CRM D0004085.A2 / Final
November 2001

A Survey of Enlisted Retention: Models and Findings

Matthew S. Goldberg

CNA
4825 Mark Center Drive • Alexandria, Virginia 22311-1850

Approved for distribution: November 2001

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 Document Control and Distribution Section at 703-824-2123.

Contents

Introduction and summary

The supply of manpower has always been a concern to the military, but this issue took on greater importance in the events leading up to the creation of the All-Volunteer Force (AVF) in 1973. In 1969, President Nixon established the President's Commission on an All-Volunteer Force, commonly known as the Gates Commission. The commission's staff papers were among the first to systematically study the supply of both enlistments and reenlistments to the military. These papers, along with concurrent literature in the professional economics journals, demonstrated that an AVF was feasible from a fiscal perspective.[1]

A variety of studies through the early and mid-1970s continued to examine the supply of reenlistments.[2] A major advance occurred during the late 1970s with development of the Annualized-Cost-of-Leaving (ACOL) model. Under this model, the primary driver of the reenlistment decision is the discounted difference between the military pay stream from reenlisting, and the civilian pay stream from leaving the military. In particular, ACOL combined all the elements of military pay (basic pay, allowances, reenlistment bonuses, retirement pay) into a single, discounted present value. Moreover, ACOL suggested a time horizon over which the military and civilian pay streams must be measured and compared. From a statistical perspective, ACOL expressed the reenlistment rate as a logit or probit function of the discounted pay difference, and possibly other regressors.

---

1 The concurrent literature includes Altman and Fechter [1], Fisher [2], Hansen and Weisbrod [3], Miller [4], and Oi [5]. These papers also made the important distinction between the fiscal cost of an AVF and the opportunity cost of diverting individuals from the civilian careers they would otherwise have pursued.

2 For example, see [6, 7, and 8].

In parallel to ACOL, Glenn Gotz and John McCall developed a dynamic-programming model of Air Force officer retention [9, 10]. Rather than specifying a single, dominant time horizon, their model allowed for probabilistic weighting of multiple time horizons. Although their model was theoretically elegant, it proved difficult to estimate given the computer hardware and software environment of the early 1980s.

ACOL remained the conventional point of departure for much of the research conducted during the 1980s and 1990s. However, considerable effort went into improving the statistical estimation of reenlistment models. That research effort took two major directions. First, panel probit models were formulated to better track the composition of cohorts making successive reenlistment decisions during their military careers. For example, those induced to reenlist by a Selective Reenlistment Bonus (SRB) might have less of a taste for military life than others who would have reenlisted even absent an SRB. These bonus-induced individuals would be less likely to remain in the military at subsequent decision points, unless the SRB were sustained. Panel probit models are designed precisely to capture the effects of cohort composition on the outcome of successive binary decisions.

The second research direction was to recognize the distinction between reenlistments (i.e., commitments for 36 or more additional months of service) and shorter extensions. Only individuals who reenlist are eligible to receive SRBs. Thus, an increase in SRB levels not only will increase the total retention rate but also will change the mix of individuals retained between those who reenlist and those who merely extend. The resulting change in the mix of commitments is clearly important for personnel planning purposes. Thus, a binary logit or probit model was replaced by a trichotomous model, such as conditional logit, multinomial logit, or nested logit.

The statistics literature tells us little about adding cohort-composition effects to trichotomous choice models. The panel probit approach and the various trichotomous logit approaches have advanced essentially independently, although some of the same researchers have applied both approaches, at one time or another, in modeling the reenlistment decision.

Other statistical problems have prompted researchers to modify or enhance the logit or probit models in various ways. First, there may be reverse causation between pay and the reenlistment rate. The goal of the analysis is to estimate the positive effects of SRB and other incentives on the reenlistment rate. However, enlisted occupations with chronically low reenlistment rates tend to be compensated with higher SRB levels. This pattern of reverse causation may lead to a downward bias in the estimated pay coefficient. At least two studies [11 and 12] have used panel data and applied a fixed-effect estimator in an effort to alleviate this source of bias.

We have already discussed the possibility that people who reenlist for an SRB might be less likely to reenlist a second time. Similarly, those who enlist for an accession bonus might be less likely to reenlist at the first-term decision point. Two studies [13 and 14] have attempted to control for the composition of the accession cohort when modeling the first-term reenlistment decision. They did so by jointly modeling survival to the first-term decision point with the outcome of that decision.

Several issues arise in computing the elasticity of the reenlistment rate with respect to military pay. The definition of "reenlistment" is complicated by a number of factors, including reenlistment eligibility and the treatment of extensions. Some studies exclude individuals declared ineligible to reenlist from the denominator of the reenlistment rate. However, the eligibility determination may be endogenous if, for example, individuals expressing a disinclination to reenlist are subsequently declared ineligible by their units. Some studies combine extensions with reenlistments, modeling total retention. Others defer their analysis of extensions, instead tracking them to learn whether they ultimately reenlist. It is difficult to compare the pay elasticities from studies that differ in their treatment of extensions.

Computation of the pay elasticity is further complicated by the definition of "military pay." Many studies measure pay in terms of ACOL or some other difference between the military and civilian pay streams. However, it is perilous to directly compute the elasticity of the reenlistment rate with respect to a pay *difference*. The elasticity, so computed, will have the same algebraic sign as the baseline pay difference. Thus, even if increased pay has a positive effect in encouraging more reenlistments, the elasticity may be zero or even negative. Instead, the model should be exercised by hypothesizing a fixed, discrete increase in military pay (e.g., $1,000). Express this increase as a percentage of baseline military pay, and divide the resulting percentage increase in the reenlistment rate by the percentage increase in military pay. This procedure estimates the arc elasticity with respect to military pay *(not* the pay difference), and is guaranteed to yield the correct algebraic sign.

At various points in time, the SRB has been paid either as a lump-sum on the date of reenlistment, or in equal annual installments over the duration of the reenlistment contract (with no indexing for inflation). To the first order of approximation, lump-sum bonuses are cost-effective if military members' discount rates exceed that of the federal government.[3] Since 1992, the Office of Management and Budget (OMB) has tied the federal government's discount rate to the market rate on Treasury bonds. Several studies have estimated the discount rates of military members. Two of these studies [12 and 15] exploited the natural experiment that occurred when the method of SRB payment switched from annual installments to lump-sum payments. The estimates of military members' real (i.e., inflation-adjusted) discount rates are in the range of 6 to 26 percent. By contrast, real Treasury rates have generally been in the range of 3 to 4 percent. Thus, lump-sum bonuses are the preferred method of payment.

