Estimating the retention effects of Continuation Pay

Jared Huff, Mikhail Smirnov, Greggory Schell, and James Grefer

April 2018

This document contains the best opinion of CNA at the time of issue.

It does not necessarily represent the opinion of the sponsor.

Distribution

DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

SPECIFIC AUTHORITY: N00014-16-D-5003 4/27/2018

Request additional copies of this document through inquiries@cna.org.

Approved by: **April 2018**

Molly McIntosh
Defense Workforce Analysis
Resource Analysis Division

This work was performed under Federal Government Contract No. N00014-16-D-5003.

Abstract

In this study, CNA uses a dynamic modeling approach to analyze the retention impacts of the recent change in the military retirement system. Our focus is on a lump-sum Continuation Pay that sailors receive in the middle of their careers. This Continuation Pay is designed to be able to offset the retention decline that results from some of the other retirement changes, and we find that it can do so. We also discuss some of the ways in which Continuation Pay is likely to interact with other Navy force-shaping policies.

This page intentionally left blank.

Executive Summary

In 2016, Congress approved a substantial change to the military retirement system. Under the old system, the minority of sailors who reached 20 years of service (YOS) became eligible for a substantial annuity upon retirement, which served as a retention "draw" in the second half of sailors' careers. Under the new system, the Navy will now offer most sailors some benefit through both direct contributions and limited matching to a Thrift Savings Plan (TSP). This is accompanied by a 20 percent decline in the retirement annuity and the introduction of a midcareer bonus, continuation pay (CP), whose intention is to retain force profiles. DOD has offered the services some latitude in how to set CP, which is the focus of this study.

To perform this analysis, we study the tradeoffs between retention, lifetime income flows, and willingness to obligate for multiple years of military service. Traditional statistical models (especially those with rudimentary estimates of expected lifetime military pay) tend to do a poor job of modelling two types of decisions in particular:

- Decisions based in part on uncertain future outcomes
- Decisions that require a tradeoff between pay and contract length

Since these two points are key to our analysis of CP, we are developing a more technical statistical model, the Dynamic Decision Model (DDM), to perform the estimation.

We use the DDM to analyze the changes in the retirement system and show that the CP options available to the Navy are more than sufficient to counteract the decline in retention due to the reduction in retirement annuity. In fact, our estimates show that a relatively modest CP can achieve the same level of cumulative retention rates at YOS 20 as existed before the Blended Retirement System (BRS). The minimum level of CP is predicted to lead to a more junior force but one that has negligible difference in steady-state force size between YCS 5 and 20.

The shift from the old retirement system to the new one is likely to have several effects, some of which may interact with other force-shaping policies. First, absent other policy changes, it results in a more junior force. Second, it incentivizes longer contracts near the YOS where CP is offered. If the obligations incurred due to CP and reenlistment bonuses can be served concurrently, the longer contracts near CP will likely increase in higher reenlistment bonus expenditures as well. This effect is likely greater if CP is offered earlier (where reenlistment bonuses are larger). The longer contracts have non-monetary implications as well—they may ease community management and have an ambiguous effect on force management writ large: if the Navy decides to downsize via a policy vehicle such as Career Waypoints, longer contracts may serve as a barrier to these goals.

This page intentionally left blank.

Contents

This page intentionally left blank.

List of Figures

This page intentionally left blank.

List of Tables

This page intentionally left blank.

Glossary

ACOL	Annualized Cost of Leaving
AVF	All-volunteer force
BRS	Blended Retirement System
CP	Continuation Pay
DDC	Dynamic Discrete Choice
DDM	Dynamic Decision Model
DOD	Department of Defense
MCRMC	Military Compensation and Retirement Modernization Commission
NDAA	National Defense Authorization Act
SRB	Selective Reenlistment Bonus
TSP	Thrift Savings Plan
YOS	Years of service

This page intentionally left blank.

Introduction

The FY 2016 National Defense Authorization Act (NDAA) introduced the Blended Retirement System (BRS) to reform military retirement. [1] BRS reduces the value of the pension, adds a defined-contribution savings plan (i.e., the Thrift Savings Plan (TSP)), and adds a continuation pay (CP). Because retirement benefits can represent a third or more of a retiree's lifetime military compensation, this is a major change.

If it performs as designed, the BRS will manage three goals that are seemingly at odds:

1. Maintain current retention profiles
2. Provide some retirement benefits to sailors who leave the Navy before 20 years of service (YOS)
3. Cost the Navy less

Two mechanisms work to simultaneously achieve these goals. First, by providing a matching TSP, the Navy allows market returns (that sailors receive on their TSP savings) to provide some of the retirement savings that would otherwise be paid by the Navy. Furthermore, it provides some retirement benefit for sailors who leave the Navy before 20 YOS. However, since this is paid to sailors regardless of whether they reach retirement, we would expect the creation of TSP and the reduction of retirement benefits to lower the retention profile, since TSP is likely to serve as much less of a draw to retirement (20 YOS).

Second, the Navy provides CP to sailors mid-way through their careers (currently at YOS 12). This exploits the fact that, based on research about sailor behavior, the Navy values money paid out in the future more than sailors do. By shifting the money to earlier in sailors' careers, the Navy can offer sailors less money without making sailors worse off.

In this study for N1T, CNA was tasked with identifying the implications of one of the recent changes to the military retirement system—the introduction of CP. CP is a bonus offered to sailors who complete a requisite YOS and who are willing to obligate for an additional four years. The bonus is intended to offset some of the decrease in the value of the pension portion of the retirement system while providing "flexibility for Service personnel managers to adjust force profiles if future manpower requirements change." [2]

In the following section, we review the literature on the military retirement system and efforts to reform it, from the Department of Defense (DOD) Authorization Act of 1981 to the FY16 NDAA. We follow this by detailing some of the theoretical considerations of CP (the focus of this study), with an emphasis on how CP will likely interact with other Navy policies. Next, we review statistical approaches to modeling retention, breaking out compensation and non-compensation factors and associated empirical estimates from the literature. We then discuss the model that we use in this study, both from a technical standpoint and in terms of the assumptions that we use to model retirement. We conclude with our results and policy implications.

The history of military retirement reform

In 1916, the military established retirement pay, informally called the final basic pay system. [3] Under that system, a retiree received an annuity equal to 2.5 percent of basic pay in his or her final year of military service, multiplied by the number of years served and indexed to average US cost of living.

The military retirement was first reformed 65 years later to address rising retirement costs to the military. The DOD Authorization Act of 1981 changed the Final Basic Pay system to a High-3 Basic Pay system, where the annuity was the average of the highest three years of basic pay instead of the final year of basic pay. [4] All servicemembers who joined prior to the reform were grandfathered under the Final Basic Pay system.

Yet pressure to cut the cost of military retirement persisted. Just five years later, the Military Retirement Reform Act of 1986 further reduced military retirement pay, lowering the annuity multiplier from 2.5 percent of basic pay to 2.0 percent, effectively reducing the pension by 20 percent. The reduction produced the desired savings for DOD, but by the late 1990s retention rates were falling for myriad reasons, including a strong civilian economy. Military leaders, also blamed the reduced rate of retention on the cut in retirement pay.

Despite a Congressional Budget Office report that found only minor effects of the 1986 reform on reenlistment [5], the Military Retirement Reform Act of 1986 was repealed in 1999, and the multiplier went back to 2.5 percent. Under a program called Redux, servicemembers who joined between 1986 and 1999 were given a choice: stay with the lower 2.0 percent multiplier and receive a $30,000 bonus at 15 years of service (YOS), or opt for the higher 2.5 percent multiplier.

More than a decade later, the FY13 NDAA established the Military Compensation and Retirement Modernization Commission (MCRMC) to review and recommend reforms to military pay and benefits. In the FY16 NDAA, Congress passed into law the BRS. On January 1, 2018, servicemembers with 12 YOS or more were grandfathered under the current retirement system, servicemembers with less than 12 YOS are allowed to choose between the legacy system and the BRS, and new accessions will be enrolled in the BRS.

Continuation Pay: Theoretical Implications

Continuation pay and additional obligation

The additional obligation associated with accepting CP is an important factor in the analysis. As a general rule, sailors value shorter obligations. This can be seen empirically both in enlistment (sailors declining longer enlistment contracts despite the presence of a bonus [6]) and reenlistments (sailors choosing less than a 6-year obligation in return for Selective Reenlistment Bonus (SRB), as seen in Navy reenlistment data). Furthermore, this is what we would theoretically expect to see: additional obligation prevents the sailor from leaving the Navy if he or she gets assigned to an undesirable command or if family circumstances change. Since the additional obligation has these negative characteristics, sailors will generally need to be compensated for it.

Noting our expectation that obligation requires compensation, the obligation that is attached to CP will have at least two effects. First, longer required obligation will result in some people reaching retirement who would have left in the presence of obligation attached to CP (since the draw of retirement will be greater at, e.g., 16 YOS than at 15 YOS). Second, the required obligation will result in some sailors declining the bonus. A lower level of CP results in a larger share of sailors who decline it.

