

Why Do Pay Elasticity Estimates Differ?

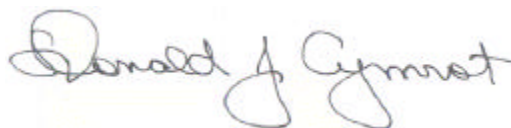
Michael L. Hansen • Jennie W. Wenger



4825 Mark Center Drive • Alexandria, Virginia 22311-1850

Approved for distribution:

March 2002

A handwritten signature in black ink that reads "Donald J. Cymrot". The signature is written in a cursive style with a large initial "D" and "C".

Donald J. Cymrot, Director
Workforce, Education and Training Team
Resource Analysis Division

This document represents the best opinion of CNA at the time of issue.
It does not necessarily represent the opinion of the Department of the Navy.

Approved for Public Release; Distribution Unlimited. Specific authority: N00014-00-D-0700.
For copies of this document call: CNA Production Services (703) 824-2123 or (703) 824-2123.

Contents

Summary	1
Background.	1
Objectives.	1
Data and methodology.	2
Findings.	3
Implications and recommendations	4
Introduction	7
Pay elasticities, reenlistment forecasts, and marginal cost	11
A baseline model of enlisted retention	13
Data	13
Composition of sample	13
Measuring reenlistment	14
Measuring pay.	15
Identifying taste for military service	16
Descriptive statistics.	19
Model specification	19
Results	22
Estimates using common alternative specifications	25
Excluding those not eligible to reenlist	26
Excluding variables used to predict civilian earnings	27
Excluding race/ethnicity and age	28
Excluding controls for fiscal year	29
Calculating separate effects for military and civilian compensation	30
Sensitivity of estimates to grouping of occupations	32
Estimating occupation-specific pay elasticities	33
Changes in the discount rate	35
Military-civilian pay ratio	37
Treatment of extensions	38
Grouped data.	39

Results using grouped data	40
Are grouped data as sensitive to changes in model specification?	41
Summary of alternative specifications.	42
Has the pay elasticity changed over time?	45
Estimating pay elasticities separately by fiscal year	45
Did Sailors respond differently to pay during the drawdown?	47
Choosing a model of reenlistment behavior	49
Goodness of fit	49
Predictions of reenlistment behavior	51
Within-sample predictions of reenlistment behavior .	53
Predictions of reenlistment behavior using test data .	56
Out-of-sample predictions of reenlistment behavior .	57
A framework for updating reenlistment models	63
Conclusion	67
Appendix A: UICs and geographical groups	69
Appendix B: Navy enlisted ratings and occupational groups . .	71
Appendix C: Coefficients from baseline model	75
Appendix D: Results of baseline model	77
Characteristics of military service	77
Paygrade	79
Sea/shore rotation	79
VSI/SSB and previous extensions	79
Differences by rating	80
Geographic differences	81
Demographic characteristics.	82
Race/ethnicity	83
Marital status	83
Other characteristics	84
Fiscal year effects	84
References	87
List of tables	91

Summary

Background

The supply of manpower has been a concern of the Navy since the creation of the All-Volunteer Force. To effectively man a volunteer force, the Navy must offer compensation that motivates men and women not only to enlist in the Navy, but to remain in the Navy past their initial commitments. A substantial body of literature examines the supply of enlistments and reenlistments and, in particular, the relationship between changes in relative compensation and changes in reenlistment behavior.

The literature often focuses on estimating the magnitude of this relationship, controlling for other factors that affect the reenlistment decision. This empirical approach produces estimates of a *pay elasticity of reenlistment*, which measures the percentage change in reenlistment associated with a 1-percent increase in pay. Even though this approach is conceptually straightforward, the previous literature has produced widely varying estimates of this relationship; because policy-makers rely on these estimates to set pay competitively, a precise understanding of the magnitude of this relationship is crucial to shaping the force.

Objectives

Given these concerns, the objective of this study is to examine the potential sources of variation in the literature. It is not clear whether the different estimates result from differences in methodology, differences in the composition of the sample, or actual changes over time in the responsiveness of enlisted personnel to pay. We examine each possibility to assess the degree to which researchers' decisions influence the estimates of the pay elasticity versus the degree to which reenlistment behavior has changed over time.

Data and methodology

We estimate the relationship between relative military compensation and enlisted retention in two stages. First, we use the personal characteristics of enlisted personnel to predict both future military compensation and future civilian earnings opportunities. Second, we use these predicted earnings, as well as additional information on enlisted personnel, to estimate the relationship between compensation and retention. We focus primarily on the Annualized Cost of Leaving (ACOL) framework, and use the dichotomous logit model when estimating the determinants of the probability of reenlistment.

We begin by estimating a baseline model of enlisted retention that reflects an empirical specification consistent with decisions made by the majority of researchers in the field. Next, we compare these estimates with those obtained using a variety of empirical specifications also found in the literature; in this way, we can trace out the effects of empirical specifications on the estimates. We then examine whether the pay elasticity has changed over time, and compare this variation over time with the variation we observe using different empirical models.

Finally, we attempt to quantify the relative success of each of these models at predicting reenlistment behavior. We make use of other data from the same time period as the data with which our baseline model was estimated, as well as more recent data, to assess the degree to which these models accurately predict reenlistment.

We use two primary sources of data when examining the relationship between relative compensation and the reenlistment decisions of Navy enlisted personnel. The first is the Enlisted Master Record (EMR) data, which we use to provide information on zone A reenlistment decisions and the demographic characteristics of the enlisted members who make these decisions. The second source of data, the March Current Population Surveys (CPS), provides information on civilian earnings opportunities. These sources include data spanning the FY87–99 time period.

Findings

Our results suggest that estimates of the pay elasticity of reenlistment are highly sensitive to the choice of empirical specification. Our baseline model generates a pay elasticity estimate of 1.5; in other words, a 1-percent increase in pay is predicted to cause a 1.5-percent increase in reenlistment. Similarly, a one-level increase in the selective reenlistment bonus (SRB) multiplier is predicted to increase reenlistment by 2.5 percentage points. Both of these estimates lie well within the range of previous estimates. The observed relationships between other explanatory variables and reenlistment are also consistent with previous research.

Alternative specifications, however, yield pay elasticities ranging from 0.4 to 2.9, a statistically significant and economically sizable range. We observe similar variation in the relationship between reenlistment bonuses and reenlistment behavior. Because we use the same data to estimate each alternative model, the variation in the estimated pay elasticities does not reflect real changes in the responsiveness to pay of Navy enlisted personnel, but rather differences in the amount of responsiveness that these models *attribute* to pay.

In contrast, there is very little variation in the pay elasticity over time, with the only significant changes occurring at the beginning and end of the drawdown. We conclude, then, that most of the variation in estimates found in the previous literature results from differences in the empirical approach of researchers, rather than from differences in the reenlistment behavior of enlisted personnel.

All of the specifications we test are defensible from different theoretical points of view, but their ability to generate accurate predictions of reenlistment behavior sheds some light on which specification should be preferred. When looking at in-sample predictions of reenlistment, models designed to predict reenlistment behavior for particular subsets of the data generate the most accurate predictions for these subsets. However, these models also do the worst job at predicting reenlistment for even a slightly different subset of the data. In general, the baseline model performs fairly well at predicting reenlistment rates for different groups of ratings.

We also use each of these specifications to predict reenlistment rates for FY00 and compare these predictions to actual reenlistment in this year. In general, the baseline model continues to be the most robust specification when predicting reenlistment. For most of our tests, reenlistment rates predicted using the baseline model are not statistically different from actual reenlistment rates. Though alternate models perform well in some tests, they perform poorly in others. We conclude, therefore, that the baseline model, with an elasticity of 1.5, provides the best “fit” of the data on Navy enlisted personnel.

Implications and recommendations

A central finding of our research is that estimates of the pay elasticity of reenlistment are highly sensitive to the choice of empirical specification. However, differences in the predictive power of these alternative specifications are smaller than differences in their pay elasticities. This seems puzzling at first: how can two models, using identical data, explain the same behavior equally well while using two very different elasticities? Also, how can these two models have similar predictive power?

The answer is that, although these models *describe* the same behavior, they differ in the degree to which they *ascribe* that behavior to pay. Models with smaller pay elasticities of reenlistment place more emphasis on other variables in their explanations of reenlistment behavior. For this reason, the ability of a reenlistment model to make accurate, out-of-sample forecasts is important in choosing among various specifications.

Our results also suggest that the pay elasticity of reenlistment has not changed markedly over time. We do find some evidence that Sailors’ responsiveness to pay was different during the drawdown period, but there is no evidence of a well-established time trend. Furthermore, the variation that we do observe over time is substantially smaller than the variation associated with different empirical specifications. Therefore, we conclude that most of the variation in the literature results from differences in empirical approach rather than differences in the behavior of Navy enlisted personnel.

Although our data do not suggest dramatic changes in behavior since the drawdown, we recommend incorporating data from future fiscal years into the model as they become available. It is likely that Sailors making reenlistment decisions are more similar to their contemporaries than to their predecessors. If so, inclusion of more recent data will only improve the ability of reenlistment models to forecast future reenlistment rates.

Finally, our analysis suggests some principles for updating and reestimating models of reenlistment behavior. Even the most accurate estimates have a *range* associated with them because of the uncertainty inherent in any forecast. In many cases, actual reenlistment rates will not lie outside the range of estimates implied by a model's predictions, even if the point estimates differ. Reestimation of a reenlistment model is necessary only when actual reenlistment rates consistently fall outside the *range* of predicted reenlistment rates.

Many factors that influence reenlistment decisions are difficult or impossible to measure. Most models attempt to capture these effects by indicating the fiscal year in which a decision was made. Even though this helps the model fit the data used in the analysis, profound shifts in factors not included in these models call for reestimation. An obvious example is a potential change in reenlistment behavior in response to the terrorist attacks in 2001. Although it is difficult to predict how Sailors will respond to these attacks, incorporating this information into a model of reenlistment behavior will help to improve subsequent forecasts.

Similarly, profound shifts in the demographic composition of the enlisted force could require modifications to models of reenlistment behavior. Such a shift does not *guarantee* a change in the relationships found in existing models; over the time period we examine, Sailors' characteristics changed dramatically yet their response to pay did not. However, large shifts in demographic composition make these changes more likely. For example, as more and more women join the Navy and face reenlistment decisions, models estimated from data consisting primarily or exclusively of male Sailors are less likely to generate accurate predictions of reenlistment behavior.

Introduction¹

The supply of manpower has been a concern of the Navy since the creation of the All-Volunteer Force. To effectively man such a force, the Navy must offer compensation that motivates people not only to enlist in the Navy, but to remain in the Navy (i.e., reenlist) throughout their careers. A substantial body of literature examines the supply of enlistments and reenlistments and, in particular, the relationship between changes in relative compensation and changes in reenlistment behavior.²

Much of this literature focuses on estimating the magnitude of this relationship, controlling for other factors that affect the reenlistment decision. This empirical approach produces estimates of the *pay elasticity of reenlistment*, which measures the percentage change in the reenlistment rate associated with a 1-percent increase in pay. A similar measure is the *SRB effect*, which measures the percentage point increase in the reenlistment rate associated with a one-level increase in reenlistment bonuses. Given reenlistment targets and these estimates, it is possible to estimate the degree to which compensation should be raised to eliminate any projected manning shortfalls.

Reference [1] documents the sizable variation in the estimates of this important relationship.³ In general, recent estimates of the pay elasticity have been notably smaller than those found in earlier research. It is not clear whether these smaller estimates result from differences in methodology across studies or actual changes in the responsiveness of enlisted personnel to pay. Understanding the causes of differences in these estimates is not simply a matter of casual empirical curiosity.

-
1. We are extremely grateful to Matt Goldberg and Amanda Kraus for their comments, insights, and suggestions.
 2. Reference [1] contains a summary of this literature.
 3. In particular, see table 2 of reference [1].

Policy-makers rely on these estimates to set reenlistment bonus levels and to forecast future reenlistment, so reliable estimates of this relationship are crucial to shaping the force.

The Assistant Deputy Chief of Naval Operations, Manpower and Personnel (N1B) requested that CNA analyze the relationship between changes in compensation and changes in reenlistment behavior of Navy enlisted personnel. Another request is for CNA to suggest a framework for updating reenlistment models when their forecasts differ from observed behavior. This study will help the Navy to better understand the relationship between compensation and reenlistment and will allow the Navy to more effectively target compensation to address manning shortfalls and retain a high-quality workforce.

We begin with a discussion of the relationships between the pay elasticity of reenlistment, reenlistment forecasts, and the marginal cost of reenlistment. A higher pay elasticity implies that, for a given increase in compensation, more people are motivated to reenlist. Consequently, different pay elasticities can lead to different forecasts of future reenlistment, as well as different costs of reenlistment. We demonstrate that these differences associated with different pay elasticities are not trivial, and that the size of the pay elasticity is an important policy parameter.

Following this discussion, we briefly describe our data and then present a “baseline” model of enlisted retention. This baseline model reflects an empirical specification that is consistent with decisions made by the majority of researchers in the field. Though our focus is on the pay elasticity of reenlistment, we also discuss the relationship between other individual characteristics and the propensity to reenlist.

The third section compares our estimates from the baseline model with those obtained using a variety of alternate empirical specifications also common in the literature. In this way, we can trace out the effects of empirical specifications on the estimates. In the next section, we examine whether the pay elasticity has changed over time, and we compare this variation over time with the variation we observe using different empirical models.

Given that all of these alternative specifications are defensible from a theoretical standpoint, the fifth section attempts to quantify the relative success of each of these models at predicting reenlistment behavior. We make use of other data from the same time period with which our baseline model was estimated, as well as more recent data, to assess the degree to which these models accurately predict reenlistment. The last main section discusses a framework for updating reenlistment models when their forecasts differ from observed behavior; the final section concludes.

Pay elasticities, reenlistment forecasts, and marginal cost

Before we analyze the variation in estimates of the pay elasticity of reenlistment, it is useful to discuss the manning and financial implications to the Navy of the size of the pay elasticity. Because elasticities measure the “responsiveness” of people to changes in pay, higher elasticities mean larger increases in reenlistment for a given pay change. Similarly, the marginal costs of achieving reenlistment targets depend heavily on the pay elasticity. We consider two examples below to illustrate how different pay elasticities affect reenlistment forecasts and the costs of reenlistment.

The pay elasticity of reenlistment allows planners to forecast the increase in reenlistment associated with an expected increase in pay. Consider, for example, a current reenlistment rate of 43 percent and an expected across-the-board increase in basic pay of 6 percent. If the pay elasticity is 1.0, a 6-percent pay raise would result in a 6-percent increase in reenlistment; in this example, the reenlistment rate would rise from 43 to 46 percent. In contrast, a pay elasticity of 3.0 would result in an 18-percent increase in reenlistment.⁴ In this case, the reenlistment rate would rise from 43 to 51 percent. These differences in reenlistment are not trivial; using FY00 data, the higher elasticity would imply over 900 *additional* reenlistments given the increase in pay. For planners trying to achieve a given endstrength, the size of the pay elasticity is important for setting recruiting targets and for anticipating the experience mix of the force.

It could be argued that, regardless of the pay elasticity, an across-the-board pay raise “buys” very little reenlistment for the money because most of the increase in pay goes to people who (1) would have

4. Both of these elasticities are within the range of estimates in the following analysis.

reenlisted even without the pay raise or (2) do not make a reenlistment decision during the fiscal year. For this reason, economists have long argued that the use of selective reenlistment bonuses (SRBs) is a more effective method by which to increase reenlistment through increases in compensation. The attractive feature of the SRB is that it is “selective,” or “targeted,” and is awarded only to individuals facing a reenlistment decision who actually agree to reenlist. In contrast, increases in basic pay are given to all enlisted Sailors.