Finally, several studies have investigated the retention effects of variables other than relative military pay. In studies specific to the Navy, the variables of interest have included the incidence of sea duty, length of deployment, time between deployments, and percentage of time spent under way while not deployed [16, 17, and 18]. The Navy studies have also estimated the SRB and other incentives required to compensate for adverse changes in these duty characteristics. A more recent study has measured additional duty characteristics and extended the analysis to all four military services [19].

---

3. Other considerations include progressive income taxation and government recoupment of lump-sum bonuses from individuals who separate during the contract period. Empirically, these factors are minor and do not change the basic conclusion.

The remainder of this report reviews each of the aforementioned methodological issues in detail. It also presents a summary of the pay elasticities estimated using the various measurement and statistical techniques. Although we cannot rationalize all of the variation in pay elasticities, we attempt to correlate the elasticities with the techniques used to estimate them.

THIS PAGE INTENTIONALLY LEFT BLANK

ACOL model

John Warner and his various collaborators developed the ACOL model in a series of papers. The initial motivation was to study a proposal by the President's Commission on Military Compensation (PCMC) to reform the military retirement system [20]. Warner also programmed a forecasting version of the model in the APL language. He distributed the model to the Navy Bureau of Personnel (BuPers) and, later, to the Office of the Assistant Secretary of Defense (Manpower, Reserve Affairs, and Logistics). BuPers started using the model to analyze manpower issues in the Navy's Program Objectives Memorandum (POM), beginning with POM 1982. By the early 1980s, the ACOL model was well known and accepted throughout the defense manpower community.

The ACOL model's first appearance in the academic literature was a 1984 paper by three of its codevelopers, Enns, Nelson, and Warner [21]. During that same year, Warner and Goldberg [18] published an application of the ACOL model in a mainstream economics journal. Parallel developments were taking place in the literature on retirement from civilian-sector jobs (e.g., Stock and Wise [22], who were apparently unaware of the ACOL model). The two strands in the literature were eventually brought together by Lumsdaine, Stock, and Wise [23] and Daula and Moffitt [24].

Economic theory suggests that individuals combine all the elements of compensation associated with any alternative into a single measure, typically the discounted present value. In our context, SRBs provide both cross-section (i.e., across military occupations) and time-series variation in discounted pay. Civilian earnings provide time-series variation, and may provide additional cross-section variation to the extent that the civilian earnings functions account for military occupation. Military pay excluding SRBs (i.e., Regular Military Compensation, or RMC) provides time-series variation but only minimal cross-section variation (to the extent that differences in promotion rates are captured).

If the three pay components (SRBs, civilian earnings, and RMC) were entered as separate regressors, their respective coefficients would almost certainly be different. RMC would probably have the least significant coefficient because RMC has the least sample variation. However, it would be wrong to conclude that increases in RMC have the smallest impact on retention. To estimate the effect of RMC more precisely, one could divide RMC by civilian earnings, thereby forming an index of relative military pay. The coefficient on this index would be driven largely by the variation in civilian earnings, but it could be used to forecast the effects of changes in RMC on retention. These forecasts would be valid as long as individuals were indifferent between an increase in RMC and an equal percentage decrease in civilian earnings.

It is even more difficult to compare the efficacy of increases in RMC versus increases in SRBs. One difference is that SRBs can be targeted to military occupations experiencing retention problems. Another difference is that SRBs have a different time dimension from RMC. SRBs represent one-time payments or, at most, a short series of annual installments. On the contrary, a given dollar increase in RMC persists for the duration of a person's military career. Thus, an increase in RMC cannot be evaluated without knowing (or at least estimating) the person's time horizon. Moreover, for those whose time horizons extend to 20 or more years of service, basic pay (the largest element of RMC) also affects their retirement annuity. Table 1 compares the time dimensions of these various elements of pay.

The ACOL approach solves the dimensionality problem by combining all the elements of compensation into a single measure. In particular, the rich sample variation in SRBs can be brought to bear in estimating the coefficient on the ACOL variable in a logit or probit choice model. The ACOL coefficient, in turn, can be used to forecast the effects of *any* change in compensation, including changes in the retirement system. Indeed, the ACOL approach was developed precisely to study the military retirement system. We will also argue, in a later section, that the ACOL approach is consistent with the results of studies that segmented compensation into multiple measures (e.g., the SRB level and an index of relative military pay).

Table 1. Elements of pay and their time dimensions

Pay element	Time dimensions
RMC (basic pay + allowances + tax advantage)	Persists over entire military career
Basic pay	Persists over entire military career Determines retirement annuity
SRBs	Lump-sum is instantaneous Annual installments over the reenlistment contract
Civilian earnings stream	Entire working life

ACOL time horizon

The ACOL approach suggests a time horizon for comparing the military and civilian discounted pay streams. However, construction of the ACOL variable requires an assumption on military members' discount rates. We will describe methods for estimating discount rates in a later section. For now, we merely report that enlisted personnel at their first-term and second-term reenlistment points appear to have real (i.e., net of inflation) discount rates of 6 to 26 percent.

To develop the ACOL variable, suppose initially that the retention decision were made solely by comparing the military and civilian discounted pay streams. Then, assuming that the pay streams could be measured precisely, we could predict *with certainty* the choice made by any individual—simply the one yielding the highest discounted pay stream.

Relaxing these assumptions gradually, suppose next that the pay streams were known exactly to the individual decision-maker, but not to the data analyst. This would be the case if the analyst were using a regression function to predict civilian earnings, yet the individual had more precise knowledge of his or her own earnings potential. In this situation, we could no longer predict an individual's choice with certainty. Instead, we could predict only the probabilities of staying or leaving for each individual.

As a further relaxation, we can recognize that a person's occupational choice depends on a comparison not only of discounted pay streams but also of the nonmonetary advantages and disadvantages of military versus civilian life. A general assumption in the literature is that the nonmonetary factors may be expressed as monetary equivalents (e.g., "I will remain in the military only if they pay me $1,000 more per year than I could earn as a civilian"). Most authors further combine the nonmonetary factors with the unmeasured portion of the pay streams, and label the result the "taste factor." Continuing the example, suppose that the same person who requires a $1,000 annual premium also knows that his or her potential civilian earnings are $500 above the regression prediction. The taste factor for this person would be the sum, $1,500. Note also that the taste factor could be negative if people prefer military life or if their potential civilian earnings are below the regression prediction.

