



INSTITUTE FOR DEFENSE ANALYSES

Implementing Dynamic Discrete Choice Models of Military Retention

Jacklyn Kambic
Mikhail Smirnov
John W. Dennis
Alan B. Gelder

September 2023
Approved for public release;
distribution unlimited.
IDA Paper P-33602
Log: H 23-000282

INSTITUTE FOR DEFENSE ANALYSES
730 E. Glebe Rd
Alexandria, VA 22305



The Institute for Defense Analyses is a nonprofit corporation that operates three Federally Funded Research and Development Centers. Its mission is to answer the most challenging U.S. security and science policy questions with objective analysis, leveraging extraordinary scientific, technical, and analytic expertise.

About This Publication

The work was conducted by the Institute for Defense Analyses (IDA) under CRP C6601. The views, opinions, and findings should not be construed as representing the official position of either the Department of Defense or the sponsoring organization.

Acknowledgments

The authors wish to thank Cullen A. Roberts, R. Abraham Holland, Valerie K. Bullock, various other members of the FFRDC community, including participants of the Dynamic Discrete Choice Workshop held at RAND, and the publications team Joseph F. Coombs, Amberlee Mabe-Stanberry, and Flo Purnell for their thoughtful review.

For More Information:

Dr. John W. Dennis III, Project Leader
jdennis@ida.org, 703-845-2166

Ms. Jessica L. Stewart, Director, SFRD
jstewart@ida.org, 703-575-4530

Copyright Notice

© 2023 Institute for Defense Analyses
730 E. Glebe Rd
Alexandria, VA 22305 • (703) 845-2000

This material may be reproduced by or for the U.S. Government pursuant to the copyright license under the clause at DFARS 252.227-7013 (Feb. 2014).

INSTITUTE FOR DEFENSE ANALYSES

IDA Paper P-33602

Implementing Dynamic Discrete Choice Models of Military Retention

Jacklyn Kambic
Mikhail Smirnov
John W. Dennis
Alan B. Gelder

This page is intentionally blank.

Executive Summary

Military manpower and personnel policy analyses can benefit greatly from the use of structural dynamic discrete choice (DDC) models. Policy makers design compensation and personnel policies with specific objectives in mind, such as meeting a retention target in a given occupation or increasing the average length of service obligation. To evaluate whether a proposed policy would achieve its objective, we need models that can tell us how service members would behave under the new policy.

Structural DDC models can do that, because they specify how compensation and other individual and career characteristics affect service members' decisions. This formulation enables researchers to predict retention outcomes under a new policy. Retention predictions can be used to assess the effect a proposed policy would have on the military's ability to meet manpower and readiness objectives.

In contrast with machine learning (ML) models that emphasize prediction, structural DDC models are designed to uncover causal relationships. In the case of military personnel policy analysis, a DDC model can be specified to identify the causal effect of a policy lever (such as a retention bonus) on service members' retention decisions. Predictive ML models and structural DDC models alike can be used to predict future outcomes under the assumption that no policy changes occur that would affect service members' decisions. However, DDC models can also provide valid predictions of future outcomes under hypothetical new policies. This capability makes them uniquely well-suited for personnel policy analysis.

Traditional estimation methods for structural DDC models have relied on dynamic programming (DP) methods, which come with some drawbacks. In particular, the computational complexity that comes from the curse of dimensionality has always been a major obstacle when using DP. Models that are easy enough to solve are generally too simplistic for applied work, which has left DDC as a mostly academic pursuit. To be useful for policy makers, models must be able to capture all of the salient features of the decisions that are being modeled.

Recent developments in approximation and econometric methods, combined with greater availability of computational power, allow us to estimate complex and realistic models. Conditional choice probability (CCP) methods can be used to incorporate unobserved taste for military service, enabling the model to identify selection out of the service on the basis of unobservable characteristics. Neural networks and other ML

algorithms can be used to approximate solutions to DP models, significantly reducing the computational time required to solve them.

This Institute for Defense Analyses paper presents a DDC model that is designed to analyze the effects of military personnel policies on service members' retention. It is intended to introduce the model to researchers who are interested in using these methods to analyze questions about military retention and to illustrate how advancements from the academic literature can be applied to improve DDC models for compensation and personnel policy analysis.

Contents

1.	Introduction	1
A.	Modeling Retention Decisions	2
B.	Stylized Model	3
C.	CCP Representation	7
2.	Estimation.....	9
A.	Data and Specifications	13
B.	Using DDC for Policy Evaluation.....	15
3.	Model Extensions	17
A.	Error Distribution	17
B.	State Variable Transitions	17
C.	Simulating Future States	18
D.	Unobserved States	18
E.	Utility Function Specification	18
F.	Choices Over Obligation Length.....	19
G.	Estimating on Short Panels.....	20
H.	Path to a Terminal Action	22
I.	Reducing the Computational Burden of the Outside Option	23
J.	Modeling the Reserve Participation Decision	26
4.	Other Considerations	29
A.	Identification of Structural Parameters.....	29
B.	Identification of the Discount Factor.....	31
C.	Estimation in Two Stages.....	31
D.	Approximation Methods.....	33
E.	Rationality and Time Preferences	33
F.	Peer Effects.....	35
	Appendix A. Recent Military Applications of DDC	A-1
	Appendix B. References	B-1
	Appendix C. Abbreviations	C-1

This page is intentionally blank.

1. Introduction

Analyses of military manpower and personnel issues can benefit greatly from the use of structural dynamic discrete choice (DDC) models. Policy makers design compensation and personnel policies with specific objectives in mind, such as meeting a retention target in a given occupation or increasing the average length of service obligation. To evaluate whether a proposed policy would achieve its objective, we need models that can tell us about how service members would behave under the new policy.

These models need to take into account not just compensation, but all aspects of the service members' careers that are relevant to the retention decision. Structural DDC models can do that. These models specify how compensation and other individual and career characteristics affect service members' decisions. This formulation enables researchers to predict retention outcomes under a new policy. Retention predictions can be used to assess the effect a proposed policy would have on the military's ability to meet manpower and readiness objectives.

DDC models have a long history in economics and are one of the main tools economists use to analyze decisions that involve intertemporal trade-offs. Two of the best known applications studied machine replacement (Rust 1987) and educational attainment (Keane and Wolpin 2001). Aguirregabiria and Mira (2010) provide a comprehensive and relatively recent survey of the academic literature. Service members' retention decisions are a natural fit for DDC models because service members compare their future military career to a potential future life as a civilian when deciding whether to stay in the military. Military personnel policy researchers have used DDC models to study retention since RAND introduced the Dynamic Retention Model (DRM) in 1984 (Gotz and McCall 1984).

Despite their usefulness, DDC models have some limitations. Rust (2019) argues that the primary obstacle to broad application of DDC models in policy and decision making is the difficulty in specifying and estimating models that are sufficiently complex to reflect real life circumstances. In this paper we focus on methods that can help overcome this limitation using techniques that make it possible to specify and estimate realistic models of individual behavior.

This Institute for Defense Analyses paper presents a structural DDC model that is designed to analyze the effects of military personnel policies on service members' retention. Our objective in writing this paper is to introduce the model to researchers who are interested in using these methods to analyze questions about military retention. As such,

we discuss technical nuances and details of model implementation that may typically be relegated to an appendix or not included at all in an empirical paper.

Section 1.B presents theory for a structural DDC model that is designed to analyze the effects of military personnel policies on service members' retention, and Section 2 discusses estimation methods. We demonstrate our preferred method for specifying and estimating a DDC model of military retention decisions, which is the conditional choice probability (CCP) estimation procedure developed by Hotz and Miller (1993) and Arcidiacono and Miller (2011). The CCP formulation simplifies the estimation procedure, enabling the use of a model that more accurately represents the real-life decisions made by service members. Consequently, it provides more credible policy analyses.

Sections 3.A through 3.F discuss standard ways to relax some of the assumptions made in the stylized model. Then, Sections 3.G through 3.J develop several model extensions and adjustments to the estimation procedure that are relevant in the context of military retention. First, we present a method for estimating the model using a short panel data set, alleviating the need to use multiple decades of historical data while still allowing us to estimate the structural utility parameters of the model. Next, we extend the model to incorporate decision points where the service member is unable to leave, such as early reenlistment or contract renegotiation. We then propose two methods for approximating the value of leaving the military that reduce the computational burden of estimation. Finally, we outline an extension to jointly model the active duty and reserve participation decisions.

A. Modeling Retention Decisions

When choosing whether to stay in the military, service members consider how their future military career compares to their potential future life as a civilian. Characterizing the service members' retention decision using a structural DDC model allows us to explicitly capture the trade-offs involved in this decision. The model specifies how service members tradeoff between current and future compensation, and how non-monetary aspects of their career and family life, such as the type of assignment and marital status, affect the retention decision.

As with most career decisions, pay is an important consideration; service members making a retention decision compare what they will earn in the military to what they could earn in the civilian labor market. Importantly, service members consider both current and deferred compensation. For example, the defined-benefit military retirement pay provides a significant incentive to stay in the service for those who are approaching retirement eligibility.

The value of staying in the military and that of leaving are defined by a utility function that represents service members' preferences. This function translates variables such as

compensation, type of assignment, and marital status into a numeric value of each available action, and the relative values of staying versus leaving determine the service member's retention probability. The structure and parameters of the utility function determine the way the variables affect service members' retention decisions. We use economic theory to make the necessary assumptions about the structure of the utility function, and then use historical data on observed retention decisions and associated variables to estimate the structural parameters of the function.

Once we have estimated the utility function that characterizes how service members make retention decisions, we can use it for counterfactual policy analyses. Counterfactual analysis considers how an outcome of interest (in this case, retention) is affected by changes in the environment in which individuals make decisions. This type of analysis is the best available method for prospective evaluations of the impacts of policies that are being considered for implementation, because it provides an estimate of the effect of a policy change while holding other factors constant.

For example, we can use the model to answer, "What would happen to retention if officers were offered a \$10,000 bonus for three additional years of obligation after their initial active duty service obligation (ADSO)?" Furthermore, we can exploit the intertemporal specification of the model to analyze policies involving different lengths of obligation, such as a bonus that incentivizes six-year versus three-year contracts. The model describes individual behavior, so the retention probabilities it produces can be aggregated flexibly to produce estimates of aggregate effects that are important to policy makers. For example, we can compute the average impacts of a policy within community, race, and gender categories to understand its effect across different groups of service members.

Importantly, structural models also allow for counterfactual analyses of policies that have no precedent in the historical data. Because we estimate the utility function describing the underlying preferences that determine service members' retention decisions, we can evaluate the impact of a change in any variable that is included in the utility function. For example, a service may want to know how effective retention bonuses would be at retaining officers at 20 years of service (YOS) long enough for them to be considered for promotion to the rank of Colonel/Captain. Even if such a bonus has never been offered at this point in the career, a structural model can provide a credible estimate of its impact before it is implemented. There are no non-structural models (aka, "reduced form") that produce valid results for these kinds of analyses.

B. Stylized Model

The goal of this section is to illustrate how the DDC model works and what the CCP estimation procedure entails. The stylized model presented here is undoubtedly a simplification of a more full-fledged retention model, but it retains the key features of the

full model, including persistent unobserved taste for military service that varies across individuals. Appropriately dealing with this unobserved heterogeneity is particularly important for making the model realistic, and it is also non-trivial from the technical perspective. The simplifications in this section allow us to illustrate how the CCP method deals with this feature of the model while avoiding particularly cumbersome notation. We discuss how to relax many of the simplifying assumptions later in the paper.

Throughout their military careers, service members periodically make decisions about whether to stay in the military or to leave. In the model, we assume that these retention decisions occur every few years, typically when a service member is nearing the completion of an existing service obligation. Each service member has a small, finite number of decision points during their military career. At each decision point, the service member chooses whether to stay in the military or to leave.

If the service member chooses to leave, they exit military service and make no further decisions. If the service member chooses to stay, they remain in the military for a few more years and then decide again. In general, the interval between decision points varies based on individual circumstances. For example, enlisted service members may have the option to choose the length of their next contract when reenlisting. However, we simplify the exposition in this section by assuming that decision points always occur at three-year intervals. Section 3.F discusses how this assumption can be relaxed.

For the purpose of this stylized model, our goal will be to use the model to set military compensation policy. Specifically, we would like to find the level of retention bonuses at each decision point that achieves a target level of retention. The target may be based on readiness requirements or some other considerations that are important to the policy makers. Note that this question includes an intertemporal component; service members may decide to retain at a decision point with relatively low compensation if they know that the compensation at the next decision point will be high. Because of these intertemporal considerations, the solution is a set of bonus amounts at different decision points that together achieve the retention target.