Interactions between CP and other force-shaping policies

As we noted earlier in this paper, CP functions as a quasi-SRB. A traditional SRB offers a "multiplier" formula to determine the size of the bonus as a function of the length of obligation:

$$SRB = multipler * monthly\ base\ pay * years\ obligated\ upon\ reenlistment \quad (1)$$

CP, on the other hand, offers a set amount in return for exactly four additional years of obligation:

$$CP = CP\ Multiplier * monthly\ base\ pay \quad (2)$$

This means that, in practice, the CP multiplier will have no more of an effect than an SRB multiplier that is a quarter its size (since CP requires an obligation of 4 years). The lowest CPs (a multiplier of 2.5), is equivalent to a small SRB (e.g., 0.625). Likewise, the maximum CP of 13 would be equivalent to an SRB (3.25) that is fairly large, especially for Zone C.

Regardless of size, one way in which CP interacts with SRBs is by creating an additional incentive to reenlist for longer obligations at the YOS where CP is offered: non-linearity: instead of all contract lengths being rewarded by the SRB multiplier, CP amplifies the value of moving from a 3-year obligation to an obligation of 4 or more-years obligation.

If CP and SRB can be accepted concurrently, CP could increase SRB payments: sailors who might otherwise have left but stay for the CP would also have the incentive to accept 4 years' worth of SRB (since the 4-year obligation has already been created by acceptance of CP). Likewise, a sailor who might have reenlisted for 3 years under a modest SRB will now have a very strong incentive to reenlist for a fourth year (since increasing the obligation from 3 to 4 years means that the sailor will get CP). Once the sailor decides to accept CP, there is no additional obligational cost to reenlisting for 4 years for SRB. As such, even if aggregate retention profiles remain unchanged, we would expect that the incentives for longer contracts due to CP would increase SRB expenditures near the YOS where CP is offered.

Finally (and relatedly), the combination of CP and SRB are likely to increase the extent to which money is paid to retain people who were going to stay anyway (often referred to as "rents").

Retention Models

With the transition from conscription to an all-volunteer force (AVF) in the 1970s, it became critical for the services to understand what motivates retention. In the decades since the inception of the AVF, retention modeling has been the subject of much analytic discourse. [7] Here, we discuss how monetary and non-monetary retention factors have been modeled and then review empirical estimates of these relationships from the literature.

Modeling monetary and non-monetary factors of retention

Retention models have historically recognized the primacy of compensation in the retention decision. Typically, compensation is grouped into three bins:

- Regular military compensation (RMC), which includes basic pay, allowances, and the tax advantage that applies to allowances
- Incentive pays, such as reenlistment and continuation bonuses, combat compensation, and sea pay
- Civilian pay (frequently relative to some of the military pay listed above)

Early retention models included each of these three types of compensation as separate regressors in the model, thereby allowing a dollar in RMC to have a different retention effect than a dollar of sea pay. However, it is not clear that this should be the case: dollars should be fungible regardless of the source. So, retention models evolved to include a single measure of compensation and to consider other factors such as differences in taste for military service and hardships associated with serving (e.g., deploying to a combat zone), giving rise to the Annualized Cost of Leaving (ACOL) model.

In the ACOL model, a servicemember will remain in the military if the difference between the future streams of military and civilian pay (discounted to the present day) is sufficient to offset the servicemember's taste for civilian life (or distaste for military life). Equation 3 shows this mathematically. Given a servicemember's personal discount rate (PDR) d and taste for civilian life v, he or she will stay in the military if for any s years of military service, the present discounted value (PDV) of expected military compensation (*M*) is greater than the PDV of expected civilian compensation (*C*) plus the dollar value of *v*:[1]:

$$\sum_{t=y+1}^{y+s} M_t * (1+d)^{y-t} > \sum_{t-y+1}^{y+s} (C_t + v) * (1+d)^{y-t} \tag{3}$$

Rearranging equation 1 into equation 4 gives us the ACOL choice function: a servicemember stays in the military if the ACOL quotient—the ratio on the left of the inequality sign—exceeds the taste for civilian life for any feasible value of *s*.

$$\left(\frac{\sum_{t-y+1}^{y+s}(M_t - C_t) * (1+d)^{y-t}}{\sum_{t=y+1}^{y+s}(1+d)^{y-t}}\right) > v \tag{4}$$

Two studies provide a rationale for a range of PDRs. First, Warner and Pleeter use data on voluntary separation payments offered during the military drawdown of the 1990s to estimate implied PDRs.[8] To incentivize separation and save money, the services offered two separation schemes: annuity payments or a lump sum payment. The break-even PDR for those two options was a PDR of 17 percent. Based on take-up rates, the authors estimated enlisted PDRs of 17.3 to 35.4 percent, and an officer PDR of around 10.4 percent.

Second, Simon, Warner, and Pleeter analyzed the option to take the $30,000 Redux bonus (in exchange for a lower annuity payment in retirement) following the repeal of the 1986 military retirement reform. [9] Using the data from Redux take-rates, the authors estimated enlisted PDRs of 7.2 percent, and officers' PDRs of 4.3 percent, much lower than their first study. By comparison, the MCRMC used middle-of-the-road PDRs: 12.7 percent for enlisted and 6.4 percent for officers.

Simon, Warner, and Pleeter argue that the estimates differ between these two studies, because the servicemembers in the first study expected to face some career uncertainty and perhaps to be more credit-constrained during the recession, and thus were more inclined to take the lump-sum payment. [9] Also of note: in both studies, there was evidence that PDRs varied with age—younger servicemembers had substantially higher estimated PDRs than mid-level or senior servicemembers. Therefore, any force-wide average PDR estimate is likely to understate the true PDR of younger personnel and overstate the true PDR for more senior personnel.

Key to modeling the effect of compensation on retention is the pay elasticity of retention, defined as a percent change in retention following a 1 percent change in pay. In the literature, estimated pay elasticities vary widely, but these estimates were produced with different empirical specifications and data samples over different periods of time. Hansen and Wenger explored the variation in pay elasticities and

[1] 1 *v* is positive when the servicemember prefers civilian life, negative if the servicemember prefers military life, and zero if the servicemember is neutral.

found that the estimates were highly sensitive to empirical specification and that pay elasticities do not change over time. Hansen and Wenger also found that using the ratio of military and civilian pay (instead of the ACOL method) produced estimates that varied greatly with minor changes in model specification. [10]

Empirical estimates of pay, benefits and retention

In this section, we review the literature on the empirically estimated retention effects of noteworthy monetary factors (reenlistment and continuation bonuses, combat compensation, and sea pay) and non-monetary factors (dangerous, numerous, or long deployments, and other aspects of quality of life).

Reenlistment and continuation bonuses

Four studies provide insight on the effects of enlisted reenlistment and officer continuation bonuses.

Hosek and Peterson analyzed the effects of enlisted SRBs on retention, holding many other factors constant. The estimation is complicated by reverse causality: when retention is low, the services increase SRBs; when retention hits desired levels, SRBs are reduced again. This reverse causality will bias the SRB coefficient toward zero in a retention model. Hosek and Peterson corrected for this, with some success, by adding a fixed effect into their model, resulting in a larger SRB effect on first-term reenlistment. [11]

Gray and Grefer addressed the reverse causation problem in a model of military physician retention by instrumenting total physician pay with RMC, which is highly correlated with total physician pay, but not directly influenced by physician retention. The authors show the instrument is statistically valid and produced substantially larger estimates of the effect of continuation bonuses than those produced in a standard fixed-effects model. [12]

Finally, Hattiangadi et al. [13], and Cox and Philips [14] examine whether SRB payments would be more effective and efficient if they were paid in a lump sum instead of as an annuity. The authors use a PDR of 20 percent, as most SRBs go to younger servicemembers who are likely to have higher-than-average PDRs. They find that although there would be a large up-front cost for the first cohorts receiving lump sums, the annual savings in the long run could be as high as $11 million (out of a roughly $150 million annual budget). The savings would come at the cost of reduced flexibility. However, in low-retention periods, lump-sum SRB expenditures could exceed the budget.