Different estimates of the relationship between SRBs and reenlistment lead to very different estimates of the costs associated with offering the SRB. Consider, for example, the Aviation Electronics Technician (AT) rating in FY99. In our data, the median SRB multiplier was 3, and 36 percent of the ATs in our sample chose to reenlist. Suppose that the Navy sets a target reenlistment rate of 42 percent for ATs; the amount by which SRBs should be increased depends completely on the relationship between changes in SRBs and changes in reenlistment. If a one-level increase in the SRB multiplier leads to a 2-percentage-point increase in reenlistment, increasing the SRB multiplier to 6 would be necessary to achieve a reenlistment rate of 42 percent. However, if a one-level increase in the SRB level leads to a 3-percentage-point increase in reenlistment, the Navy could achieve a 42-percent reenlistment rate by increasing the SRB multiplier to 5.⁵ A one-level difference in SRBs might appear small, but the costs of reaching the reenlistment target are more than *40 percent* higher with the smaller relationship between SRBs and reenlistment. When assessing the cost-effectiveness of increases in SRBs, then, the size of the relationship between pay and reenlistment is a vital piece of information to policy-makers.

5. Again, both are within the range of estimates found in the literature.

A baseline model of enlisted retention

A notable feature of the empirical literature on reenlistment is the sizable variation in estimates of the pay elasticity of reenlistment. Because a focus of our research is to understand the underlying causes of the differences in these estimates, we begin by specifying a baseline model. We then compare estimates from a variety of alternative empirical specifications to those from the baseline model. In this way, we can trace out the effects of different empirical specifications on the estimates. In our baseline model, we make decisions that are consistent with those made by the majority of researchers in the field.

Data

Reference [2] discusses in detail the composition of our sample, definitions of key variables of interest, and assumptions about the economic environment faced by the Sailor. In this section, we briefly outline the fundamental choices on sample selection that we made for our analysis.

Composition of sample

For FY87 to FY99, we focus on the first, nonobligated, long-term decision of the enlisted Sailor.⁶ Consistent with the bulk of previous empirical research, we focus only on male Sailors and exclude those who work in the nuclear field or are not rated by the time of the first reenlistment decision (GENDETs). We exclude most data “outliers” by focusing on those in paygrades E-3 to E-6 and aged 19 to 40.⁷

6. See [2] for a detailed explanation of what is meant by the “first, non-obligated, long-term decision.”

7. Ninety-seven percent of our sample is 30 or younger. Excluding people over 30 does not materially affect our results.

Given the concerns of [3], that “many of the actions leading to ineligibility are, in fact, a manifestation of the decision not to continue” with military service, we include Sailors classified as ineligible to reenlist in our baseline analysis.⁸ Following a discussion of our empirical results, we examine the degree to which estimates of the pay elasticity of reenlistment are altered by a different treatment of Sailors not eligible to reenlist.

To assess the relative success of various models at predicting actual reenlistment behavior, we randomly split our sample of enlisted Sailors and use one-half to generate our estimates of the pay elasticity of reenlistment. In a subsequent section, we use the other half of the sample to compare actual and predicted reenlistment.

Measuring reenlistment

Reference [2] argues that the difference between unconditional extensions (a continuation of the current enlistment contract for at least 24 months) and reenlistments (a signing of a new enlistment contract) is often one of semantics. Rather, it is the *length* of the reenlistment or extension that is responsible for any meaningful distinction. People who reenlist or extend for at least 36 months are entitled to receive any SRB for which they are eligible; those who reenlist or extend for less than 36 months are not.

In our baseline specification, then, we define our dependent variable as those who reenlist or extend for at least 3 years. Note, however, that this dichotomous framework makes no distinction between people who separate and those who reenlist or extend for fewer than 3 years.

Although the empirical literature has no consistent, widely accepted definition of “reenlistment,” our baseline definition is similar to that used by a number of other researchers. For example, both [4] and [5] include all those making unconditional extensions with those who formally reenlist. In contrast, [3] uses a slightly different definition, excluding people from the data until they make a decision

8. By definition, all who are “ineligible to reenlist” at the time of the reenlistment decision separate from the Navy.

to separate or formally reenlist. Our definition, then, falls somewhere between these two approaches.

These differences in definitions do affect a sizable proportion of the data. In our sample, 87 percent of all unconditional extensions and 10 percent of all reenlistments are fewer than 36 months in length; these make up about 8 percent of our entire sample. In our analysis of the sensitivity of estimates to changes in specification, we address whether estimates of the pay elasticity of “reenlistment” are sensitive to changes in the definition of “reenlistment.”

Measuring pay

For the majority of our analysis, we make use of the Annualized-Cost-of-Leaving (ACOL) framework to estimate the relationship between relative military compensation and reenlistment. In an ACOL model, “pay” is the discounted difference between expected military compensation (if a person were to reenlist) and expected civilian compensation (if a person were to leave the Navy). The theoretical advantage of the ACOL framework is that it reveals the (person-specific) time horizon over which relative military compensation should be calculated.⁹

Most of the literature has focused on the three largest components of military pay: basic pay, allowances for subsistence and housing, and retirement pay [1]. In addition, measures of military pay used in reenlistment models usually include selective reenlistment bonuses, which are often the policy tool being examined. We follow the literature in our computation of military compensation and predict future military compensation based on predictions of promotion opportunities, future dependency status, and retirement pay.¹⁰

We make use of the March Current Population Surveys (CPS) to estimate civilian earnings opportunities for Navy enlisted personnel. Using these data, we estimate log earnings regressions, controlling

9. See [1] for an excellent summary of the ACOL framework.

10. See [2] for a more detailed discussion of the computation of military compensation.

for race, educational attainment, and experience. We estimate these regressions separately for each year, and for technical and nontechnical occupations.¹¹ The estimated coefficients allow us to predict civilian earnings for each enlisted person.

Finally, we use the predicted military and civilian compensation to compute the ACOL variable for each person at each future year. In an ACOL framework, the maximum of these values is used in estimation of the relationship between pay and reenlistment.¹²

Identifying taste for military service

As [2] recognizes, selecting appropriate proxies for “taste for military service” is critical to the estimation of an ACOL model. In our model, we include several variables that reflect both a relative taste for military service and different earnings opportunities in civilian labor markets (e.g., race/ethnicity). We also identify a few variables that are reasonable proxies of taste for military service and do not substantively affect earnings opportunities.

Expected sea/shore rotation

Sea duty is one of the central features of a career in the Navy and a characteristic on which many researchers and policy-makers have focused. The conventional wisdom is that Sailors have a relative preference for shore duty over sea duty, and that Sailors with a disproportionate amount of sea duty are those most likely to separate from the Navy.¹³ Most researchers, however, have found a *positive* relationship between current sea duty and the propensity to reenlist.¹⁴ It has been hypothesized that this reflects an expectation of an imminent rotation to shore duty and not a relative preference for sea duty; indeed, references [5] and [8] estimate a significant, negative relationship

11. In defining technical and nontechnical occupations, we follow [6].

12. See [2] for a detailed discussion of the actual construction of the ACOL variable.

13. For example, see [5] for a discussion of the relationship of sea duty, expected sea duty, and retention.

14. For example, see [7].

between the amount of second-term sea duty and the level of first-term reenlistment. We follow this literature and try to classify Sailors by the type of duty to which they can expect to rotate in their next term.

Our data indicate whether enlisted personnel are currently on sea or shore duty. We assume that Sailors currently serving at sea will rotate to shore on completion of their sea tours, and that those currently on shore duty will rotate to sea duty. A comparison of a person's prescribed sea (shore) tour length (PST) with the number of months currently served in that duty allows us to calculate the expected expiration of the current tour.

For Sailors currently on sea duty, we classify those who expect to spend at least the next 12 months on sea duty as facing sea duty; for Sailors currently on shore duty, those who expect to rotate to sea duty within the next year are also classified as facing sea duty. If a person on sea (shore) duty is expected to spend less (more) than the next 12 months on sea (shore) duty, he is classified as facing shore duty.

Our data include some Sailors who have already spent more time on sea/shore duty than their PST indicates they should. This could occur for a variety of reasons. For example, Sailors could remain at sea simply because their tours extend beyond their PST. Similarly, personnel may serve back-to-back sea or shore tours that are not differentiated in our data. Rather than make assumptions about their individual expectations, we classify them as currently being on a "long" sea/shore tour and estimate the effect of these tours on reenlistment decisions.

Finally, Sailors in certain ratings expect to spend very little time on sea duty. For example, those in the Cryptology ratings serve tours within and outside the United States but do not have the traditional sea/shore rotation. For these people, we include a variable to indicate that a person's rating is characterized by "nontraditional" rotations.¹⁵

15. This variable is time-dependent because some ratings have switched over time between a traditional and a nontraditional sea/shore rotation.

Current duty station

A person's taste for military service may also be affected by the environment in which he serves. It is likely that some billets are preferred over others for such reasons as the civilian opportunities available for one's spouse, the quality of public services, the geographic location, or even the weather.¹⁶ To measure the effect of these surroundings, we match the Unit Identification Code (UIC) where the Sailor is currently stationed to the state in which each UIC is located. In most cases, the characteristics that affect taste for military service probably vary by state. However, for some states with a large number of UICs, we subdivide the states into multiple geographic locations.¹⁷

VSI/SSB eligibility

To support the drawdown, the Navy used two programs—primarily from FY92 to FY95—designed to *encourage* separation: the Voluntary Separation Incentive (VSI) and Special Separation Benefit (SSB) programs.¹⁸ Although no one in zone A was ever directly eligible for VSI/SSB, our data indicate whether zone A Sailors were in a rating or community for which VSI/SSB was offered to more senior enlisted personnel. We control for these programs in our estimation because their use reflects a period during which the Navy's attitudes toward reenlistment were different, particularly for certain ratings. If a person belongs to a rating or community that is targeted by VSI/SSB, it is likely that the Sailor's expectations about future promotion differ

16. It is possible that one might observe a *negative* relationship between being stationed at a "desirable" UIC and the propensity to reenlist if being stationed at a "desirable" UIC increases the likelihood of being rotated to an "undesirable" UIC. Although this is similar to the argument for including measures of *expected* sea duty rather than current sea duty, we have no reliable method of determining the UIC at which a person will be stationed if he chooses to reenlist.

17. One example is Virginia. We hypothesize that living and serving in Norfolk may be fundamentally different from living and serving in the Washington, DC, area of northern Virginia. See appendix A for a complete listing of UIC-state matches and an explanation of how some UICs were grouped.

18. Reference [9] discusses the VSI/SSB programs in greater detail.

in a way that is not captured by our specification of expected military compensation.

Previous conditional extension

Some Sailors execute conditional extensions (lasting less than 24 months) at the end of their initial obligation. Reference [10] shows that these Sailors are much more likely to reenlist than those without a history of conditional extensions. Therefore, we include a variable indicating whether a person executed such an extension before his zone A reenlistment decision as a measure of relative taste for military service.

Descriptive statistics

Table 1 presents the means (or, where appropriate, the proportion of our sample with each characteristic) of the variables used in our estimation. For purposes of clarity, variables have been grouped to indicate the particular aspect of a Sailor's environment that they are hypothesized to reflect.

Model specification

To estimate the effect of military compensation on reenlistment, we make use of a standard logit regression model.¹⁹ Other, more sophisticated models have been used in previous studies; for our baseline model, however, we chose a commonly used econometric framework.²⁰

In addition to a measure of relative military compensation in models of reenlistment behavior, variables used by previous researchers can be separated into three general categories: variables that affect military compensation, variables that affect civilian earnings opportunities, and those that reflect a relative preference for military service. These categories are not always mutually exclusive; for example, both military and civilian compensation vary by fiscal year in virtually all specifications. Most of the previous literature, however, is divided

19. For a detailed explanation of the logit model, see [11].

20. Use of a probit regression model generates virtually identical results.

Table 1. Descriptive statistics

Variable	Mean	Variable	Mean
Reenlistment rate	0.32	UIC^a	
Economic data		Northern (coastal) California ^b	0.07
ACOL (\$1,000)	3.68	Southern (coastal) California ^b	0.23
Unemployment rate (state)	5.92	Hawaii ^b	0.04
Characteristics of military service		Coastal Virginia ^b	0.26
E-3 ^b	0.14	Washington State ^b	0.06
E-4 ^b	0.61	South Carolina ^b	0.04
E-5 ^b	0.25	Northern (coastal) Florida ^b	0.08
E-6 ^b	0.01	Other ^b	0.21
Length of service (months)	49.4	Personal characteristics	
Previous extension ^b	0.16	Married ^b	0.40
Eligible to reenlist ^b	0.98	Number of children	0.32
Rating eligible for VSI/SSB ^b	0.15	Age	23.8
Nontraditional rotations ^b	0.04	AFQT	58.5
Long sea tour ^b	0.10	White ^b	0.74
Long shore tour ^b	0.06	Black ^b	0.16
Expected sea tour ^b	0.50	Hispanic ^b	0.07
Expected shore tour ^b	0.31	Other ethnicity ^b	0.03
Rating group^c		Fiscal year	
SEABEE Construction ^b	0.02	FY87 ^b	0.07
Non-SEABEE Construction ^b	0.01	FY88 ^b	0.10
Marine Engineering ^b	0.16	FY89 ^b	0.10
Ship Maintenance ^b	0.06	FY90 ^b	0.10
Aviation Maintenance ^b	0.16	FY91 ^b	0.10
Aviation Ground Support ^b	0.07	FY92 ^b	0.11
Media ^b	0.01	FY93 ^b	0.09
Logistics ^b	0.09	FY94 ^b	0.08
Administration ^b	0.04	FY95 ^b	0.06
Data Systems ^b	0.05	FY96 ^b	0.05
General Seamanship ^b	0.11	FY97 ^b	0.06
Health Care ^b	0.07	FY98 ^b	0.05
Cryptology ^b	0.03	FY99 ^b	0.04
Ordnance Systems ^b	0.05		
Communications/Sensor ^b	0.03		
Weapons Systems/Control ^b	0.06		

a. Appendix A contains a complete listing of the locales that make up each category.

b. Proportion with this characteristic is presented.

c. Appendix B contains a complete listing of the individual ratings that make up each category.

over which of these variables “should” be included in models of reenlistment behavior.²¹

Despite an apparent consensus that variables affecting compensation can (and likely do) influence reenlistment decisions above and beyond their effects through pay, the decision by some researchers to exclude these variables seems to reflect concern over eliminating variation in the pay variable.²² Given our relatively detailed data and construction of expected military and civilian pay, there is a considerable amount of variation in our measure of relative military compensation. In our baseline specification, then, we include variables used to predict military and civilian compensation as measures of taste for military service. With minor exceptions, our measures of taste are quite similar to those used by [3].²³

Because the logit model estimates a nonlinear relationship between the explanatory variables and the probability of reenlistment, the interpretation of the coefficients is not straightforward. To facilitate interpretation of the results, we calculate and present two different statistics that measure the relationship between changes in military compensation and changes in reenlistment behavior. The first is the *pay elasticity of reenlistment*, which measures the percent change in reenlistment associated with a 1-percent increase in basic pay. The second is the *SRB effect on reenlistment*, which measures the percentage-point change in reenlistment associated with a 1-level increase in the SRB multiplier.

21. See [2] for a discussion of the econometric and methodological issues.

22. For example, [8] states that “attempts to include [education level, mental group and race] as well as ACOL were unsuccessful due to the multicollinearity between these variables and ACOL” (page 31). In contrast, [3] includes race, education level, and AFQT score in its reenlistment equation but does not cite multicollinearity problems.

23. Rather than total number of dependents, we include marital status and number of children separately. Also, we include geographic location and expected sea duty ([3] uses data on Army enlisted personnel, so expected sea duty is not a relevant variable).