Let $i \in \{1, \dots, N\}$ index service members and $t \in \{1, \dots, T\}$ index time periods, which we assume to be years. Each service member will make at least one decision in the model and will continue to make decisions every three time periods as long as they remain in the military. Let $d \in \{1, \dots, D_i\}$ index decision points, and let t_d be the time period in which decision point d occurs. D_i may differ across individuals; a service member who chooses to leave at the end of their initial obligation will have only one decision point, while someone who stays until retirement will have multiple.

At decision point d , service member i chooses an action $a_{i,t_d} \in A = \{\text{stay}, \text{leave}\}$ based on the expected utility associated with each action. A service member's utility $u_{i,t}$ in each time period is a function of that individual's state, which contains both observed

variables and persistent unobserved heterogeneity. Observed variables are denoted by $x_{i,t}$ and while the service member is in the military they include variables such as gender, family status, military occupational specialty (MOS), and military compensation.

If the service member has left the military then $x_{i,t}$ includes only civilian compensation. We follow the Heckman and Singer (1984) finite mixture framework to model persistent unobserved heterogeneity: we assume there are a small number of distinct types of individuals with different latent preferences for serving in the military. We assume that the unobserved type, denoted by $s_i \in \{1, \dots, S\}$, is constant over time. For convenience we will sometimes denote the entire state space as $z_{i,t} = (x_{i,t}, s_i)$.

Some observed state variables are likely to change over time. The state at some future period $\tau > t$ may depend on whether the service member chose to stay in service or to leave at time t . For example, a service member who remains in the military may convert to a different MOS, which would change their military career and their potential civilian earnings. Let $f(x_{i,\tau}|x_{i,t}, a)$ denote the probability that observed variables take on values $x_{i,\tau}$, conditional on them having values $x_{i,t}$ and taking action a at time t .

In particular, $f(x_{i,d+1}|x_{i,d}, \text{stay})$ is the distribution of observed variables at the next decision point (three years in the future), conditional on the service member choosing to stay in service at time t . This will represent what the service member expects to happen during their next tour, including events like promotions, marriage, having children, and changes in the civilian economy.

To streamline notation, we define an aggregate utility function, suppress the individual subscript i , and let z_d denote the state at time period t_d for the remainder of this section. Service members generally consider their entire expected future military career and compare it to their expected civilian career when making the retention choice. The model captures this forward-looking nature of the decision problem. The function $U(z_d, a)$ represents the total expected utility between decision point d and the next decision point, conditional on choosing action a .

If $a = \text{stay}$, U is a sum over the next three years of utility that the service member expects to receive while in the military, before reaching the next decision point. If $a = \text{leave}$, this is the last decision in the model, and U is the sum of the remainder of expected lifetime utility. Service members discount utility they expect to receive in the future according to an exponential discount factor $\beta \in (0,1)$.

$$U(z_d, a) = \begin{cases} \sum_{\tau=t}^{t+2} \beta^{\tau-t} \mathbb{E}[u(z_\tau)|z_d, \text{stay}] & \text{if } a = \text{stay} \\ \sum_{\tau=t}^T \beta^{\tau-t} \mathbb{E}[u(z_\tau)|z_d, \text{leave}] & \text{if } a = \text{leave} \end{cases} \quad (1)$$

In addition to the utility from state variables, the value of each action is affected by a choice-specific shock. This shock represents circumstances that are unique to the service member, independent of the state variables, and not persistent over time. Perhaps the service member received an unusually good civilian job offer, or maybe their next assignment is at the location where their best friend is stationed. Service members know the value of the current shocks, but we as researchers do not. Service members know that future decision points will also have choice-specific shocks, but do not know what they are ahead of time. Denote the shock associated with action a as $\epsilon_d(a)$. We assume it is independently distributed Type 1 Extreme Value and is additively separable from the rest of the utility associated with that action.

The timing of the model is as follows. Consider a service member who needs to make a retention decision at t_d . At the start of period t_d the service member observes the realized values of the current state $z_d = (x_d, s)$ and the utility shocks ϵ_d . Next, the service member chooses $a_d \in A = \{\text{stay}, \text{leave}\}$ and receives the current period utility associated with a_d , $u(z_d, a_d) + \epsilon_d(a_d)$. Finally, the state transitions to z_{d+1} according to $f(x_{d+1}|x_d, a_d)$ and the next period begins. The service member collects discounted period utility associated with the state until the next decision period and then repeats the process.

We assume that service members chose the best action from their point of view, taking into account current utility from that action as well as the effect of the action on their future utility. Formally, service members choose a sequence of actions $\mathbf{a}^* = \{a_1^*, \dots, a_D^*\}$ that maximizes expected present discounted lifetime utility. Recalling that U is the expected discounted sum of utility from a single decision, we now write the service members' optimization problem as

$$\mathbf{a}^* = \underset{\mathbf{a}}{\operatorname{argmax}} \mathbb{E} \left[\sum_{d=1}^D \beta^{3(d-1)} [U(z_d, a_d) + \epsilon_d(a_d)] \right]. \quad (2)$$

It is easier to work with the decision problem using value functions. We denote by $\bar{V}(z_d)$ the value associated with state z_d at the beginning of time period t_d . This is the expected discounted sum of current and future utility, before ϵ_d are realized, and conditional on taking the optimal action, both in this period and in the future, given that the current state is z_d . Recall that the state is a collection of observed variables x_{i,t_d} together with the unobserved taste for military service s_i , so the value function gives the total payoff associated with the optimal sequence of actions $\mathbf{a}^*$ for a service member with a specific combination of observed variables and taste for service.

To understand the value of a specific action, define the alternative-specific value function $v(z_d, a)$ as the value of taking action a in state z_d , and then taking the optimal action at all future points. Because there are no decision points after taking the leave action, $v(z_d, \text{leave})$ is just the sum of the remainder of expected lifetime utility. For the stay

action, it is the utility associated with that action, plus the utility associated with taking the optimal action at the next decision point, which we defined above as $\bar{V}$. The alternative-specific value functions can be written as

$$v(z_d, a) = \begin{cases} U(z_d, \text{stay}) + \beta^3 \mathbb{E}[\bar{V}(z_{d+1})] & \text{if } a = \text{stay} \\ U(z_d, \text{leave}) & \text{if } a = \text{leave} \end{cases} \quad (3)$$

To estimate the model, we need to connect this value function representation of the problem with observable outcomes. Specifically, we need to know what the value functions imply about the probability of observing a service member stay versus leave. Our parametric assumption regarding the distribution of the idiosyncratic shocks, ϵ , implies that the choice probabilities follow the standard multinomial logit form. The CCP of choosing action $a = \text{stay}$ is given in (4); the probability of choosing to leave takes the same form.

$$p(\text{stay}|z_d) = \frac{\exp(v(z_d, \text{stay}))}{\exp(v(z_d, \text{stay})) + \exp(v(z_d, \text{leave}))} \quad (4)$$

C. CCP Representation

Equation (4) for $p(a|z_d)$ may look like a standard multinomial logit, but it cannot yet be estimated directly. We have defined the CCPs in terms of $v(z_d, a)$, but $v(z_d, \text{stay})$ is defined in (3) in terms of $\bar{V}(z_{d+1})$, which generally lacks a closed form. The inversion property from Hotz and Miller (1993) and the properties of the Type 1 Extreme Value distribution provide a closed form for the value function in terms of the alternative-specific value and CCP for an arbitrary action $\tilde{a} \in A$ (and the Euler's constant γ).

$$\bar{V}(z_d) = v(z_d, \tilde{a}) - \ln(p(\tilde{a}|z_d)) + \gamma \quad (5)$$

Essentially, (5) tells us that the value associated with a specific state can be expressed as the value of taking any specific action, adjusted for the probability of taking that action. The following example may provide some intuition. Suppose we know the value a service member might expect if they leave the service at 18 YOS because we have estimated their likely civilian earnings. We also know that a service member deciding to leave at 18 YOS is extremely unlikely; practically no service member voluntarily leaves at that point, because they will be eligible for the military pension in just two years. Equation (5) tells us that we can adjust the known value of leaving by the very low probability of actually leaving to arrive at the (much higher) value of being in the service at 18 YOS. We will use this intuition to construct our value functions in a way that makes the model simple to estimate.

To make the best use of the closed form representation in (5), we need to choose $\tilde{a}$ such that $v(z_d, \tilde{a})$ is known. We have assumed that leaving the service is a terminal decision and $v(z_d, \text{leave}) = U(z_d, \text{leave})$, which involves no further decision points and

actions. Therefore, $v(z_d, leave)$ is known (up to structural parameters); we can calculate it for any z_d . We set $\tilde{a} = leave$ and rewrite the alternative-specific value function for staying using the CCP representation for $\bar{V}$.

$$v(z_d, stay) = U(z_d, stay) + \beta^3 \mathbb{E}[U(z_{d+1}, leave) - \ln(p(leave|z_{d+1})) + \gamma] \quad (6)$$

Equation (6) implies that the value of staying can be broken down into the current utility from staying plus the utility of leaving at the next decision point adjusted for the probability of actually leaving at the next decision point. We now have representations of both $v(z_d, stay)$ and $v(z_d, leave)$ that do not involve the future value function. Note that (6) includes the CCP for leaving at the next decision point, $p(leave|z_{d+1})$. If there were no persistent unobserved heterogeneity across individuals, then these CCPs could be computed directly from the data and the multinomial logit in (4) could be used to estimate the model directly. In our example, retention decisions are affected by the service member's taste for service, so the CCPs have to be estimated jointly with the parameters of the utility function, a process which we turn to in the next section.

2. Estimation

To estimate the model we use the CCP method, which was first described by Hotz and Miller (1993) and extended to include persistent unobserved heterogeneity by Arcidiacono and Miller (2011). The CCP method exploits the mapping between CCPs and alternative-specific value functions, shown in (6). This avoids solving the dynamic programming problem for every value of the model parameters, which can quickly become computationally infeasible.

The CCP method is particularly attractive when the model includes terminal actions, as our model does with the leave action. For the purposes of exposition, we first describe a CCP estimator of our model without unobserved taste for military service. We then show how the unobserved taste can be included in the model using the well-known expectation-maximization (EM) algorithm.

Our goal is to estimate the structural parameters θ of the service members' utility function $U(z_d, a; \theta)$. We can then use the utility function to evaluate the impact of different policies on the retention decisions of service members. Our data consist of a sample of N service members observed for D_i decision points. At each decision point, we observe the relevant variables x_d and action a_d for each service member. For now, assume there is no difference in the taste for military service between service members so that $z_d = x_d$. We assume the discount factor β is known, and the idiosyncratic error ϵ is independently distributed Type 1 Extreme Value across alternatives, as before. We again simplify the exposition of the problem by assuming that decisions occur at three-year intervals.

Recall that the CCPs that link the observed data to the model are given by the multinomial logit form, and the alternative-specific value functions for these CCPs can be expressed in a closed form, as discussed in the previous section:

$$p(stay|z_d) = \frac{\exp(v(z_d, stay))}{\exp(v(z_d, stay)) + \exp(v(z_d, leave))} \quad (7)$$

$$v(z_d, stay) = U(z_d, stay; \theta) + \beta^3 \mathbb{E}[U(z_{d+1}, leave; \theta) - \ln(p(leave|z_{d+1})) + \gamma] \quad (8)$$

$$v(z_d, leave) = U(z_d, leave; \theta) \quad (9)$$

Equation (8) shows that the value of remaining in the military can be summarized by the current utility of staying ($U(z_d, stay; \theta)$), the expected utility of leaving the military at the next decision point ($U(z_{d+1}, leave; \theta)$), and the probability of leaving at the next decision point ($p(leave|z_{d+1})$). The key insight of Hotz and Miller (1993) is that

$p(\text{leave}|z_{d+1})$ can be estimated separately from θ . We do not need to know the structural parameters of the utility function to estimate the relationship between the observed variables and the CCPs. If we observe all of the variables in z_d , as we have assumed for now, then estimation can be done in two easy stages.

In this case, all of the components of $v(z_d, a)$ other than $U(z_d, a; \theta)$ are estimated in a first stage. Reintroducing the subscript i to enable indexing over individuals, the CCPs and transition functions can be estimated using bin estimators:

$$\hat{p}(\text{leave}|z) = \frac{\sum_i \sum_d \mathbb{1}[a_{i,d} = \text{leave}] \mathbb{1}[z_{i,d} = z]}{\sum_i \sum_d \mathbb{1}[z_{i,d} = z]} \quad (10)$$

$$\hat{f}(z'|z, a) = \frac{\sum_i \sum_d \mathbb{1}[z_{i,d+1} = z'] \mathbb{1}[z_{i,d} = z] \mathbb{1}[a_{i,d} = a]}{\sum_i \sum_d \mathbb{1}[z_{i,d} = z] \mathbb{1}[a_{i,d} = a]} \quad (11)$$

Alternatively, these functions can be estimated using flexible logits. For the CCPs in particular, it is important to estimate them as precisely as possible in the first stage, since they play a crucial role in our formulation of $v(z_{i,d}, \text{stay})$ in the second stage.