Combat compensation

Lien and McIntosh examined the retention effect of two forms of combat compensation: hostile fire pay and combat-zone tax exclusion (CZTE). It is not possible to cleanly separate the retention effects of combat compensation from the retention effects of other aspects of combat (e.g., exposure to danger and time away from family). However, the authors exploit aggregate differences in hazardous combat deployments by YOS, service, and location (Global War on Terror deployments versus deployments to other combat areas), in an attempt to isolate combat deployments that are more hazardous than others with no meaningful differences in combat compensation. For active servicemembers, the authors found a negative retention effect of combat compensation on more hazardous deployments—perhaps the compensation was insufficient to balance the negative aspects of the more-hazardous deployments. For reservists, they found a positive retention effect of mobilizing and deploying, relative to mobilizing and not deploying—perhaps reflecting a combined response to combat compensation and a stronger connection to the mission. [15]

Dangerous, numerous, or long deployments

At least three studies examine the retention effects of deployments that are dangerous, numerous, or especially long. First, during Operation Iraqi Freedom and Operation Enduring Freedom, members of the Explosive Ordnance Disposal (EOD) units were deployed frequently and reenlistment rates were low. To evaluate the extent to which these two observations correlated, Lien and Gregory used data on CZTE and imminent danger pay (IDP) as proxies for time deployed on a dangerous mission. The authors found a small but statistically significant negative effect of receiving IDP for first-term sailors and no effect for second- and third-term sailors. Because it is impossible to disentangle the offsetting effects of dangerous duty and additional pay, the estimated effects of the dangerous deployments (which we would presume would negatively affect retention) are likely to be positively biased (toward zero) due to the presence of additional pay. [16]

Second, Hosek and Martorell examined the effects of the number of deployments on self-reported willingness to stay in the military, and the actual effects of deployments on retention. Using data from 10 Status of Forces Surveys (SOFS) of active duty servicemembers linked to administrative personnel and pay files, the authors found that the self-reported deployment intensity and stress were increasing across time, reducing intentions to reenlist among first- and second-term enlisted servicemembers. From the personnel and pay data, however, deployments were shown to have mostly no or positive effects on retention decisions for most first-term and second-term servicemembers. The exception was for second-term servicemembers during the height of the second Iraq War, when deployments were appreciably longer. [17]

Third, Lien et al. examine the effect of frequency and length of deployments on retention of Marines during the Iraq War. Using data from 2004 to 2007, they find that adding 100 days of deployment time has no significant retention effects on first-term Marines overall or on Marine Corps officers in general. However, long deployments decreased retention for enlisted Marines without dependents. [18]

Quality of life

In 2003, Kraus et al. conducted the Navy's first choice-based conjoint (CBC) survey on pay and benefits. A CBC survey allows servicemembers to identify relative preferences among a list of cash payments and non-cash benefits, and determine sailors' willingness to trade off one for another. With CBC, researchers can estimate a range of monetary equivalent retention effects that could follow a reduction or increase of one or more of these benefits. The authors discovered that sailors place a high value on noncash benefits. For example, sailors in the survey claim to value choice of duty location as roughly equivalent to a 5 to 6 percent pay increase. [19]

Grefer et al. conducted a statistical analysis of the effect of commissaries on retention, in light of the inherent reverse causality between retention and commissary use. Assuming the correlation is positive in both directions, this would result in coefficients that overstate the retention effect. Still, the results did not reveal a statistically significant relationship. In a separate analysis, the authors broadened the approach by constructing an index of all on-base benefits—including the commissary, base exchanges, fitness centers, and community centers—and found large, positive results, suggesting that any retention effects come from overall use of on-base benefits, rather than just the commissary. [20]

Modeling Sailor Decisions in a Dynamic Framework

In this section, we briefly summarize the technical aspects of the DDM. A more complete discussion can be found in the Appendix.

We model sailors as making a decision in the last year of their current obligation to leave the Navy ($d = 0$), take a short-term extension for up to 2 years ($d = 1{,}2$), or reenlist for up to 6 additional years ($d = 3{,}4{,}5{,}6$). Sailors make this decision by comparing the relative payoffs from each of the possible decisions. Importantly, when considering the payoffs of the decisions to in the Navy, the sailors consider the payoffs of the entire potential Navy career, rather than just the duration of their next contract.

When making the decision a sailor knows the current “state of the world”, summarized by the vector x. This includes all relevant information about the world, such as the state of the civilian economy, as well as information about the sailor, such as rate, time in grade, family status, etc. The sailor does not know the future state of the world, as she does not have perfect foresight. What she does know is the probability of certain events happening. For example, she knows that the economy tends to revert to its mean level after periods of boom and recession, and that young single sailors tend to get married in their 20s. So while she does not know exactly what will happen to her over the next few years, she is able to make probabilistic forecasts and incorporate information from them into her retention decision.

The relative payoffs from the sailors’ decisions are summarized by a utility function $u(\cdot)$ that captures the relative benefits and costs of remaining in the Navy for one year. The value of $u(\cdot)$ depends on the state of the world x, including the monetary compensation, foregone civilian opportunities, as well as nonmonetary factors, such as family status. Sailors consider their probabilistic lifetime utilities in the civilian sector and in the Navy and chose the option that gives them the highest total value. Sailors that are different from one another, or who make decisions at different points in time, have a different value of utility from being in the Navy and thus make systematically different decisions about their future.

Importantly, sailors also know that they will make good decisions in the future and incorporate this into their decision-making process. While this may seem like a straightforward assumption, it results in an empirical strategy that is quite complex: the ACOL models that we discussed in the previous section, for example, assume instead that the sailors expect to get the *average* payoff from staying in the Navy, and as a result these models were able to generate estimates at a time when computational power did not permit the estimation of models that did make this assumption. This assumption also has notable impacts on estimates when compared to the ACOL model and other naive “lifetime income” models: sailors value remaining in the military more, but also dislike longer contracts more, than simpler models would suggest.

Armed with these forecasts of the future, sailors make the decisions that benefit them the most (i.e., has the highest value of utility). Some of the forecasted series of events are much more likely to happen than others, and sailors weight the utility associated with each forecast according to its probability. This is a sailor’s *expected* utility. Sailors also values today’s utility more than future utility, and thus *discount* the utility of events that are far in the future. Thus sailors pick the decision that maximizes their *expected present discounted value* of utility, given the function $u(\cdot)$, and the vector x, and their preferences about the future. The next section provides more detail on these two objects and their relation to data and model results.

Empirical details

The previous section describes the retention decisions sailors have to make. They make it by maximizing the expected present discounted value of their utility. To make this calculation, they consider the current state of the world, x, the probabilistic evolution of this state of the world over time, and the utility function $u(\cdot)$. In this section, we describe these three components of a sailor’s problem, and the estimation strategy we use for finding a model of sailor behavior that most closely matches the data.

The vector x is a combination of sailor-specific information and a summary of the civilian economy. The sailor-specific information includes demographic information (race, ethnicity, gender), family status (marital status, number of children), and Navy-specific information (recruit quality, paygrade, time in grade, EMC, sea or shore assignment, YOS). Additionally, the vector x includes the information about the fiscal year that sailors are making their decisions, and information on the SRBs that are available to them during their decision windows. Together with their paygrades, these are used to calculate the total monetary compensation from each potential new contract.

The function $u(\cdot)$ is a linear combination of x and a vector of parameters θ. Our specification normalizes the utility from leaving the Navy to zero. The utility of remaining in the Navy depends on the sailor’s total compensation, marital status, the number of children, rank, gender, race, and ethnicity, the EMC group, the state of the civilian economy, and whether the new contract will take the sailor to 20 YOS. The functional form allows, for example, female sailors to be affected differently by marital status and the number of children than are male sailors. The complete specification is available along with the estimates of parameters in the Appendix.

Two components of x require special attention. The first is EMC group, which we include to capture different patterns in contract length across enlisted communities. For example, the intelligence community has very few 3- and 4-year contracts, but a lot of 5-year contracts relative to other communities, even when other differences (such as SRBs) are taken into account. These EMC group variables allow the model to be flexible enough to match the observed retention decision patterns across communities. More importantly, however, they make it possible to correctly and precisely estimate the parameters on compensation and retirement variables. Models that omit controls for EMC group are likely to conflate compensation with EMC group since SRBs vary widely by EMC.

The second important component of x is the measure of the civilian economy. Economic theory would expect the sailor to use a forecast of the civilian economy to make her decisions. Unfortunately, developing a sophisticated forecast of the civilian economy is well outside of the scope of this project.[2] We use, instead, the CNA economic index, which combines a number of different measures of the civilian economy into a single measure. [21] The index does include forward-looking measures, such as treasury interest rates, so it does incorporate some information about the future, but it is not, by itself, a full forecast of the economy. In essence, our current model implementation limits the amount of information the sailor has about the economy to its current state, excluding any economic forecasts from consideration.

The final component of the utility function is the probability that sailors will chose to leave the Navy at the end of their *next* contract, P_0. We make a number of assumptions that guarantee that the model sufficiently summarizes the future payoffs the sailor can expect from staying in the Navy. For details, please see the Appendix. While the mathematical derivation is involved, the intuition for this is straight-forward. Consider sailors who are thinking about reenlisting for 4 years. If they expects that at the end of these next 4 years they are likely to leave the Navy (a large P_0), then they know that the Navy career options that *follow* the 4-year reenlistment are not very attractive relative to the civilian opportunities. Conversely, if sailors know that after the 4-year reenlistment they are very likely to stay in the Navy (a small P_0), they know that the Navy options at that point in time will be good relative to the civilian ones. This is how we summarize and incorporate the sailor's expectations about her future decisions into the model.

