	regular military compensation
ROTC	Reserve Officer Training Corps
S&I	special and incentive
SD	standard deviation
SRB	selective retention bonus
TIG	time-in-grade
TIS	time-in-service
WEX	Work Experience File
YOS	year(s) of service

References

Aguirregabiria, Victor, and Pedro Mira, "Dynamic Discrete Choice Structural Models: A Survey," *Journal of Econometrics*, Vol. 156, No. 1, May 2010.

Asch, Beth J., and James Hosek, *Military Compensation: Trends and Policy Options*, RAND Corporation, DB-273-OSD, 1999. As of June 17, 2008: http://www.rand.org/pubs/documented_briefings/DB273/

Asch, Beth J., James Hosek, Jennifer Kavanagh, and Michael G. Mattock, *Retention, Incentives, and DoD Experience Under the 40-Year Military Pay Table*, RAND Corporation, RR-1209-OSD, 2016. As of July 26, 2023: https://www.rand.org/pubs/research_reports/RR1209.html

Asch, Beth J., James Hosek, and Michael G. Mattock, *A Policy Analysis of Reserve Retirement Reform*, RAND Corporation, MG-378-OSD, 2013. As of July 26, 2023: https://www.rand.org/pubs/monographs/MG378.html

Asch, Beth J., James Hosek, and Michael G. Mattock, *Toward Meaningful Compensation Reform: Research in Support of DoD's Review*, RAND Corporation, RR-501-OSD, 2014. As of July 26, 2023: https://www.rand.org/pubs/research_reports/RR501.html

Asch, Beth J., James Hosek, Michael G. Mattock, David Knapp, and Jennifer Kavanagh, *Workforce Downsizing and Restructuring in the Department of Defense: The Voluntary Separation Incentive Payment Program Versus Involuntary Separation*, RAND Corporation, RR-1540-OSD, 2016. As of July 26, 2023: https://www.rand.org/pubs/research_reports/RR1540.html

Asch, Beth J., James R. Hosek, Michael G. Mattock, and Christina Panis, *Assessing Compensation Reform: Research in Support of the 10th Quadrennial Review of Military Compensation*, RAND Corporation, MG-764-OSD, 2008. As of July 26, 2023: https://www.rand.org/pubs/monographs/MG764.html

Asch, Beth J., Richard Johnson, and John T. Warner, *Reforming the Military Retirement System*, RAND Corporation, MR-748-OSD, 1998. As of June 17, 2008: http://www.rand.org/pubs/monograph_reports/MR748/

Asch, Beth J., David Knapp, and Michael G. Mattock, *The Effects of Public Sector Retirement Plan Reform on Workforce Retention: Evidence from Teachers in Three States,* RAND Corporation, WR-A816-2, 2022. As of January 11, 2023: https://www.rand.org/pubs/working_papers/WRA816-2.html

Asch, Beth J., Michael G. Mattock, and James Hosek, *A New Tool for Assessing Workforce Management Policies Over Time: Extending the Dynamic Retention Model,* RAND Corporation, RR-113-OSD, 2013. As of July 26, 2023: https://www.rand.org/pubs/research_reports/RR113.html

Asch, Beth J., Michael G. Mattock, and James Hosek, *The Federal Civil Service Workforce: Assessing the Effects on Retention of Pay Freezes, Unpaid Furloughs, and Other Federal-Employee Compensation Changes in the Department of Defense*, RAND Corporation, RR-514-OSD, 2014. As of July 26, 2023: https://www.rand.org/pubs/research_reports/RR514.html

Asch, Beth J., Michael G. Mattock, and James Hosek, *Reforming Military Retirement: Analysis in Support of the Military Compensation and Retirement Modernization Commission*, RAND Corporation, RR-1022-MCRMC, 2015. As of July 26, 2023: https://www.rand.org/pubs/research_reports/RR1022.html

Asch, Beth, Michael Mattock, and James Hosek, *The Blended Retirement System: Retention Effects and Continuation Pay Cost Estimates for the Armed Services*, RAND Corporation, RR-1887-OSD/USCG, 2017. As of July 26, 2023: https://www.rand.org/pubs/research_reports/RR1887.html

Asch, Beth J., Michael G. Mattock, and James Hosek, *Effects of the Blended Retirement System on United States Army Reserve Participation and Cost*, RAND Corporation, RR-2591-A, 2019. As of February 4, 2020: https://www.rand.org/pubs/research_reports/RR2591.html