With these pre-computations complete, $v(z_{i,d}, a)$ is known up to θ and we can use (7) to form the likelihood:

$$l(a_{i,d}|z_{i,d}; \theta) = p(a|z_{i,d}; \theta). \quad (12)$$

In some applications, it may be convenient to normalize the utility of the outside option $U(z_{i,d}, \text{leave})$ to zero. This normalization makes sense when individuals have an outside option that essentially has them opt out of the process being modeled (e.g., decision to not participate in the labor force or to not purchase a particular good). In our model the outside option is for service members to leave the military and participate in the civilian economy by working in an appropriate job.

Because service members' civilian incomes likely vary based on their occupation and experience in the military, we keep utility parameterized as $U(z_{i,d}, \text{leave}; \theta)$. Hotz and Miller (1993) prove that the maximum likelihood estimates (MLE) $\hat{\theta}$ are consistent and asymptotically normal with $\sqrt{N}$ convergence under their regularity assumptions, and can be estimated using a logit or any other M-estimator.

The downside of the Hotz and Miller (1993) CCP estimator is the assumption that all variables in $z_{i,d}$ are observed in the data. Both intuition and prior research suggest that there are important and persistent differences between service members in the way they view military service. We now show how unobserved state variables can be integrated into the CCP estimation procedure. Typically, the unobserved information is grouped together under the moniker “taste for military service,” although nothing in the model limits what this variable is capturing. Persistently good or bad civilian job prospects, the belief that

military service is not compatible with family life, and many other unobservable factors may all be captured here.

We reintroduce a permanent unobserved component to the state such that $z_{i,d} = (x_{i,d}, s_i)$, where $x_{i,d}$ are the observed state variables and s_i is persistent unobserved heterogeneity, which we refer to as the taste for military service. As mentioned earlier, we follow Heckman and Singer (1984) and assume that the taste for military service can be captured by a fixed number of discrete types $s_i \in \mathcal{S} = \{1, \dots, S\}$.

For now, we make a number of important assumptions about how taste for military service can interact with other parts of the model. First, we assume that service members' taste for military service does not change over time. Next, we assume that the taste for military service does not affect the transitions of observed variables, meaning $f(x_{i,d+1}|x_{i,d}, a_{i,d}, s_i) = f(x_{i,d+1}|x_{i,d}, a_{i,d})$. With these two assumptions, we can estimate $\hat{f}(x_{i,d+1}|x_{i,d}, a_{i,d})$ as in (11). Finally, we assume that the taste for military service does not affect the utility of working in the civilian sector, so $U(x_{i,d}, s_i, \text{leave}) = U(x_{i,d}, \text{leave})$.

Due to the presence of unobserved taste for military service, it is no longer possible to estimate CCPs in the first stage because $\hat{p}(\text{leave}|x_{i,d}, s_i)$ is conditioned on the unobserved type s_i . Let $\pi(s)$ denote the distribution of types in the population. This distribution can be conditional on $x_{i,1}$, but for now we assume that every service member has the same underlying type distribution. We integrate this distribution out of the likelihood to arrive at the MLE:

$$\{\hat{\theta}, \hat{\pi}\} = \underset{\theta, \pi}{\operatorname{argmax}} \sum_{i=1}^N \ln \left[\sum_{s_i=1}^S \pi(s_i) \prod_{d=1}^D p(a_{i,d}|x_{i,d}, s_i; \theta) \right]. \quad (13)$$

Direct estimation of the MLE is complicated by the fact that the likelihood is not additively separable in θ and π . Arcidiacono and Miller (2011) demonstrate how the expectation-maximization (EM) algorithm can be adopted to compute the MLE. They present three different algorithms that iterate between estimating $\pi(s)$, $\hat{p}(a|x_{i,d}, s_i)$, and θ and converge to the MLE solution. These estimators retain the $\sqrt{N}$ -consistency and asymptotic normality of the full MLE estimates under regularity conditions in Arcidiacono and Miller (2011). We discuss two of these estimators here, and return to their “two-stage” estimator later in the paper.

In order to compute an estimate of $\hat{p}(a|x_{i,d}, s_i)$ we need an estimate of the probability distribution of s_i . This distribution is conditional on all of the decisions of service member i , denoted by the vector $\mathbf{a}_i$. Intuitively, service members who repeatedly choose to stay in the military likely have a higher taste for service than those who leave at an early career point. We can use Bayes' rule to write this probability distribution as:

$$\hat{q}_i(s) = Pr\{s_i = s | \mathbf{a}_i, \mathbf{x}_i; \hat{\pi}\} = \frac{\hat{\pi}(s) \prod_d p(a_{i,d} | x_{i,d}, s)}{\sum_{s'} \hat{\pi}(s') \prod_d p(a_{i,d} | x_{i,d}, s')} \quad (14)$$

The overall population probability distribution for each type can be found by summing all of the individual probabilities:

$$\hat{\pi}(s) = \frac{1}{N} \sum_i \hat{q}_i(s) \quad (15)$$

Given an estimate of the type distribution $\hat{q}_i(s)$, we can compute the CCPs. Specifically, if we take the types as given, then we are back to a setting where all relevant state variables are observed. Because (14) gives a probability distribution of types, we use these $\hat{q}_i(s)$ as weights in a bin estimator:

$$\hat{p}(\text{leave} | x_{i,d}, s_i) = \frac{\sum_i \sum_d \hat{q}_i(s) \mathbb{1}[a_{i,d} = \text{leave}] \mathbb{1}[x_{i,d} = x]}{\sum_i \sum_d \hat{q}_i(s) \mathbb{1}[x_{i,d} = x]} \quad (16)$$

With an estimate of $\hat{p}(\text{leave} | x_{i,d}, s_i)$ we can now estimate θ . We take $\hat{q}_i(s)$ as given again and use it as population weights in an estimator where the types are observed. The estimator for θ is

$$\hat{\theta} = \underset{\theta}{\operatorname{argmax}} \sum_i \sum_s \sum_d \hat{q}_i(s) \ln[p(a_{i,d} | x_{i,d}, s_i, \hat{p}(\text{leave}), \theta)] \quad (17)$$

Note that this estimator is not the likelihood given in (13). The two estimators have the same first order conditions, and so are maximized at the same $\hat{\theta}$. The estimator given in (17) is easier to use in the EM framework, is globally concave under our assumptions, and is simple to compute. It converges to the MLE under standard regularity conditions.

The estimation proceeds by iterating on equations (14), (15), (16), and (17) until convergence. The first three steps are computationally straight-forward. The maximization over θ treats the unobserved taste for military service as known, so each step in the algorithm is approximately as computationally expensive as estimating one model without the unobserved taste. While this can still be computationally burdensome, it is significantly faster than a full-solution method that requires solving the dynamic problem for every trial value of parameters.

The second estimator in Arcidiacono and Miller (2011) uses a different approach to obtain $\hat{p}(\text{leave} | x_{i,d}, s_i)$. For any given estimate of $\hat{\theta}$ and $\hat{p}(\text{leave})$, the model will generate $p(\text{leave} | x_{i,d}, s_i, \hat{p}(\text{leave}), \hat{\theta})$. This suggests that we can replace the empirical bin estimator in (16) with model predictions $\hat{p}(\text{leave} | x_{i,d}, s_i) = l(\text{leave}, x_{i,d}, s_i, \hat{p}, \hat{\theta})$. If the bin estimator can be thought of as getting $\hat{p}(\text{leave})$ “from the data” then we can think of this method as obtaining it “from the model.” In practice, the bin estimator may be difficult to implement if some values of the variables are relatively rare (or if some x are continuous), and so updating $\hat{p}(\text{leave})$ from the model may be preferred.

A. Data and Specifications

To estimate the model, we need individual-level data on retention decisions. Each observation represents a single retention decision point, and includes information about the alternatives available to the service member at that point and which alternative was chosen. Observations for the same service member are linked using a unique identifier, which allows us to use information across multiple decision points to estimate the service member's taste for military service. In addition to observed decisions, we use data on both military and civilian labor markets to fully characterize the alternatives available to the service member.

The military data include variables that affect the utility that a service member receives from staying in the military. These variables form the observable state $x_{i,t}$, though not all observable variables enter the utility function directly. For example, military compensation is a key variable of interest, and is included in the utility function. Deciding how to calculate military compensation is up to the researcher; typically, it is either based on published policy documents or estimated as a function of service members' observed characteristics. Certain characteristics, such as YOS and pay grade, are necessary to correctly calculate military compensation even though they may not directly influence service members' utility.

In addition to military compensation, expected wages in the civilian labor market are an important state variable. Military data only contain information on service members who have not yet left the military, so we do not observe a service member's civilian pay once they leave. We make assumptions about how the observable variables translate into expected civilian pay, so we can treat it as an "observed" state variable.

For example, assume that civilian wages are a function of the service member's career field, experience, and education level. We can estimate a wage equation using civilian data from the American Community Survey (ACS) or the Annual Social and Economic Supplement of the Current Population Survey (CPS-ASEC), and use this equation to predict civilian wages for the service members in our military data.

Spouse wages may also be a component of the state space, particularly if the retention decision is a household, rather than an individual, decision. If the utility function is nonlinear in income, then we need to carefully consider whether individual or household income is the relevant state variable. For example, suppose $u(\text{income}) = \theta \ln(\text{income})$; this imposes a decreasing marginal utility of income.

In this case, spouse wages shift household income up, and reduce the marginal utility from the service member's income, suggesting that married service members may be less responsive to compensation incentives than single service members. Additionally, there is evidence that spouses of service members earn less than spouses of veterans; if household

income is the relevant variable for utility, the effect of the retention decision on the spouse's wages should also be taken into account.

Service members' utility from being in the military may be affected by variables other than pay, such as their marital status, their duty station location, and the likelihood of deployment during the next tour, as well as the duration and nature of deployment, which may itself depend upon persistent geopolitical shocks. The set of variables that enter the utility function should include anything that the researcher is interested in using during policy evaluation. In addition, it should include variables that:

1. Directly affect the utility that a service member expects to receive from staying in the military, and
2. Are correlated with the variables of interest for policy evaluation

Here, the variable of interest is compensation, so we endeavor to include variables in the utility function that directly affect utility and are correlated with military or civilian compensation. This includes variables that are correlated with differences in compensation across individuals (e.g., gender) as well as those that are correlated with changes in compensation for the same person over time (e.g., pay grade). Failing to include utility variables that are correlated with the variables of interest for policy evaluation may lead to omitted variable bias in the estimated coefficients of interest.

Service members in the model are forward-looking, so the utility function captures the future utility that they expect to receive, conditional on the retention decision they make. Accurate predictions of future state variables are a key part of the model because they represent service members' beliefs about future utility. Expectations of promotions, conversions to different communities, changes in location, marriage, having children, and other factors play a role in retention decisions.

To be useful in the utility function, a variable must be predictable (from the perspective of the researcher) at least to some degree. We define a variable to be predictable to the researcher if the data contains information that facilitates construction of a prediction that is more accurate than a random guess. For example, service members may have information about the quality of housing or the command climate at the units they expect to go to, and this knowledge affects their retention decisions. Unfortunately, housing quality and command climate are typically unobservable to the researcher. Effects like these would therefore have to be captured via other observable variables in the utility function, or with the idiosyncratic shocks.

B. Using DDC for Policy Evaluation

Estimating a structural DDC model of retention makes it possible to evaluate the effect of counterfactual personnel or compensation policy changes on retention. The most direct way to assess the effect of a proposed policy change is to modify the state variables in x in existing data to reflect the anticipated effect of the policy, and then predict new choice probabilities using the structural model.

This method of policy evaluation is often called a “counterfactual” policy evaluation; it answers the retrospective question, “What decisions would service members have made in the past, if the hypothetical policy had been in effect?” Often, this is not the relevant question of interest. The true question might be, “What will service members do in the future if the hypothetical policy is put into place?” To the extent that the service member population and their career and family circumstances in the future resemble those of the past, the counterfactual approach can provide valuable insight. However, if the salient features of the environment or the population are different, we can simulate new data to provide a better answer to this question.

Retention bonuses and other compensation policies are often of interest because these are expected to have a direct effect on retention decisions. However, structural DDC models can be used to evaluate a wide variety of personnel policies, including changes to promotion/advancement rates or to the assignment process. In general, the model can assess hypothetical policies that meet three criteria:

1. The policy affects observable state variables.
2. The affected state variables could be accurately predicted before the policy change.
3. The policy has a predictable effect on the state variables.

Condition 1 requires that policy operates on at least one observable state variable. An easy example is a retention bonus; it increases military compensation for at least some service members if they decide to stay. But not all policies that may affect retention do so through observable variables. Service members may value access to high-quality on-base childcare, and a policy that increases the availability of childcare may increase retention.