Suppose, for the moment, that a person currently in year of service (YOS) t is contemplating only two choices: remain in the military for an additional s years, or leave immediately. He or she will remain in the military if:

$$\sum_{j=t+1}^{t+s} M_j \times (1+r)^{t-j} \quad > \quad \sum_{j=t+1}^{t+s} (C_j + v) \times (1+r)^{t-j} \, , \tag{1}$$

where M_j is military pay (including any SRBs) in YOS j, C_j is potential civilian pay in the same year, and v is the taste factor. Note that the taste factor is assumed to be time-invariant.[4] Equivalently, the person will remain in the military if:

$$\text{ACOL}_s \equiv \frac{\sum_{j=t+1}^{t+s} (M_j - C_j) \times (1+r)^{t-j}}{\sum_{j=t+1}^{t+s} (1+r)^{t-j}} > v. \tag{2}$$

As its name suggests, the ACOL variable is simply the annualized (or annuitized) difference between the military and civilian pay streams. Put differently, a stream of s pay differences, each equal to ACOL_s, has the same discounted value as the pay stream $\{(M_j - C_j),\ j = t+1, \ldots, t+s\}$, namely, the numerator of the previous expression for ACOL_s.

Now considering *all* possible horizons $\{s = 1, 2, 3, \ldots\}$, the person will remain in the military if there is at least one horizon over which ACOL exceeds the taste factor. Mathematically, this condition is equivalent to having the maximum ACOL greater than the taste factor:

$$\text{Max}_s\{\text{ACOL}_s\} > v. \tag{3}$$

Conversely, the individual will leave the military immediately if there is no horizon over which ACOL exceeds the taste factor. Mathematically, this condition is equivalent to having the maximum ACOL less than the taste factor:

$$\text{Max}_s\{\text{ACOL}_s\} < v. \tag{4}$$

---

4 Inequality (1) is written so that potential civilian pay depends on calendar year (equivalently, the person's age), but not on the length of his or her military career (i.e., not upon the value of s). This assumption can be relaxed, at the expense of some additional terms that measure the gain or loss in potential civilian pay from continued military service. The ACOL expression under this relaxation is found in [11] or [23].

Thus, the maximum ACOL summarizes all of the information on pay streams necessary to predict a person's retention decision. Earnings further than s^* years into the future (where s^* is the horizon that maximizes ACOL) need not be considered. This result is impressive because earnings beyond s^*, even when discounted, need not be negligible numerically; yet the retention decision can be made without considering them.

"Optimality" of the ACOL time horizon

The horizon s^* is sometimes called the "optimal horizon," but this nomenclature is misleading. It seems to imply that, among all possible horizons that involve remaining in the military at least one additional year, the horizon s^* is the most preferred. However, some simple counterexamples disprove this conjecture.[5]

Suppose the only two possible career lengths involve staying for one additional year ($s = 1$) or two additional years ($s = 2$). Suppose further that the military/civilian pay differences are \$2,000 in the first year and \$1,000 in the second year. If the discount rate is 10 percent, the ACOL values are $ACOL_1$ = \$2,000 and $ACOL_2$ = \$1,524. The optimal horizon over which ACOL is maximized is $s = 1$ year. Thus, the person will stay in the military for *some* duration if the taste factor is less than the maximum ACOL, or \$2,000. Yet he or she would prefer to stay for two additional years, rather than just one, if the taste factor is sufficiently small (or negative). Specifically, having already stayed for one additional year, the person would prefer to stay for the second

---

5 In one of many published examples of the misleading use of the term "optimal horizon," Gotz [25, p. 266] states that, "associated with [*the* ACOL variable] is a known optimal future quitting date." Black, Moffitt, and Warner [26, p. 270] agree with Gotz on this point: "the ACOL model assumes that the individual picks a single optimal date of leaving some time in the future." A rare correct statement is found in Mackin et al. [27, p. C-5]: "Note that the ACOL measure should be considered an index describing the financial incentive to stay at least one more year. The horizon associated with the maximum ACOL is *not necessarily* the optimal leaving point" [emphasis added].

year as well if the taste factor is less than the military/civilian pay difference in that year, \$1,000. Figure 1 illustrates this situation.

Figure 1. First counterexample to optimality of ACOL time horizon

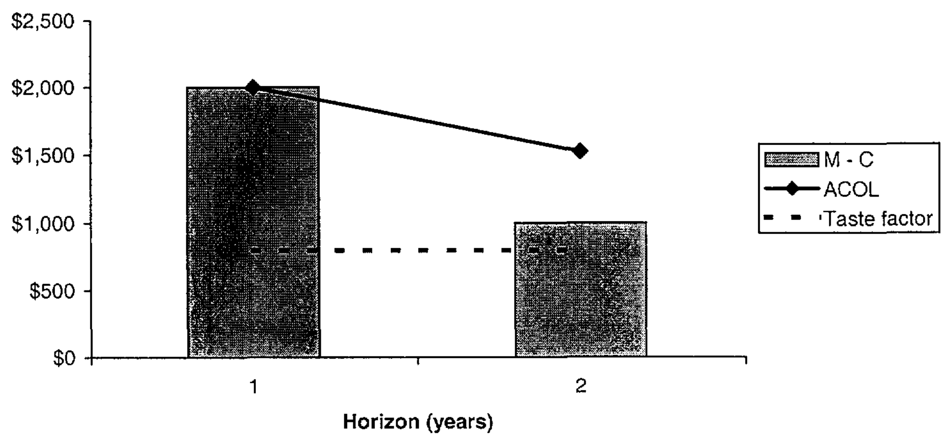

Conversely, suppose the military/civilian pay differences are \$1,000 in the first year and \$2,000 in the second year. In this case, the ACOL values are $ACOL_1$ = \$1,000 and $ACOL_2$ = \$1,476. The optimal horizon over which ACOL is maximized is now $s = 2$ years. Yet the individual would prefer to leave the military after just one additional year, rather than two, if the taste factor is sufficiently large. Specifically, having already stayed for one additional year, the individual would prefer to leave before the second additional year if the taste factor is greater than the pay difference in that year, \$2,000. It remains true that, because the taste factor exceeds the maximum ACOL, the individual would most prefer to leave the military immediately. Our point, however, is that among the various career lengths that involve staying, the so-called optimal horizon is not necessarily the most preferred. We show this situation in figure 2.