Results

For purposes of clarity, we focus here only on the relationship between compensation and reenlistment. Note, however, that our baseline model includes all the variables listed in table 1. A complete listing of the coefficients for all variables included in the model can be found in appendix C; a complete discussion of these results can be found in appendix D.

Table 2 presents our estimates of two key statistics that are of central interest in the literature on the supply of military manpower: the pay elasticity of reenlistment and the effect of a one-level increase in the SRB multiplier. The coefficient on relative military compensation is positive and statistically significant, suggesting that increases in military pay do lead to increases in reenlistment. Specifically, our results indicate that a 1-percent increase in basic pay leads to a 1.5-percent increase in reenlistment.²⁴ Furthermore, a one-level increase in the SRB multiplier is associated with an increase in the reenlistment rate of 2.5 percentage points.

Table 2. Pay effects—baseline model

Effect	Estimate
Pay elasticity of reenlistment	1.5 percent ^a
One-level increase in SRB multiplier	2.5 percentage points ^a

a. Zero lies outside the 99-percent confidence interval for this estimate.

It is useful to compare these estimates with those found in the previous literature on Navy enlisted personnel. Reference [1] summarizes

24. We calculate the pay elasticity of reenlistment using the approach suggested in [1]. First, we use the predicted coefficients to estimate the average predicted probability of reenlistment. Second, we increase basic pay by 1 percent and find the new maximum ACOL value for each person. Third, we use this new value to estimate a new, average predicted probability of reenlistment using the original coefficients from the regression. Finally, we compare this new prediction to the original prediction to establish the percentage change in reenlistment.

not only the theoretical literature on compensation and enlisted retention, but also estimates of pay elasticities from various studies.²⁵ Of those that deal with Navy enlisted personnel, estimates of the pay elasticity of reenlistment range from 0.8 to 3.4, with most of the estimates between 1.2 and 2.2. Our estimated pay elasticity of 1.5, then, lies well within the range of previous estimates. Our estimate of the relationship between SRBs and reenlistment is slightly larger than the “rule of thumb” presented in [1], which states that a one-level SRB increase raises the reenlistment rate by about 2 percentage points.

25. See, in particular, table 2 in reference [1].

Estimates using common alternative specifications

With these baseline estimates of the pay elasticity of reenlistment and the SRB effect on reenlistment, we now assess the degree to which these estimates change when using alternative specifications commonly found in the literature. Reference [1] argues that much of the variation in previous estimates is probably caused by differences in the average reenlistment rate found in different samples. Specifically, [1] demonstrates that pay elasticity estimates are extremely sensitive to the “point of evaluation” (i.e., the average reenlistment rate used to calculate the pay elasticity) and that even small differences in reenlistment rates across samples can almost triple estimates of the pay elasticity. Therefore, [1] argues that pay elasticities are “relatively unstable.” In contrast, SRB effects are relatively constant for a wide range of reenlistment rates and are therefore more stable. This latter finding is consistent with the relative uniformity in estimates of the SRB effect from one study to the next.

It is important to note, however, that differences in pay elasticities due to differences in model specification are *not* caused by changes in the average reenlistment rate. For the bulk of the specifications we consider below, the sample used with each model is identical. This implies that the underlying reenlistment rate is also the same from one specification to the next. Differences in our estimates of the pay elasticity are truly differences in estimates of the responsiveness to pay and are not artifacts of the functional form of the pay elasticity. In other words, holding constant the average reenlistment rate allows us to assess the degree to which different specifications lead to different predictions of the relationship between pay and reenlistment for the same set of individuals.

This section is descriptive in the sense that we discuss the extent to which estimates vary across specifications. We make no specific

recommendations about which models are most appropriate; instead, we address that subject in the next two sections.

Excluding those not eligible to reenlist

The literature is divided over whether to include ineligibles in models of enlisted retention. On one hand, the empirical literature is concerned with modeling the relationship between changes in military compensation and changes in the *voluntary* supply of labor to the Navy. Consequently, many researchers have excluded ineligible personnel from their analyses, concluding that those who are not eligible to reenlist are not making a voluntary labor supply decision.²⁶ On the other hand, [3] argues persuasively that most of the reasons for being declared “ineligible to reenlist” are the direct result of actions, or lack of actions, by the enlisted Sailor and, therefore, reflect a voluntary decision to separate from the Navy.

As table 1 indicates, only about 2 percent of our sample of enlisted personnel are considered “ineligible to reenlist” by the Navy at the time of their reenlistment decision. Because these people are not eligible to remain in the Navy at the end of their initial obligation, all of them separate at the time of their reenlistment “decision.”

Given that these ineligibles constitute a small percentage of our entire sample, it is not surprising that, when we exclude them from our analysis, estimates of the relationship between changes in compensation and changes in retention remain unaltered. This alternate specification yields an estimated pay elasticity of 1.5; similarly, our estimates indicate that an increase in SRBs by one level leads to increases in reenlistment of about 2.6 percentage points.

As [2] indicates, the percentage classified as ineligible to reenlist was dramatically different during the later years of the drawdown (FY94-95) compared with the periods before and after the drawdown (see figure 1 in [2]). Given these stark differences in eligibility during the drawdown, it is possible that models estimated using data primarily from the drawdown era could be more sensitive to the inclusion/exclusion of those classified as ineligible to reenlist. With our data,

26. Examples can be found in [5] and [7].

however, a model using data only from the drawdown yields identical pay elasticities, regardless of whether one excludes those ineligible to reenlist.²⁷

We offer three plausible explanations for this lack of variation in estimates of the pay elasticity of reenlistment. First, the sample of ineligible Sailors is so small, even during the drawdown, that it is unlikely to have a substantive influence on any estimates. Second, it is possible that we have “adequately controlled for demand differences” across Sailors, which [3] argues is the source of identical estimates when including or excluding ineligible. Finally, it is possible that the true relationship between compensation and retention is identical for all Sailors, regardless of reenlistment eligibility.

Excluding variables used to predict civilian earnings

Most researchers agree that variables affecting civilian compensation influence reenlistment decisions above and beyond their effects through pay. For example, it is well known that civilian earnings vary substantially among different ethnic groups. It is possible, however, that different ethnic groups have different relative preferences for military service. If this is the case, these variables will have two theoretically distinct effects on reenlistment: an effect due to different civilian opportunities, and an effect due to different preferences for military service. Consequently, many researchers have simultaneously used individual characteristics both to predict earnings and as separate explanatory variables in the reenlistment equation.

There are arguments, however, for *excluding* these variables from models of reenlistment behavior. From an econometric standpoint, estimates of the effect of pay on reenlistment are more precise if there are large differences in compensation in the data.²⁸ Further,

27. Using FY92–95 data, our baseline model generates a pay elasticity of 1.5, with an SRB effect on reenlistment of 2.2 percentage points.

28. As an illustration, consider the extreme case where there is no difference in compensation from one person to the next. If all have identical pay, it would be impossible to measure the effect of pay on reenlistment decisions, even if there truly is an effect.

the bulk of the variation in relative compensation comes from differences in civilian pay; compared with civilian pay, there is little variation in military compensation. Several researchers, have chosen to exclude these individual characteristics from the reenlistment equation to preserve variation in the key variable of interest.²⁹

If these characteristics are excluded, the implicit assumption is that differences in reenlistment behavior for different demographic groups result *entirely* from differences in civilian opportunities. This is explicitly argued in [7], which states that this approach “does not assume that these demographic variables are unimportant,” but that “differences in reenlistment behavior among these demographic groups are implicitly attributed to differences in their civilian opportunities.” This is also consistent with the interpretations of regression results in previous studies. For example, [5] attributes a negative relationship between ability and reenlistment to the “stronger civilian opportunities” of high-ability individuals, rather than inherently different preferences for military service.³⁰ By comparing these alternate specifications, we assess the degree to which excluding these variables from the reenlistment equation alters estimates of the pay elasticity of reenlistment.

Excluding race/ethnicity and age

As table 3 shows, including race/ethnicity and age only in predictions of civilian compensation significantly affects estimates of both the pay elasticity and the effect of SRBs on reenlistment. Our new estimates of both the pay elasticity and the effect of a one-level increase in SRBs are 40 percent lower than those implied by our baseline model.

These sizable differences are consistent with the observations of previous researchers that models using individual-level data are sensitive to the inclusion/exclusion of these variables [12]. Note, however,

29. For example, see [8].

30. Even though [5] does not use ability to predict civilian earnings, the interpretation of the relationship between ability and reenlistment is that, if one could use ability to predict civilian opportunities, ability would have no additional effect on reenlistment.

that these differences do not suggest that these variables *should* be excluded or included; they merely indicate that such a decision affects estimates of the pay elasticity of reenlistment.

Table 3. Pay effects—excluding race/ethnicity and age

Effect	Estimate
Pay elasticity of reenlistment	0.9 percent ^a
One-level increase in SRB multiplier	1.5 percentage points ^a

a. Zero lies outside the 99-percent confidence interval for this estimate.

Excluding controls for fiscal year

Excluding variables that indicate the fiscal year in which Sailors make their decisions is different from excluding controls for ethnicity and age because both civilian *and* military compensation vary by fiscal year. There are stronger theoretical reasons for including controls for fiscal year than there are for including race and ethnicity.³¹ However, table 4 shows that excluding controls for fiscal year along with race and age has very little additional effect on estimates of a Sailor’s responsiveness to pay; the pay effects are quite similar to those in table 3.

Table 4. Pay effects—excluding race/ethnicity, age, and fiscal year

Effect	Estimate
Pay elasticity of reenlistment	1.0 percent ^a
One-level increase in SRB multiplier	1.6 percentage points ^a

a. Zero lies outside the 99-percent confidence interval for this estimate.

31. Fiscal year variables capture any changes in a Sailor’s environment not already identified in our model.

Calculating separate effects for military and civilian compensation

Although ascribing differences between demographic groups entirely to differences in civilian opportunities is extreme, it is defensible from a theoretical standpoint. For example, it is difficult to think of a plausible argument for why ethnic groups might have different preferences for military service, holding constant relative civilian opportunities. It is possible that civilian opportunities are not fully captured by models of predicted civilian compensation; in theory, however, a perfectly specified model of civilian earnings would completely incorporate differences among demographic groups.

Unfortunately, our model of civilian earnings is *not* perfectly specified. Using data from the March CPS, we estimate log earnings regressions for each person in our sample. Our specification of these regressions is similar to many that appear in the labor economics literature, but it leaves out a number of factors that are likely to be correlated with earnings. For example, we do not include such factors as industry and occupation of employment because we do not know the industry/occupation in which each Sailor's "best civilian alternative" lies. Also, it is highly unlikely that the functional form we use to estimate civilian earnings is the actual function that determines the relationship between civilian characteristics and civilian earnings. If our civilian earnings function is specified with error, a reliance on that functional form to identify the effect of earnings is highly tenuous.³²

In contrast, our military pay variable is a relatively accurate representation of the pay a Sailor expects to receive in the Navy. The largest component of military pay is basic pay, which is completely deter-

32. The coefficient on civilian earnings in the reenlistment equation is identified solely from the functional form chosen to estimate civilian earnings when the variables used to predict earnings are included in the reenlistment equation. For example, if we had estimated a linear relationship between civilian earnings and Sailors' characteristics, it would not be possible to separately identify the effect of civilian earnings and the additional impact of these characteristics on reenlistment. Because civilian earnings enter our baseline model as part of ACOL, any errors in predicted civilian earnings potentially affect our ACOL estimates.

mined by rank and years of service. This is a particular functional form, but we *know* the relationship between these variables and use it in our construction of military compensation. Our estimates of future dependency status and promotion probabilities are not perfect, but the overall variance in military pay is fairly small; any differences between our measures of expectations and the actual expectations of the Sailor will probably make only small differences in the discounted future pay stream.

Our ACOL variable, therefore, is formed from the difference of two predicted variables. Because we believe the errors in expected civilian compensation to be substantially larger than those in expected military compensation, we also use an alternative specification that *separates* military pay from civilian pay and enters them as separate regressors. The maximum ACOL value is still calculated in the same fashion as in the baseline model; the only difference is that, once this value is found, the military and civilian components of this ACOL value are entered separately in the regression. The theoretical foundation of ACOL implies that this should make no difference in estimates of the effect of pay. If the model is properly specified, the coefficient on ACOL should equal the coefficient on military pay, which should be equal (in absolute value) to the coefficient on civilian pay because one dollar of military pay is assumed to have the same value as one dollar of civilian pay. Given the aforementioned problems with expected civilian compensation, however, it is possible that the coefficients might differ. Separating civilian pay from military pay allows us to focus on the unbiased effect that changes in military compensation have on reenlistment behavior.

When entering military and civilian pay separately, table 5 shows that our estimate of the pay elasticity almost doubles from 1.5 to 2.8, whereas the effect of a one-level increase in the SRB multiplier is about 25 percent smaller than that found using our baseline specification.

Table 5. Pay effects—separate effects of military and civilian pay

Effect	Estimate
Pay elasticity of reenlistment	2.8 percent ^a
One-level increase in SRB multiplier	1.9 percentage points ^a

a. Zero lies outside the 99-percent confidence interval for this estimate.

Furthermore, when we separate military and civilian pay, the coefficient on military pay is about *four times* as large as that on civilian pay; a statistical test of this difference allows us to reject, with 99-percent confidence, the hypothesis that these two coefficients are equal in absolute value. This finding is consistent with the notion that misspecification of predicted civilian compensation biases estimates of the pay elasticity when using the traditional ACOL coefficient. Separating the military and civilian components of ACOL allows us, in principle, to isolate the effect of changes in military compensation from any bias due to misspecified civilian earnings. The dramatic differences in our estimates of the pay elasticity suggest that this is not a trivial distinction.

Sensitivity of estimates to grouping of occupations

There is a possibility that the data will show a “reverse causation” between pay and reenlistment because enlisted occupations with chronically low reenlistment rates are typically given higher SRB levels. In this sense, low reenlistment “causes” higher compensation, even if Sailors respond positively to increases in pay. When using grouped data, this reverse causation would decrease the observed effect between increases in compensation and (positive) changes in reenlistment.

Although this problem is more common with grouped data, reference [1] suggests that reverse causation may also occur when using individual-level data if Sailors alter the timing of their reenlistment decisions to take advantage of expected increases/decreases in SRB levels. Furthermore, if an enlisted occupation is relatively small, it is possible that individual decision-makers in these occupations are not truly “price-takers.”

As a solution, [1] suggests using a fixed-effects estimator, controlling for each enlisted rating. According to [1], these fixed effects are “intended to capture permanent deviations between that occupation’s reenlistment rate and the overall sample average.” While our baseline model controls for broad occupational categories, there is significant variation in both reenlistment and SRB multipliers within many of these groups. This suggests that fixed-effects for individual

ratings, rather than for broad occupational categories, can further refine our estimates of the pay elasticity. Therefore, we assess the degree to which our estimates change when including a complete set of controls for a person’s rating.

Our results are consistent with the hypothesis that there is some downward bias in our baseline estimates. As table 6 indicates, including a complete set of controls for a person’s rating generates a pay elasticity of 2.7, which is about 80 percent higher than the baseline estimate. The effect of a one-level increase in the SRB multiplier is also dramatically larger.

Table 6. Pay effects—fixed effects for each rating

Effect	Estimate
Pay elasticity of reenlistment	2.7 percent ^a
One-level increase in SRB multiplier	4.4 percentage points ^a

a. Zero lies outside the 99-percent confidence interval for this estimate.

Alternatively, it is possible that differences across ratings in the propensity to reenlist, holding pay constant, were not adequately captured in our baseline specification using broad occupational categories. If this were the case, these differences would be attributed to differences in compensation or job conditions across ratings rather than to the ratings themselves. In either case, these results suggest that controlling for a person’s rating leads to substantial changes in the estimates of the relationship between pay and reenlistment.