However, unless the historical availability of childcare is observed in the data, we cannot capture the role of childcare in the model. It is crucially important to consider what policies the model will be used to evaluate, and ensure that the utility function specification includes variables that capture the policies’ effect. If a policy affects state variables that are not modeled (whether they are observable or not), the model-predicted retention decisions cannot capture the full effect of the policy.

Condition 2 requires that the affected state variables can be accurately predicted using available data. In the model, retention decisions depend on the service member’s

expectations about the future in the future utility terms. For a variable to affect the retention decision, the service members have to be able to predict, at least to some extent, the future values of the variable. If location assignment was unpredictable before the policy change, then it could not have affected the behavior of service members, and we cannot estimate their preferences for one location versus another.

Even if location assignment is perfectly predictable after the policy change, the model cannot predict the impact of this policy on retention because the data before the policy change were not informative about service members' location preferences. More generally, if the policy affects something that service members could not or did not consider when making previous retention decisions, then the model estimated using the past data will not be useful in evaluating the effect of this policy.

Finally, Condition 3 requires that the policy have a predictable effect on the state variables. To perform a counterfactual policy evaluation, we simulate the effect of the policy on the state variables in the utility function. A bonus that is targeted to certain service members based on observed characteristics, such as MOS and pay grade, is easy to simulate and evaluate.

However, a bonus that is offered to service members based on a recommendation from their supervisor would not be easy to simulate. We do not know who would be recommended or why, and while we may try to make some assumptions, they are unlikely to fully capture the supervisors' decision process. In this case, the policy has an unpredictable effect on the state variable; we cannot reliably tell who will be offered higher compensation under this policy. The model may provide predictions of different scenarios under different assumptions, but these would need to be interpreted with the appropriate degree of caution.

Retaining service members with the right set of skills and expertise is important because they fill critical jobs in the military and maintain a capable and ready military force. Policies that affect retention therefore affect the inventory of service members with different skills, the job match quality of service members, and the readiness of military units. Understanding the connection between retention decisions and these additional metrics requires additional post-processing.

Inventory projection models take predictions from retention models and simulate changes in the personnel inventory and unit assignment of service members to evaluate the downstream effects of a retention policy. Because the retention model operates at an individual level, predicted retention decisions can be aggregated along any dimension necessary to work with an inventory projection model.

3. Model Extensions

The stylized model of retention presented in the previous section has the essential features of a retention model, but it also makes a lot of simplifications that may not be appropriate in specific applications. The CCP estimation framework reduces the computational burden of estimation to the point where we can construct models that are rich enough to reflect real-life conditions under which retention decisions are made.

This section outlines several extensions that allow us to provide more accurate and credible estimates of the impact of personnel and compensation policies on retention. Some of these extensions are straight-forward and are largely covered in Arcidiacono and Miller (2011). We briefly discuss these first, before turning to topics that are specific to the military retention application.

A. Error Distribution

The stylized model assumed that ϵ are distributed Type 1 Extreme Value, leading to standard multinomial logit specification. In general, CCPs can be numerically calculated for any Generalized Extreme Value (GEV) distribution, but this is typically unattractive on computational grounds. In the special case of a nested logit model, the CCPs have an analytic solution. Nested logit models may be attractive when there is reason to believe that ϵ shocks are not independent across alternatives.

For example, if there are multiple stay options, such as shorter and longer contracts, the shocks associated with the stay actions may be correlated among themselves, but be independent of the shock for the leave action. The nested logit specification would imply that if one of the stay actions were removed, a service member who previously chose that action would be more likely to switch to a different stay action, all else equal. This extension is particularly important in settings where service members have choices over the characteristics of their next contract and assignment.

B. State Variable Transitions

Some of the assumptions we made about the transition function may be relaxed at the cost of additional computational burden. State variables may transition differently depending on the action taken; for example, service members who choose to commit to a longer obligation may receive systematically different assignments. Transitions of observed variables may depend on the unobserved taste for military service.

For example, we may hypothesize that service members with a high taste for military service are more likely to get promoted. The unobserved taste for military service itself may also change over time. In both cases, the transition function $f(x_{d+1}, s_{d+1} | a_d, x_d, s_d)$

has to be estimated as part of the EM algorithm. In practice, identifying the relationship between transitions and the unobserved taste for military service requires repeated observations of the same service member, but most service members leave the military at the first or second decision point. This may be a plausible and desirable extension in applications dealing with specific communities where the identification is possible.

C. Simulating Future States

Most theoretical DDC papers assume that the state space is discrete, allowing for an easy construction of $f(z_{t+k}|z_t, a_t)$ as a Markov transition matrix. In practice, some variables in z may be continuous. Additionally, in military retention applications, the state space is usually large, complex, and sparse enough that it is impractical to directly use the full Markov transition matrix to calculate expected utility.

Instead, we numerically approximate expectations over z_{t+k} using simulations, a process which is easily extended to continuous variables. With the CCP methodology, we need to simulate z_{t+k} only until the next decision point, which is usually only a few years in the future. Using simulations to approximate expected utility enables us to capture more complexity in the utility function than would otherwise be possible.

D. Unobserved States

In the stylized model, we assumed that taste for military service is the only relevant unobserved variable. In general, there could be multiple unobserved state variables that influence the transitions $f(z_{t+k}|z_t)$ and utilities $u(z_t, a)$ in different ways. For example, we may think that taste for military service directly affects service members' utility, but does not affect transitions, while unobserved ability affects the promotion probability, but does not directly affect utility. Identifying this two-dimensional unobserved state variable may be possible, depending on the available data and how each variable is assumed to affect utility and transitions.

Note that treating ability as an unobserved state variable is not as simple as it may seem. Ability is likely related to the civilian wage; service members with high ability and good promotion prospects may also have better job prospects in the civilian labor market. It can be difficult to make assumptions a priori about which unobserved variable is affecting which part of the model, and mis-specification may confound interpretation. However, this extension may be useful if the policy question at hand is directly related to unobserved variables such as ability.

E. Utility Function Specification

We often represent utility $u(z_{i,t}, a)$ as a linear function of the state variables. However, u does not need to be linear, and the aggregate utility U does not have to be a

sum of annual utilities. In some cases, economic theory may provide guidance about the shape of the utility function. A linear specification is popular and easy to interpret, but it is possible to estimate other variations; examples include hyperbolic absolute risk aversion and isoelastic utility functions. The additional structural parameters that come with these specifications increase the computational cost of estimation, but are possible to implement given our estimation procedure outlined above.

A simple and useful way to extend the utility specification is to add variables that only appear at the decision period and not every year. Recall that U captures the utility associated with a decision, conditional on the current state. Some parts of the utility are annual events or flows such as income, but it may also include per-decision components. The utility function must be time-separable across decisions (i.e., we must be able to separate $U(z_d, a_d)$ from $U(z_{d+1}, a_{d+1})$).

However, it does not need to be separable across time periods within $U(z_d, a_d)$; this is an unusual feature that stems from the fact that service members do not have an opportunity to make decisions every year. In fact, our specification in the stylized model already includes one component that is per-decision: the idiosyncratic error term. Some other examples of potential per-decision utility components are:

- An indicator for whether a service member will deploy during the next tour
- A one-time utility for leaving the service and for retiring

We can incorporate information about the recent past in a similar way. We could include data on whether a service member recently returned from a deployment, the command climate at their last unit, or other variables. These do not directly fit into the framework of utility based on future state variables, but they may provide important information about a service member's retention probability. Care must be taken to interpret model parameters, especially when there is a mixture of per-period utility flow and per-decision utility. However, in some cases, this flexibility can be used to capture important aspects of a service member's retention decision.

F. Choices Over Obligation Length

In the stylized model, all stay decisions incurred three additional years of obligation. In practice, service members often have the ability to choose how long they will stay in the military. Retention bonuses, such as Selective Reenlistment Bonuses (SRBs), frequently offer higher amounts of money for longer obligations, so incorporating a choice of obligation length is an important generalization relative to the simple stay/leave model.

The fundamental change from this generalization is that the service member has control over when the next decision point $d + 1$ will occur. Decision points might not occur at regular intervals; a service member may choose to extend for one year, then reenlist for

six years, then leave. Not all stay options may be available at every decision point; a short one-year extension may be available at one decision point, but not at another one. For the remainder of the paper, we will use a to index available actions and to represent the length of the obligation associated with each action. For example, action $a = 3$ represents the decision to stay for three more years, while $a = 0$ represents the decision to leave.

We can accommodate this generalization via an aggregate utility function U . For a stay action, $a > 0$, U sums expected utility over a periods. If the service member chooses the leave action, $a = 0$, there are no further decision points, and the function aggregates the remainder of expected lifetime utility.

$$U(z_t, a) = \begin{cases} \sum_{\tau=t}^{t+a-1} \beta^{\tau-t} \mathbb{E}[u(z_\tau)|z_t, a] & a > 0 \\ \sum_{\tau=t}^T \beta^{\tau-t} \mathbb{E}[u(z_\tau)|z_t, 0] & a = 0 \end{cases} \quad (18)$$

The rest of the changes are equally straight-forward. The continuation value $\bar{V}(z_{d+1})$ in the alternative-specific value function now depends on the length of obligation associated with the stay action a_d . Choosing $a_d > 0$ implies that the next decision point $d + 1$ occurs a_d periods in the future. Therefore, the state at decision point $d + 1$ may be different for each $a_d \in A$. Additionally, the continuation value must be appropriately discounted to reflect when the next decision point will occur.

We again use the terminal property of the leave decision $a = 0$ to obtain the closed form representation of the alternative-specific value functions. There are no changes to the estimation procedure. The CCPs and their components take the following form:

$$p(a|z_d) = \frac{\exp(v(z_d, a))}{\sum_{a' \in A} \exp(v(z_d, a'))} \quad (19)$$

$$v(z_{t_d}, a > 0) = U(z_{t_d}, a) + \beta^a \mathbb{E}[U(z_{d+1}, 0) - \ln(p(0|z_{d+1}))] + \gamma \quad (20)$$

$$v(z_{t_d}, 0) = U(z_{t_d}, 0) \quad (21)$$

G. Estimating on Short Panels

We previously assumed that every service member is observed starting with the first decision point, $d = 1$. For this assumption to be true, we would need a panel that covers entire careers of the service members in our sample. Given that military careers can last in excess of 30 years, this is a burdensome requirement. Furthermore, the behavior of service members who entered the military three decades ago is unlikely to represent in a

satisfactory manner the behavior of service members who have recently entered service. Young service members today have different attitudes and experiences and face a different military and civilian environment than service members in prior generations.

We can make some assumptions that allow us to estimate the model using a short panel. Our approach takes inspiration from the Keane and Wolpin (2001) simulation method for dealing with the initial conditions problem. If the first period of our data contains a mid-career service member, we can infer that this service member made at least one decision to stay in the past.

With this in mind, we can use the observed data to backwards simulate the salient features of this service member’s career, and then infer the past decisions the service member must have made to reach the point where they are observed at the beginning of our dataset. Doing this requires an assumption that the distribution of unobserved types $\pi(s|x_{i,1})$ is invariant over time. While this assumption is not innocuous, it allows us in practice to estimate the model with only one stay-leave decision per individual – as few as four years of data.

Our method proceeds as follows. Assume we have observations of all service members who make decisions over a short time period, such as four or five years. Some service members make their first retention decisions during this time period, and other service members at a later point in their careers. Based on observable characteristics such as MOS, we assign to service members the modal decisions and timing of those decisions that would result in the state we observe them in during the first period of our data.

For example, if we see a service member for the first time at 14 YOS finishing a six-year contract, we would look at other service members in their MOS at 8 YOS (when the service member must have started their current contract) and find that most of them are finishing a four-year contract. We would then assign to this service member two “unobserved” decisions: the first at 4 YOS to stay for four years and the second at 8 YOS to stay for six years.

We use the unobserved decisions to help us correctly estimate the selection process based on the unobserved taste for military service, but we do not want to use these decisions in the estimation of the structural parameters. We include unobserved decisions in our calculation of $q_i(s)$ in (14) so we can correctly condition the distribution of types for a service member in the middle of the career on their full history of decisions to not leave the service.

We do not fully characterize the unobserved decisions in the model, as doing so would require more assumptions regarding the circumstances faced by the service members in the past. Instead, we simply match the unobserved decisions to observed decisions based on a subset of observable characteristics, such as career field and YOS, and use the probabilities from the observed decisions to update $q_i(s)$. We then use these $q_i(s)$ as weights in the

likelihood estimation, but include in the likelihood only the decisions we do observe. Although we have to make the additional assumption on π for this method to work, it allows us to estimate the model using relatively recent decisions of service members. We believe this is a good trade-off because estimating utility function parameters using recent decisions mitigates the possible effects of changes in the true parameters over time.