Intuitively, a comparison of the $ACOL_s$ values among the various horizons $\{s = 1, 2, 3, \ldots\}$ cannot determine the optimal leaving date because ACOL does not account for the taste factor, only the relative earnings. An individual may choose to remain until later, despite a decreasing sequence of ACOL values, because he or she has a net preference for military life (i.e., a sufficiently small taste factor).

Conversely, a person may choose to leave sooner despite an increasing sequence of ACOL values, because the taste factor is overwhelmingly large.

Figure 2. Second counterexample to optimality of ACOL time horizon

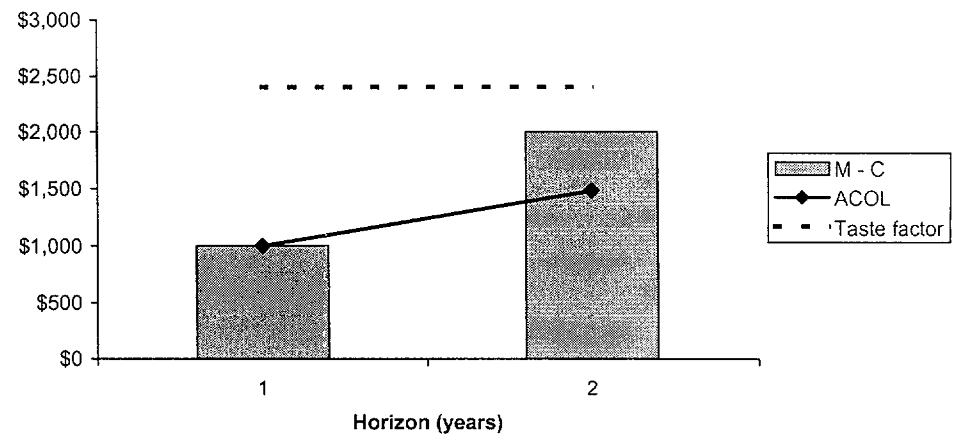

Daula and Moffitt [24] pointed out that, even if the taste factor is identically zero, the optimal horizon that maximizes ACOL may differ from the horizon that maximizes the discounted present value of earnings. Returning to the first example, suppose that military earnings are \$10,000 in both years. With the stated differentials, civilian earnings are \$8,000 in the first year and \$9,000 in the second year. The discounted present values (again using a 10-percent discount rate) are \$16,182 for leaving immediately, \$18,182 for staying one additional year and then leaving, and \$19,091 for staying two additional years. In this example, ACOL is maximized at $s = 1$, yet the discounted present value of earnings is maximized at $s = 2$. With the assumed zero taste factor, the individual would prefer to stay for the second year in order to maximize discounted earnings. He or she would be undeterred by the decline in ACOL values from $ACOL_1 = \$2,000$ to $ACOL_2 = \$1,524$.

As a technical matter, the ACOL calculation truncates the military and civilian earnings streams after s years. However, the discounted present value of earnings is calculated through a predetermined horizon—in practice, through an individual's entire working life, or even longer if retirement pay is considered. Because it is truncated, ACOL is *not* a monotonic transformation of the discounted present value over the predetermined horizon. Thus, the two expressions could easily achieve their respective maxima at different values of s.

None of these arguments vitiate the use of maximized ACOL to predict the individual's retention decision (although we will soon consider some different arguments against the ACOL approach). But the arguments do militate against labeling as "optimal" the horizon over which ACOL is maximized.

Statistical estimation of the ACOL model

If the distribution of the taste factor across decision-makers is normal, the probability of staying in the military follows a probit model. If the distribution of the taste factor is logistic, the probability of staying follows a logit model. Both of these models take the form of S-curves, so that the estimated probability of staying increases up to a limit of 1.0 as conditions become more conducive to staying (e.g., as relative military compensation increases). Conversely, the probability of staying decreases to a limit of 0.0 as conditions become less conducive to staying. When properly calibrated, the probit and logit S-curves are virtually indistinguishable, although the logit model is somewhat simpler mathematically and easier to compute. Software is readily available to estimate both models.

The logit and probit models allow for the introduction of additional regressors, apart from the maximum ACOL, that help explain the retention decision. For example, the retention rate has been found to vary directly with the civilian unemployment rate. The retention rate is also related to personal characteristics, such as marital status, race, education, and mental group.

The older studies estimated first-term and second-term retention models completely independently of each other. Many studies used grouped data, but even studies that used individual (panel) data made no allowance for correlation over time in the taste factor for a given person. We will argue later that disregard for intertemporal correlation likely led to upward-biased estimates of the coefficient on the ACOL variable. As we will see, the ACOL–2 model imposes a permanent/transitory error structure in an effort to avoid this source of bias.

Independent of the ACOL developments, David Wise and his various collaborators developed an essentially equivalent model in their research on retirement from civilian-sector jobs. In particular, they independently discovered the "maximum ACOL" condition (our equation 3).[6] Operationally, the only difference from ACOL is that Wise specified a first-order autocorrelation (AR1) error structure when estimating sequential retention decisions using panel data.

Interestingly, for a time Wise seemed unaware of the connection between ACOL and his own research on civilian retirement. He was the discussant on Warner and Solon's [14] paper at an Army retention conference. Although the proceedings were published in 1991, the conference actually took place in 1989, at which time Wise must have been working on his paper with James Stock that would be published in 1990. Yet Wise [28, p. 278] made the following comment on Warner and Solon, indicating his apparent lack of familiarity with the ACOL concept:

> the ACOL variable should be explained briefly in [Warner and Solon's] paper. The authors refer the reader to explanations presented in other project reports. But the variable plays a key role in the analysis; several of the other variables that are included make little sense if the reader does not understand what the ACOL variable is supposed to capture.

The two strands in the literature were finally brought together by Lumsdaine, Stock and Wise [23], some 3 years after the Army retention conference; further developments were contained in Daula and Moffitt [24].

---

6 Stock and Wise [22], equations 2.12 through 2.14 on p. 1162; or Lumsdaine, Stock, and Wise [23], equation 10 on p. 27.

Panel probit models

Critics of the ACOL approach point to its poor treatment of the dynamics of retention over a person's military career. The ACOL values often increase over one's career, as fewer years remain until retirement and the discounted value of the retirement annuity dramatically increases. According to a strict interpretation of the ACOL approach, anyone who stayed at the first decision point would certainly stay at all subsequent decision points because the taste factor is assumed time-invariant yet the financial incentive to stay (as measured by the ACOL value) increases with time. As an empirical matter, however, we know that retention rates at the second and third decision points are significantly below 1.0.