Asch, Beth J., Michael G. Mattock, James Hosek, and Shanthi Nataraj, *Assessing Retention and Special and Incentive Pays for Army and Navy Commissioned Officers in the Special Operations Forces*, RAND Corporation, RR-1796-OSD, 2019. As of October 17, 2022: https://www.rand.org/pubs/research_reports/RR1796.html

Asch, Beth J., Michael G. Mattock, James Hosek, and Patricia K. Tong, *Capping Retired Pay for Senior Field Grade Officers: Force Management, Retention, and Cost Effects*, RAND Corporation, RR-2251-OSD, 2018. As of May 25, 2018: https://www.rand.org/pubs/research_reports/RR2251.html

Asch, Beth J., Michael G. Mattock, and Patricia K. Tong, *Analysis of a Time-in-Grade Pay Table for Military Personnel and Policy Alternatives*, RAND Corporation, RR-A369-1, 2020. As of October 17, 2022: https://www.rand.org/pubs/research_reports/RRA369-1.html

Asch, Beth J., Michael G. Mattock, Patricia K. Tong, and Jason M. Ward, *Increasing Efficiency and Incentives for Performance in the Army's Selective Reenlistment Bonus (SRB) Program*, RAND Corporation, RR-A803-1, 2021. As of October 7, 2022: https://www.rand.org/pubs/research_reports/RRA803-1.html

Asch, Beth J., and John T. Warner, *A Policy Analysis of Alternative Military Retirement Systems*, RAND Corporation, MR-465-OSD, 1994a. As of August 28, 2019: https://www.rand.org/pubs/monograph_reports/MR465.html

Asch, Beth J., and John T. Warner, *A Theory of Military Compensation and Personnel Policy*, RAND Corporation, MR-439-OSD, 1994b. As of April 28, 2020: https://www.rand.org/pubs/monograph_reports/MR439.html

Asch, Beth J., and John T. Warner, "A Theory of Compensation and Personnel Policy in Hierarchical Organizations with Application to the United States Military," *Journal of Labor Economics*, Vol. 19, No. 3, July 2001.

Ausink, John, and David A. Wise, "The Military Pension, Compensation, and Retirement of U.S. Air Force Pilots," in David A. Wise, ed., *Advances in the Economics of Aging*, University of Chicago Press, 1996.

Blundell, Richard, Eric French, and Gemma Tetlow, "Retirement Incentives and Labor Supply," in John Piggott and Alan Woodland, eds., *Handbook of the Economics of Population Aging*, Vol. 1, Elsevier, 2016.

Calkins, Avery, Michael G. Mattock, Beth J. Asch, Ryan Schwankhart, and Tara L. Terry, *Army Aviation Special and Incentive Pay Policies to Promote Performance, Manage Talent, and Sustain Retention*, RAND Corporation, RR-A2234-1, 2023. As of October 2, 2023: https://www.rand.org/pubs/research_reports/RRA2234-1.html

Cooley, T. F., "Calibrated Models," *Oxford Review of Economic Policy*, Vol. 13, No. 3, Autumn 1997.

Daula, Thomas, and Robert Moffitt, "Estimating Dynamic Models of Quit Behavior: The Case of Military Reenlistment," *Journal of Labor Economics*, Vol. 13, No. 3, July 1995.

Dempster, A. P., N. M. Laird, and D. B. Rubin, "Maximum Likelihood from Incomplete Data via the EM Algorithm," *Journal of the Royal Statistical Society, Series B (Methodological)*, Vol. 39, No. 1, 1977.

Gotz, Glenn A., and John McCall, *A Dynamic Retention Model for Air Force Officers: Theory and Estimates*, RAND Corporation, R-3028-AF, 1984. As of June 17, 2008: http://www.rand.org/pubs/reports/R3028/

Hosek, James, Beth J. Asch, and Michael G. Mattock, *Toward Efficient Military Retirement Accrual Charges*, RAND Corporation, RR-1373-A, 2017. As of July 27, 2023: https://www.rand.org/pubs/research_reports/RR1373.html

Hosek, James, David Knapp, Michael G. Mattock, and Beth J. Asch, "Incentivizing Retirement: An Analysis of Cash Retirement Incentives for Chicago Teachers," *Educational Researcher*, Vol. 52, No. 2, March 2023.