H. Path to a Terminal Action

In the stylized model, we assumed that the individual has an opportunity to leave at every decision point and used the value of leaving at the following decision point ($d + 1$) to obtain a closed form for $\bar{V}(z_{d+1})$. In some cases, an option to leave may not be available at every retention decision point. For example, some service members are offered an extension in the middle of their contract, and the available options are to accept the extension or to reject it and finish their current obligation.

In cases where the leave action is not available at $d + 1$, we can instead use the value of leaving at $d + 2$ to calculate $\bar{V}(z_{d+1})$. To see this, let $\tilde{a}_1 > 0$ be an arbitrary action at $d + 1$ and $\tilde{a}_2 = 0$ be the leave action at $d + 2$. We can then write the continuation value $\bar{V}(z_{d+1})$ as follows. First, since the terminal action is assumed to be unavailable at $d + 1$, at we replace it with $\tilde{a}_1$ when applying (5).

$$\bar{V}(z_{d+1}) = v(z_{d+1}, \tilde{a}_1) - \ln(p(\tilde{a}_1|z_{d+1})) + \gamma$$

Next, we apply (3) to replace $v(z_{d+1}, \tilde{a}_1)$ as usual.

$$\bar{V}(z_{d+1}) = U(z_{d+1}, \tilde{a}_1) + \beta^{\tilde{a}_1} \mathbb{E}[\bar{V}(z_{d+2})] - \ln(p(\tilde{a}_1|z_{d+1})) + \gamma$$

However, since $\tilde{a}_1$ is not a terminal action, this step now introduces a continuation value $\bar{V}(z_{d+2})$. Therefore, we apply (5) again to expand $\bar{V}(z_{d+2})$. Because a terminal action is assumed to be available at $d + 2$, this yields closed form for $\bar{V}(z_{d+1})$.

$$\bar{V}(z_{d+1}) = U(z_{d+1}, \tilde{a}_1) + \beta^{\tilde{a}_1} \mathbb{E}[U(z_{d+2}, 0) - \ln(p(0|z_{d+2})) + \gamma] - \ln(p(\tilde{a}_1|z_{d+1})) + \gamma$$

$\bar{V}$ is used to represent the continuation value, so it depends on the expected values of state variables at $d + 1$. Calculating $U(z_d, \tilde{a}_1)$ and $U(z_{d+1}, 0)$ requires the expectation over state variables at $d + 2$. Without a terminal action available at $d + 1$, the model now relies on state variables at $d + 2$, which is further in the future. More generally, we need to compute or simulate the expected state variables at the next decision point where the leave option is available, however far in the future this may be.

Additionally, we may want to avoid using $p(0|z_d)$ to identify $\bar{V}$ in some cases even when the leave option is available. Specifically, if $p(0|z_d)$ is near zero, then the value associated with z_d may not be well identified. For example, the probability of voluntarily leaving the military within a couple of years prior to retirement eligibility is exceptionally low. Consider a service member whose decision point d takes place when they have 16 YOS, whose available stay action at this point is to remain in service for three additional

years. Denote the state at this point as z_{16} , and the state at decision point $d + 1$ as z_{19} because it would take place at 19 YOS. Applying (3) and (5) as usual, the value of staying at 16 YOS is given by

$$v(z_{16}, 3) = U(z_{16}, 3) + \beta^3 \mathbb{E}[U(z_{19}, 0) - \ln(p(0|z_{19})) + \gamma]$$

However, suppose that $p(0|z_{19}) = 0.01$. The CCP estimation method uses model-predicted CCPs to iteratively update $\hat{p}(0|z_d)$ during the estimation process. These model-predicted probabilities are the result of a multinomial logit specification, which requires dramatic differences in values to produce such small CCPs. It is therefore likely that $\hat{p}(0|z_{19}) > p(0|z_{19})$ in cases like this. Suppose $\hat{p}(0|z_{19}) = 0.03$, which may not seem like a particularly large discrepancy. However, note that $\ln(0.01) = -4.6$ while $\ln(0.03) = -3.5$.

At very low probabilities, a small change in the CCPs can lead to a large difference in the continuation value, which in this example would bias $v(z_{16}, 3)$ and therefore $p(3|z_{16})$ downward. Because the likelihood is calculated by comparing observed decisions to the estimated probabilities of those decisions, this also biases the likelihood contributions of stay decisions made at YOS 16.

We can mitigate this issue by choosing a different arbitrary action $\tilde{a}_{19}$ at the 19 YOS decision point that is more representative of the actions of service members. Suppose that many service members making retention decisions at YOS 19 choose to stay for two years; we can then use $\tilde{a}_{19} = 2$ instead of a terminal action when expanding $v(z_{16}, 3)$.

$$\begin{aligned} v(z_{16}, 3) &= U(z_{16}, 3) + \beta^3 \mathbb{E}[v(z_{19}, 2) - \ln(p(2|z_{19})) + \gamma] \\ &= U(z_{16}, 3) + \beta^3 \mathbb{E}[U(z_{19}, 2) + \beta^2 \mathbb{E}[\bar{V}(z_{22})] - \ln(p(2|z_{19})) + \gamma] \end{aligned}$$

This introduces a continuation value $\bar{V}(z_{22})$ where previously there was only a utility function. However, if service members making retention decisions at YOS 22 often choose to leave, we can use a terminal action at that decision point to obtain a closed form solution.

$$\begin{aligned} v(z_{16}, 3) &= U(z_{16}, 3) \\ &\quad + \beta^3 \mathbb{E}[U(z_{19}, 2) + \beta^2 \mathbb{E}[U(z_{22}, 0) - \ln(p(0|z_{22})) + \gamma] - \ln(p(2|z_{19})) + \gamma] \end{aligned}$$

This adds to the computational burden to calculating $v(z_d, a)$ at some decision points, but it avoids identification of parameters using the sharp changes in the natural log function for values near zero.

I. Reducing the Computational Burden of the Outside Option

Using the terminal property of the leave action requires the utility of that action to be known. $U(z_d, 0)$ must be estimated for all possible states z_d from which a service member may choose to leave the military. Some applications normalize this value to zero for all z_d , but this normalization is not appropriate in our case. Current and future compensation is one of the most important factors in a retention decision. Expected civilian wages can differ

substantially based on the service member's occupation, training, and experience. Failing to account for this variation in wages may lead the model to incorrectly associate differences in retention probabilities with other related variables.

Calculating $\mathbb{E}[U(z_d, 0)]$ exactly for all z_d is impractical. Income depends on occupation, experience, non-work income, investment returns, tax credits, and many other factors. Our goal is to incorporate those factors that account for the largest predictable variation in expected income.

We can use data on a service member's MOS, training, and certifications to match them to likely civilian occupations. Data on the service member's current YOS and pay grade can be used to estimate military retirement pay, if any. In general, we estimate a reduced form function for civilian income (typically only wage income) using civilian data, and make assumptions about state variable transitions that allow us to calculate a service member's hypothetical lifetime wage trajectory in the civilian world.

This calculation is straight-forward but due to the size of the state space for z_d it is too computationally intensive to be feasible at each step of the estimation. Therefore, the utility of leaving is usually restricted to be a continuous and monotonic function of expected income; it can then be precomputed and re-used during estimation. We propose two potential simplifications to reduce the complications involved in estimating $U(z_d, 0)$.

First, it may be possible to cancel a portion of the value of the terminal action. This would reduce the computational burden of calculating the value of the terminal action, and may make it more feasible to relax other assumptions. Suppose that at some future time period, the distribution of *relevant* state variables conditional on choosing to leave at d is the same as the distribution of state variables conditional on choosing to stay at d and leave at $d + 1$. That is, let

$$f(\tilde{z}_{t_d+\rho} | z_{t_d}, a_d = 0) = f(\tilde{z}_{t_d+\rho} | z_{t_d}, a_d = y > 0, a_{d+1} = 0) \quad (22)$$

where $\tilde{z}$ contains only the state variables relevant to the value of the terminal action. This is similar to finite dependence, which we discuss in more detail in Section 3.J.

For example, if both d and $d + 1$ occur before the service member is eligible for retirement, then the YOS and pay grade at the time of separation is arguably not relevant because there is no military retirement pay to consider. If one or both decision points occurs after retirement eligibility, the annual military retirement pay could be added as a variable in $\tilde{z}$ to remove the dependence on YOS and pay grade.

If we can argue that (22) holds for some $\rho > y$, we can cancel part of the continuation value by splitting up $U(z_d, 0)$ in the alternative-specific value functions as follows:

$$\begin{aligned}
v(z_d, y) &= U(z_d, y) + \beta^y \mathbb{E}[U(z_{d+1}, 0) - \ln(p(0|z_{d+1})) + \gamma] \\
&= U(z_d, y) + \beta^y \sum_{\tau=t_{d+1}}^T \beta^{\tau-t_{d+1}} \mathbb{E}[u(z_\tau)|a_{d+1} = 0] - \beta^y \mathbb{E}[\ln(p(0|z_{d+1}))] + \beta^y \gamma \\
&= U(z_d, y) - \beta^y \mathbb{E}[\ln(p(0|z_{d+1}))] + \beta^y \gamma \\
&\quad + \beta^y \sum_{\tau=t_{d+1}}^{t_d+\rho-1} \beta^{\tau-t_{d+1}} \mathbb{E}[u(z_{t_{d+1}+k})|a_{d+1} = 0] \\
&\quad + \beta^y \sum_{\tau=t_d+\rho}^T \beta^{\tau-t_{d+1}} \mathbb{E}[u(z_\tau)|a_{d+1} = 0]
\end{aligned}$$

Noting that $t_d + y = t_{d+1}$ and therefore $y + \tau - t_{d+1} = \tau - t_d$ and rewriting expected utility as the integral with respect to the distribution of relevant state variables, the last term becomes:

$$\sum_{\tau=t_d+\rho}^T \beta^{\tau-t_d} \int u(\tilde{z}_\tau) f(\tilde{z}_\tau|z_{t_d}, a_d = y, a_{d+1} = 0) d\tilde{z}_\tau \quad (23)$$

Next, we expand the value of leaving at t_d .

$$\begin{aligned}
v(z_d, 0) &= U(z_d, 0) \\
&= \sum_{\tau=t_d}^T \beta^{\tau-t_d} \mathbb{E}[u(z_\tau)|a_d = 0] \\
&= \sum_{\tau=t_d}^{t_d+\rho-1} \beta^{\tau-t_d} \mathbb{E}[u(z_\tau)|a_d = 0] + \sum_{\tau=t_d+\rho}^T \beta^{\tau-t_d} \mathbb{E}[u(z_\tau)|a_d = 0]
\end{aligned}$$

Note that the last term can be rewritten as:

$$\sum_{\tau=t_d+\rho}^T \beta^{\tau-t_d} \int u(\tilde{z}_\tau) f(\tilde{z}_\tau|z_{t_d}, a_d = 0) d\tilde{z}_\tau \quad (24)$$

If (22), then (23) is equal to (24) and we are left with a closed form for $v(z_d, a') - v(z_d, a)$ (and therefore $p(a|z_d)$) that does not require us to calculate the expected lifetime utility of the terminal action past $t_d + \rho$.

If we do not believe (22) holds for some reasonable value of ρ , it may still be possible to simplify the calculation of the outside option. We propose replacing the sums in (23) and (24) with a parametric function $h(z_{t_d}, a_d)$ that captures the most important differences in $\tilde{z}_{t_d+\rho}$ that result from taking different actions at t_d . For example, suppose that the service member is at 17 YOS at t_d ; leaving at this point will result in the service member receiving no military retirement benefit, while staying for at least three years will result in eligibility

for retirement. In this case, a significant difference in future state variables persists for the rest of the individual's life, and there is no ρ such that (22) plausibly holds.

Instead of canceling the time periods from $t_d + \rho$ to T , we could summarize the relevant differences using $h(z_{t_d}, a) = \theta^r \mathbb{1}(YOS_{t_d} + a \geq 20)$. This function uses a simple retirement eligibility indicator; for $a = 0$ the indicator is equal to 1 if the individual is eligible to retire at the current decision point, while for $a > 0$ the indicator represents retirement eligibility if the individual were to leave at the next decision point.

Substituting $h(z_{t_d}, a)$ in for the last term in $v(z_d, a)$, we now have a closed form for $v(z_d, a') - v(z_d, a)$ that again requires us to calculate only ρ periods of transitions and utility past the current decision point. A single indicator for retirement eligibility may be too simple, but the principle holds. $U(z_d, 0)$ is not required to be a sum of annual utilities, so we can significantly reduce the computational burden of calculating $U(z_d, 0)$ for all z_d by using a parametric function to approximate part of the value.