To develop a second criticism, consider a person who would have left the military after one term of service except for the lure of an SRB. This person has a larger taste factor (i.e., a larger distaste for military life) than others who would have stayed even absent an SRB. Unless the SRB is sustained, bonus-induced people are less likely to remain in the military at the second and subsequent decision points.

As an example, suppose the person had a taste factor of $2,000 and a first-term baseline ACOL of $1,000, but was offered an SRB that raised ACOL to $3,000. This person would stay through the first decision point because ACOL ($3,000, including the SRB) exceeds the taste factor ($2,000). However, the same individual would leave at the second decision point unless a sustained SRB or other compensation incentive raised ACOL above the baseline value of $1,000 to some value exceeding the (time-invariant) taste factor of $2,000. By contrast, a non-bonus-induced person would stay at the second decision point absent any compensation incentives. The latter individual, by definition, had a taste factor less than the baseline ACOL value of $1,000. This person would stay at the second decision point because the taste factor is time-invariant whereas ACOL tends, if anything, to increase as retirement approaches.

We see that the second-term reenlistment rate depends on the circumstances under which a person survived the first-term reenlistment decision. In an effort to capture this effect, Warner and Simon [29] included the lagged first-term ACOL value in a model to predict the second-term reenlistment rate. Along similar lines, Goldberg and Warner [30] include the lagged first-term SRB multiple in the second-term reenlistment model. The effect of lagged SRB was marginally significant with an unexpected positive sign for one occupational group (Electronics), and highly significant with the expected negative sign for one other occupational group (Non-electronics). Despite the names of these two groups, they are *not* mutually exhaustive. Goldberg and Warner's taxonomy contained six other occupational groups, for which the lagged SRB effect was statistically insignificant.

ACOL–2 model

The ACOL–2 model was an attempt to improve on ad hoc inclusion of lagged variables in second-term reenlistment models. Black, Hogan, and Sylwester [31] used the ACOL–2 model to predict retention decisions of Navy enlisted personnel. Black, Moffitt, and Warner [32] applied the model to retention decisions of Department of Defense (DoD) civilian employees. The ACOL–2 model was further developed in a dialogue between the latter authors and Glenn Gotz [25], and in a subsequent paper on Army reenlistments by Daula and Moffitt [24].

The ACOL–2 model follows a long tradition in the literature on panel data. Specifically, the taste factor for each person is decomposed into (a) a permanent component, constant over time through all decision points, and (b) a transitory component, randomly varying over time from one decision point to another. This permanent/transitory structure has several advantages. First, the retention rate is no longer predicted as 1.0 at the second and third decision points. Returning to the example above, the person who stayed at the first decision point might choose to leave at the second decision point, if the transitory component of the taste factor were sufficiently positive. Several events, such as an unusually arduous tour of duty or failure to receive an expected promotion, could "sour" a person at the second decision point. This effect might offset the general tendency for ACOL to increase over the individual's career, causing him or her to leave the military at the second decision point.

Simply pooling retention data from several decision points, without imposing a permanent/transitory structure, would lead to an upward-biased estimate of the ACOL coefficient. We have noted both the general tendency for ACOL to increase over an individual's career, and the general tendency for retention rates to increase (though not all the way to 1.0). Suppose that the first- and second-term data were pooled, but the two decisions for each person were treated as statistically independent. Then the entire increase in retention rates would be attributed to the increase in ACOL, leading to a large ACOL coefficient. In fact, however, part of the increase in retention rates results from the early departure from the sample of people with a stronger distaste for the military. Put differently, the ACOL coefficient would pick up not only the effect of changes in relative compensation on a fixed population, but also changes in the population composition itself. This phenomenon, known as "unobserved heterogeneity," leads to biased coefficient estimates.

Note that unobserved heterogeneity would not lead to any bias in the ACOL coefficient estimated from a single cross-section of first-term reenlistment decisions. Nor would there be any bias if data were pooled on first-term reenlistment decisions made by different cohorts of individuals in consecutive fiscal years. Instead, the bias arises from the failure of the simple ACOL model to adequately track a cohort (or cohorts) of individuals through successive decision points. Thus, the bias would be manifest in simple ACOL models only when applied at the second-term (or later) decision points.

The ACOL–2 model avoids the problem of unobserved heterogeneity by explicitly tracking the permanent taste distribution as a given cohort advances through successive decision points. At each decision point, the main forcing variable is again the maximum ACOL over all possible horizons. Suppose, for example, that the first reenlistment decision occurs in 1990 after 4 years of service, and the second reenlistment decision occurs in 1994 after 8 years of service. Then the first-term reenlistment decision is driven by the maximum ACOL over the horizons of staying 1 additional year up to 26 additional years (assuming mandatory retirement after 30 years of service). For the second-term reenlistment decision, ACOL is recomputed over the horizons of staying 1 additional year up to 22 additional years. Both ACOL values are computed using data from the fiscal years in which the respective decisions were made (e.g., a person's first-term decision might be modeled using the military and civilian wages that prevailed in 1990, but then the second-term decision would be modeled using the wages that prevailed in 1994). Thus, the model captures not only a person's progression through a fixed military pay table but also any growth over time in the military pay table or in civilian wages.[7] The ACOL–2 model also allows additional regressors, such as the civilian unemployment rate. This variable, too, is measured contemporaneously with the decision years, thus capturing additional information on trends in the civilian economy.

Black, Moffitt, and Warner [32] applied the ACOL–2 model to separation decisions of DoD civilian employees. Because estimation of the ACOL–2 model requires numerical integration of the multivariate normal density, they achieved a considerable computational efficiency by adopting a likelihood-factorization technique previously developed by Butler and Moffitt [33]. Glenn Gotz [25] wrote a comment on Black, Moffitt, and Warner, to which they immediately responded. Some of Gotz's points apparently spurred Robert Moffitt and his various collaborators to further improve on the ACOL–2 formulation.

---

7 For example, in their study of Navy enlisted retention, Black, Hogan, and Sylwester [31] reported that the sample average ACOL value doubled (in constant dollars) from the first-term to the second-term decision point. The average ACOL value nearly doubled again from the second-term to the third-term decision point.

In his comment, Gotz [25, p. 266] makes the following statement:

> Recall that associated with [the ACOL variable] is a known optimal future quitting date [*sic*], $t + s^*$....By construction of [Black, Moffitt and Warner's] model, any reduction in civil service pay more than s^* years from t [i.e., beyond the "optimal future quitting date"] will have absolutely no effect on the predicted quit rate at t.