Estimating occupation-specific pay elasticities

Reference [8] argues that working conditions vary significantly from one Navy occupation to the next. Similarly, it is likely that civilian opportunities also vary by Navy rating. Given these differences, it is not surprising that reenlistment rates vary by occupation as well (table 7).

Given these differences, it is possible that the relationship between changes in compensation and changes in reenlistment also differ by

occupation. Our baseline model allows us to observe how reenlistment rates vary by occupation while controlling for other observable characteristics, but it does not reveal any differences in elasticities. Therefore, some researchers have chosen to estimate models of reenlistment separately for broad occupational categories [3, 8].

Table 7. Pay effects—occupation-specific elasticities

Rating	Reenlistment rate	Pay elasticity	SRB effect
SEABEE Construction	0.29	0 ^a	0 ^a
Non-SEABEE Construction	0.28	0 ^a	0 ^a
Marine Engineering	0.29	2.8 ^b	4.2 ^b
Ship Maintenance	0.25	3.8 ^b	6.1 ^b
Aviation Maintenance	0.32	0.7 ^b	1.3 ^b
Aviation Ground Support	0.27	0.5 ^b	0.6 ^b
Media	0.26	1.9 ^c	2.1 ^c
Logistics	0.34	3.3 ^b	4.8 ^b
Administration	0.34	3.1 ^b	4.1 ^b
Data Systems	0.39	1.5 ^b	3.2 ^b
General Seamanship	0.28	3.2 ^b	4.5 ^b
Health Care	0.31	2.6 ^b	3.7 ^b
Cryptology	0.46	0.2 ^c	0.6 ^c
Ordnance Systems	0.43	0.3 ^b	0.7 ^b
Communications/Sensor	0.31	2.0 ^b	3.1 ^b
Weapons Systems/Control	0.41	1.3 ^b	2.8 ^b

a. No significant effect of pay on reenlistment.

b. Zero lies outside the 99-percent confidence interval for this estimate.

c. Zero lies outside the 95-percent confidence interval for this estimate.

As table 7 demonstrates, this approach suggests that individuals in different occupations do respond differently to changes in compensation. We estimate separate regressions for each of the occupation groups listed in column 1; columns 2 through 4 display each group's actual reenlistment rate, estimated pay elasticity of reenlistment, and estimated SRB effect on reenlistment, respectively.

Estimates of the pay elasticity of reenlistment range from 0.2 (Cryptology) to 3.8 (Ship Maintenance); estimates of a one-level increase

in the SRB multiplier exhibit similar variation.³³ The implication is that estimates of the pay elasticity of reenlistment are extremely sensitive to the choice of ratings being studied in the analysis.

Comparing these results to those found in [8], which uses data from FY74 to FY78, we see that some groups exhibit relatively large SRB effects during both time periods (e.g., Logistics, Health Care), whereas others consistently have relatively small relationships between pay and reenlistment (e.g., Media, Ordnance Support). However, there are also many differences between our results and those found in [8]. For example, [8] finds that Sailors in Ship Maintenance have a relatively low pay elasticity; we find their pay elasticity to be among the highest of any group. In contrast, we find relatively low pay elasticities for the Cryptology and Aviation Maintenance ratings; reference [8] finds the opposite.

There could be several reasons for the differences between these two sets of results. First, it is possible that relative working conditions and civilian opportunities changed significantly for these groups of ratings from the 1970s to the 1990s. Second, [8] estimates a relatively sparse model, including only ACOL, marital status, and the unemployment rate as covariates.

Changes in the discount rate

Construction of an ACOL framework requires an assumption about enlisted personnel's discount rates. Several studies have attempted to estimate the discount rate of enlisted personnel; these results, summarized in [1], typically range from 4 to 17 percent. The most recent estimates discussed in [1], however, place the discount rate between 26 and 37 percent for enlisted personnel [13]. Our baseline model assumed a discount rate of 20 percent, a relatively conservative estimate. To assess the degree to which changing the discount rate influences our estimates of the pay elasticity, we reestimate the baseline model, first assuming a discount rate of 10 percent, and then one of 30 percent.

33. Both SEABEE and Non-SEABEE Construction ratings groups have insignificant relationships between pay and reenlistment.

An across-the-board pay raise of 1 percent increases pay for all Sailors and at all future years of service by the same percentage. Changing the discount rate alters the discounted present value of this pay raise, but it does so in the same manner for all individuals. Because elasticities measure relationships in percentage terms, we hypothesize that the choice of discount rate will not affect estimates of the pay elasticity. A one-level increase in the SRB multiplier, however, does not increase pay by the same amount for all individuals; therefore, we expect the choice of discount rate to affect these estimates more than those of the pay elasticity.

As table 8 indicates, a reduction in the discount rate to 10 percent does not result in a significant change in the estimated pay elasticity. In fact, our estimate remains unchanged at 1.5. Yet our estimate of the effect of a one-level increase in SRB multipliers declines dramatically, from 2.5 to only 0.9 percentage point.

Table 8. Pay effects—10-percent discount rate

Effect	Estimate
Pay elasticity of reenlistment	1.5 percent ^a
One-level increase in SRB multiplier	0.9 percentage point ^a

a. Zero lies outside the 99-percent confidence interval for this estimate.

Similarly, table 9 presents our estimates after raising the discount rate to 30 percent. With this specification, the pay elasticity remains unchanged, and the effect of a one-level increase in SRB levels increases from 2.5 to 3.3 percentage points.

Table 9. Pay effects—30-percent discount rate

Effect	Estimate
Pay elasticity of reenlistment	1.5 percent ^a
One-level increase in SRB multiplier	3.3 percentage points ^a

a. Zero lies outside the 99-percent confidence interval for this estimate.

These changes in the discount rate do dramatically change the point at which the maximum ACOL is realized. As one might expect, increases in the discount rate increase the value of current compensation relative to future compensation. With a 10-percent discount rate, ACOL is maximized for virtually all enlisted personnel at the point at which retirement benefits become available. As the discount rate rises, more and more Sailors have a maximum ACOL earlier in their careers. These differences are more pronounced for those offered a zone A SRB because the (nondiscounted) value of current compensation is relatively higher.

Military-civilian pay ratio

Specifying relative military compensation as a ratio of contemporaneous military and civilian pay is a common approach in previous literature [4, 5, 7, 14].³⁴ Unlike ACOL, this approach is not firmly grounded in economic theory but can be characterized as an ad hoc specification of the economic environment faced by the Sailor. An advantage is that calculating a military-civilian pay ratio is significantly easier than constructing an ACOL framework. Furthermore, much of the previous literature that uses such a pay ratio finds estimates of the pay elasticity of reenlistment similar to those found in the ACOL literature [5, 14]. Given these similarities, the traditional thinking has been that, while ACOL has theoretical advantages and is more flexible when modeling alternative changes in compensation, both frameworks generate similar estimates of pay elasticities. Recent estimates, however, have cast doubt on these assumptions [7].

Our assessment of models using a pay ratio rather than ACOL is that these specifications are among the most unstable of all that we examine. Depending on the exact specification, models that exclude variables used to predict civilian compensation yield estimates of the pay elasticity of reenlistment between 0.4 and 1.0.³⁵ These are significantly

34. Typically, this pay ratio is expressed as current (annual) military pay divided by current (annual) civilian pay, or as a relative pay index that measures growth over time. A person's zone A SRB multiplier is traditionally entered as a separate regressor.

35. Using a specification similar to that found in [7] yields an elasticity of 0.4, which is extremely close (0.4 vs. 0.37) to that found in [7].

lower than the estimates we obtain in an ACOL framework. When we include variables used to predict civilian earnings, the estimated pay elasticity is actually *negative*, although relatively small in magnitude. Without any formal theoretical foundation for a military-civilian pay ratio, these specifications are inferior to the more traditional ACOL approach, given these huge differences in parameter estimates.

Treatment of extensions

Models of reenlistment behavior that examine two distinct choices of the enlisted Sailor (e.g., remain in the Navy vs. leave to pursue civilian employment) are the most common in the literature. All of these approaches, however, necessarily collapse three conceptually distinct choices (e.g., reenlistment, unconditional extension, and separation) into two separate decisions. Depending on the focus of previous studies, the literature has combined reenlistments and extensions, combined extensions and separations, or simply excluded extensions altogether. Our baseline model follows most of the literature and combines extensions and separations, focusing on “long-term commitments” to the Navy as our variable of interest.³⁶

Moving from a dichotomous framework to a trichotomous specification is not trivial when using ACOL. Reference [15] shows that, in a trichotomous framework, there is not a “unique theoretically correct time horizon.”³⁷ Rather, the author argues that, if there is relatively little growth in expected military and civilian earnings over time, examination of relative compensation over a 4-year time horizon is a reasonable approximation. We follow this reasoning and use such a time horizon to assess the degree to which our estimates are influenced by our treatment of unconditional extensions. We make use of a multinomial logit framework and analyze extensions, reenlistments, and separations as three distinct choices.

36. In theory, increases in the SRB multiplier will have different effects on the propensities to reenlist and execute an unconditional extension. Because SRBs are available only to those who reenlist, an increase in the SRB multiplier is expected to increase reenlistments and decrease extensions. The efficacy of the SRB program is often a focus, so extensions typically are not combined with reenlistments.

37. See appendix A of [15].

With this specification, we estimate a pay elasticity of reenlistment of 2.9. This is extremely similar to the elasticity that we find when calculating separate effects for military and civilian compensation (see table 5).³⁸ The effect of a one-level increase in SRB multipliers is similar to that observed in [7]. Despite a sizable increase in reenlistments, our estimates suggest that over one-half of this increase is because of a *decrease* in the number of extensions. The multinomial logit estimation, then, indicates an increase in retention of about 2.9 percentage points, which is significantly larger than the estimates in table 5. Therefore, although treating extensions, reenlistments, and separations as distinct choices has little effect on estimates of the pay elasticity of reenlistment, it does have a sizable impact on the effect of increases in SRBs.

Grouped data

Much of the past research on reenlistment used grouped data. Although data on individual Sailors were available, this research relied on earnings data for veterans, which were not available at the individual level.³⁹ At least one researcher suggests that the use of grouped vs. individual data could explain some of the variance in elasticity estimates [1]. To test this hypothesis, we use our individual-level data to create grouped data, and then reestimate our baseline model using the new data set.

The literature provides no clear guide on how estimates will vary when using grouped data instead of individual-level data. Both [17] and [18] demonstrate that the use of grouped data could result in biased coefficients, a loss of efficiency, and/or heteroskedasticity. Reference [19] argues that, in some cases, aggregated data (such as data

38. In the multinomial logit framework, we calculate separate effects for military compensation if one chooses to reenlist, military pay if one chooses to extend, and civilian compensation if one chooses to separate. The pay elasticity of reenlistment is calculated by estimating the percentage increase in the probability of reenlistment associated with a 1-percent increase in *both* basic pay if one chooses to reenlist and basic pay if one chooses to extend.

39. References [4] and [16] are examples of research using grouped data.

grouped at the state level) may be *preferable* to individual data, although this argument holds only under limited circumstances.⁴⁰ As [19] shows, however, it is extremely difficult to predict the sign or magnitude of any bias that results from using grouped data. Reference [16] cites [20], which demonstrates that the bias can be minimized by using groups that are as narrowly defined as possible.⁴¹

Early researchers did not have the luxury of deciding between individual and grouped data. They simply grouped their data in the same manner as the veterans' earnings data were compiled. We follow the spirit of this earlier literature and group our data by three covariates: fiscal year, paygrade, and rating. Each "observation" in the data set, then, represents a different combination of these three variables.

For each observation, we calculate the average of each independent variable we use in the baseline model. For categorical variables, this approach gives us the proportion of each group with a given characteristic.⁴² Finally, we calculate the maximum ACOL value for each observation in the same manner as we did with our individual data, rather than taking the average of each Sailor's ACOL. This strategy reflects a desire to replicate the efforts of earlier research as closely as possible.

Results using grouped data

Given the strong possibility of reverse causation between pay and reenlistment when using grouped data [1], our "baseline model" for grouped data includes fixed effects for each individual rating. Our comparison with results using individual-level data, then, is the model estimated with rating fixed-effects (table 6).

40. Grouped data are preferable to individual-level data when there is substantial measurement error and reason to believe that grouped-level data are significantly more accurate.

41. This result makes intuitive sense. Individual-level data and grouped data are identical if the "group" is narrowly defined to be the "individual."

42. For variables with discrete values, such as SRBs, we calculate the median rather than the mean. This does not affect our estimates, but it allows us to make calculations using values actually observed in the data (e.g., while some can receive an SRB of 1.0 and the median Sailor can receive an SRB of 0, no one receives a mean SRB of 0.6).

Table 10 presents both the pay elasticity and SRB effect that we estimate using grouped data.⁴³ A comparison with table 6 reveals that the pay elasticity of reenlistment is about 30 percent smaller when using grouped data. Similarly, the effect of a one-level increase in the SRB multiplier is about 40 percent smaller than our estimate when using individual-level data. Although both models predict a positive relationship between changes in compensation and changes in reenlistment, the predicted relationship is substantively smaller when using grouped data.

Table 10. Pay effects—grouped data

Effect	Estimate
Pay elasticity of reenlistment	1.8 percent ^a
One-level increase in SRB multiplier	3.1 percentage points ^a

a. Zero lies outside the 99-percent confidence interval for this estimate.

In general, the observed relationships between explanatory variables and reenlistment rates are similar with grouped and individual-level data. While there are some differences in magnitude, the coefficients on all of the strongest predictors of reenlistment are of the same sign when using both types of data.

Are grouped data as sensitive to changes in model specification?

Reference [12] notes that, in the author’s experience, estimates using grouped data are *less* sensitive to the inclusion/exclusion of variables used to predict civilian earnings, although no theoretical explanation is provided. Our results are consistent with the observations of [12]. When excluding race/ethnicity and age, for example, the pay elasticity of reenlistment drops slightly, from 1.8 to 1.7, and the SRB effect drops from 3.1 to 2.9 percentage points. This stands in contrast to the 40-percent declines when using individual-level data. In general, changes in model specification result in smaller changes in estimates

43. We estimate all grouped-data models by generalized least squares (GLS) regression. See [21] for details.

of the pay elasticity of reenlistment. This does not necessarily advocate the use of grouped data over individual-level data, but it does suggest that the choice of specification is more critical when using individual-level data.

Summary of alternative specifications

Our analysis of commonly used specifications, summarized in table 11, suggests that a great deal of variation in existing estimates of pay elasticity arises from the use of different models of reenlistment behavior. The estimates that we obtain from these theoretically defensible specifications span the range of previous estimates summarized in [1].

Table 11. Summary of alternative specifications

Model	Pay elasticity
Baseline	1.5
Exclusion of those ineligible to reenlist	1.5
Exclusion of race/ethnicity and age	0.9
Exclusion of race/ethnicity, age, and fiscal year	1.0
Separate variables for military and civilian pay	2.8
Fixed-effects for each rating	2.7
Occupation-specific elasticities	0.2 - 3.8 ^a
Different discount rates	1.5
Military-civilian pay ratio	0.4
Separating reenlistments, extensions, and separations	2.9
Grouped data	1.8

a. Separate elasticities are calculated for each occupation group; therefore, a range of elasticities is presented.

We use the same data to estimate each alternative model, so these differences in the estimated pay elasticity do not reflect changes in the responsiveness to pay of Navy enlisted personnel. Rather, these differences in estimates represent shifts in the magnitude of responsiveness

that these models *attribute* to pay. For example, as explanatory variables are removed from the baseline model, their effect on reenlistment is picked up by the remaining variables in the model with which they are correlated. If these variables are correlated with pay, at least part of the relationship between this explanatory variable (e.g., race) and reenlistment will be ascribed to pay instead.