J. Modeling the Reserve Participation Decision

In addition to expanding the set of stay options, we may want to understand more about service members' decisions to participate in the reserves after leaving active duty. The baseline model assumes that leaving active duty is a terminal action, after which service members do not make any more decisions in the dynamic model. Here, we consider two ways to relax this assumption to accommodate the reserve participation decision.

First, suppose that active duty service members can choose to join the reserves or become a civilian upon leaving active duty, and each year service members in the reserve component choose whether to continue in the reserves or become a civilian. Additionally, suppose that they cannot rejoin the active or reserve components once they leave to become a civilian.

In this example, we might have separate value functions for service members in the active and reserve components. Let v^A be the alternative-specific value function for stay actions; let R represent the choice to join the reserves and C be the choice to become a civilian. Choosing to leave to become a civilian is still a terminal action, so we can use this property to obtain a closed form for the continuation value just as before.

$$v(z_d, a \in A) = U(z_d, a) + \beta^a \mathbb{E}[U(z_{d+1}, C) - \ln(p^A(C|z_{d+1})) + \gamma]$$

Note that the state z_d must now track whether the service member is in the active or reserve components because this may influence the current period utility and the probability of leaving to become a civilian ($p(C|z_{d+1})$). We add a superscript on the CCP to emphasize this distinction. We can also use the terminal action property to write the alternative-specific value of choosing to join (or stay in) the reserves.

$$v(z_d, R) = U(z_d, R) + \beta \mathbb{E}[U(z_{d+1}, C) - \ln(p^R(C|z_{d+1})) + \gamma]$$

Participation decisions for reservists are assumed to be annual, so we do not need to use a to discount the continuation value or aggregate utility. The value of the civilian option is unchanged from the baseline model.

$$v(z_d, C) = U(z_d, C)$$

The value of being in state z_d prior to choosing an action is equal to the maximum value of the available actions. Therefore, the choice probabilities take the same form as in the baseline model, for $a \in A_A = \{A, R, C\}$.

$$p^A(a|z_d) = \frac{\exp(v(z_d, a))}{\sum_{a' \in A_A} \exp(v(z_d, a'))}$$

In our example, reserve component service members do not have the option to re-enter the active component, so the set of available actions is $A_R = \{R, C\}$.

$$p^R(a|z_d) = \frac{\exp(v(z_d, a))}{\sum_{a' \in A_R} \exp(v(z_d, a'))}$$

Depending on the service member population or the policy question of interest, it might be undesirable to assume that a service member can never rejoin the reserves after becoming a civilian. If we include the reserve participation decision in the same model as we did earlier in this section, we can no longer use the terminal action property to obtain a closed form for $\bar{V}$.

However, we can split the reserve participation decision into a separate second model, and use a property called finite dependence to solve it. Actions a and a' are defined to be ρ -period dependent if there exist sequences of actions starting with a and a' that lead to the same distribution of state variables ρ periods in the future.

For example, consider the decision of a current reservist. They can continue as a reservist next year or become a civilian; because military status (active, reserve, civilian) is part of the state, these two actions lead to different states at $d + 1$. Comparing the sequences of actions $\{R, C\}$ and $\{C, C\}$, note that the state at $d + 2$ is different even if the action at $d + 1$ is C in both cases because years of service is an important component of the state space.

Sequences $\{R, C\}$ and $\{C, R\}$ lead to the same years of service at $d + 2$, but still differ in the military status upon entering $d + 2$. However, sequences $\{R, C, C\}$ and $\{C, R, C\}$ satisfy finite dependence: the distribution of state variables at $d + 3$ is the same in both cases. We can iterate with equation (5) to expand the alternative-specific value function as follows:

$$\begin{aligned}
v^R(z_d, R) &= U(z_d, R) + \beta \mathbb{E}_d[\bar{V}^R(z_{d+1})] \\
&= U(z_d, R) + \beta \mathbb{E}_d[v(z_{d+1}, C) - \ln(p^R(C|z_{d+1})) + \gamma] \\
&= U(z_d, R) + \beta \mathbb{E}_d \left[[U(z_{d+1}, C) + \beta \mathbb{E}_d[\bar{V}^C(z_{d+2})]] - \ln(p^R(C|z_{d+1})) + \gamma \right]
\end{aligned}$$

Iteration of this expansion yields the following equation for the value of choosing R at d , using the sequence of actions $\{R, C, C\}$.

$$\begin{aligned}
v^R(z_d, R) = U(z_d, R) &+ \beta \mathbb{E}_d[U(z_{d+1}, C) - \ln(p^R(C|z_{d+1})) + \gamma] \\
&+ \beta^2 \mathbb{E}_d[U(z_{d+2}, C) - \ln(p^C(C|z_{d+2})) + \gamma] \\
&+ \beta^3 \mathbb{E}_d[\bar{V}^C(z_{d+3})]
\end{aligned}$$

Similarly, we can write the value of choosing C at d using the sequence $\{C, R, C\}$.

$$\begin{aligned}
v^R(z_d, C) = U(z_d, C) &+ \beta \mathbb{E}_d[U(z_{d+1}, R) - \ln(p^C(R|z_{d+1})) + \gamma] \\
&+ \beta^2 \mathbb{E}_d[U(z_{d+2}, C) - \ln(p^R(C|z_{d+2})) + \gamma] \\
&+ \beta^3 \mathbb{E}_d[\bar{V}^C(z_{d+3})]
\end{aligned}$$

The continuation value $\beta^3 \mathbb{E}_d[\bar{V}^C(z_{d+3})]$ is the same in both cases (the individual is currently a civilian and has one additional year of reserve service compared to z_d). This portion of the continuation value will cancel when we take the difference between the two alternative-specific values to calculate the choice probabilities. The finite dependence assumption has now allowed us to obtain closed forms for the choice probabilities in the reserve participation decision model.

In the active duty retention model, we can treat reserve/civilian as a single terminal action, just as we did before in the motivating example and the baseline model. This may require some changes to the value that we calculate for $U(z_d, C)$, but the structure of this model is fundamentally unchanged.

4. Other Considerations

We have described a structural DDC model of military retention and a number of possible extensions that fit into the simple CCP estimation framework. There are some other considerations that are relevant to determining whether the structural DDC approach is appropriate in a given application, and what we may need to keep in mind if we chose to implement it.

The first consideration is identification, which we briefly mentioned earlier in the paper. It is important to understand what we can and cannot learn from the data that we have and what assumptions allow us to do so. It is also necessary to consider alternative estimation strategies and approximations, as even with the CCP method the computational requirements for estimating these models can become burdensome. Finally, testing the degree to which our assumptions about rationality and time preferences – and the presence or absence of peer effects – change our results may be an important part of applying a DDC model in certain situations.

A. Identification of Structural Parameters

We are particularly concerned with identification of the structural parameters of the model. The usefulness of the model in counterfactual policy analyses depends on our ability to recover the true parameters that determine individual behavior. As such, it is important to understand what assumptions and what variation in the data provide the information that we use to estimate our model. There are two practical concerns here: first are the assumptions that we need to make to ensure identification justifiable, and second, do we have the kinds of data we need to identify the parameters we are interested in. In practice, we specify structural forms for our models while considering the limitations implied by the non-parametric identification results to ensure we do not rely on the structural form assumptions beyond what is necessary.

The key identification result for DDC models is given by Magnac and Thesmar (2002). These models are not generically non-parametrically identified. In the absence of unobserved heterogeneity, identification requires specifying the distribution of unobserved preference shocks, the discount factor, and the value of one reference alternative. Semi-parametric identification results describing the case where the payoff function is known up to a set of parameters and the preference shock distribution is non-parametric are available in Buchholz, Shum, and Xu (2021). We make the assumptions required for identification by specifying the distribution of ϵ , the value of β , and normalizing the intercept in the utility of the outside option to zero. These assumptions guarantee that the parameters of the utility functions are uniquely identified from the data.

Identification in the presence of unobserved heterogeneity is more difficult. Heckman and Singer (1984) show that mixture models are not generically identified, resulting in inconsistent estimators of the structural parameters. This makes it crucially important to correctly non-parametrically identify the distribution of unobserved heterogeneity. Fox et al. (2012) give conditions for the identification of the mixed logit model and Kasahara and Shimotsu (2009) give the conditions for general dynamic models.

In both cases, identification relies on sufficient variation in the characteristics of alternatives that induces variation in the actions of individuals based on the heterogeneity. In the dynamic case studied by Kasahara and Shimotsu (2009), a panel of three periods is required for this variation to identify the distribution of unobserved heterogeneity.

This leaves the practical concern of what sort of parameters can be identified by the available data. An important consideration is that DDC models are forward-looking, so only parameters associated with “forecastable” variables can be estimated.¹ Consider, for example, the question of whether service members prefer a specific location over another.

In order to estimate this preference, the researcher must forecast the transition of locations, and the forecasted locations must capture meaningful variation in observed outcomes. That is, expected outcomes must be different for some service members versus others. The estimated $\hat{f}(x_{d+1}|x_d, a_d)$ shows us whether the expectation of a particular variable varies across service members in a meaningful way. After estimating these transitions, we can use them to understand what specifications and parameters we can expect our data to be informative about.

The following question often comes up: how can we separately identify the effect of a covariate x on the per-period utility $u(x)$ and its effect on the initial distribution of unobserved types $\pi(x)$? Let us try to offer some intuition through an example.

Suppose x is gender and we want to know whether women (who retain at lower rates than men) receive an additional disutility from being in the military or are more likely to be of the type s that has a low taste for military service. If the effect of gender is via the per-period utility, then we would expect to see women retain at lower rates than men at all points in the career. If the effect is via the unobserved types, we would expect more women to leave at the first opportunity and the retention rates to eventually equalize. This distinction is important because it helps us understand the behavior and the potential impacts of policies on service members of different genders.

¹ By forecastable, we mean that the data must contain information that facilitates construction of a prediction that is more accurate than a random guess.

B. Identification of the Discount Factor

Though we typically specify the value of the discount factor, in theory it is possible to estimate the discount factor jointly with the other structural parameters. Identification of the exponential discount factor is addressed by Abbring and Daljord (2020). This paper applies an exclusion restriction on utilities to extend the identification strategy of Magnac and Thesmar (2002). Under the required exclusion restriction, there exist known values a , a' , z , z' , and c such that

$$U(z, a) = U(z', a') + c, \quad (25)$$

where either $a \neq a'$ or $z \neq z'$, or both. If this exclusion restriction is satisfied, a finite set of discount factors is identified.

The authors show that the set of discount factors can be characterized by a set of moment conditions that depend only on choice probabilities and transition probabilities. Calculating the identified set of β values requires the full Markov matrix of state transitions, which implies that the state must be discrete. The size of the identified set of β may be reduced by imposing additional exclusion restrictions. In addition, models with single-action finite dependence have no more than ρ values in the identified set of β .

To impose the exclusion restriction in (25), we would need to select two states z and z' , actions a and a' , and assume a value for the utility differential c . However, given the size of the state space and the kinds of information that it may contain in retention applications, these assumptions are difficult to justify. The size of the state space also makes working with the full Markov transition matrix difficult, and suggests that the set of β values identified by the moment condition is likely to be large. Therefore, we typically do not attempt to estimate β even though it fits within the Arcidiacono and Miller (2011) estimation strategy. Instead, we can estimate the model and compare the counterfactual predictions under different assumptions on β .²

C. Estimation in Two Stages

Even when using the CCP estimation strategy, estimating complex DDC models can be a computational challenge. One promising approach is to use a two-stage estimation procedure that deals with unobserved heterogeneity in the first stage and allows us to vary the utility specification in the second stage. This two-stage implementation of the Arcidiacono and Miller (2011) estimator recovers the CCPs and the distribution of unobserved heterogeneity in the first stage and the structural parameters in the second stage.

² In addition to different values for the exponential discount factor, we can estimate the model and make counterfactual predictions under alternative discount structures such as present-biased time preferences.

The second stage allows for an easy way to test multiple structural specifications and can be used with MLE, simulation, or other estimation methods. This method requires the CCPs and the unobserved types to be identified from the data without the assumptions imposed by the structural model. Discussions of different conditions that lead to identification can be found in Arcidiacono and Jones (2003), Kasahara and Shimotsu (2009), and Fox et al. (2012).

The first stage of the estimation iterates on (14), (15), and (16) until convergence. This is essentially a repeated application of the “E” step from the EM algorithm. In practice, we follow Arcidiacono et al. (2016) and specify a reduced form equation for $\hat{p}(a|x_{i,d}, s_i)$, say $\tilde{p}(a|x, s)$. This is not our structural model because it does not impose any of the intertemporal or optimality constraints.

We use this reduced form $\tilde{p}$ to impose the condition that observationally equivalent service members, meaning those with the same x , should behave in the same way. If they do not, then the difference is attributed to the unobserved heterogeneity s . In cases where the unobserved type affects an additional observed outcome, such as if the transitions $\hat{f}(x_{d+1})$ depend on s , the structural parameters of this outcome equation can be estimated in this step as well.