Gotz's statement is too severe. When simulating a policy change, knowledgeable users of the ACOL model always recompute the sequence of ACOL values and locate the new maximum ACOL value. Consider, for example, an increase in military retirement pay, and suppose that the individual's horizon was initially 4 years ahead ($t + 4$). Gotz's statement implies that the horizon would remain fixed at $t + 4$ and, thus, the increase in retirement pay would have no effect on retention. In fact, the horizon might easily move out to year 20, so that retirement pay now enters the calculation and affects retention.[8]

Figure 3 illustrates this situation for a first-term decision-maker. In the base case, ACOL is maximized over the horizon of a 4-year reenlistment. The prospect of retirement pay after 20 years causes a jump in the ACOL value to nearly $4,500 at YOS 20, but that value still lies below the maximum ACOL of $5,000. Now consider an increase in the present value of retirement pay, equal to $100,000 when discounted to the date of retirement. The ACOL value jumps to almost $7,000 at YOS 20, so the ACOL horizon now encompasses the 20-year retirement point. The increase in the maximum ACOL from $5,000 to $7,000 provides a substantial retention incentive, even though the underlying change in compensation takes place beyond the initial ACOL horizon.

Paradoxically, Gotz and his collaborators had already recognized this point 5 years earlier, although, like many others, they misinterpreted the ACOL horizon as the planned leave point. According to Fernandez, Gotz, and Bell [34, p. 16]:

---

8 Indeed, recalculation of the ACOL horizon was included in the forecasting version of the ACOL model developed by John Warner in the early 1980s.

the calculated ACOL for any particular decision point reflects a specific horizon, the planned leave point [*sic*] for the marginal individual. Changes in earnings beyond that horizon generally do not affect the [maximum] ACOL value, and so cannot change the model's retention predictions for earlier decision points. Only an increase in military earnings (or decrease in potential civilian earnings) large enough to move the horizon outward can have any effect.

Figure 3. Example of shift in ACOL time horizon

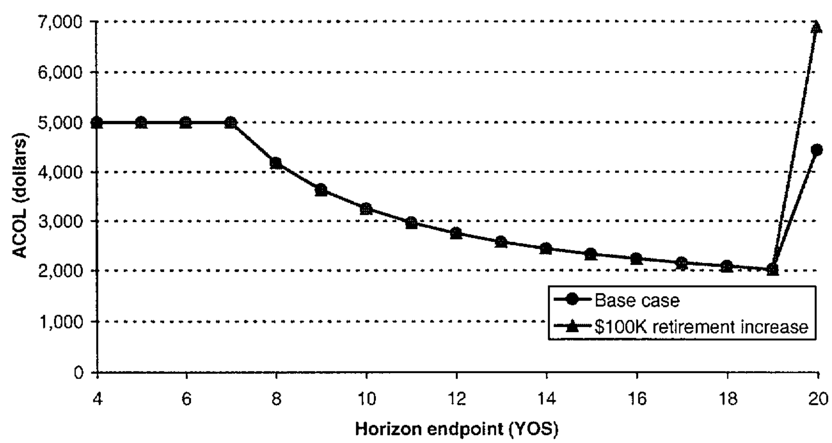

It was clearly the intention of Black, Moffitt, and Warner [32, pp. 258-259] that the maximum ACOL be recalculated after a policy change:

> To incorporate [the effects of a policy change] a new set of [ACOL] values must be calculated and a [maximum] selected for each individual in the file. The recalculated [maximum ACOL] is then inserted into the quit model, along with the other variables and their respective parameters, to obtain a simulated pattern of quit rates.

Other authors, such as Daula and Moffitt [24, p. 520], recognized the need to recalculate the maximum ACOL after a policy change, though again mislabeling the ACOL horizon as "optimal":

> To construct the...ACOL forecasts...would require recalculating optimal leaving dates [*sic*] at every date in the future (each of which requires rechecking all possible future leaving dates at each future date).

Dynamic-programming models

Along with John McCall, Glenn Gotz had developed a dynamic-programming model of Air Force officer retention [9, 10]. Their approach was particularly well suited to modeling officer retention because it offered the individual an opportunity to leave the military during every future year. Although military officers certainly face minimum service requirements, their mid-career commitments are usually less rigid than the typical 4-year terms served by enlisted personnel. Gotz and McCall were also very careful in modeling alternative promotion paths, capturing the adverse retention effect of being passed over for promotion.

Unfortunately, Gotz and McCall's formulation was computationally intensive, especially given the computer hardware and software environment of the early 1980s. They were able to estimate only three model parameters: the mean and standard deviation of the permanent taste factor, and the standard deviation of the transitory taste factor (the latter factor has a mean of zero by assumption). In particular, they did not estimate the effects of other regressors, such as the unemployment rate or various personal characteristics. Nor did they estimate the discount rate, which they fixed a priori. Finally, they were unable to estimate the standard errors of the three model parameters.[9]

Moffitt and his collaborators took some lessons from Gotz and went on to develop a dynamic-programming model of their own. Their approach was crystallized in an impressive paper by Daula and Moffitt [24]. Recall that the simple ACOL model summarizes the military and civilian pay streams with a single discounting calculation over the

9 A simple approximation was developed by Warner [35, pp. 27-28], who fit the three model parameters to the cross-sectional survival profile (by term of service) that prevailed in the Navy enlisted force in FY 1979. Using a grid search, Warner estimated the mean permanent taste factor as $2,800 (in FY 1979 dollars), the standard deviation of the permanent taste factor as $3,500, and the standard deviation of the transitory taste factor as $4,500. However, Warner reported that his objective function was extremely flat, so that many alternative sets of parameter values fit the data about equally as well.

dominant optimal horizon. The ACOL–2 model tracks individuals through time, using contemporaneous pay streams to update the ACOL calculation at each decision point. Thus, under ACOL–2 there is a single, dominant horizon at the first-term decision point; a single (generally different) dominant horizon at the second-term decision point; and so on. These calculations are illustrated in figure 4, where the dominant horizon shifts from YOS 7 when evaluated at the first-term decision point to YOS 20 when reevaluated at the second-term decision point.