Having said this, there are theoretical justifications for each of the specifications that we have examined in this section. Consequently, economic theory does not provide clear guidelines for choosing one model over another. The choice of an appropriate model is a crucial one, however, because the pay elasticity of reenlistment varies significantly from one specification to the next. After looking at how the pay elasticity has changed over time, we turn to the relative *performance* of each model to assess which specification yields accurate predictions of reenlistment behavior.

Has the pay elasticity changed over time?

Reference [1] documents the sizable variation in the estimates of the pay elasticity of reenlistment. In general, estimates presented in recent studies have been notably smaller than those found in previous research. It is not clear whether these smaller estimates are the result of differences in methodology or actual changes in the responsiveness of enlisted personnel to pay. The previous section examined the degree to which changes in specification of the reenlistment model could explain changes in the magnitude of the pay elasticity. We now examine changes in the responsiveness to pay over time.

Estimating pay elasticities separately by fiscal year

To assess whether the pay elasticity has changed substantively over our sample period, we modify our baseline model to include separate ACOL variables for each fiscal year. Table 12 presents these estimates of both the pay elasticity and the effect of a one-level increase in the SRB multiplier, calculated separately for each fiscal year. This approach yields pay elasticities ranging from 1.3 (FY99) to 2.0 (FY97), and SRB effects ranging from 2.1 (FY92) to 3.2 percentage points (FY96 and FY97). Inspection of table 12 suggests that some of the largest pay elasticities are observed during the drawdown years and that the largest SRB effects are found at the end of the drawdown.

Statistical tests of the underlying coefficients allow us to reject (with 99-percent confidence) the hypothesis of no differences in the effect of pay across fiscal years.⁴⁴ In other words, we can be certain that

44. It is possible that the variation in estimates of the pay elasticity in table 12 results from changes in the “point of evaluation” because reenlistment rates vary by fiscal year [1]. However, the statistical test of significant differences by fiscal year is based on the estimated coefficients in the model and not the pay elasticities themselves. These significant differences in coefficient estimates imply that the differences in estimates of the pay elasticity are truly differences in the responsiveness to pay.

there are statistically significant differences in the responsiveness of Sailors to changes in military compensation over the FY87–99 time period. When comparing each fiscal year to the previous year, however, there appear to only be two distinct changes in the responsiveness to pay. The pay elasticity increases by almost 30 percent from FY92 to FY93, and then decreases by about 30 percent from FY97 to FY98. For all other years, the estimated pay elasticity is statistically indistinguishable from the previous year’s estimate.

Table 12. Pay effects—calculated separately by fiscal year^a

Fiscal year	Pay elasticity ^b	SRB effect ^b
1987	1.5	2.4
1988	1.6	2.3
1989	1.5	2.6
1990	1.5	2.7
1991	1.4	2.4
1992	1.4	2.1
1993	1.8	2.6
1994	1.9	2.6
1995	1.9	3.0
1996	1.9	3.2
1997	2.0	3.2
1998	1.4	2.4
1999	1.3	2.3

a. Note that the two data columns are highly correlated, suggesting a close relationship between changes in the pay elasticity of reenlistment and changes in the effect of increases in SRB multipliers.

b. Zero lies outside the 99-percent confidence interval for all these estimates.

Although the pay elasticity clearly declined from FY97 to FY98, the data also suggest a slightly lower pay elasticity at the end of our sample period than at the beginning. Comparing our estimates in the years before the drawdown (FY87–92) with those at the end of our sample (FY98–99) reveals statistically significant (although perhaps not

economically significant) declines.⁴⁵ There is some evidence, then, that the pay elasticity has declined, both in recent years and over the FY87–99 period.

The variation in pay elasticity estimates over this time period, however, is much smaller than the variation we observe when making use of different specifications of the reenlistment model. Furthermore, most of this variation occurred around the drawdown period, a time that was not considered in most of the previous literature on compensation and reenlistment behavior. It is likely, therefore, that differences in estimates found in the previous literature are being driven by differences in methodology rather than just differences in the time period being studied.

Did Sailors respond differently to pay during the drawdown?

Our estimates of the pay elasticity of reenlistment suggest that the responsiveness of Sailors to pay was substantively different during the drawdown. In fact, these higher elasticities imply that enlisted Sailors who made reenlistment decisions during the drawdown were *more* responsive to pay than either their immediate predecessors or their successors.

It is possible, however, that our baseline model is not well specified for decisions made during the drawdown; if so, estimates of the pay elasticity of reenlistment may not be reliable for this period. The drawdown was, by definition, a time when the Navy was deliberately trying to *reduce* the size of the enlisted force. The sharp increase in the number of Sailors classified as “ineligible to reenlist” and the decline in the proportion of Sailors who held a rating at their first reenlistment point reflect the desire of the Navy to actively manage the size and composition of the enlisted force [2]. If our model does not accurately control for these demand constraints, it is possible that these differences are being captured by our ACOL variable.

45. We reestimate our model to produce a single pay elasticity for FY87–92 (1.4) and for FY98–99 (1.3).

Similarly, our predictions of expected future military compensation are based on the contemporaneous (at the time of the reenlistment decision) distribution of individuals across paygrades, within each rating, for each year of service. Although this is a reasonable assumption for a steady-state environment, the expectations of individual Sailors during the drawdown were probably quite different from what our model would predict. This misspecification of expected military compensation will undoubtedly influence estimates of the pay elasticity of reenlistment for the drawdown period. Without a rigorous way to improve on our predictions, however, it is difficult to assess the degree to which this shortcoming of our model affects our estimates.

Choosing a model of reenlistment behavior

The previous sections document the sizable variation in estimates of the pay elasticity of reenlistment that arise under a variety of empirical specifications. Given that all of these specifications are defensible from different theoretical points of view, questions as to which model is “preferred” and, similarly, which is the “true” pay elasticity of reenlistment are natural ones.

The econometric properties of each of the models that we examine imply that each estimate of the pay elasticity “best” describes the data being used. Therefore, we choose to turn to different criteria to assess which specification is preferred as a model of reenlistment behavior of Navy enlisted personnel. We use two different criteria—statistical measures of “goodness of fit” and a direct comparison of actual and predicted reenlistment rates. Although both of these criteria offer similar guidance about the preferred econometric specification, we argue that a comparison of predicted and actual reenlistment behavior is a more appropriate measure for the Navy.

Goodness of fit

Goodness of fit refers to the ability of an econometric model to explain variation in the variable of interest (in our case, variation in reenlistment behavior). In other words, a “good” model will include regressors (e.g., age, unemployment rates, and relative military compensation) that are strong predictors of the variable of interest. Several measures of goodness of fit indicate how much of this variation can be explained by the independent variables within the model.

First, the “R-squared” statistic measures the proportion of the variation that is explained by the model; the higher the R-squared, the greater the explanatory power of the model. Unfortunately, a simple comparison of this metric for each specification does not always lead a researcher to the “best” model. Models estimated with grouped

data, for example, will generally have a higher R-squared than models using individual-level data. This does not necessarily imply that a model with grouped data is superior; rather, more of the variation in reenlistment rates can be explained with grouped data simply because these data have dramatically less variation than individual-level data. Consequently, selecting a preferred model involves more than comparing a simple measure of explained variation.

Furthermore, the R-squared measure for nonlinear models always increases as the number of explanatory variables increases. The implication is that, if we focus solely on this metric, the preferred model is one that includes the most independent variables in the analysis. Of course, the inclusion of explanatory variables that are strong predictors of reenlistment behavior is desired; including variables with weak explanatory power, however, decreases the precision of the other parameters of interest (e.g., the pay elasticity).

Because of the drawbacks associated with this metric, we turn to an alternative measure of goodness of fit. The “weighted sum of squared residuals” (SSR) has several advantages over the R-squared measure. It is not directly tied to the number of explanatory variables in the model. Also it considers the total amount of variation, so it does not necessarily favor models that use data with little variation. The SSR measures the difference between actual and predicted behavior for each observation in the data; more accurate models will have a smaller differential and, therefore, a lower SSR.⁴⁶

Use of this metric indicates that, of all the specifications we examine, our model including fixed effects for each rating has the best “fit” to the data with which it is estimated. At best, however, this is a tenuous recommendation of the fixed-effects model. A comparison of this “best” model with this metric’s “worst” model (excluding variables used to predict civilian pay) reveals only a 0.3-percent difference in goodness of fit.

46. See [21] for a complete discussion of this metric and of alternate measures of goodness of fit. A weighted SSR measure has properties similar to those of the “adjusted R-squared” measure (a measure often used to judge the fit of linear regression models).

In summary, although measures of goodness of fit can suggest a preferred econometric specification, we caution against relying solely on any goodness-of-fit measure. In fact, [17] suggests that there may be a trade-off between maximizing the amount of variation explained by a model and attaining “good” parameter estimates. Because our purpose is to select a model that accurately predicts the reenlistment behavior of Navy enlisted personnel and their response to specific policies (e.g., an increase in reenlistment bonuses), we prefer the model(s) with “good” parameter estimates to one that happens to have a good “fit” for the data that is used in the analysis.

Predictions of reenlistment behavior

Our preferred measure of a model’s predictive power is a comparison of the percentage of Sailors that choose to reenlist with the reenlistment rate predicted by the model. Ultimately, a model of reenlistment behavior is used by the Navy to forecast reenlistment rates, given characteristics of Navy enlisted personnel and their economic environment. The success or failure of a given model, from the Navy’s perspective, will be determined by its ability to accurately predict reenlistment behavior.

Simply comparing the predicted average reenlistment rate with the actual reenlistment rate found in the data used to estimate the model yields no clues about the relative performance of a model. With logistic regression, this predicted reenlistment rate for the sample will *always* match the actual reenlistment rate. Therefore, alternative specifications that make use of the same data all yield the same predictions of the average reenlistment rate. This does not mean, however, that all logistic models are equally good.

Instead, we use our parameter estimates from each model to predict reenlistment rates on *different* data. There is no guarantee that the predicted and actual reenlistment rates will be equal. In fact, the difference between these two rates provides us with a measure of the degree to which a model accurately predicts reenlistment behavior.

We have three distinct sources of data that we can use to test the predictive power of the models that we have estimated. First, we can

make use of a *subset* of the data used in the preceding analysis to predict reenlistment rates for this group. Though predicted and actual reenlistment rates are equal *on average* in the data used in the original analysis, the econometric properties of logistic regression do not prevent overprediction or underprediction of reenlistment for subsets of the data. Furthermore, this type of regression places fewer restrictions on the *degree* to which these predictions deviate from the sample average. For this analysis, we focus on a few specialties considered to be “critical ratings” to assess the relative performance of each model at predicting reenlistment for Sailors in these occupations.

Second, we originally split our sample of enlisted Sailors randomly into two parts and used one-half to generate our estimates of the models presented in previous sections.⁴⁷ We now use the other half of this sample (referred to as the “test data”) to test the predictive power of our baseline model and some of the alternative specifications. In doing so, we can examine whether our models “fit” well (rather than being fitted to specific statistical anomalies in the data) and can test our models using data from the same time period over which the models were originally estimated. With these data, we can compare both predicted and actual reenlistment rates for the sample as a whole, as well as for the ratings examined in the original data.

Finally, we use data from FY00 to test our empirical models and examine the predictive power of various specifications for reenlistment decisions made *outside* the time period over which the models were estimated. Again, we can evaluate our alternate models both on their ability to predict overall reenlistment rates and on the relative accuracy of their predictions for a few, critical ratings. All three approaches give us a sense of the relative predictive power of these different models of reenlistment behavior.

47. We used a random-number generator to split the sample in half. Not surprisingly, the descriptive statistics of each subsample are virtually identical.

Within-sample predictions of reenlistment behavior

Reference [7] examines the relationship between manning shortfalls in several Navy enlisted ratings and the relative earnings of Sailors in these occupations. The author finds that, while highly technical ratings have some of the highest levels of military compensation, occupational differentials are substantially smaller in the military than in the civilian sector. Consequently, these ratings with the highest civilian earnings opportunities also have some of the most severe manning problems in the Navy. Reference [7] identifies three ratings—Aviation Electronics Technician (AT), Electronics Technician (ET), and Fire Control Technician (FC)—as highly technical ratings with significant manning shortages.

Given the Navy’s increasing reliance on a more skilled workforce [22], we begin to assess the relative performance of each of our reenlistment models by comparing actual and predicted reenlistment rates for Sailors in these critical ratings. To quantify the relative efficacy of these models, we measure the percentage difference in predicted and actual reenlistment rates: a positive value indicates that the model predicts higher reenlistment than is actually observed (an “overprediction”), whereas negative values imply that the model “underpredicts” reenlistment. To distinguish between “historical” and “contemporaneous” goodness of fit, we examine predictions both for the average reenlistment rate over the FY87–99 time period, as well as for the most recent data used in our analysis (FY99).

As table 13 suggests, most of our models *underpredict* reenlistment for these ratings over the FY87–99 time period. For example, our baseline model predicts a reenlistment rate that is 4.1 percent lower than what we actually observe. These differences between actual and predicted reenlistment rates are statistically significant; even given the range of estimates associated with the prediction, we can be confident that the baseline model underpredicts reenlistment for these ratings. Similarly, all three models concerned with the variables used to predict civilian earnings underpredict reenlistment. In contrast, the grouped-data model and the model that distinguishes extensions, reenlistments, and separations both slightly overpredict reenlistment in the AT, ET, and FC ratings; in fact, the actual reenlistment rate lies well within the range implied by each of these predictions.

Table 13. Within-sample predictions of reenlistment behavior—
AT, ET, and FC ratings

Model	Level of overprediction (percentage)	
	FY87–99	FY99 only
Baseline	-4.1 ^a	-0.7
Exclusion of race/ethnicity and age	-4.9 ^a	-5.5 ^a
Exclusion of race/ethnicity, age, and fiscal year	-4.8 ^a	-9.4 ^a
Separate variables for military and civilian pay	-3.7	4.1 ^a
Fixed-effects for each rating	0	13.0 ^a
Separating reenlistments, extensions, and separations ^b	0.6	21.5 ^a
Separate elasticities by year	-4.2 ^a	-1.5
Grouped data	0.8	5.6

a. The actual reenlistment rate lies outside the 95-percent confidence interval of the predicted reenlistment rate.

b. For comparability, the percentage difference in predicted and actual reenlistment rates is presented.

Although the fixed-effects model exactly predicts reenlistment in these ratings over the FY87–99 time period, this does *not necessarily* imply that the model is superior. Using fixed-effects for each rating *guarantees* that predicted and actual reenlistment rates by rating will be identical. When we focus on predictions for FY99 only, table 13 suggests that the fixed-effects model does very poorly at predicting reenlistment in these ratings, with an overprediction of 13 percent. The results for FY99 imply that the baseline model does the best at predicting reenlistment, with only a negligible difference between predicted and actual reenlistment rates. The model that calculates separate effects of military and civilian compensation appears to be the least stable; depending on the period for which reenlistment is predicted, the model can underpredict or overpredict reenlistment in these ratings.⁴⁸

48. Despite the sizable overpredictions of the grouped data model, we cannot conclude that predicted reenlistment rates are significantly different from actual reenlistment rates. This is because of the relative imprecision of estimates using grouped data and should not be interpreted as a reason to use grouped-data models.

As an alternate subset of the data, we also focus on ratings that experienced changes from FY99 to FY00 on their SRB multipliers. The changes in the reenlistment bonus within these ratings suggest that the Navy was concerned about reenlistment in these ratings, and it is likely that the Navy had projections of the amount of reenlistment that would result from these changes in SRBs. Focusing on these ratings allows us to assess how well each reenlistment model does at predicting reenlistment in these ratings.