Our estimation procedure is broken down into two steps. In the first step, we estimate P_0 and the sailors' forecasts about the future. The probabilities of leaving, P_0, are estimated using a logistic regression with all of the elements of x as controls. To build the sailors' forecasts we estimate transition probabilities for marital status, number of

[2] We use AR1 and AR2 models to estimate the evolution of the economic index over time. These models do no better at explaining the sailors' retention decisions than the static value of the index.

children, and paygrade using frequency estimators. The probabilities of 1) marriage and divorce depend on age and gender, 2) the number of children on age, gender, and marital status, and 3) promotion on paygrade, time in grade, and EMC group. These first-stage estimates are then treated as data for the purposes of the second stage of the estimation. The literature on dynamic models stresses the importance of accuracy of these first-stage estimates, and future work will focus on further developing this part of the estimation routine.

In the second step we estimate the parameters of the utility function $u(\cdot)$, which uses a maximum likelihood algorithm. For any proposed set of parameters θ, the algorithm computes the likelihood of observing the decisions we see in the data. Using a gradient ascent algorithm, we maximize this likelihood over the full space of parameter values. One model parameter cannot be estimated as part of this estimation procedure. The discount factor, β, that determines how much sailors value the present vs the future, has to be set to a fixed value before the estimation. We estimate models with $\beta = 0.88$ and report counterfactual policy evaluations for the latter.[3] Additional work is needed to explore the time preferences of sailors in greater detail to help inform dynamic models of sailor behavior.

[3] The typical value in economics literature is 0.95. Meanwhile, prior research has estimated a discount factor of 0.88 for sailors.

Modeling the Retirement Change

In work supporting the Military Compensation and Retirement Modernization Commission (MCRMC), RAND used a dynamic modeling approach to estimate the effect of the change in the retirement system on retention profiles. RAND's Dynamic Retention Model (DRM) functions similarly to the model described in the previous section—sailors make decisions based on today's circumstances, expectations about future circumstances, and expectations about how they will deal with these future circumstances. RAND estimates a discount factor of 0.88 to 0.90—that is, sailors value \$1 in one year equivalent to between \$0.88 and \$0.90 today. Because our effort represents the first step of a dynamic framework that will be built upon later, our model omits some of the secondary factors modeled in the DRM (e.g., whether to affiliate with the Reserve Component).

Our assumptions and RAND's assumptions notably diverge in their treatment of the TSP contributions. In reality, the DOD contributes 1 percent of a sailor's basic pay into an account (where it accumulates market returns). The DOD also matches sailor contributions to the same account over a limited range: 1 to 1 for the first 3 percent of basic pay and then 0.5 to 1 for the 4th and 5th percent. This matching represents a *100 percent initial return* to sailors' investments. The caveat to the TSP contributions is that, under normal circumstances, they cannot be withdrawn until the sailor reaches the age of 59.5.

This delayed ability to withdraw funds ends up being a complicating factor. Commonly assumed discount factors for servicemembers (including those used by RAND) are larger than average market returns; that is, based on typical assumptions, market returns cannot keep up with the extent to which sailors discount future income. The difference between the discount rate and average market returns is typically between 5 and 7 percent per year—this means that each year that a sailor's contributions (and DOD match) are inaccessible but gaining market returns, the sailor is losing 5 to 7 percent of the value compared to spending the money now. Since market returns fall below sailor discount rates, sailors would (under these assumptions) prefer to not invest. The initial match from the military provides more of an incentive to invest, but it may not be able to overcome the preference for current use of income for all sailors (especially younger sailors who have longer to wait until TSP funds can be withdrawn).

Despite these considerations, we see people investing in TSP (including servicemembers and other people whose discount rates are likely similar to enlisted servicemembers), which means they are made better off by contributing. The question, then, is how to model TSP contributions that we expect to (according to their discount rates) make sailors worse off.

RAND's solution was to model DOD TSP contributions as contributions that happen completely separately from anything that sailors contribute. This means that sailors can be modeled as not contributing funds—consistent with their calculated discount rates—but still be modeled as receiving benefits from the TSP program. In doing so, they mitigate the modeling difficulty of TSP contributions that sailors would

We take a different approach. We assume that sailors already invest some of their salary into savings for retirement apart from what the Navy will provide if they reach YOS 20. The DOD's automatic contribution of 1 percent of basic pay allows sailors to shift some of the money that they would be paying into their retirement account into money that they can spend today (without reducing their retirement savings). As such, in our model, TSP results in a 1 percent increase in basic pay each year.

Having decided how to model TSP, we model the remainder of the factors that affect retention. Sailors' utility function is linear in the parameters of interest. While the sailor is in the Navy utility is determined by pay, a variable to capture whether the sailor is at sea or ashore, and a variable that captures the disutility of being in the Navy compared to the civilian economy. The marginal utilities of being at sea or ashore and the disutility of being in the Navy vary with gender, education, rating, and whether the sailor is married or has children.

We do not have data on the civilian earnings of sailors who leave the Navy. Therefore, the civilian earnings are modeled as a function of sailor education, AFQT, rating, paygrade at separation, the value of the CNA economic index, and the military to civilian pay ratio. Sailors who receive retirement pay receive it in addition to their estimated civilian earnings[4].

Sailors decide whether to remain in the Navy and for how long (1 to 6 years). Any SRBs that are available come into effect for retaining for 3 or more years. After making this decision, the sailor does not make another decision until the new obligation has been met. There is no way for sailor to leave the Navy of their own volition mid-contract to pursue civilian opportunities, regardless of how attractive they may be.

Most of the variables in the state space are fixed or evolve deterministically.[5] There are three main groups of variables that evolve stochastically: sailor paygrade (and thus pay) and assignment, sailor family situation, and the civilian economy.[6] Paygrade and sea or shore assignment evolve stochastically given rating, paygrade, and time in

[4] We do not estimate likely civilian retirement payments in this iteration of the model

[5] Sailor gender, race, ethnicity, education, AFQT, and rating are fixed. Sailor age, YOS, time to EAOS, time in grade, and any accumulated retirement benefits evolve deterministically.

[6] One of the model assumptions is that all state space variables are discrete. For this reason, the measures of the civilian economy are discretized into a number of bins. Simulations can allow continuous state space variables, but they tend to be computationally expensive.

grade. SRBs are determined stochastically by rating given the state of the civilian economy. Sailor family situation (marriage, children) evolves stochastically as a function of sailor gender and age.

We assume exogenous attrition every period, based on rating and YOS. Every period a sailor is in the Navy carries some probability he or she will separate from the Navy and enter the civilian workforce for an exogenous reason. Additionally, some sailors will hit high-year tenure and be forced out of the Navy.

We also begin the model at the first reenlistment decision (and thus do not model attrition before 4 years of service) Due to the fairly large number of variables for a dynamic model,[7] we use a representative 10 percent sample of sailor decisions to form our estimates.

[7] One of the commonly noted downsides of dynamic modeling is that such models cannot contain a large number of variables due to computational and data limitations. The computational limitation means that we faced a tradeoff between more variables with a smaller sample and fewer variables with a larger sample. We chose the former; eliminating the computational burden is the next step in our model development process.

Results

In this section, we summarize our results on the retention effects of different levels of CP. We focus on four scenarios:

- The lower retention annuity without any additional compensation
- The lower retention annuity with TSP and a low CP (2.5 multiplier)
- The lower retention annuity with TSP and a moderate CP (7.5 multiplier)
- The lower retention annuity with TSP and a high CP (12.5 multiplier)

In this relatively early stage of development for the DDM, we model career retention conditional on remaining in the military to the first reenlistment decision. As such, first-term attrition is excluded from the model. Furthermore, our focus is on sailors' preretirement decisions, since CP seems to be targeted at ensuring that sufficient sailors reach retirement eligibility at YOS 20.

Also note that, unlike the RAND analysis which focused on "steady state" results, the DDM estimates the short- and medium-term effects of these changes on actual sailors in our database. Thus, in the graphs which follow, our retention profiles under different CP schemes represent our estimates about how historical sailors would have responded to the BRS instead of the historical retirement system.

As we noted in a previous section, one of the difficult parts of the analysis is determining how to treat sailors' valuations of TSP. In our baseline model, we assume that sailors value TSP contributions equivalently to how they would view a 1 percent increase in pay. This assumption results in post-BRS retention profiles that are lower than those under the previous retirement system. We vary this assumption about TSP valuation, which notably changes baseline retention patterns. However, the estimated effect of CP on these retention patters is fairly consistent, regardless of the estimated size of the TSP valuation.

Figure 1 shows retention profiles under the base assumptions—in particular, we look at the cumulative rate of retention between starting at YOS 5 (relative to YOS 4). The solid black line shows cumulative retention under the legacy retirement system, while the grey line shows cumulative retention after the 20 percent cut in retirement annuity (and before TSP or CP are introduced). The red lines show cumulative retention with TSP and low, medium, and high values of CP (where darker lines correspond to higher CP values).