Figure 4. Example of recalculation of dominant time horizon

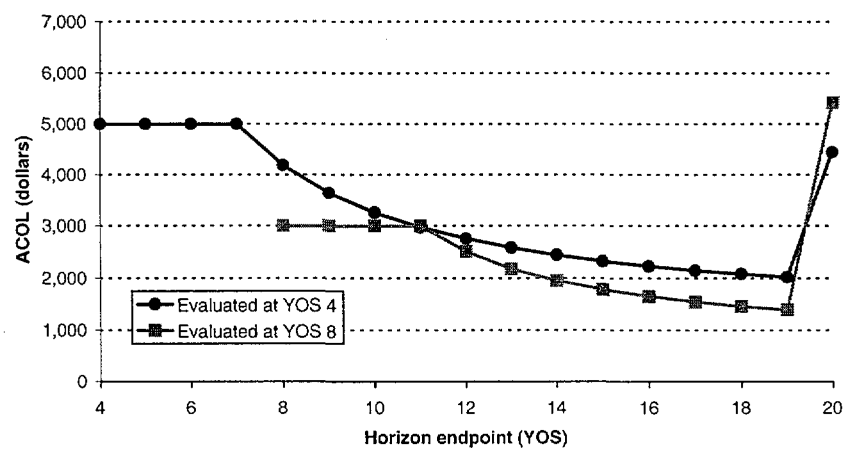

By contrast, at any particular decision point, Daula and Moffitt probabilistically weight the discounted pay differences over *all* future leaving points. Thus, there is no longer a single, dominant horizon.[10] In addition, Daula and Moffitt were more careful in their specification of the error terms than had been Black, Moffitt, and Warner [32]. Finally, they estimated their model by embedding the dynamic program inside the panel probit approach of Butler and Moffitt [33].

---

10 The equivalence between dynamic programming and probabilistic weighting in this context had previously been established by Warner [35]. Further theoretical developments along these lines are found in Hotz and Miller [36].

Daula and Moffitt [24] touted the ease with which their estimates were computed: "we show that dynamic retention models are considerably less difficult to estimate than [the] literature implies" (p. 500); "estimation of the model in this form is not difficult...no difficult calculations are involved" (p. 503); and "since the single-period model is not overly burdensome itself, its multiple evaluation [using panel data] is still well within the power of modern computational facilities" (p. 507). However, they later conceded that estimation took about 450 CPU minutes per iteration, and six or seven iterations per model run (p. 514). Thus, each model run took about 48 hours—hardly an improvement over Gotz and McCall.

For comparison purposes, Daula and Moffitt also estimated the ACOL–2 model using the bivariate probit technique.[11] Interestingly, they report that the log-likelihood value is slightly better for the ACOL–2 model than for their dynamic-programming model. In light of the computational difficulty of the latter (notwithstanding the authors' statements to the contrary), the ACOL–2 model becomes an extremely compelling alternative.

[11] As Daula and Moffitt correctly point out, multivariate probit is equivalent to Butler and Moffitt's panel probit technique. The latter was developed primarily for long panels spanning three or more decision points, to avoid numerical integration of the trivariate (or higher order) normal density. These days, both techniques are available in the LIMDEP package developed by Econometric Software, Inc. (www.limdep.com). In fact, LIMDEP is advertised as being able to estimate the multivariate probit model with up to 20 correlated decisions, though one must be skeptical about the computational speed of such high-dimensional models. Also, it should be possible to program the panel probit model in PROC NLMIXED of SAS.

THIS PAGE INTENTIONALLY LEFT BLANK

Conditional logit models

More detailed models partition the event "staying" into reenlistments and extensions. Reenlistments are defined as commitments to stay in the military for 36 months or longer, whereas extensions are defined as commitments to stay for fewer than 36 months. The distinction between reenlistments and extensions is clearly important for personnel planning purposes. There are also behavioral differences, because only those who reenlist are eligible to receive SRBs. We would expect an increase in the SRB level to increase the total probability of staying. Underlying that effect, we would expect an increase in the SRB level to reduce the probability of extending but to increase the probability of reenlisting by a larger magnitude.

Various models are available to estimate the three probabilities of reenlisting, extending, or leaving. One approach, the conditional logit model, was pursued by Goldberg and Warner [30] and Goldberg [11]. These authors collected data on reenlistment, extension, and separation rates in cells defined by fiscal year, Navy enlisted rating, and years of service (in the range of 3 to 6 years). They computed a discounted pay stream associated with each of the three choices for the "typical" sailor in each cell. In particular, the pay stream associated with reenlistment contained the SRB, whereas the pay stream associated with extension did not. Their models contained background variables, including the civilian unemployment rate, marital status (i.e., percentage married in each cell), race, education, and mental group. They estimated coefficients from which one can compute the marginal effect of each background variable on the three choice probabilities.

Goldberg and Warner also estimated a single pay coefficient, interpretable as the "marginal utility of income." Using this coefficient, one can compute the reallocation of the three choice probabilities in response to a change in the discounted pay stream associated with one or more of the three choices. For example, a change in the SRB level affects only the pay stream associated with reenlistment (which we denote as M), but affects all three choice probabilities as follows:

$$\partial P_R/\partial M = b\, P_R\, (1 - P_R),$$
$$\partial P_E/\partial M = -\, b\, P_E\, P_R, \qquad (5)$$
$$\partial P_L/\partial M = -\, b\, P_L\, P_R,$$

where b is the pay coefficient and P_R, P_E, and P_L are the respective probabilities of reenlisting, extending, and leaving.

Hogan and Black [37, p. 41] opine that,

> The conditional logit model...is a poor choice in the analysis of extensions versus reenlistments because it constrains reenlistment bonuses to reduce extensions by the same percentage that it reduces losses.

Their statement of this mathematical property of the conditional logit model is correct; in terms of percentage changes:

$$(\partial P_E/\partial M)/P_E = -\, b\, P_R = (\partial P_L/\partial M)/P_L. \qquad (6)$$

Hogan and Black argue that reenlisting and extending are closer substitutes than are reenlisting and leaving. If that were the case, an increase in the SRB level would draw more reenlistments from those who otherwise would have extended, rather than from those who otherwise would have left. Thus, one might prefer an alternative model with the following mathematical property:

$$(\partial P_E/\partial M)/P_E < (\partial P_L/\partial M)/P_L < 0. \qquad (7)$$

Logit models with correlated taste factors

Alternative models, satisfying the Hogan and Black critique, may be formulated by returning to the theoretical underpinnings of occupational choice. For this purpose, we change the notation slightly so that each choice has its own taste factor. Thus, v_R is the monetary equivalent of the nonmonetary factors associated with reenlisting; v_E

and v_L are defined similarly for extending and leaving. The single "taste factor" in the earlier discussion would be interpreted as $v = v_L - v_R$, the net preference for civilian life.