Table 14 presents the level of overprediction of reenlistment in these ratings for each of our models, both for the entire sample period and for FY99 in particular. Over the FY87–99 period, every model except that which excludes race/ethnicity, age, and fiscal year predicts reenlistment rates that are statistically indistinguishable from what we actually observe. When looking at FY99, only this same model predicts reenlistment that is significantly different from actual reenlistment. Though all other models predict rates that are lower than what is actually observed, they are not significantly different from actual reenlistment rates.

Table 14. Within-sample predictions of reenlistment behavior—
ratings with change in SRB multiplier from FY99 to FY00

Model	Level of overprediction (percentage)	
	FY87–99	FY99 only
Baseline	2.6	-1.6
Exclusion of race/ethnicity and age	4.0	-1.1
Exclusion of race/ethnicity, age, and fiscal year	4.0 ^a	-6.7 ^a
Separate variables for military and civilian pay	1.3	-2.2
Fixed-effects for each rating	0	-4.1
Separating reenlistments, extensions, and separations ^b	-2.0	-3.0
Separate elasticities by year	2.6	-0.7
Grouped data	3.6	-1.3

a. The actual reenlistment rate lies outside the 95-percent confidence interval of the predicted reenlistment rate.

b. For comparability, the percentage difference in predicted and actual reenlistment rates is presented.

In conclusion, tables 13 and 14 suggest that the relative performance of each model depends heavily on the particular subset of the data being examined. Models that are explicitly designed to focus on a subset of the data (e.g., fixed-effects models) do the best job at predicting reenlistment for that particular subset; however, these models also appear to do the worst job at predicting reenlistment for even a slightly different subset of the data. In general, the baseline model performs fairly well at predicting reenlistment rates for different groups of ratings.

Predictions of reenlistment behavior using test data

Although our original data set is extremely large, one potential area of concern in choosing between models is that a model may fit statistical anomalies in the data very well but not have good predictive power over data from a different sample. This is less likely to be a problem with large data sets than with smaller ones, but guarding against this possibility is one of our motivations for estimating our original models with only one-half of the sample.

For each model, then, we use the coefficients estimated using the original data set to predict the probability of reenlistment for Sailors in the test data. Then, we compare the predicted reenlistment rate with the actual reenlistment rate observed in the test data.

Table 15 suggests that each model does a very good job at predicting reenlistment in our test data; for most of these models, the predicted reenlistment rate is within 1 percent of the actual reenlistment rate.⁴⁹ Furthermore, with each model, the predicted reenlistment rate is not statistically different from the actual reenlistment rate. This suggests that our econometric models do a good job, on average, of predicting a general relationship between our explanatory variables and reenlistment behavior over the FY87–99 time period.

49. Note that the average reenlistment rate in the test data is about 1 percent higher (although not statistically different) than in the original sample. Given these slight differences, it is not surprising that our models predict reenlistment that is 1 percent lower in the test data.

Table 15. Predictions of reenlistment behavior—test data

Model	Level of overprediction (percentage)
Baseline	-0.6
Exclusion of race/ethnicity and age	-0.7
Exclusion of race/ethnicity, age, and fiscal year	-0.7
Separate variables for military and civilian pay	-0.6
Fixed-effects for each rating	-1.3
Separating reenlistments, extensions, and separations ^a	-0.6
Separate elasticities by year	-0.6
Grouped data	1.3

a. For comparability, the percentage difference in predicted and actual reenlistment rates is presented.

Using the test data, we also predict reenlistment behavior for the subsets of the data presented in tables 13 and 14. These results (not shown) are very similar to those we obtain with the original sample, and confirm that our models are not being driven by any statistical anomalies present in the original data.

Finally, we also reestimate each model using the test data. This approach generates a set of coefficients that can be compared to those estimated with our original data. For the most part, there are only minor differences in parameter estimates, and these differences are well within the confidence intervals implied by each estimate. In other words, we cannot reject the hypothesis that each data set generates identical parameter estimates.

Out-of-sample predictions of reenlistment behavior

A primary focus of our research is to allow the Navy to forecast reenlistment rates with as much accuracy as possible. This is necessary both to achieve endstrength and to adequately budget for the resources needed to achieve reenlistment targets. For this reason, we believe that the most meaningful measure of how models perform stems from their ability to make out-of-sample predictions. There-

fore, we examine predictions of future reenlistment using estimates of prior relationships between Sailor characteristics and reenlistment decisions. In principle, this is how reenlistment models are actually used. In this section, we use data from FY00 to test each model, and compare actual reenlistment rates with those predicted by each specification.⁵⁰

Data

The demographic characteristics of Sailors, the UICs at which they are stationed, and the economic environment faced by these Sailors all changed significantly between FY87 and FY00. Table 16 lists descriptive statistics that reflect both changes in relative military compensation and changes in reenlistment behavior over this time period. To examine differences between our original data and our out-of-sample data, we list data separately for the FY87–FY99, FY99, and FY00 time periods.

Table 16. Descriptive statistics

Variable	FY87–99	FY99	FY00
Reenlistment rate	0.32	0.34	0.43
Unconditional extension ^a	0.08	0.07	0.05
Economic data			
ACOL (\$1,000) ^b	3.68	4.33	4.71
SRB multiplier ^c	0	1	1.5
Eligible for SRB > 0 ^a	0.48	0.61	0.73
Eligible for SRB > 3 ^a	0.08	0.15	0.21

a. Proportion with this characteristic is presented.

b. The average ACOL value is presented.

c. The median SRB multiplier is presented.

Table 16 indicates a substantial increase in the reenlistment rate over this time period; in FY00, 43 percent of all Sailors chose to reenlist, compared with an average of 32 percent over the FY87–99 period.

50. We choose our sample for FY00 using the same restrictions as we did for our original data.

This change is more dramatic when compared with the FY99 reenlistment rate of 34 percent. Although this increase in the reenlistment rate is partially offset by a decrease in the proportion that execute unconditional extensions, the proportion of Sailors who choose not to separate from the Navy (reenlistments plus extensions) rises from FY99 to FY00.

These changes are not surprising when examining the SRBs offered to individual Sailors in FY99 and FY00. By FY00, almost 75 percent of all Sailors were eligible for an SRB, a dramatic increase from FY99. Similarly, the proportion eligible for a relatively large SRB (a multiplier greater than 3) had risen to about 21 percent.

There were also changes in the demographic composition of Sailors between FY87 and FY00 (not shown). By FY00, both the marital and ethnic composition of the junior enlisted force changed in step with overall societal trends. The geographic distribution of Sailors also changed dramatically over this time period, probably as a direct result of the drawdown and base closures in the 1990s.

In summary, our data for FY00 are significantly different from the FY87–99 data. Rather than a cause of concern, however, these differences present an *ideal* test of our reenlistment models. Models of reenlistment behavior measure the relationship between changes in several explanatory variables and changes in reenlistment. With stark changes over time in the explanatory variables that we use in our analyses, the data provide a powerful test of our models to assess whether the associated changes in reenlistment behavior that we observe are consistent with the changes predicted by our reenlistment model.

Predictions of reenlistment behavior

Table 17 contains our out-of-sample predictions of reenlistment behavior in FY00 for each of the empirical specifications that we examine. As table 17 shows, most of the models do a fairly good job at predicting reenlistment. The baseline model, the model that excludes race/ethnicity and age, the model that includes fixed-effects for each rating, and the one that calculates separate pay elasticities by year all predict reenlistment rates that are not statistically different from what is actually observed.

Table 17. Out-of-sample predictions of reenlistment behavior—FY00

Model	Level of overprediction (percentage)
Baseline	2.1
Exclusion of race/ethnicity and age	0.7
Exclusion of race/ethnicity, age, and fiscal year	-4.3 ^a
Separate variables for military and civilian pay	13.6 ^a
Fixed-effects for each rating	2.6
Separating reenlistments, extensions, and separations ^b	22.6 ^a
Separate elasticities by year	1.8
Grouped data	6.2

a. The actual reenlistment rate lies outside the 95-percent confidence interval of the predicted reenlistment rate.

b. For comparability, the percentage difference in predicted and actual reenlistment rates is presented.

None of these models predict FY00 reenlistment as accurately as they predict reenlistment in the test data. This is not surprising, however, for two reasons. First, in-sample predictions will always be more accurate than out-of-sample predictions because the models are estimated using in-sample data. Second, most of our models include fiscal year variables intended to reflect any differences from one fiscal year to the next not captured by other explanatory variables. When using these coefficients to make predictions for FY00, we are assuming that these unobserved differences are identical in FY99 and FY00. Our predictions for FY00, then, will not be accurate if there are changes in reenlistment behavior due to factors other than what we included in our model.

As an additional test, we also predict reenlistment rates for the subsets of ratings identified in tables 13 and 14. As table 18 shows, the baseline model does the best job at predicting reenlistment rates in these critical ratings. As before, our models tend to produce more accurate predictions for the entire sample than for particular subsets of ratings. One likely reason for this is the relative imprecision of our civil-

ian wage estimates, which are distinguished by occupation only by separating “technical” and “nontechnical” ratings. It is likely that this distinction understates civilian earnings opportunities in the highly technical ratings, which reduces the precision of our predictions.⁵¹

Table 18. Out-of-sample predictions of reenlistment behavior—subsets of ratings

Model	Level of overprediction (percentage)	
	ATs, ETs, and FCs	Ratings with new SRBs in FY00
Baseline	-0.5	-2.5
Exclusion of race/ethnicity and age	-4.9 ^a	-4.3 ^a
Exclusion of race/ethnicity, age, and fiscal year	-8.4 ^a	-8.8 ^a
Separate variables for military and civilian pay	17.2 ^a	9.0 ^a
Fixed-effects for each rating	8.9 ^a	-1.0
Separating reenlistments, extensions, and separations ^b	37.2 ^a	21.0 ^a
Separate elasticities by year	-1.3	-2.5
Grouped data	4.1	3.4

a. The actual reenlistment rate lies outside the 95-percent confidence interval of the predicted reenlistment rate.

b. For comparability, the percentage difference in predicted and actual reenlistment rates is presented.

In general, then, the baseline model, with a pay elasticity of 1.5, appears to be the most robust specification when predicting reenlistment. For most of our tests, both in- and out-of-sample, reenlistment

51. Reference [7] matches a number of technical or critical ratings to specific civilian occupations and finds that average earnings in these jobs are over \$40,000 per year for civilians of similar age and education as our first-term Sailors. This figure is significantly larger than predicted civilian wages for Navy enlisted personnel in technical ratings. Therefore, we suggest caution when using these models to predict reenlistment behavior of Sailors in these highly technical ratings.

rates predicted using the baseline model are not statistically different from actual reenlistment rates. Although other models perform well in some tests, they perform poorly in others. We conclude, therefore, that the baseline model provides the best “fit” of the data on Navy enlisted personnel.

A framework for updating reenlistment models

Given a preferred model of reenlistment behavior, a remaining issue is when to rely on historical predictions of key parameters of interest versus when to reestimate the model with more recent data. In other words, it is natural to question when models of reenlistment behavior should be updated to more accurately capture the responsiveness of today's enlisted personnel to changes in their environment. As the previous section suggests, because these models are principally used to forecast reenlistment behavior, updating them when their forecasts differ from observed behavior would seem to be a reasonable approach.

Our analysis suggests a few principles that can serve as guidance in the use and reestimation of a model of reenlistment behavior. First, while these models are designed to forecast behavior as accurately as possible, any forecast has a *range* of estimates associated with it. Although a given forecast can appear different from what is actually observed, actual reenlistment rates often fall within the range of estimates implied by the forecast. In other words, for relatively small differences, one cannot be certain whether this difference reflects changes in the responsiveness of Sailors to changes in their environment or the uncertainty inherent in any econometric model. This suggests that, when reenlistment rates consistently fall within the range associated with a model's forecast, there is not necessarily a need to reevaluate or reestimate a model of reenlistment behavior.

Second, just as accurate predictions of the overall reenlistment rate are a priority for the Navy, so too are accurate predictions for critical skills or those with chronic manning shortfalls. If a model continues to perform well, on average, but systematically overpredicts or underpredicts reenlistment for these key ratings, reestimating the model with more recent data, or even reevaluating the choice of model used to predict this behavior, is in order.

Third, our examination of changes in the pay elasticity over time suggests that there has been relatively little change in this key parameter. The most substantive changes that we do observe, however, occurred during the drawdown period, when the environment faced by Navy enlisted personnel changed in a fundamental way that no model could have forecast. This suggests that models of reenlistment behavior may need to be reevaluated in response to profound shifts in policy, societal changes, or events that directly affect military personnel. It is difficult to predict, for example, *how* reenlistment behavior might change after the terrorist attacks in 2001, but understanding that behavior *might* change, and incorporating this information into a model of reenlistment behavior, will help to improve subsequent forecasts.

Absent any of these concerns, there is nothing “wrong” with adding more recent data to more accurately estimate these econometric models. Additional data can only improve the precision of these models, and recent data are more likely than historical data to reflect the responsiveness of today’s Sailor to changes in his or her environment. As data for future fiscal years become available, incorporating them into a model of reenlistment behavior will only improve the forecasts of the model for subsequent planning purposes.

Another possible adjustment to models of reenlistment behavior is to exclude earlier data. For example, our results suggest that behavior of enlisted personnel differed during the drawdown. This suggests that models excluding data from that time period might produce more accurate estimates than models that include decisions made during the drawdown. To test this hypothesis, we predict FY00 reenlistment rates using coefficients obtained from models that exclude the drawdown period.⁵² These results, shown in table 19, indicate that the models that exclude the drawdown do a slightly better job at predicting FY00 reenlistment than the baseline model (using all data

52. Although there is no single definition of the drawdown, we choose two alternate definitions. The first is the FY92–95 period, during which VSI/SSB were available to more senior enlisted personnel. The second is the FY94–97 period, in which we observe substantively different responsiveness to pay.

from FY87 to FY99). However, these differences are extremely small and *not* statistically significant.⁵³ In other words, excluding the draw-down slightly improves forecasts of reenlistment behavior, but not substantively.

Table 19. Predictions of FY00 reenlistment rates—baseline model and exclusion of various subperiods

Model	Level of overprediction (percentage)
Baseline	2.1
Excluding drawdown	
FY92–95	1.4
FY94–97	0.3
Excluding FY87–95	6.0
Excluding FY87–97	6.3
Excluding FY87–98	5.2

Although excluding the drawdown results in marginal improvements in predicted reenlistment rates, table 19 also cautions against excluding *all* earlier data. Regardless of the exact definition of the draw-down, using only post-drawdown data tends to *overpredict* reenlistment behavior; these overpredictions are up to three times as large as the overpredictions of the baseline model. Consistent with our evidence that the pay elasticity has changed very little over time, table 19 suggests that using data over a longer time period provides a better estimate of reenlistment rates. This may be because using data from a longer time period provides more variation in the variables that help explain reenlistment behavior.

53. None of the predictions in table 19 are statistically different from actual reenlistment rates. The difference, however, between a 0.3-percent overprediction and a 6.3-percent overprediction is not a trivial one. Without a statistical distinction on which to rely, we focus instead on the differences in the point estimates.

These results for our particular sample do *not* suggest, however, that historical data will *always* improve forecasts of current reenlistment behavior. At some point, it is possible that historical relationships between a Sailor's characteristics or environment and reenlistment behavior will change as the population of enlisted Sailors changes. This is particularly likely as women, a growing proportion of the enlisted force, make proportionately more of the reenlistment decisions.

In other words, it is possible that forecasts of reenlistment at some point in the future will be improved by focusing on more recent data and on populations not considered in this analysis. A periodic examination of existing models of reenlistment behavior to assess their performance with and without historical data, then, is a useful test to ensure the reliability of a model's forecasts.

Conclusion

Our results suggest that estimates of the pay elasticity of reenlistment are highly sensitive to the choice of empirical specification. Our baseline model generates a pay elasticity estimate of 1.5 and a one-level increase in the SRB multiplier is predicted to increase reenlistment by 2.5 percentage points. Both of these estimates lie well within the range of previous estimates.