Hosek, James, Michael G. Mattock, and Beth J. Asch, *A Wage Differential Approach to Managing Special and Incentive Pay*, RAND Corporation, RR-2101-OSD, 2019. As of February 4, 2020: https://www.rand.org/pubs/research_reports/RR2101.html

Hosek, James, Michael G. Mattock, C. Christine Fair, Jennifer Kavanagh, Jennifer Sharp, and Mark E. Totten, *Attracting the Best: How the Military Competes for Information Technology Personnel*, RAND Corporation, MG-108, 2004. As of July 27, 2023: https://www.rand.org/pubs/monographs/MG108.html

Hosek, James, Shanthi Nataraj, Michael G. Mattock, and Beth J. Asch, *The Role of Special and Incentive Pays in Retaining Military Mental Health Care Providers*, RAND Corporation, RR-1425-OSD, 2017. As of July 27, 2023: https://www.rand.org/pubs/research_reports/RR1425.html

Hotz, V. Joseph, and Robert A. Miller, "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, Vol. 60, No. 3, July 1993.

Knapp, David, Beth J. Asch, and Michael G. Mattock, *Public Employee Retention Responses to Alternative Retirement Plan Design: South Carolina Teachers and State Public Employees*, RAND Corporation, WR-A816-1, 2021. As of October 7, 2022: https://www.rand.org/pubs/working_papers/WRA816-1.html

Knapp, David, Beth J. Asch, Michael G. Mattock, and James Hosek, *An Enhanced Capability to Model How Compensation Policy Affects U.S. Department of Defense Civil Service Retention and Cost*, RAND Corporation, RR-1503-OSD, 2016. As of July 27, 2023: https://www.rand.org/pubs/research_reports/RR1503.html

Knapp, David, Kristine M. Brown, James Hosek, Michael G. Mattock, and Beth J. Asch, *Retirement Benefits and Teacher Retention: A Structural Modeling Approach*, RAND Corporation, RR-1448-RC, 2016. As of July 27, 2023: https://www.rand.org/pubs/research_reports/RR1448.html

Knapp, David, James Hosek, Michael G. Mattock, and Beth J. Asch, *Exploring Voluntary Retirement Incentives for Teachers: Effects on Retention and Cost in Chicago Public Schools*, RAND Corporation, WR-1249, 2018. As of July 18, 2018: https://www.rand.org/pubs/working_papers/WR1249.html

Knapp, David, James Hosek, Michael G. Mattock, and Beth J. Asch, “Predicting Teacher Retention Behavior: Ex Ante Prediction and Ex Post Realization of a Voluntary Retirement Incentive Offer,” *Economics of Education Review*, Vol. 93, April 2023.

Marrone, James V., Michael G. Mattock, Beth J. Asch, and Hannah Acheson-Field, *Payment of the Full Rate of Special and Incentive Pays to Members of the Reserve Components*, RAND Corporation, RR-A669-1, 2022. As of October 7, 2022: https://www.rand.org/pubs/research_reports/RRA669-1.html

Mattock, Michael G., and Jeremy Arkes, *The Dynamic Retention Model for Air Force Officers: New Estimates and Policy Simulations of the Aviator Continuation Pay Program*, RAND Corporation, TR-470-AF, 2007. As of July 27, 2023: https://www.rand.org/pubs/technical_reports/TR470.html

Mattock, Michael G., and Beth J. Asch, *An Initial Look at the U.S. Air Force Aviation Professional Pay Proposal*, RAND Corporation, PE-309-AF, 2019. As of February 4, 2020: https://www.rand.org/pubs/perspectives/PE309.html

Mattock, Michael G., Beth J. Asch, Avery Calkins, and Daniel Schwam, *Civilian Cyber Workers in the U.S. Department of Defense: Demographics, Retention, and Responsiveness to Training Opportunities*, RAND Corporation, RR-A730-3, 2022. As of October 7, 2022: https://www.rand.org/pubs/research_reports/RRA730-3.html

Mattock, Michael G., Beth J. Asch, and James Hosek, *Making the Reserve Retirement System Similar to the Active System: Retention and Cost Estimates*, RAND Corporation, RR-530-A, 2014. As of July 27, 2023: https://www.rand.org/pubs/research_reports/RR530.html