Figure 1. Career retention profiles

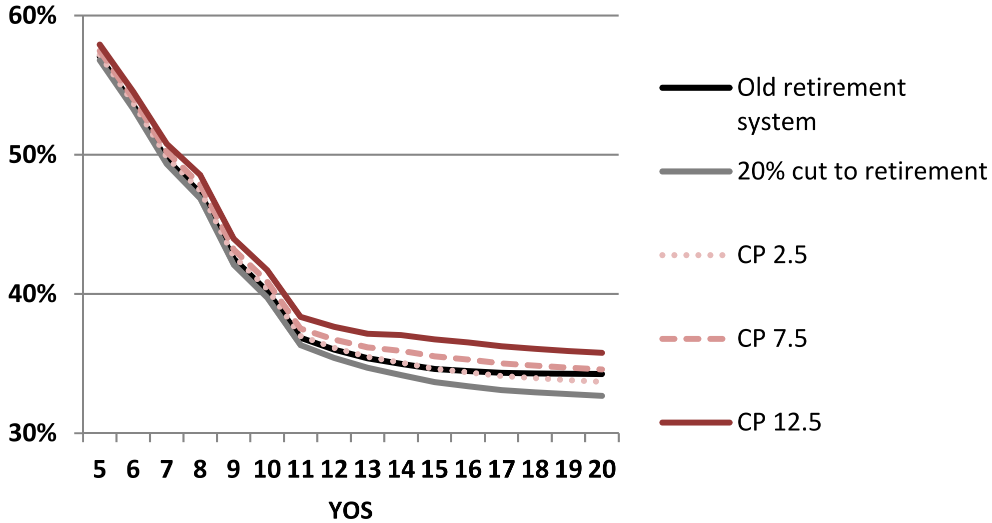

Source: CNA calculations

The cut to retirement lowers cumulative retention, particularly after YOS 10 (by YOS 20, the difference in cumulative retention is 3.7 percentage points). This makes sense: at lower YOS, the reduction in retirement is less valuable (since the amount of time until sailors would receive retirement is longer and thus discounted more heavily). The lowest level of CP (2.5 times monthly basic pay) is not sufficient to keep cumulative retention at pre-BRS levels. However, moderate and high amounts are sufficient, with CP multipliers of 7.5 and 12.5 being more than enough to make up for the retention loss at all YOS between 4 and 20. Furthermore, the "optimal" CP calculated by RAND [22] of a multiplier between 4.2 and 5.2 falls very much in the range that we would expect to be roughly cumulative retention neutral at YOS 20.

Since the point of CP is to make up the gap between the black and grey lines in Figure 1, we look more precisely at the level of the gap and the extent to which each of the three CP levels covers or fails to cover it. In Table 1 , the cumulative retention gap at each YOS is shown in the second column. The third, fourth, and fifth columns show the percent of that gap that the CP amount covers. Numbers larger than 100 percent mean that cumulative retention at that YOS is higher than it was before the change in the retirement system.

Table 1. CP and career cumulative retention

YOS	Cumulative retention gap (percentage points)	CP 2.5	CP 5.5	CP 12.5
5	0.7	147.7%	253.9%	422.5%
6	0.7	151.8%	259.9%	431.8%
7	0.8	153.0%	265.1%	443.6%
8	1.0	148.8%	262.1%	442.7%
9	1.1	142.6%	256.3%	437.9%
10	1.2	137.1%	247.3%	423.8%
11	1.4	124.9%	224.1%	383.6%
12	1.5	120.8%	220.7%	372.3%
13	1.8	114.9%	213.3%	354.1%
14	2.0	111.9%	216.0%	358.2%
15	2.3	101.0%	196.1%	324.4%
16	2.5	91.0%	175.5%	288.6%
17	2.8	81.3%	155.2%	253.8%
18	3.1	75.1%	141.7%	230.8%
19	3.4	69.0%	130.2%	212.2%
20	3.7	64.3%	120.8%	196.5%

Source: CNA calculations

As the table shows, even the lowest amount of CP (combined with TSP) leads to cumulative retention that is at or above pre-BRS levels through YOS 15. This falls to about two-thirds of the gap by YOS 20. As noted above, both the moderate and high levels of CP lead to higher cumulative retention through YOS 20 than was the case pre-BRS. This effect is especially strong for the large CP, where the positive effect of the CP is roughly twice as large (at YOS 20) as the negative effect of the cut in the retirement annuity. In all cases, the retention gains (relative to the gap caused by the BRS) are largest at lower YOS—this suggests a shift toward a more junior workforce based on changes in retention patterns. We also calculated the impact on total inventory between YOS 5 and 20 under the different levels of CP. Our estimates suggest that, in the steady state, the total available inventory falls by less than one tenth of one percent with a CP multiplier of 2.5. Multipliers that match retention to YOS 20 have larger inventories than existed before the BRS changes.

We also examined the impact of BRS on groups of ratings to identify substantive community-level changes that may be masked by the aggregate changes. However, our analysis predicted similar changes across communities: the total inventory in each falls within a quarter of a percent above or below pre-BRS levels. As BRS opt-in and opt-out decisions become available, this is one area that is worth revisiting: different opt-in rates across ratings would imply different discount rates, which (in turn) would affect our estimates of the impact of BRS on the rating.

We also model three additional scenarios for robustness that we compare against the Navy's current policy of a 2.5 CP at YOS 12:

- 5 CP at YOS 12
- 2.5 CP at YOS 8
- 2.5 CP with a 3-year obligation (instead of a 4-year obligation)

The CP of 5 essentially replicated cumulative retention under the legacy system with higher retention before YOS 20. A CP of 2.5 at YOS 8 outperforms a CP of 2.5 at YOS 12: we estimate that cumulative retention across sailors' career is higher for the former. This is likely due to the CP at YOS 8 serving as a larger draw than it does at YOS 12. Likewise, were a shorter obligation an option for the Navy, CP of 2.5 with a 3-year obligation outperforms CP of 2.5 with a 4-year obligation if both are placed at YOS 12. Here, the increased willingness of sailors to accept CP (and thus increased draw of CP) offsets the shorter contractual "lock-in" for those who accept.

As we noted in an earlier section, there are theoretical interactions between CP and SRB. Some of these results show how the two are at least partially substitutes for each other. Regardless of when in the career CP is offered, it is predicted to result in a more junior force absent changes in SRBs. SRBs, on the other hand, can be timed in such a way to preserve the seniority of the inventory. Furthermore, instead of paying a bonus to every sailor who retains to a specific YOS, SRBs can be used to target low-retention ratings, which results in more cost-effective retention.

In addition to changes in cumulative retention, we also note some changes in average obligation accepted when reenlisting. As noted earlier, CP could lead to higher SRB payments. In Figure 2, we show that this is especially likely to happen in the years directly surrounding CP—especially for larger values of CP, there is a marked increase in the average obligation assumed near CP, although this is coupled with slightly lower obligation in the preceding years. While not shown here, it is worth noting that the same effect would likely to occur if CP was moved to YOS 8 instead of YOS 12. Since SRBs are typically higher at YOS 8 than at YOS 12, we would anticipate the increase in SRB expenditures to rise by more.

Figure 2. The average contract length of sailors who do not leave at the initial decision point

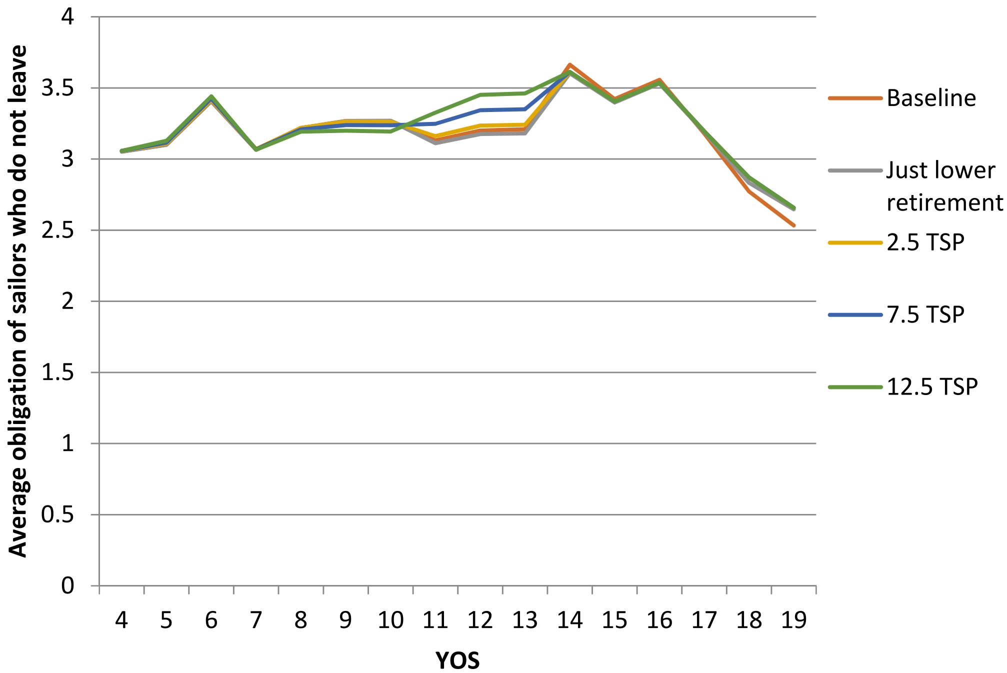

Summary, Policy Implications, and Way Forward

In this paper, we described the development of the DDM, which allows us to better model the behavior of sailors in various circumstances under which traditional statistical models of retention struggle. This behavior includes the following decisions:

- Decisions based in part on uncertain future outcomes
- Decisions that require a tradeoff between pay and contract length

Both of these decisions are applicable to analysis of the changes in the BRS: sailors (in theory) base their decisions on what they think their Navy careers and civilian opportunities will look like in the future. Furthermore, giving up the ability to leave the Navy in return for a bonus (as with SRB and CP) is costly, as it means that the sailor is forgoing the option to leave the Navy if, for instance, a high-paying civilian job becomes available.