Alternative specifications, however, yield pay elasticities ranging from 0.4 to 2.9, a statistically significant range that spans the variation found in the literature. We observe similar variation in the relationship between reenlistment bonuses and reenlistment behavior. Because we use the same data to estimate each alternative model, these differences in the pay elasticity do not reflect changes in the responsiveness to pay of Navy enlisted personnel, but rather shifts in the magnitude of responsiveness that these models *attribute* to pay.

In contrast, there is very little variation in the pay elasticity over time, with the only significant changes occurring during the drawdown. We conclude, then, that most of the variation in estimates found in the previous literature results from differences in the empirical approach of researchers, rather than from differences in the reenlistment behavior of enlisted personnel.

Though all of these specifications are defensible from different theoretical points of view, their ability to generate accurate predictions of reenlistment behavior sheds some light on which specification should be preferred. When looking at in-sample predictions of reenlistment, models designed to predict reenlistment behavior for particular subsets of the data generate the most accurate predictions for these subsets. However, these models also do the worst job at predicting reenlistment for even a slightly different subset of the data. In general, the baseline model performs fairly well at predicting reenlistment rates for different groups of ratings.

We also use each of these specifications to predict reenlistment rates for FY00 and compare these predictions to actual reenlistment in this year. In general, the baseline model continues to be the most robust specification when predicting reenlistment. For most of our tests, reenlistment rates predicted using the baseline model are not statistically different from actual reenlistment rates. Other models perform well in some tests but poorly in others. We conclude, therefore, that the baseline model, with a pay elasticity of 1.5, provides the best “fit” of the data on Navy enlisted personnel.

Appendix A: UICs and geographical groups

To explore the influence of a person’s surroundings on the reenlistment decision, we first matched each Sailor’s UIC to the state in which the UIC is located. Next, we subdivided states with a large number of UICs (and a large number of people in each UIC) to assess whether there was significant variation within large states. In our regression analyses, we include dummy variables that indicate whether a person is stationed at one of the seven largest geographic regions at the time of the reenlistment decision. About 79 percent of all Sailors in our data are stationed in one of these geographic regions; the remainder are stationed in areas smaller than the smallest region that we differentiate (Hawaii). Table 20 lists the UICs associated with each of these large geographic regions.

Table 20. Major geographic regions and their locale codes

Geographic region	Locale codes
Northern (coastal) California	KLS, KMI, KPS, LAL, LCE, LCP, LFE, LFO, LMI, LMN, LMO, LOA, LSC, LSI, LSJ, LSU
Southern (coastal) California	KAN, KCC, KCV, KFB, KIM, KLB, KLO, KSB, KSC, KSD, KSN, KSR, KUE, LCL, LTF, SNI
Northern (coastal) Florida	GCF, GEG, GJK, GML, GMY, GPE, GPS
Hawaii	BPI, QKI, QTH
South Carolina	GBE, GCE, GCR, GGS, GPC
Coastal Virginia	FCF, FDA, FDN, FFT, FHM, FLI, FNN, FNO, FOA, FPO, FWI, FWO, FYK
Washington State	MAN, MAS, MBR, MEV, MII, MKE, MLS, MMV, MOA, MOS, MSE, MSP, MTA, MWI

Appendix B: Navy enlisted ratings and occupational groups

This appendix lists the Navy enlisted ratings found in each occupational group used in our analysis. This classification is similar to that used by [8], with a few exceptions. First, because [8] uses data from the 1970s, some of today's ratings did not exist or were excluded from the data. We have placed these ratings in the most appropriate groups created by [8]; for example, the Gas Turbine Systems ratings have been grouped with other Marine Engineering ratings. Second, the duties associated with some ratings have changed significantly because of the increased use of technology in today's Navy. Consequently, some ratings are more similar today than they were in the 1970s; for example, we have placed the ET and FC ratings in the same occupational group.

1. SEABEE Construction

- Builder (BU), Construction Electrician (CE), Construction Mechanic (CM), Engineering Aid (EA), Steelworker (SW), Utilitiesman (UT)

2. Non-SEABEE Construction

- Constructionman (CN), Equipment Operator (EO)

3. Marine Engineering

- Boiler Technician (BT), Electrician's Mate (EM), Engineman (EN), Gas Turbine Systems Technician - Electrical, Mechanical (GSE, GSM), Interior Communications Electrician (IC), Machinist's Mate (MM)

4. Ship Maintenance

- Damage Controlman (DC), Hull Maintenance Technician (HT), Instrumentman (IM), Molder (ML), Machinery Repairman (MR), Opticalman (OM), Pattern Maker (PM)

5. Aviation Maintenance

- Air Traffic Controller (AC), Aviation Machinist's Mate (AD), Aviation Electrician's Mate (AE), Aviation Structural Mechanic - Safety Equipment, Hydraulics, Structures (AME, AMH, AMS), Aviation Ordnanceman (AO), Aviation Fire Control Technician (AQ), Aviation Electronics Technician (AT)

6. Aviation Ground Support

- Aviation Boatswain's Mate - Launching and Recovery Equipment, Fuels, Aircraft Handling (ABE, ABF, ABH), Aerographer's Mate (AG), Aviation Support Equipment Technician -- Electrical, Hydraulics and Structures, Mechanical (AS, ASE, ASH, ASM), Aviation Antisubmarine Warfare Operator (AW), Aviation Maintenance Administrationman (AZ), Parachute Rigger (PR)

7. Media

- Draftsman Illustrator (DM), Journalist (JO), Lithographer (LI), Photographer's Mate (PH)

8. Logistics

- Aviation Storekeeper (AK), Disbursing Clerk (DK), Mess Management Specialist (MS), Ship's Serviceman (SH), Storekeeper (SK)

9. Administration

- Legalman (LN), Master-at-Arms (MA), Postal Clerk (PC), Personnelman (PN), Religious Program Specialist (PR), Yeoman (YN)

10. Data Systems

- Data Processing Technician (DP), Data Systems Technician (DS), Radioman (RM)

11. General Seamanship

- Boatswain's Mate (BM), Operations Specialist (OS), Quartermaster (QM), Signalman (SM)

12. Health Care

- Dental Technician (DT), Hospital Corpsman (HM)

13. Cryptology

- Cryptologic Technician - Administration, Interpreter / Linguist, Maintenance, Communications, Collection, Technical (CTA, CTI, CTM, CTO, CTR, CTT), Intelligence Specialist (IS)

14. Ordnance Systems

- Fire Control Technician -- Ballistic Missile, Gun Fire Control (FT, FTB, FTG), Gunner's Mate -- Guns, Missiles, Technician (GM, GMG, GMM, GMT), Mineman (MN), Missile Technician (MT), Sonar Technician - Submarine (STS), Torpedoman (TM), Weapons Technician (WT)

15. Communications / Sensor

- Electronics Warfare Technician (EW), Ocean Systems Technician - Analyst, Maintenance (OTA, OTM), Sonar Technician - Surface (STG)

16. Weapons Systems / Control

- Electronics Technician (ET), Fire Control Technician (FC)

Appendix C: Coefficients from baseline model

Table 21 lists the coefficients and standard errors for each variable in our baseline model, estimated using a standard logistic regression. The fourth column presents the probability that the sample coefficient is equal to zero, and is used to determine the statistical significance of each estimate. For example, a probability less than 0.01 means that zero lies outside the 99-percent confidence interval for this estimate.

Table 21. Logit results - baseline model

Independent variable	Coefficient	Standard error	Probability coefficient equals zero
ACOL (\$1,000)	0.1162989	0.0024833	0.000
E-3	-0.8885403	0.0183556	0.000
E-5	0.4619234	0.0138606	0.000
E-6	0.4948727	0.0574651	0.000
Expected sea tour	-0.2918802	0.0129481	0.000
Long shore tour	-0.4164279	0.0236652	0.000
Long sea tour	-0.1706097	0.0196064	0.000
Nontraditional rotation	-0.1048387	0.0389739	0.007
Rating eligible for VSI/SSB	-0.1937269	0.0214721	0.000
Previous extension	0.6192732	0.0146074	0.000
Length of service (months)	-0.0449512	0.0007288	0.000
SEABEE Construction	-0.1916408	0.040682	0.000
Non-SEABEE Construction	-0.5039881	0.0803791	0.000
Ship Maintenance	-0.2445109	0.0267851	0.000
Aviation Maintenance	0.3697221	0.0198309	0.000
Aviation Ground Support	-0.4105172	0.025317	0.000
Media	-0.4241226	0.0383129	0.000
Logistics	-0.1931014	0.0231311	0.000
Administration	-0.2513335	0.0289934	0.000
Data Systems	0.0463619	0.0268935	0.085
General Seamanship	-0.6239253	0.0224232	0.000

Table 21. Logit results - baseline model (continued)

Independent variable	Coefficient	Standard error	Probability coefficient equals zero
Health Care	0.2817947	0.0271484	0.000
Cryptology	0.0387455	0.0458419	0.398
Ordnance Systems	0.238134	0.0254851	0.000
Communications / Sensor	0.1941102	0.0350754	0.000
Weapons Systems / Control	0.5526943	0.0254265	0.000
Northern (coastal) California	-0.8267596	0.024168	0.000
Southern (coastal) California	-0.5725326	0.0156568	0.000
Hawaii	-0.2376489	0.0269174	0.000
Coastal Virginia	-0.4736201	0.0151719	0.000
Washington State	-0.5818533	0.0239897	0.000
South Carolina	-0.2092629	0.0265656	0.000
Northern (coastal) Florida	-0.3609583	0.0208369	0.000
Black	0.2198179	0.0178799	0.000
Hispanic	-0.2674819	0.0219017	0.000
Other ethnicity (non-white)	0.2724968	0.0284857	0.000
Married	0.5226453	0.0117708	0.000
Number of children	0.1667898	0.008379	0.000
Age	0.1138637	0.002811	0.000
AFQT	-0.0007174	0.0003025	0.018
Unemployment rate	0.0306778	0.0038016	0.000
FY88	-0.0561674	0.026136	0.032
FY89	0.074045	0.0256196	0.004
FY90	0.0894276	0.0255508	0.000
FY91	0.139277	0.0252867	0.000
FY92	0.1373262	0.0268643	0.000
FY93	-0.1026117	0.0268614	0.000
FY94	0.0548417	0.0315602	0.082
FY95	0.2219484	0.0335539	0.000
FY96	0.1564371	0.0306878	0.000
FY97	0.0376345	0.0297543	0.206
FY98	0.0874207	0.032043	0.006
FY99	0.1165215	0.0332619	0.000
Constant	-1.718971	0.074743	0.000

Appendix D: Results of baseline model

For clarity, we separate variables into different categories and present their marginal effects in separate subsections below. Note, however, that these results all come from the regression presented in appendix C.

Because the logit model estimates a nonlinear relationship between the explanatory variables and the probability of reenlistment, the interpretation of the coefficients is not straightforward. To facilitate interpretation of the results, we calculate and present the “marginal effects” rather than the underlying coefficients.

For variables that take on a wide range of values (unemployment rates, age, etc.), the marginal effect measures the percentage-point change in the probability of reenlistment for a unit change in one of the independent variables, holding the other independent variables constant. For sets of variables that indicate the status of a Sailor (e.g., marital status), the marginal effect measures the percentage-point difference in the probability of reenlistment (holding all else constant) for individuals with a given characteristic, relative to an excluded group. For example, a marginal effect for marital status of 0.10 implies that, for two otherwise identical Sailors, the probability of reenlistment is 10 percentage points higher for the married than for the single one.

Characteristics of military service

Table 22 presents our estimates of the effects of various characteristics of military service on reenlistment. The second column displays the average predicted probability associated with each characteristic; for example, a predicted probability of 0.32 for E-4s implies that, if all Sailors were E-4 at the time of the reenlistment decision, predicted reenlistment would be 32 percent. The third column presents the marginal effects, which measure the percentage-point difference in

the predicted probability of reenlistment for Sailors with a given characteristic, relative to a reference group. For example, a marginal effect for E-3s of -0.15 implies that, for two otherwise identical Sailors, the probability of reenlistment is 15 percentage points lower for the E-3 than for the E-4. Finally, the fourth column lists the percentage difference in reenlistment rates between individuals with a given characteristic and those of the reference group. Again, a 15-percentage-point difference between predicted reenlistment rates of E-3s and E-4s implies that reenlistment for E-3s is about 47 percent lower than that of E-4s, an extremely large difference.

Table 22. Characteristics of military service—baseline model

Independent variable	Predicted probability	Marginal effect ^a	Percentage change
Paygrade			
<i>E-4^b</i>	0.32		
E-3	0.17	-0.15	-47
E-5	0.42	0.10	31
E-6	0.42	0.10	31
Sea/shore rotation			
<i>Expected shore tour^a</i>	0.36		
Expected sea tour	0.30	-0.06	-17
Long shore tour	0.28	-0.08	-22
Long sea tour	0.33	-0.03	-8
Nontraditional rotation	0.34	-0.02	-6
VSI/SSB			
<i>Rating is ineligible^a</i>	0.33		
Rating is eligible	0.29	-0.04	-12
Previous extensions			
<i>No previous extension^a</i>	0.30		
Previous extension	0.43	0.13	43
Other characteristics			
Length of service (months)	0.32	-0.01	-3

a. Zero lies outside the 99-percent confidence interval for this estimate.

b. Italicized variables indicate the reference group with which the effects of other variables are compared.

Paygrade

The marginal effects for the set of paygrade variables suggest that the propensity to reenlist increases with promotion. The marginal effects are measured relative to the reenlistment rate of E-4s; although E-3s have dramatically lower reenlistment than E-4s, both E-5s and E-6s have much higher reenlistment.

Because paygrade is one of the primary determinants of military compensation, these marginal effects measure the relationship between paygrade and reenlistment above and beyond their effect through pay. These effects could reflect, for example, a better “job match” [14], unobserved ability, or differences between ratings even after controlling for ratings groups [5].

Sea/shore rotation

Table 22 also suggests that expectations about future sea duty strongly affect one’s reenlistment decision. Sailors who expect to rotate to sea duty have reenlistment rates that are 17 percent lower than those who expect to rotate to shore duty. Similarly, those who have been on a shore tour longer than their PST are also significantly less likely to reenlist; this is consistent with the notion that these Sailors will probably rotate to sea duty in the near future.

Sailors currently on long sea tours are less likely to reenlist than those who expect to rotate soon to shore duty. It is possible that these lower reenlistment rates are caused by an inordinate amount of time spent at sea during one’s first term, although this is not directly testable. In general, however, these differences in reenlistment by expected sea/shore rotation are consistent with our hypothesis that the amount of sea duty is an important factor in the reenlistment decision of the enlisted Sailor.

VSI/SSB and previous extensions

Finally, table 22 also indicates that other variables thought to reflect a relative preference for military service do indeed have significant effects on reenlistment behavior. Consistent with the findings of [10], Sailors with a history of conditional extensions are 43 percent more

likely to reenlist than those who have not had a previous extension.⁵⁴ Our results also indicate that individuals in ratings offered VSI/SSB at the time of their reenlistment decision are about twelve percent less likely to reenlist than other Sailors. This is consistent with a decrease in the promotion opportunities available to these individuals due to the drawdown.⁵⁵

Differences by rating

Table 23 presents differences in reenlistment by occupation groups. These differentials are calculated holding constant everything else that is included in our model; therefore, they reflect differences in reenlistment above and beyond differences due to promotion opportunities, sea/shore rotation, geographic location, or SRBs. As table 23 indicates, substantial differences in reenlistment behavior persist by occupation group after controlling for these other factors. The Weapons Systems/Control group (i.e., the ET and FC ratings) has the highest predicted reenlistment over the FY87-99 period, while the General Seamanship ratings (BM, OS, QM, and SM ratings) have the lowest.