The first stage provides us with an estimate of the CCPs, conditional both on observed and unobserved variables, the population distribution of types $\hat{\pi}(s)$, and the individual-specific probability distribution of types $\hat{q}_i(s)$. At this point we treat these parameters as known and proceed to the estimation of the structural parameters in the second stage. Any CCP estimator that works without unobserved heterogeneity can now be used to estimate θ . We form the likelihood just as described in the main part of the paper and estimate $\hat{\theta} = \arg\max_{\theta} \sum_i \sum_s \sum_d q_i(s) \ln l(a_{i,d}|x_{i,d}, s_i, \tilde{p}(0); \theta)$.

This simplicity of the second stage allows us to explore many potential utility functions $U(x, s, a)$. We can consider more variables and functional forms than typically possible in a one-stage estimation. If desirable, it should be possible at this point to incorporate variable selection, regularization, and other functional approximation techniques.

A further additional benefit of the two-stage approach is the possibility of using simulation methods in the second stage. When finite dependence requires more than one period, the form of the likelihood function can become unwieldy. Bajari, Benkard, and Levin (2007) propose a forward simulation approach based on the differences in conditional value functions.

Recall from (5) that $\bar{V}(x, s) = v(x, s, \tilde{a}) - \ln(p(\tilde{a}|x, s)) + \gamma$. This implies that the difference between any two conditional value functions can be written as $v(x, s, \tilde{a}) - v(x, s, \tilde{b}) = \ln(p(\tilde{b}|x, s)) - \ln(p(\tilde{a}|x, s))$. These form the basis of moment conditions

for a Generalized Method of Moments (GMM) estimator. If finite dependence does not hold, then this estimator is still feasible using simulations until the end of the decision process or until the discount rate renders future values insignificant.

D. Approximation Methods

The CCP method of estimating DDC models significantly reduces this computational challenge, but is still subject to the usual curse of dimensionality. An alternative approach is to use approximations to specify the value functions or to evaluate them at a small grid of points.

There are two broad categories of approximation methods that have been developed to deal with models that require a large state space. The first attempts to approximate the value function using a lower-dimension representation. Examples of this approach include the sieve value function iteration of Arcidiacono et al. (2012) and the neural network approximations of Norets (2012). Related to this approach are methods of interpolating the value function, such as the well-known method of Keane and Wolpin (1994) and the sparse grid method from Brumm and Scheidegger (2017). Taking the approximation idea one step further, Bernal and Keane (2010) use a quasi-structural method to approximate the decision rules from a structural model using a reduced form method.

The other method attempts to reduce the number of times the value function needs to be evaluated. Keane and Wolpin (2001) propose a simulation method that can be combined with CCP methods to reduce the number of outcome histories that need to be considered. Bajari, Benkard, and Levin (2007) adapt this method to a setting with dynamic games. This approach is particularly attractive in cases where the model exhibits finite dependence. If finite dependence is satisfied after k periods, then only k periods of simulations are needed.

E. Rationality and Time Preferences

A common critique of economic models, particularly by those outside the field, is that the models assume individuals act rationally. It is true that some amount of rationality on the part of the individuals is required for the models to make sense, but this assumption can generally be summed up as “individuals have preferences and act on them.” Specifically, economic rationality boils down to two properties of individual preferences: completeness (existence of a preference ranking over alternatives), and transitivity (preservation of the ranking of those alternatives) (Mas-Colell et al. 1995).

Service members in our models may, for example, believe that they will make more money in the civilian sector or that a promotion will improve their quality of life, whether this is true or not. Strictly speaking, we are estimating the utility that service members expect to receive from different characteristics of their careers. We do not expect or impose complete rationality from them.

One particular testable assumption of rationality is that individuals are forward-looking. Service members in the model understand and take into account the likely impact of their actions today on their future careers. A higher discount rate, hyperbolic discounting, or a model of myopic behavior can describe a situation where service members do not care as much (or at all) about the future.

Arcidiacono, Sieg, and Sloan (2007) test models of drinking and smoking with different degrees of rationality and find that decisions are better explained by models where individuals are rational. We can perform similar comparisons to evaluate whether the degree of rationality and forward-looking in our models is consistent with the behavior of service members we observe in the data.

Fang and Wang (2015) discuss estimation of quasi-hyperbolic time preferences. Quasi-hyperbolic discounting is one way to represent present-biased preferences; it adds a present-bias discount factor δ to the standard exponential discounting framework.³ The present value of lifetime utility is then given by:

$$u(z_t) + \delta \sum_{k=t+1}^T \beta^{k-t} u(z_k)$$

Individuals are said to be sophisticated if they know that $\delta < 1$ and make decisions accordingly. Sophisticated individuals understand that their present bias will affect their future decisions, and adjust their continuation value accordingly. Those who believe that their future decisions will be made under time-consistent preferences are said to be naïve. Naïve individuals believe that their present-bias discount factor is $\tilde{\delta} = 1$. Individuals may also be partially naïve; they may understand that they are present-biased but believe that they will be less present-biased in the future ($\delta < \tilde{\delta} < 1$).

The exclusion restriction that gives identification of quasi-hyperbolic preferences requires that utility is time-separable and that there exists a variable in the state space that affects transition probabilities but does not affect current utility. Suppose states $z_t = z$ and $z_t = z'$ differ only in the exclusion restriction variable. Then the current period utilities are equal; $u(z, a) = u(z', a)$. If $p(a|z) \neq p(a|z')$, then the difference in choice probabilities can be attributed to differences in expected future utility. The extent to which differences in expected future states affect current decisions can be used to identify the time preference. Translating this exclusion restriction into our model, we would require a state variable such that

³ Because we use β as the exponential discount factor, this notation is the opposite of the standard (β, δ) notation used in most papers with quasi-hyperbolic discounting.

$$U(z, a) = \sum_{\tau=t}^{t+a-1} \beta^{\tau-t} \mathbb{E}[u(z_\tau)|z, a] = \sum_{\tau=t}^{t+a-1} \beta^{\tau-t} \mathbb{E}[u(z_\tau)|z', a] = U(z', a)$$

for all $a \in A$, and

$$f(z_{t+k}|z, a) \neq f(z_{t+k}|z', a)$$

for some $a \in A$ and $k \geq a$.

F. Peer Effects

It is reasonable to wonder if the decisions and attitudes of peers play a role in service members' retention decisions. If everyone else in your cohort plans to leave, then you may be more likely to leave as well. If these peer effects are present, then policy may be able to induce a shift from a bad equilibrium in which everyone wants to leave because everyone else is leaving to a good equilibrium where everyone wants to stay because everyone else is staying. Unfortunately, the identification of peer effects is notoriously tricky due to the "reflection problem" articulated by Manski (1993): when group formation is endogenous, the effects of peers are difficult or impossible to separate from selection into the group.

In a military setting we may observe that service members in a specific field are highly likely to retain. Is this due to the peer effect or are service members who like the military more likely to select into this field? We may be able to test peer effects in settings where groups are formed exogenously, such as service member assignment to specific units.

Additionally, Bramoullé, Djebbari, and Fortin (2009) show that peer effects may be identified through social networks, and we could use data on which service members served with other service members to create such a social network for a military cohort. Dennis et al. (2021) and Eliezer et al. (2021) provide more information on constructing professional networks of service members. Finally, Davezies, d'Haultfoeuille, and Fougère (2009) indicate how to identify these effects using the variation in group size and the variance of outcomes.

Most of the applied research examining peer effects has focused on educational settings. We are unaware of any study of military retention that incorporates peer effects, but we strongly recommend that future studies examine the effects of peers on retention and other personnel issues among service members.

This page is intentionally blank.

Appendix A. Recent Military Applications of DDC

The military's detailed historical personnel data provides a unique opportunity for estimating dynamic discrete choice (DDC) models of individual decisions. However, whether due to data access restrictions or some other constraint, the academic community outside of Federally Funded Research Development Centers (FFRDCs) has largely stayed away from studies of the military. In recent years, the only significant applications of dynamic discrete choice models to study military career decisions have been at FFRDCs.

Since the 1980s, RAND has applied their Dynamic Retention Model (DRM) to analyze retention incentives across the DOD. In general, the DRM assumes that military retention decisions are influenced only by compensation and an individual's "taste" for military service. DRM studies often assess potential compensation policies by simulating the survival curve that is expected to occur under the policy in steady state.

Recently, CNA built the Dynamic Decision Model (DDM) to estimate the retention effects of both monetary and non-monetary personnel policies. The DDM includes many demographic and career variables as state variables, and is typically used to predict retention effects of potential policies among specific subgroups of service members.

In 2015, IDA developed the Military Career Analysis Model (MCAM), which includes a modified version of the DRM called the Observed-Performance Dynamic Retention Model (OPDRM). The goal of this model was to combine a promotion model, a dynamic discrete choice retention model, a life-cycle cost model, and an accession model into a single wholistic model for studying military careers.

A. RAND DRM

The model and estimation strategy developed by Mattock and Arkes (2007) introduces the framework for RAND's modern applications of the DRM. The authors note that at the time, "the DRM is not widely used to analyze manpower policy questions, in part because of the computational complexity of the model." However, with the advances in computing power that have taken place over the last two decades, at least 15 RAND studies have estimated or applied the DRM since 2013.

The model incorporates a nested stay decision that allows the service member to choose an obligation length and receive a retention bonus that depends on the obligation

chosen. The authors have data on military participation but not actual decisions; that is, the data indicate whether an individual is in the Active Component or the Reserve component, but do not contain information on retention bonus contracts or remaining obligation. Therefore, the estimation strategy is to maximize a likelihood that represents the probability of observing a service member staying for a given number of years.

The most significant step toward the current DRM occurs in Asch, Mattock, and Hosek (2013). This study introduces a nested leave option that models the decision whether to participate in the Reserves after leaving active duty. The model also includes distinct taste parameters for Active and Reserve service. These taste parameters are assumed to be distributed according to a bivariate normal distribution across the population of service members; the distribution of taste for service is assumed constant across entry cohorts and its parameters are estimated jointly with the other structural parameters.

Because the DRM does not include any demographic or career characteristics, the taste for service is assumed to capture all non-monetary utility variation across individuals. Finally, the model incorporates a switching cost to capture the fact that service members cannot easily choose to leave active duty before the end of their initial service commitment.

This version of the DRM was used over the next several years to estimate the effects of retention bonuses in various officer communities. Mattock et al. (2016) estimate the model using data on Air Force pilots; Mattock et al. (2019) use the estimated model to assess cost-effectiveness of retention versus accession. These studies are specifically concerned with pilot retention in an economy where civilian demand for pilots is growing, and they carefully model the demand for and wages of pilots employed by major airlines.

In addition to switching costs incurred if the service member leaves before completing their initial active duty service commitment or total service obligation, the model includes switching costs for moving to the reserves from either the active component or the civilian sector. Hosek et al. (2017) apply the same model to study the retention effects of special pays for officers in mental health care professions across the services, and Asch et al. (2019) use it to study retention bonuses in the Army and Navy Special Operations Forces.

More recent versions of the DRM often assume that service members make annual decisions, and do not include nested stay decisions for obligations beyond one year. Asch, Mattock, and Hosek (2019) are interested in reserve participation under the new Blended Retirement System, and extend the model to differentiate between participation in the Army Reserve and the National guard. They estimate the model using data on enlisted service members and officers that entered the Army in 1990-1991.

Asch, Mattock, and Tong (2020) estimate the DRM in order to simulate the retention effects of a time-in-grade (rather than a time-in-service) pay table and evaluate the impact on the ability distribution of the force. The model is extended to include pay grade as a

stochastic state variable; expected military pay now depends on YOS and pay grade.⁴ Promotion opportunities are assumed to occur at known intervals throughout a service member's career; an individual who fails to promote never receives another chance.

All other state variables are deterministic, and there are only two possible values for next year's pay grade at any decision point. Therefore, the continuation value is simply the weighted average of the continuation values associated with these two states. To assess the performance incentives associated with Selective Reenlistment Bonuses (SRBs) for enlisted Soldiers, Asch et al. (2021) again estimate the DRM with pay grade as a state variable. This study makes several assumptions about how innate ability and effort affect the timing and probability of promotion, and simulate the effects of increasing SRBs on the ability distribution and effort choices of the force.

B. CNA DDM

The DDM was first developed by Huff et al. (2018) to evaluate whether the Blended Retirement System (BRS) Continuation Pay can offset the anticipated decline in Navy enlisted retention due to the other aspects of BRS. The model assumes that sailors make retention decisions at the end of each contract, and is estimated on observed choices at these identified decision points. At each decision point, a sailor has seven alternatives available: leave the Navy, or sign a new contract for one to six additional years of obligation.