We use the model to show that the range of CP available to the Navy is more than capable of counteracting the retention lost due to the lower BRS retirement annuity. We also note possible interactions between CP and other policies, the consequences of which are worth consideration:

- Will sailors be able to serve obligation from CP and SRB concurrently?
    - If so, an increase in SRB expenditures absent a reduction in SRB is likely.
- How will the CP obligation be treated when the Navy is trying to downsize?

Note that this analysis serves as an initial demonstration of the DDM— the Navy faces many problems with similar implications. For instance, how much do high SRBs paid to junior sailors necessitate higher SRBs when these sailors are more senior (i.e., do sailors who stay for monetary reasons continue to need monetary reasons to stay)? Will the timing of CP (and the change in the retirement system writ large) disproportionately affect the retention of women and minorities and, if so, how can this be counteracted? These issues will be explored in a follow-on study.

Appendix: Dynamic Decision Model

This is a description of the dynamic discrete choice (DDC) framework. We apply this framework to sailor retention in the Dynamic Decision Model (DDM). Recent surveys of the DDC literature by Aguirregabiria and Mira [23] and Arcidiacono and Ellickson [24] provide more details. There are a number of areas in which the DDM departs significantly from the rest of the DDC literature, to which we pay special attention.

General framework

Time is discreet and indexed by t. Individuals are indexed by i. Individuals have preferences over a sequence of states of the world from $t = 0$ to $t = T$. The state of the world has two components: a vector of state variables $s_{i,t}$ that is known to the individual and a decision $d_{i,t} \in \mathrm{D_t} = \{0,1,\ldots,J\}$. Individuals have preferences that can be represented by a period utility function $U(d_{i,t}, s_{i,t})$. Utility is time separable, so preferences over a sequence of states of the world can be represented as $\sum_{j=0}^{T} \beta^j\, U(d_{i,t+j}, s_{i,t+j})$, where $\beta \in (0,1)$ is the time-invariant discount factor. State variables $s_{i,t}$ evolve according to a Markov process $F(s_{i,t+1}|d_{i,t}, s_{i,t})$ that is known to the individual. Every period the individual makes a decision $d_{i,t} \in D_t$ to maximize the expected utility:

$$E(\sum_{j=0}^{T} \beta^j\, U(d_{i,t+j}, s_{i,t+j})|d_{i,t}, s_{i,t}) \tag{5}$$

DDM note: In the DDM, individuals do not necessarily have the option to make a decision every time period. Sailors sign multi-year contracts and are, in general, unable to freely leave the Navy until the end of the contract. This is atypical in the DDC literature, where most models assume the choice set D to be constant over time and across individuals. Some models index the choice set by t as we have done, but we have not been able to find any application that models variation in the choice set over time. In the DDM, not only does the choice set vary over time, but it varies *as a result of the individuals' decisions.* This is the most important departure of the DDM from the DDC literature and we will emphasize where it becomes particularly important.

The solution of the dynamic programming (DP) problem described above is characterized by the Bellman equation for the value functions $\mathrm{V}(s_{i,t})$:

$$V(s_{i,t}) = \max_{d \in D_t}\{U(d, s_{i,t}) + \beta \int V(s_{i,t+1}) dF(s_{i,t+a}|d, s_{i,t})\} \tag{6}$$

An individual's optimal decision rule is $\delta(s_{i,t}) = \operatorname{argmax}_{d \in D_t} v(d, s_{i,t})$, *where* $v(d, s_{i,t})$ *are choice-specific value functions:*

$$v(d, s_{i,t}) \equiv U(d, s_{i,t}) + \beta \int V(s_{i,t+1}) dF(s_{i,t+a}|d, s_{i,t}) \tag{7}$$

The vector of state variables $s_{i,t}$ *can be divided into two parts,* $s_{i,t} = (x_{i,t}, \epsilon_{i,t})$. *The first part,* $x_{i,t}$, *is observed by both the individual and the econometrician, while the second part,* $\epsilon_{i,t}$, *is observed only by the individual.*

Denote by θ *the vector of structural parameters of the model. In order to compute any estimation criterion, such as a likelihood or GMM conditions, it is necessary to solve for the individuals' optimal decision rule as a function of structural parameters:* $\delta(x_{i,t}, \epsilon_{i,t}, \theta)$. *Most of the difficulty in estimating DDC models comes from this step, as solving the DP problem for each trial value of* θ *is computationally burdensome. We use an estimation procedure, described in detail below, which avoids solving the DP problem at the expense of additional assumptions and loss of asymptotic efficiency. Asymptotic properties of the estimators discussed throughout are in the context of fixed* T *and* $N \to \infty$.

Assumptions

We make the following assumptions in the DDML

1. Additive separability: $U(d, x_{i,t}, \epsilon_{i,t}) = u(d, x_{i,t}) + \epsilon_{i,t}(d)$.
2. Distribution of unobservable: $\epsilon_{i,t} \sim \text{iid } G_\epsilon(\epsilon)$, where G_ϵ is the type 1 extreme value distribution.
3. Discrete state variables: $x_{i,t} \in X = \{x^{(1)}, x^{(2)}, \ldots, x^{(|X|)}\}$, $|X| < \infty$.
4. Conditional independence: $CDF(x_{i,t+1}|d_{i,t}, x_{i,t}, \epsilon_{i,t}) = F(x_{i,t+1}|d_{i,t}, x_{i,t})$.

DDM note: Additive separability implies that the marginal utility of the observed variables does not depend on the unobservables. The model cannot capture the relationship between unobserved heterogeneity and differences in the marginal utility of sea or shore assignments, compensation, or other variables. An extension of the model can allow for permanent unobserved heterogeneity between individuals, but even that version rules out transitory shocks to the marginal utility of observable variables.

Conditional independence and independent and identically distributed (iid) distribution of unobservables have two major implications. First, unobserved shocks to the utility of every decision are transitory. Second, the unobserved shocks can affect the evolution of the state variables only through their effect on the individuals' decisions. This assumption rules out the possibility that these shocks capture sailor quality or some other attribute that may affect promotion, retention, billet assignment, or other outcomes. To the degree that these are unobserved qualities are persistent, they can be captured by permanent unobserved heterogeneity in an extended version of the model.

The assumptions of conditional independence and iid unobservables imply that the solution of the DP problem is fully characterized by the integrated value function $\bar{V}(x_{i,t}) \equiv \int V(x_{i,t}, \epsilon_{i,t}) dG_\epsilon(\epsilon_{i,t})$. The integrated value function is the unique solution of the integrated Bellman equation:

$$\bar{V}(x_{i,t}) = \int \max_{d \in D_t} \{u(d, x_{i,t}) + \epsilon_{i,t}(d) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1}) f_x(x_{i,t+1} | d, x_{i,t})\} dG_\epsilon(\epsilon_{i,t})$$

The choice-specific value functions can be rewritten as $v(d, s_{i,t}) = v(d, x_{i,t}) + \epsilon_{i,t}(d)$:

$$v(d, x_{i,t}) = u(d, x_{i,t}) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1}) f_x(x_{i,t+1} | d, x_{i,t}) \tag{8}$$

The optimal decision rule can be rewritten as $\delta(x_{i,t}, \epsilon_{i,t}) = arg \max_{d \in D_t} \{v(d, x_{i,t}) + \epsilon_{i,t}(d)\}$. The conditional choice probability (CCP) of observing decision d given a vector of state variables $x_{i,t}$ and a vector of structural parameters θ, $P(d_{i,t} | x_{i,t}, \theta)$, is the optimal decision rule integrated over the unobserved error terms:

$$\begin{aligned} P(d | x, \theta) &= \int I\{\delta(x, \epsilon; \theta) = d\} dG_\epsilon(\epsilon) \\ &= \int I\{v(d, x_{i,t}) + \epsilon_{i,t}(d) > v(d', x_{i,t}) + \epsilon_{i,t}(d') \; \forall \, d' \\ &\qquad \neq d\} dG_\epsilon(\epsilon_{i,t}) \end{aligned} \tag{9}$$

The assumption that G_ϵ is the type 1 extreme value distribution gives closed-form solutions for the integral in the integrated Bellman equation and the CCPs. In this case the full model has the following form:

$$P(d | x_{i,t}, \theta) = \frac{exp(v(d, x_{i,t}))}{\sum_{j=0}^{J} exp(v(j, x_{i,t}))} \tag{10}$$

$$v(d, x_{i,t}) = u(d, x_{i,t}) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1}) f_x(x_{i,t+1} | d, x_{i,t}) \tag{11}$$

$$v(d, x_{i,t}) = u(d, x_{i,t}) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1}) f_x(x_{i,t+1} | d, x_{i,t}) \tag{12}$$

$$\bar{V}(x_{i,t}) = log(\sum_{d=0}^{J} exp\{u(d, x_{i,t}) + \beta \sum_{x_{i,t+1}} \bar{V}\,(x_{i,t+1}) f_x(x_{i,t+1}|d, x_{i,t})\}) \tag{13}$$

DDM note: The complexity (and usefulness) of DDC models comes from the fact that individuals take into account that they will act optimally in the future. This complexity only comes into play in periods where the individual has a decision to make. In the DDM sailors do not make a decision every period, and so what happens to them between the period when they make a decision and the next chance they have to make another decision can be rolled into the deterministic part of utility.