Mattock, Michael G., Beth J. Asch, James Hosek, and Michael Boito, *The Relative Cost-Effectiveness of Retaining Versus Accessing Air Force Pilots*, RAND Corporation, RR-2415-AF, 2019. As of February 4, 2020: https://www.rand.org/pubs/research_reports/RR2415.html

Mattock, Michael G., Beth J. Asch, James Hosek, Christopher M. Whaley, and Christina Panis, *Toward Improved Management of Officer Retention: A New Capability for Assessing Policy Options,* RAND Corporation, RR-764, 2014. As of July 27, 2023: https://www.rand.org/pubs/research_reports/RR764.html

Mattock, Michael, James Hosek, and Beth J. Asch, *Reserve Participation and Cost Under a New Approach to Reserve Compensation*, RAND Corporation, MG-1153-OSD, 2012. As of July 27, 2023: https://www.rand.org/pubs/monographs/MG1153.html

Mattock, Michael G., James Hosek, and Beth J. Asch, *Policies for Managing Reductions in Military End Strength: Using Incentive Pays to Draw Down the Force*, RAND Corporation, RR-545-OSD, 2016. As of July 27, 2023: https://www.rand.org/pubs/research_reports/RR545.html

Mattock, Michael G., James Hosek, Beth J. Asch, and Rita T. Karam, *Retaining U.S. Air Force Pilots When the Civilian Demand for Pilots Is Growing*, RAND Corporation, RR-1455-AF, 2016. As of July 27, 2023: https://www.rand.org/pubs/research_reports/RR1455.html

Neal, Radford M., and Geoffrey E. Hinton, "A View of the EM Algorithm That Justifies Incremental, Sparse, and Other Variants," in Michael I. Jordan, ed., *Learning in Graphical Models*, MIT Press, 1998.

Office of the Under Secretary of Defense for Personnel and Readiness, Directorate of Compensation, *Selected Military Compensation Tables*, 1980–2018.

Piessens, Robert, Elise Doncker-Kapenga, Christoph W. Überhuber, and David K. Kahaner, *Quadpack: A Subroutine Package for Automatic Integration*, Vol. 1, Springer Science & Business Media, 1983.

Rennane, Stephanie, Beth J. Asch, Michael G. Mattock, Heather Krull, Douglas C. Ligor, Michael Dworsky, and Jonas Kempf, *U.S. Department of Defense Disability Compensation Under a Fitness-for-Duty Evaluation Approach*, RAND Corporation, RR-A1154-1, 2022. As of October 17, 2022: https://www.rand.org/pubs/research_reports/RRA1154-1.html

Robbert, Albert A., Michael G. Mattock, Beth J. Asch, John S. Crown, James Hosek, and Tara L. Terry, *Supplemental Career Paths for Air Force Pilots: A Warrant Officer Component or an Aviation Technical Track?* RAND Corporation, RR-2617-AF, 2018. As of February 4, 2020: https://www.rand.org/pubs/research_reports/RR2617.html

Robbert, Albert A., Patricia K. Tong, and Chaitra M. Hardison, *Retention of Enlisted Maintenance, Logistics, and Munitions Personnel: Analysis and Results,* RAND Corporation, RR-A546-1, 2022. As of February 24, 2023: https://www.rand.org/pubs/research_reports/RRA546-1.html

Stafford, Darlene E., and Henry S. Griffis, *A Review of Millennial Generation Characteristics and Military Workforce Implications*, CNA, CRM D0018211.A1/Final, May 2008.

Tong, Patricia K., Michael G. Mattock, and Beth J. Asch, *Cost-Benefit Analysis of Special and Incentive Pays for Career Enlisted Aviators*, RAND Corporation, RR-A189-1, 2021. As of October 17, 2022: https://www.rand.org/pubs/research_reports/RRA189-1.html

Tong, Patricia K., Michael G. Mattock, Beth J. Asch, James Hosek, and Felix Knutson, *Modeling Career Enlisted Aviator Retention in the U.S. Air Force*, RAND Corporation, RR-3134-AF, 2020. As of October 17, 2022: https://www.rand.org/pubs/research_reports/RR3134.html

Train, Kenneth, *Discrete Choice Methods with Simulation*, 2nd ed., Cambridge University Press, 2009.

Under Secretary of Defense for Personnel and Readiness, *Military Compensation Background Papers: Compensation Elements and Related Manpower Cost Items: Their Purpose and Legislative Background*, 8th ed., U.S. Department of Defense, July 2018.