The value of each alternative depends on the utility the individual expects to receive and a continuation value representing the opportunity to re-optimize in the future. In addition to expected future military pay, the model includes several non-monetary state variables (e.g., race, gender, marital status) directly in the utility function. These state variables, along with those that determine pay, are simulated so that the model accounts for uncertainty in outcomes such as future promotion probability and timing. The model is estimated using the CCP estimation strategy described in Arcidiacono and Miller (2011), though it does not yet include permanent unobserved heterogeneity across individuals.

An extended version of the DDM is presented in Levy et al. (2020), where it is used to assess the effectiveness of retention incentives among officers in Navy medical specialties. The model is extended to include heterogeneity in taste for service across individuals as a mixture of discrete types, following Heckman and Singer (1984).

Since this study is concerned with officers, who do not require enlistment contracts to remain in the service, retention decisions are assumed to occur once per tour or at the

⁴ Previous versions of the DRM use average pay by YOS.

end of a service obligation.⁵ The updated model also incorporates a careful treatment of the officer's expected civilian wage, which accounts for the high wages of medical professionals as well as the wide variation across occupations.

C. IDA OPDRM

Doyle (2015) developed a multi-part model called the Military Career Analysis Model (MCAM); the objective was to combine several related aspects of personnel modeling into one unified framework. The MCAM included a dynamic discrete choice retention model based on the DRM, called the Observed-Performance Dynamic Retention Model (OPDRM). The OPDRM removes the taste parameter found in the DRM and replaces it (conceptually) with a time-varying measure of individual match quality based on the observed time to promote to the current paygrade.

Having removed the persistent unobserved heterogeneity, the model is estimated using the CCP method of Hotz and Miller (1993), which significantly reduces the computational burden of estimation relative to the DRM. MCAM was used to simulate future force profiles for a subsequent IDA study on the feasibility of a proposed training policy.

D. Model Comparisons

Evaluating the impact of compensation and personnel policies on retention decisions across the military career requires a model that can predict how future service members will respond to a prospective policy. Until recently, models of retention decisions focused either on long-run equilibrium effects or on very near-term projections. The DRM is an example of the former, typically requiring a complete panel of data to estimate long-run impacts of changes in compensation. The DDM is an example of a near-term approach, using fewer years of data to predict retention outcomes in the immediate future.

This paper develops two methodological advances to combine the advantages of both approaches. The extended model provides a means to incorporate information from unobserved early-career retention decisions, alleviating the need for a complete panel. The same principle can be applied to estimate a model of non-retention career decisions, such as BRS enrollment, jointly with retention decisions.

This model also adjusts for extremely low probabilities of certain decisions – like choosing to leave immediately before becoming eligible for military retirement benefits – to reduce bias in estimates of the structural parameters of interest. These advances enable

⁵ Service obligations include the initial obligation, obligation incurred as a result of additional medical education or training, and obligation associated with a retention bonus contract.

the model to make credible near-term and long-run equilibrium predictions about retention outcomes under alternative compensation and personnel policy scenarios.

This page is intentionally blank.

Appendix B. References

- Abbring, Jaap H., and Øystein Daljord. “Identifying the Discount Factor in Dynamic Discrete Choice Models.” *Quantitative Economics* 11, 2 (2020): 471–501. <https://onlinelibrary.wiley.com/doi/full/10.3982/QE1352>.
- Aguirregabiria, Victor, and Pedro Mira. “Dynamic Discrete Choice Structural Models: A Survey.” *Journal of Econometrics* 156, 1 (2010): 38–67. <https://www.sciencedirect.com/science/article/pii/S0304407609001985>.
- Arcidiacono, Peter, Esteban Aucejo, Arnaud Maurel, and Tyler Ransom. “College Attrition and the Dynamics of Information Revelation.” Vol. Working Paper 22325 (2016). NBER Working Paper Series. <https://www.nber.org/papers/w22325>.
- Arcidiacono, Peter, Patrick J. Bayer, Federico A. Bugni, and Jonathan James. “Approximating High-Dimensional Dynamic Models: Sieve Value Function Iteration.” *NBER Working Paper*, no. w17890 (2012). <https://www.nber.org/papers/w17890>.
- Arcidiacono, Peter, and John Bailey Jones. “Finite Mixture Distributions, Sequential Likelihood, and the EM Algorithm.” *Econometrica* 71, 3 (2003): 933–46. <https://doi.org/10.1111/1468-0262.00431>.
- Arcidiacono, Peter, and Robert A. Miller. “Conditional Choice Probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity.” *Econometrica* 79, 6 (2011): 1823–67. <https://doi.org/10.3982/ECTA7743>.
- Arcidiacono, Peter, Holger Sieg, and Frank Sloan. “Living Rationally Under the Volcano? An Empirical Analysis of Heavy Drinking and Smoking.” *International Economic Review* 48, 1 (2007): 37–65. <https://doi.org/10.1111/j.1468-2354.2007.00417.x>.
- Asch, Beth J., Michael G. Mattock, and James Hosek. “A New Tool for Assessing Workforce Management Policies over Time: Extending the Dynamic Retention Model.” Santa Monica, CA: RAND Corporation, 2013. https://www.rand.org/pubs/research_reports/RR113.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.” Santa Monica, CA: RAND Corporation, 2019. https://www.rand.org/pubs/research_reports/RR1796.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.” Santa Monica, CA: RAND Corporation, 2020. 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.” Santa Monica, CA: RAND Corporation, 2021. https://www.rand.org/pubs/research_reports/RRA803-1.html.

Asch, Beth J., Michael G. Mattock, and James Hosek. “Effects of the Blended Retirement System on United States Army Reserve Participation and Cost.” Santa Monica, CA: RAND Corporation, 2019. https://www.rand.org/pubs/research_reports/RR2591.html.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. “Estimating Dynamic Models of Imperfect Competition.” *Econometrica* 75, 5 (2007): 1331–70. <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0262.2007.00796.x>.

Bernal, Raquel, and Michael P. Keane. “Quasi-Structural Estimation of a Model of Childcare Choices and Child Cognitive Ability Production.” *Journal of Econometrics* 156, 1 (2010): 164–89. <https://www.sciencedirect.com/science/article/pii/S0304407609002140>.

Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin. “Identification of Peer Effects Through Social Networks.” *Journal of Econometrics* 150, 1 (2009): 41–55. <https://www.sciencedirect.com/science/article/pii/S0304407609000335>.

Brumm, Johannes, and Simon Scheidegger. “Using Adaptive Sparse Grids to Solve High-Dimensional Dynamic Models.” *Econometrica* 85, 5 (2017): 1575–1612. <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA12216>.

Buchholz, Nicholas, Matthew Shum, and Haiqing Xu. “Semiparametric Estimation of Dynamic Discrete Choice Models.” *Journal of Econometrics* 223, 2 (2021): 312–27. <https://www.sciencedirect.com/science/article/pii/S0304407620303274>.

Davezies, Laurent, Xavier d’Haultfoeuille, and Denis Fougère. “Identification of Peer Effects Using Group Size Variation.” *The Econometrics Journal* 12, 3 (2009): 397–413. <https://doi.org/10.1111/j.1368-423X.2009.00296.x>.

Dennis, John W., Nathaniel T. Cleaves, Heidi C. Reutter, Dina Eliezer, and J. Brandon Carter. *Understanding and Assessing Unit Cohesion: Technical Supplement*. IDA Paper P-20508. Alexandria, VA: Institute for Defense Analyses, 2021.

Doyle, Colin M. *The Military Career Analysis Model: Theory, Methodology, Estimation, and Application of a Unified Model of Military Personnel Accession, Retention,*

Promotion, and Life-Cycle Cost. IDA Paper P-5264. Alexandria, VA: Institute for Defense Analyses, 2015.

Eliezer, Dina, Joseph F. Adams, Amy A. Alrich, Dave I. Cotting, Nathaniel T. Cleaves, Jay Dennis, Edward P. Donnelly, et al. *Understanding and Assessing Unit Cohesion at Multiple Organizational Levels*. IDA Paper P-20508. Alexandria, VA: Institute for Defense Analyses, January 2021.

Fang, and Wang. “Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting, with an Application to Mammography Decisions.” *International Economic Review* 56, 2 (2015): 565-596.

<https://onlinelibrary.wiley.com/doi/abs/10.1111/iere.12115>.

Fox, Jeremy T., Kyoo il Kim, Stephen P. Ryan, and Patrick Bajari. “The Random Coefficients Logit Model Is Identified.” *Journal of Econometrics* 166, 2 (2012): 204–12. <https://www.sciencedirect.com/science/article/pii/S0304407611001655>.

Gotz, Glenn A., and John McCall. “A Dynamic Retention Model for Air Force Officers: Theory and Estimates.” Santa Monica, CA: RAND Corporation, 1984. <https://www.rand.org/pubs/reports/R3028.html>.

Heckman, James, and Burton Singer. “A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data.” *Econometrica* 52, 2 (1984): 271–320. <http://www.jstor.org/stable/1911491>.

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.” Santa Monica, CA: RAND Corporation, 2017. 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.” *The Review of Economic Studies* 60, 3 (1993): 497–529. <https://www.jstor.org/stable/2298122>.

Huff, Jared, Mikhail Smirnov, Gregory Schell, and James Grefer. “Estimating the Retention Effects of Continuation Pay.” Arlington, VA: Center for Naval Analyses, 2018. https://www.cna.org/archive/CNA_Files/pdf/drm-2018-u-017177-final.pdf

Kasahara, Hiroyuki, and Katsumi Shimotsu. “Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices.” *Econometrica* 77, 1 (2009): 135–75. <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA6763>.

Keane, Michael P., and Kenneth I. Wolpin. “The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence.” *The Review of Economics and Statistics*, (1994): 648–72. <https://www.jstor.org/stable/2109768>.

Keane, Michael P., and Kenneth I. Wolpin. “The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment.” *International Economic Review* 42, 4 (2001): 1051–1103. <https://www.jstor.org/stable/826985>.

Levy, Robert, Jacklyn Kambic, Jessica Oi, and Mikhail Smirnov. “Evaluating Navy Medicine’s Accession and Retention Incentives for Officers in Critical Readiness Specialties.” Arlington, VA: Center for Naval Analyses, 2020.

Magnac, Thierry, and David Thesmar. “Identifying Dynamic Discrete Decision Processes.” *Econometrica* 70, 2 (2002): 801–16. <https://onlinelibrary.wiley.com/doi/10.1111/1468-0262.00306>.

Manski, Charles F. “Identification of Endogenous Social Effects: The Reflection Problem.” *The Review of Economic Studies* 60, 3 (1993): 531–42. <https://www.jstor.org/stable/2298123>.

Mas-Colell, Andreu, Whinston, Michael D. and Green, Jerry R. *Microeconomic Theory*. New York: Oxford University Press, 1995.

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.” Santa Monica, CA: RAND Corporation, 2007. https://www.rand.org/pubs/technical_reports/TR470.html.

Mattock, Michael G., Beth J. Asch, James Hosek, and Michael Boito. “The Relative Cost-Effectiveness of Retaining Versus Accessing Air Force Pilots.” Santa Monica, CA: RAND Corporation, 2019. https://www.rand.org/pubs/research_reports/RR2415.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.” Santa Monica, CA: RAND Corporation, 2016. https://www.rand.org/pubs/research_reports/RR1455.html.

Norets, Andriy. “Estimation of Dynamic Discrete Choice Models Using Artificial Neural Network Approximations.” *Econometric Reviews* 31, 1 (2012): 84–106. <https://www.tandfonline.com/doi/full/10.1080/07474938.2011.607089>.

Rust, John. “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher.” *Econometrica: Journal of the Econometric Society*, (1987): 999–1033. <https://www.jstor.org/stable/1911259>.

Rust, John. “Has Dynamic Programming Improved Decision Making?” *Annual Review of Economics* 11, (2019): 833–58. <https://www.annualreviews.org/doi/abs/10.1146/annurev-economics-080218-025721>.

Appendix C. Abbreviations

Term	Definition
ACS	American Community Survey
ADSO	Active Duty Service Obligation
BRS	Blended Retirement System
CCP	Conditional Choice Probability
CPS	Current Population Survey
CPS-ASEC	Annual Social and Economic Supplement of the CPS
DDC	Dynamic Discrete Choice
DDM	Dynamic Decision Model
DOD	Department of Defense
DP	Dynamic Programming
DRM	Dynamic Retention Model
EM	Expectation Maximization
FFRDC	Federally Funded Research Development Center
GEV	Generalized Extreme Value
IDA	Institute for Defense Analyses
MCAM	Military Career Analysis Model
MLE	Maximum Likelihood Estimator
MOS	Military Occupational Specialty
OPDRM	Observed-Performance DRM
SRB	Selective Reenlistment Bonus
YOS	Years of Service

This page is intentionally blank.