For example, assume that a sailor is making a decision to extend his or her contract for 4 years. The choice-specific value function of this decision is as follows:

$$v(d = 4, x_{i,t}) = u(d, x_{i,t}) + \sum_{j=1}^{3} \beta^j \sum_{x_{i,t+j}} u(x_{i,t+j})\, f_x(x_{i,t+j+1}|d, x_{i,t+j}) + \beta^4 \sum_{x_{i,t+4}} \bar{V}\,(x_{i,t+4}) f_x(x_{i,t+4}|d, x_{i,t}) \tag{14}$$

Crucially, the part of this equation that describes what happens up until $t + 4$ does not depend on the integrated value functions. It is simply the expected discounted flow of utility that the sailor receives while he or she waits until the next opportunity to make a decision. There is still uncertainty about what this flow of utility will be, since the state variable transitions are stochastic, but the DP problem does not need to be solved at every time period since the sailor is not making a decision every time period.

The notation above is cumbersome, so for now we will continue to use the standard DDC notation that implies that each individual makes a decision every time period. This notation is more general, since the choice set D_t depends on t and it is simply degenerate in the DDM for many time periods. It is important to keep in mind, however, that the model only has to be solved for a subset of the time periods when the sailor is making a decision.

Estimation with the DDM

The goal of estimation is to find parameters θ that maximize some kind of estimation criterion, such as a likelihood or GMM conditions. Although direct estimation that uses a full information maximum likelihood is possible, it is computationally burdensome because to obtain the likelihood of a particular decision at any trial value of θ it is necessary to solve the DP problem. The two-step estimation method described here is drawn from Hotz and Miller [25] and avoids having to solve the DP problem.

The key to this estimation method is the fact that there exists a decision in the choice set that terminates the dynamic problem. After an individual makes this terminal decision the problem stops being dynamic and no further decisions are made on the individual's part. Let this terminal decision be $d = 0$. Begin with the choice-specific value function:

$$v(d, x_{i,t}) = u(d, x_{i,t}) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1}) f_x(x_{i,t+1}|d, x_{i,t}) \tag{15}$$

Note that the one-period-ahead integrated value functions can be rewritten as follows:

$$\bar{V}(x_{i,t+1}) = \sum_{d=0}^{J} P(d|x_{i,t+1})[v(d, x_{i,t+1}) + e(d, x_{i,t+1})]$$

Where $e(d, x_t)$ is the expectation of $\epsilon_t(d)$ conditional on x_t and on d being the optimal decision: $e(d, x_t) = E(\epsilon_t(d)|x_t, \delta(x_t, \epsilon_t) = d)$. Now rewrite $\bar{V}(x_{i,t+1})$ in terms of differences in choice-specific value functions between every choice and the terminal choice:

$$\begin{aligned}\bar{V}(x_{i,t+1}) = v(0, x_{i,t+1}) &+ \sum_{d=1}^{J} P(d|x_{i,t+1})[v(d, x_{i,t+1}) - v(0, x_{i,t+1})] \\ &+ \sum_{d=1}^{J} P(d|x_{i,t+1})\, e(d, x_{i,t+1})]\end{aligned} \tag{16}$$

Hotz and Miller show that both differences in choice-specific value functions $v(d, x_{i,t+1}) - v(0, x_{i,t+1})$ and $e(d, x_{i,t+1})$ can be written as functions of transition probabilities and CCPs. This leads to the two-step estimation procedure wherein the first step transition probabilities and CCPs are recovered, and then, in the second step, are taken as given in order to estimate the utility parameters.

DDM note: In the DDM, the terminal decision is the decision to leave the Navy. After this point in the model, the sailors make no further decisions and cannot return to the Navy. These sailors continue to receive a payoff every period based on the state of the economy and any accumulated pension benefits, but the payoff is out of the sailors' control. This flow of payoffs takes the form

$\sum_{j=t}^{T} \beta^{j-t} \sum_{x_{i,j}} u(x_{i,j}) f_x(x_{i,j+1}|x_{i,j})$ and does not depend on the integrated value functions.

In the DDM, the sailors do not make a decision every time period. Still, it is possible to express the integrated value function in terms of the differences in choice-specific value functions by treating the flow of utility between decision periods as part of $u(d, x_{i,t})$. The integrated value function is then expressed as a function of CCPs and the differences in choice-specific value functions at the next decision point. We have verified that the DDM specified here still meets the assumptions of the Hotz and Miller inversion theorem described below.

To write $e(d, x_t)$ as a function of choice probabilities, note that $\{\delta(x_t, \epsilon_t) = d\}$ is equivalent to $\{v(d, x_t) + \epsilon_t(d) > v(d', x_t) + \epsilon_t(d') \; \forall \; d' \neq \text{d}\}$. Now rewrite $e(d, x_t)$ as follows:

$$e(d, x_t) = E(\epsilon_t(d)|x_t, v(d, x_t) + \epsilon_t(d) > v(d', x_t) + \epsilon_t(d') \; \forall \; d' \neq d) \tag{17}$$

$$= \frac{1}{P(d|x_t)} \int \epsilon_t(d) I\{\epsilon_t(d') - \epsilon_t(d) < v(d, x_t) - v(d', x_t) \forall d' \neq d\} \, dG_\epsilon(\epsilon_t) \tag{18}$$

Note here that $e(d, x_t)$ depend only on the CCPs, G_ϵ, and a vector of differences in choice-specific value functions $\tilde{v}(x_t) \equiv \{v(d, x_t) - v(d', x_t) : d \in D_t\}$. The CCPs also depend only on G_ϵ and $\tilde{v}(x_t)$:

$$P(d|x_t) = \int I\{\epsilon_t(d') - \epsilon_t(d) < v(d, x_t) - v(d', x_t) \forall d' \neq d\} \, dG_\epsilon(\epsilon_t) \tag{19}$$

Hotz and Miller show that this relationship between $P(d|x_t)$ and $\tilde{v}(x_t)$ is invertible. Taking the inverse, and plugging it into the equation for $e(d, x_t)$ we have $e(d, x_t)$ as a function of CCPs and G_ϵ only. Under the assumption that G_ϵ is the type 1 extreme value distribution this function has a closed-form solution: $v(d, x_t) - v(d', x_t) = \log P(d|x_t) - \log P(d'|x_t)$ and $e(d, x_t) = \gamma - \log P(d|x_t)$, where γ is the Euler constant. These expressions depend only on CCPs.

It is now possible to form an expression for the integrated value function as a function of CCPs. Take the expression for the integrated value function where choice-specific value functions are expressed relative to the terminal decision:

$$\bar{V}(x_{i,t+1}) = v(0, x_{i,t+1}) + \sum_{d=1}^{J} P(d|x_{i,t+1})[v(d, x_{i,t+1}) - v(0, x_{i,t+1})] + \sum_{d=1}^{J} P(d|x_{i,t+1}) \, e(d, x_{i,t+1})] \tag{20}$$

By using the fact that $v(d, x_t) - v(0, x_t) = \log P(d|x_t) - \log P(0|x_t)$ and $e(d, x_t) = \gamma - \log P(d|x_t)$ it is possible to obtain a simple expression for this integrated value function:

$$\bar{V}(x_{i,t+1}) = v(0, x_{i,t+1}) - log\,P(0|x_t) + \gamma \tag{21}$$

This expression has an intuitive interpretation. The value of being in a given state x_{t+1} is the sum of the value from making the terminal choice $v(0, x_{i,t+1})$, an adjustment term for the fact that this choice may not be optimal $-\log P(0|x_t)$, and the mean of the distribution of the error term γ.

Plugging in the new expression for the integrated value functions, we have an expression for the choice-specific value functions:

$$\begin{aligned} v(d, x_{i,t}) = {} & u(d, x_{i,t}) \\ & + \beta \sum_{x_{i,t+1}} [v(0, x_{i,t+1}) - log\,P(0|x_t) \\ & + \gamma] f_x(x_{i,t+1}|d, x_{i,t}) \end{aligned} \tag{22}$$

At this point this function depends only on CCPs, the transition probabilities $f_x(x_{i,t+1}|d, x_{i,t})$, and parameterized utility functions. The value of the terminal choice $v(0, x_{i,t+1})$ does not have any associated future value components. Therefore, given x, β, P, f_x, θ, it is possible to calculate the choice-specific value functions and the optimal decision rules that are needed to evaluate the estimation criterion.