A variety of different factors could account for these differences. The two most likely candidates are differences in civilian opportunities and differences in working conditions/environment. As [2] discusses, we estimate civilian earnings separately for “technical” and “nontechnical” ratings, which is a fairly broad classification.⁵⁶ Although it is possible that estimating civilian opportunities separately by rating would eliminate any remaining variation by occupation group, [7] addresses the difficulties associated with this strategy and concludes that there are only a few ratings for which reliable estimates can be obtained.

54. Reference [10] finds reenlistment rates that are 36 percent higher over the FY87-96 period.

55. As [2] discusses, the ACOL variable is estimated using the paygrade distribution at the time of the reenlistment decision. For Sailors making reenlistment decisions during the drawdown, it is likely that actual promotion opportunities are quite different from what we observe for their predecessors.

56. Reference [6] describes this classification of ratings in greater detail.

Table 23. Differences by rating—baseline model

Rating group ^a	Predicted probability	Marginal effect	Percentage change
<i>Marine Engineering^b</i>	0.33		
SEABEE Construction	0.29	-0.04 ^c	-12
Non-SEABEE Construction	0.24	-0.09 ^c	-27
Ship Maintenance	0.28	-0.05 ^c	-15
Aviation Maintenance	0.40	0.07 ^c	21
Aviation Ground Support	0.25	-0.08 ^c	-24
Media	0.25	-0.08 ^c	-24
Logistics	0.29	-0.04 ^c	-12
Administration	0.28	-0.05 ^d	-15
Data Systems	0.34	0.01	3
General Seamanship	0.22	-0.11 ^c	-33
Health Care	0.38	0.05 ^c	15
Cryptology	0.33	0.00	0
Ordnance Systems	0.37	0.04 ^c	12
Communications / Sensor	0.37	0.04 ^c	12
Weapons Systems / Control	0.44	0.11 ^c	33

a. See appendix B for a detailed explanation of these rating groups.

b. This italicized variable indicates the reference group with which the effects of other variables are compared.

c. Zero lies outside the 99-percent confidence interval for this estimate.

d. Zero lies outside the 90-percent confidence interval for this estimate.

Similarly, it is possible that these differences are caused by differences in working conditions. Though we have controlled for geographic location and sea/shore rotation in our model, there are likely differences both in working conditions and even ship type for these occupation groups. Without additional data, it is not possible to identify the source of these differences from one occupation to the next; the inclusion of these variables in the model, however, serves to capture the effects of any of these differences in environment or civilian opportunities on reenlistment decisions.

Geographic differences

Differences by geographic location are presented in table 24; again, these differentials are calculated holding constant all other factors in

our model. Marginal effects for the seven largest geographic regions are calculated relative to all other UICs; the fact that all marginal effects are negative implies that reenlistment in the regions with the most Sailors are lower than the average reenlistment rate of smaller UICs.

Table 24. Geographic differences—baseline model

Region	Predicted probability	Marginal effect ^a	Percentage change
<i>Other</i> ^b	0.40		
Northern (coastal) California	0.24	-0.16	-40
Southern (coastal) California	0.29	-0.11	-28
Hawaii	0.35	-0.05	-13
Coastal Virginia	0.31	-0.09	-23
Washington State	0.29	-0.11	-28
South Carolina	0.36	-0.04	-10
Northern (coastal) Florida	0.33	-0.07	-18

a. Zero lies outside the 99-percent confidence interval for this estimate.

b. This italicized variable indicates the reference group with which the effects of other variables are compared.

These differences from one UIC to the next reflect any differences in Sailor characteristics not already captured in our model, be they differences in civilian opportunities or in work environment or other unobservable characteristics. One possible explanation for the negative marginal effects is that they reflect perceptions about *future* UICs. In other words, Sailors stationed at these UICs may be less likely to reenlist because they expect that a future assignment (after reenlistment) would move them to another, less desirable UIC.

Demographic characteristics

Table 25 presents our estimates of the effects that various demographic characteristics have on reenlistment. Most of these characteristics are used to predict military and/or civilian compensation, so the marginal effects in table 25 represent their effect on reenlistment behavior beyond the effect they have through pay.

Table 25. Demographic characteristics—baseline model

Independent variable	Predicted probability	Marginal effect	Percentage change
Race/ethnicity			
<i>White^a</i>	0.32		
Black	0.36	0.04 ^b	13
Hispanic	0.27	-0.05 ^b	-16
Other ethnicity	0.37	0.05 ^b	16
Marital status			
<i>Not married^a</i>	0.28		
Married	0.38	0.10 ^b	36
Other characteristics			
Number of children	0.32	0.03 ^b	9
Age	0.32	0.02 ^b	6
AFQT	0.32	0.00 ^c	0
Unemployment rate	0.32	0.01 ^b	3

a. Italicized variables indicate the reference group with which the effects of other variables are compared.

b. Zero lies outside the 99-percent confidence interval for this estimate.

c. Zero lies outside the 95-percent confidence interval for this estimate.

Race/ethnicity

The marginal effects for the set of race/ethnicity variables suggest that the propensity to reenlist varies dramatically by race. Because race is one of the primary determinants of expected civilian compensation in our model, the differences in table 25 reflect relative preferences for Navy service across ethnic groups. It is also possible, however, that they reflect civilian earnings opportunities not captured in our model of compensation, or other demographic characteristics for which we do not (and cannot) control.

Marital status

Married Sailors are significantly more likely to reenlist than unmarried Sailors, even when controlling for differences in current and expected military compensation.⁵⁷ This is a result that is commonly

57. Married Sailors receive higher military compensation than unmarried Sailors because housing allowances vary by dependency status. Our compensation estimates recognize that many who are not married at the time of the first reenlistment decision eventually do marry [2].

found in the empirical literature. It is interesting to note that the percentage difference in reenlistment by marital status is slightly smaller than that found in [7], whose estimates of military compensation do not vary by dependency status. This is consistent with the hypothesis that the marginal effect on marital status in our baseline model reflects a relative preference for military service above and beyond any differential in compensation. Alternatively, the positive coefficient on marital status could reflect additional differences in compensation not present in our model.⁵⁸

Other characteristics

Other demographic characteristics have significant effects on the probability of reenlistment. Those with more children have a higher propensity to reenlist, as do older Sailors.⁵⁹ Also, people from states with high unemployment rates are significantly more likely to reenlist than those from states with lower unemployment rates. We hypothesize that the unemployment rate in a person's home state is negatively correlated with his civilian opportunities; in this context, the relationship between the unemployment and reenlistment rates confirms that civilian opportunities are an important consideration in the reenlistment decision.

Fiscal year effects

Finally, table 26 presents the marginal effects for each fiscal year relative to FY87. These coefficients represent any differences in reenlistment over time that are not caused by changes in any other independent variable in the model. Although there are certainly significant differences in reenlistment from one fiscal year to the next, there does not appear to be any discernible trend over the FY87-99 period.

58. For example, the tax advantage may be larger for married people, or there may be differences in expected civilian compensation not captured in our estimates.

59. Although the coefficient on AFQT is statistically larger than zero, its magnitude is so small that the marginal effect is essentially zero.

Table 26. Fiscal year effects—baseline model

Fiscal year	Predicted probability	Marginal effect	Percentage change
1987 ^a	0.31		
1988	0.30	-0.01 ^b	-3
1989	0.32	0.01 ^c	3
1990	0.33	0.02 ^c	7
1991	0.34	0.03 ^c	10
1992	0.34	0.03 ^c	10
1993	0.29	-0.02 ^c	-7
1994	0.32	0.01 ^d	3
1995	0.35	0.04 ^c	13
1996	0.34	0.03 ^c	10
1997	0.32	0.01	3
1998	0.33	0.02 ^c	7
1999	0.34	0.03 ^c	10

a. The italicized variable indicates the reference group with which the effects of other variables are compared.

b. Zero lies outside the 95-percent confidence interval for this estimate.

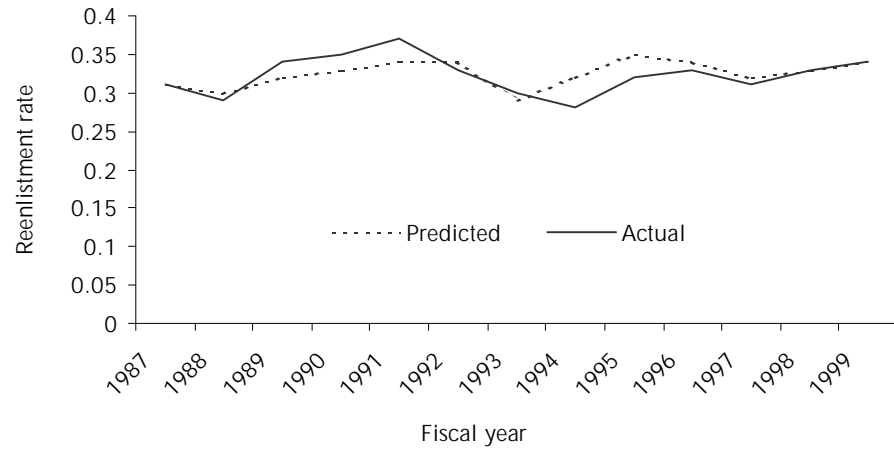
c. Zero lies outside the 99-percent confidence interval for this estimate.

d. Zero lies outside the 90-percent confidence interval for this estimate.

It is interesting to compare these predictions with the actual reenlistment rates observed over the same time period. Figure 1 presents, for each year from FY87 to FY99, both the predicted and actual reenlistment rate for our sample of Navy enlisted personnel. These predictions are calculated holding all other factors in our model constant; differences between predicted and actual reenlistment, then, are the result of differences in other independent variables from one fiscal year to the next.⁶⁰ In general, there is less variation in predicted reenlistment from one year to the next than in the actual reenlistment rates over this time period; this implies that differences in the characteristics of Navy enlisted personnel are responsible for some of the variation over time.

60. If one calculated predicted reenlistment rates by fiscal year, the predicted reenlistment rate would equal the actual reenlistment rate. Holding all other factors constant and calculating the marginal effects allows us to observe the effect that these other variables have on reenlistment.

Figure 1. Predicted and actual reenlistment rates—baseline model



References

- [1] Matthew S. Goldberg. *A Survey of Enlisted Retention: Models and Findings*, Nov 2001 (CNA Research Memorandum D0004085.A2)
- [2] Michael L. Hansen and Jennie W. Wenger, with David Gregory and Lynda G. Houck. *Construction of an ACOL Framework: Theoretical Guidelines and Practical Considerations*, Jul 2001 (CNA Research Memorandum D0004380.A1)
- [3] D. Alton Smith, Stephen D. Sylwester, and Christine M. Villa. "Army Reenlistment Models," in Curtis L. Gilroy, David K. Horne, and D. Alton Smith (eds.), *Military Compensation and Personnel Retention: Models and Evidence*, Feb 1991 (United States Army Research Institute for the Behavioral and Social Sciences)
- [4] James R. Hosek and Christine E. Peterson. *Reenlistment Bonuses and Retention Behavior*, Mar 1985 (RAND Report R-3199-ML)
- [5] Martha E. Shiells and Joyce S. McMahon. *Effects of Sea Duty and Advancement on First-Term Retention*, Jun 1993 (CNA Research Memorandum 92-205)
- [6] Linda C. Cavalluzzo and Jeremy A. Arkes. *Comparisons of Enlisted and Civilian Pay*, Aug 1998 (CNA Annotated Briefing 98-88)
- [7] Michael L. Hansen. *Compensation and Enlisted Manning Shortfalls*, Sep 2000 (CNA Research Memorandum D0001998)
- [8] John T. Warner and Matthew S. Goldberg. "The Influence of Non-Pecuniary Factors on Labor Supply: The Case of Navy

- Enlisted Personnel." *Review of Economics and Statistics*, 66 (1), Feb 1984: 26-35
- [9] Stephen L. Mehay and Paul F. Hogan. "The Effect of Separation Bonuses on Voluntary Quits: Evidence from the Military's Downsizing." *Southern Economic Journal*, 1998: 127-139
- [10] David B. Nickerson, David Gregory, and Henry S. Griffis. *Analysis of Conditional Extensions*, Aug 1998 (CNA Annotated Briefing 98-63)
- [11] G. S. Maddala. *Limited-Dependent and Qualitative Variables in Econometrics*. Cambridge: Cambridge University Press, 1983
- [12] John T. Warner. "Comment on Reenlistment Decision Models: Implications for Policy Making," in *Army Manpower Economics*, 1986 (Boulder, CO: Westview Special Studies in Military Affairs)
- [13] John T. Warner and Saul Pleeter. "The Personal Discount Rate: Evidence from Military Downsizing Programs." *American Economic Review* 91(1), Mar 2001, 33-53
- [14] Timothy W. Cooke, Alan J. Marcus, and Aline O. Quester. *Personnel Tempo of Operations and Navy Enlisted Retention*, Feb 1992 (CNA Research Memorandum 91-150)
- [15] Matthew S. Goldberg. "Compensation and Retention in the U.S. Navy." Chicago: The University of Chicago, Mar 1984
- [16] Matthew S. Goldberg and John T. Warner. *Determinants of Navy Reenlistment and Extension Rates*, Dec 1982 (CNA Research Contribution 476)
- [17] William H. Greene. *Econometric Analysis*. Upper Saddle River, NJ: Prentice-Hall, Inc., 1993
- [18] Henri Theil. *Principles of Econometrics*. New York: John Wiley & Sons, Inc., 1971

- [19] David Card and Alan B. Krueger. "Labor Market Effects of School Quality: Theory and Evidence." NBER Working Paper 5450, Feb 1996
- [20] Yoel Haitovsky. *Regression Estimation from Grouped Data*. London: Griffin, 1973
- [21] Takeshi Amemiya. "Qualitative Response Models: A Survey." *Journal of Economic Literature* 19, Mar 1981, 1483-1536
- [22] Martha E. Koopman and Heidi L. Golding. *Optimal Manning and Technological Change*, Jul 1999 (CNA Research Memorandum 99-59)

List of tables

Table 1.	Descriptive statistics	20
Table 2.	Pay effects—baseline model	22
Table 3.	Pay effects—excluding race/ethnicity and age	29
Table 4.	Pay effects—excluding race/ethnicity, age, and fiscal year.	29
Table 5.	Pay effects—separate effects of military and civilian pay.	31
Table 6.	Pay effects—fixed effects for each rating.	33
Table 7.	Pay effects—occupation-specific elasticities	34
Table 8.	Pay effects—10-percent discount rate	36
Table 9.	Pay effects—30-percent discount rate	36
Table 10.	Pay effects—grouped data	41
Table 11.	Summary of alternative specifications	42
Table 12.	Pay effects—calculated separately by fiscal year.	46
Table 13.	Within-sample predictions of reenlistment behavior—AT, ET, and FC ratings	54
Table 14.	Within-sample predictions of reenlistment behavior—ratings with change in SRB multiplier from FY99 to FY00.	55
Table 15.	Predictions of reenlistment behavior—test data.	57

Table 16.	Descriptive statistics	58
Table 17.	Out-of-sample predictions of reenlistment behavior—FY00	60
Table 18.	Out-of-sample predictions of reenlistment behavior—subsets of ratings	61
Table 19.	Predictions of FY00 reenlistment rates—baseline model and exclusion of various subperiods	65
Table 20.	Major geographic regions and their locale codes	69
Table 21.	Logit results - baseline model.	75
Table 22.	Characteristics of military service—baseline model	78
Table 23.	Differences by rating—baseline model.	81
Table 24.	Geographic differences—baseline model	82
Table 25.	Demographic characteristics—baseline model	83
Table 26.	Fiscal year effects—baseline model.	85