At this point, it is necessary to estimate the conditional choice probabilities $P(d|x)$ and the transition probabilities f_x. Consistent estimates of transition probabilities can be obtained using a maximum likelihood estimator (MLE) that maximizes the partial likelihood $\sum_{i,t} \log f_x(x_{i,t+1}|d_{i,t}, x_{i,t}, \theta_f)$. CCPs can be obtained using nonparametric regression methods for $P(d|x) = E(I\{d_{i,t} = d|x_{i,t} = x\})$, with the method depending on the amount of data that is available. For the rest of this document, we treat the CCPs and transition probabilities as known.

To make clear the dependence of the constructed choice-specific value functions on the estimated CCPs and transition probabilities, we write $v(d, x_{i,t}|\theta, \hat{P}, \widehat{f_x})$. The Hotz and Miller GMM estimator solves for the θ that is a solution to the following moment conditions:

$$\sum_{i=1}^{N} \sum_{t=1}^{T} H(x_{i,t}) \times \begin{matrix} I\{d_{i,t} = 1\} - \dfrac{exp\,v(1, x_{i,t}|\theta, \hat{P}, \widehat{f_x})}{\sum_{d=0}^{J} exp\,v(d, x_{i,t}|\theta, \hat{P}, \widehat{f_x})} \\ \vdots \\ I\{d_{i,t} = J\} - \dfrac{exp\,v(J, x_{i,t}|\theta, \hat{P}, \widehat{f_x})}{\sum_{d=0}^{J} exp\,v(d, x_{i,t}|\theta, \hat{P}, \widehat{f_x})} \end{matrix} = 0 \tag{23}$$

Here, $H(x_{i,t})$ is a matrix of instruments with dimension $\dim(\theta) \times J$. The main advantage of this estimator is that the DP problem does not need to be solved even once to obtain estimates of the structural parameters. Additionally, once transition probabilities and CCPs are obtained, a number of different model specifications can be tried with minimal additional computational difficulty. The Hotz and Miller inversion theorem works for many distributions of G_ϵ, not just the type 1 extreme value, but, in the case of other distributional assumptions, the expressions for $v(d, x_t) - v(d', x_t)$ and $e(d, x_t)$ do not have closed-form solutions, and so numerical evaluation of these objects is necessary to form the optimal decision rules.

Aguirregabiria and Mira [26] propose a specific version of the Hotz and Miller GMM estimator that is asymptotically equivalent to the partial MLE of the full solution of the DDC model. This estimator has a surprisingly simple form:

$$Q(\theta, \hat{P}, \hat{f}_x) = \sum_{i=1}^{N} \sum_{t=1}^{T} log \frac{exp\, v(d_{i,t}, x_{i,t} | \theta, \hat{P}, \hat{f}_x)}{\sum_{d=0}^{J} exp\, v(d, x_{i,t} | \theta, \hat{P}, \hat{f}_x)} \tag{24}$$

Although asymptotically equivalent to the partial MLE, Monte Carlo experiments show that in finite samples, this estimator—along with all Hotz and Miller two-stage estimators—can have larger bias. To alleviate this problem, it is possible to “update” the initial CCPs with the new ones predicted by the model. Suppose we have θ^0 that maximize $Q(\theta, \hat{P}, \hat{f}_x)$. We can use these θ to form new CCPs according to the model. Using these new CCPs and the estimates of the transition probabilities, we can form new value functions and solve for new parameters θ^1. This process is repeated until the estimates of the parameters converge. Although all of these estimates have the same asymptotic variance, the small sample performance of the estimator can be substantially improved as a result of the iterations.

References

[1] *National Defense Authorization Act for Fiscal Year 2016*.

[2] Maldon Jr, Alphonso, Larry L Pressler, Stephen E Buyer, Dov S Zakheim, Michael R Higgins, Peter W Chiarelli, Edmund P Giambastiani, JR Kerrey, and Christopher P Carney. 2015. *Report of the Military Compensation and Retirement Modernization Commission.*

[3] Breth, Bruce R. 1998. *An historical analysis and comparison of the military retirement system and the federal employee retirement system.*

[4] Henning, Charles A, ed. 2011. *Military retirement reform: A review of proposals and options for congress*, LIBRARY OF CONGRESS WASHINGTON DC CONGRESSIONAL RESEARCH SERVICE.

[5] Office, Congressional Budget. 1999. *The Effects of the Military Retirement Reform Act of 1986 on Midcareer Retention.*

[6] Huff Jared, M., K. Pinelis Yevgeniya, and W. Wenger Jennie. 2013. *Adjusting First-Term Contract Lengths in the Navy: Implications and Recommendations (U).* DRM-2013-U-004794-Final

[7] Goldberg, Matthew, S. 2001. *A Survey of Enlisted Retention: Models and Findings (U).* CRM D0004085.A2 Final.

[8] Warner, John T, and Saul Pleeter. 2001. "The personal discount rate: Evidence from military downsizing programs." *American economic review* 91 (1): 33-53.

[9] Simon, Curtis J, John T Warner, and Saul Pleeter. 2015. "Discounting, cognition, and financial awareness: New evidence from a change in the military retirement system." *Economic Inquiry* 53 (1): 318-334.

[10] Hansen, Michael L., and Jennie W. Wenger. 2002. *Why Do Pay Elasticity Estimates Differ? (U).*

[11] Hosek, James R, and Christine E Peterson. 1985. *Reenlistment Bonuses and Retention Behavior.*

[12] Gray, Bradley M, and James E Grefer. 2012. "Career earnings and retention of US military physicians." *Defence and Peace Economics* 23 (1): 51-76.

[13] Hattiangadi, Anita U., Deena Ackerman, Maria Kimble, and Aline O. Quester. 2004. *Cost-Benefit Analysis of Lump Sum Bonuses for Zone A, Zone B, and Zone C Reenlistments: Final Report (U).* CRM D0009652.A4 1Rev.

[14] Cox, Gerald, and Shannon M. Phillips. 2012. *The Costs and Benefits of Converting to Lump-Sum SRBs (U).*

[15] Lien, Diana, S., Molly F. McIntosh, and Darlene E. Stafford. 2011. *Combat Compensation and Continuation in the Active and Reserve Components (U).*

[16] Lien, Diana S., and David Gregory. 2009. *Deployment Tempo and Reenlistment Within the EOD Community (U).*

[17] Hosek, James, and Paco Martorell. 2009. *How Have Deployments During the War on Terrorism Affected Reenlistment?* : Rand Corporation.

[18] Lien, Diana S., Aline O. Quester, and Robert W. Shuford. 2008. *Marine Corps Deployment Tempo and Retention from FY04 through FY07 (U).* CRM D0018757.A2 Final.

[19] Kraus, Amanda, Diana Lien, S., and Bryan Orme, K. 2003. *The Navy Survey on Reenlistment and Quality of Service: Using Choice-Based Conjoint to Quantify Relative Preferences for Pay and Nonpay Aspects of Naval Service (U).* CRM D0008416.A2 Final.

[20] Grefer, James E., Diana S. Lien, Kletu S. Lawler, and Erin M. Rebhan. 2010. *Determining the Effects of the Military Commissary Benefit on Recruitment and Retention (U).*

[21] Pinelis, Yevgeniya K., and Jared M. Huff. 2014. *The Economy and Enlisted Retention in the Navy: Volume I: Main Report (U).* CNA. DRM-2014-U-007301-Final.

[22] Asch, Beth J, Michael G Mattock, and James Hosek. 2015. *Reforming military retirement: Analysis in support of the Military Compensation and Retirement Modernization Commission.*

[23] Aguirregabiria, Victor, and Pedro Mira. 2010. "Dynamic discrete choice structural models: A survey." *Journal of Econometrics* 156 (1): 38-67.

[24] Arcidiacono, Peter, and Paul B. Ellickson. 2011. "Practical Methods for Estimation of Dynamic Discrete Choice Models." *Annual Review of Economics* 3 (1): 363-394.

[25] Joseph Hotz, and Robert A. Miller. 1993. "Conditional Choice Probabilities and the Estimation of Dynamic Models." *The Review of Economic Studies* 60 (3): 497-529.

[26] Aguirregabiria, Victor, and Pedro Mira. 2002. "Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models." *Econometrica* 70 (4): 1519-1543.

CNA

This report was written by CNA's Resource Analysis Division (RAD).

RAD provides analytical services—through empirical research, modeling, and simulation—to help develop, evaluate, and implement policies, practices, and programs that make people, budgets, and assets more effective and efficient. Major areas of research include health research and policy; energy and environment; manpower management; acquisition and cost; infrastructure; and military readiness.

DRM-2018-U-017177-Final

CNA is a not-for-profit research organization
that serves the public interest by providing
in-depth analysis and result-oriented solutions
to help government leaders choose
the best course of action
in setting policy and managing operations.

Nobody gets closer—
to the people, to the data, to the problem.

www.cna.org ● 703-824-2000

3003 Washington Boulevard, Arlington, VA 22201