McFadden [38] showed that the conditional logit model arises when the taste factors are independent across choices, each with an extreme-value distribution. It is not as well appreciated that the conditional logit model also arises when the taste factors have Gumbel's multivariate logistic distribution, with correlations of 0.5 between each pair of taste factors.[12] In either case, Hogan and Black's critique comes into greater focus. Suppose that reenlisting and extending are indeed closer substitutes than either of the other two pairs of choices. If so, the correlation between the taste factors for reenlisting and extending should be larger than for the other two pairs of choices. For example, because reenlisting and extending are more similar, an individual who requires an above-average premium for reenlisting rather than leaving should also require an above-average premium for extending rather than leaving. In other words, the taste factors for these two choices should have a particularly high correlation. However, the conditional logit model implicitly assumes equal correlations between all three pairs of choices.[13]

---

12 This result is due to Goldberg [11, pp. 80-81]; Gumbel's multivariate logistic distribution is described in Johnson and Kotz [39, pp. 291-293]. The oft-cited converse to McFadden's theorem states that, *if* the taste factors are independent and the choice probabilities are of the logit form, then the taste factors are extreme-value distributed. However, the latter result does *not* rule out correlated taste factors.

13 Another expression of the difficulty with the conditional logit model is the "independence of irrelevant alternatives." In our example, the relative probability of extending versus leaving depends on only the background variables and the discounted pay streams associated with these two choices. The relative probability does *not* depend on the pay stream associated with reenlisting. In particular, it does not depend on the SRB level. Yet a person who extends could subsequently choose to reenlist and thereby receive an SRB. Thus, for many, the SRB level is an important determinant of the decision to extend versus leave.

Nested logit model

McFadden's [40, 41] nested logit model allows for unequal correlations. Under this model, the taste factors associated with reenlisting and extending have a bivariate extreme-value distribution, with a correlation coefficient that is free to vary in the range of 0 to +1. The taste factor associated with leaving has a univariate extreme-value distribution and is independent of the other two taste factors. In the special case where the correlation coefficient equals zero, the three taste factors are all independent extreme-value distributed, and the conditional logit model results. If the correlation coefficient is positive, however, the probability equations differ from those of the conditional logit model. In particular, the probability equations contain the correlation coefficient as a free parameter.

During the mid-1980s, Goldberg and Warner attempted to apply the nested logit model to grouped data on first-term reenlistment decisions in the Navy. Goldberg and Warner never published their results because they could not achieve convergence to reasonable parameter estimates. Nor could Mackin et al. [27] using microdata on individual Navy sailors.

Even when using microdata, there are two approaches to estimating the nested logit model. The first proceeds in two stages: (1) a logit model is estimated among individuals who stay, to predict probability of reenlisting versus extending, and (2) another logit model is estimated to predict the probability of staying (i.e., either reenlisting or extending) versus leaving. However, the second stage is not a standard logit model. Instead, it contains an additional variable, known as the "inclusive value," that must be constructed based on the results of the first stage. To avoid model failure due to multicollinearity, the inclusive value must be computed from at least some variables that are absent from the second-stage (stay/leave) model. That is, there must be some variables that drive the reenlist/extend decision but *not* the stay/leave decision. Mackin et al. opine that, because individual decision-makers ultimately compare all three choices simultaneously, the required identifying variables simply do not exist.[14]

The second approach to estimating the nested logit model is full-information maximum-likelihood. This approach has only recently become available using commercial software.[15] It would be interesting to apply this approach to the retention decision, to determine whether it circumvents the problem of multicollinearity.

Another expedient was attempted by Warner [42], using grouped data on first-term and second-term reenlistment decisions in the Marine Corps. Warner estimated sequential logit models, but simply omitted the inclusive value from the second-stage model. These sequential logit models had good explanatory power and produced reasonable estimates of the pay elasticities. However, it is not known what joint distribution of the taste factors, if any, would yield the sequential logit probability equations (without the inclusive value). Thus, Warner's approach, though pragmatic, does not have strong theoretical underpinnings.

---

[14] Theoretically, the nested logit model is identified because the inclusive value is a nonlinear construct, thus not perfectly predictable from any linear combination of second-stage regressors. As a practical matter, however, the degree of nonlinearity may not be adequate to identify the model.

[15] LIMDEP version 7.0 includes this feature in the module *NLOGIT* version 2.0.

THIS PAGE INTENTIONALLY LEFT BLANK

Multinomial logit models

Recall that Goldberg and Warner [30] and Goldberg [11] computed a discounted pay stream associated with each of the three choices, and estimated a single pay coefficient interpretable as the "marginal utility of income." Thus, their models contain terms of the form $b\,M_R$, $b\,M_E$, and $b\,M_L$, where M_R, M_E, and M_L are the respective pay streams. An alternative approach is to enter the pay variables in the same manner as the background variables. Recall that a background variable, such as the unemployment rate, affects the probabilities of all three choices. Three separate coefficients are estimated, from which one can compute the effect of a change in the unemployment rate on each choice probability. Similarly, one could enter a pay variable, such as the SRB multiple or dollar amount, as a background variable, and compute its effect on each choice probability. Thus, the alternative model would contain terms of the form b_R SRB, b_E SRB, and b_L SRB.

In the econometrics literature, logit models in which the coefficients are fixed across choices, but the regressors vary, are known as "conditional logit models." By contrast, logit models in which the regressors are fixed across choices, but the coefficients vary, are known as "multinomial logit models."[16]

The multinomial logit model satisfies the Hogan and Black [37] critique and breaks the "independence of irrelevant alternatives." Under the multinomial logit model, the relative probability of extending versus leaving is sensitive to the pay stream associated with reenlisting,

---

16 The term "conditional logit model" is unfortunate because it is not clear, in any statistical sense, which variable is conditional on which other variable. Nor is it clear why one model is "conditional" and the other (multinomial) model is, presumably, "unconditional." More recently, Greene [43] has suggested the terminology "discrete choice model" to replace "conditional logit model."

particularly the SRB level. The extension and separation rates change by (possibly) different percentages in response to an SRB increase:

$$(\partial P_E/\partial \mathrm{SRB})/P_E - (\partial P_L/\partial \mathrm{SRB})/P_L = b_E - b_L. \tag{8}$$

This difference is precisely the coefficient on SRB in an extend/leave log-odds model, and is a free parameter that may be of either algebraic sign (not necessarily zero).

Hosek and Peterson [12] and Lakhani and Gilroy [44] estimated multinomial logit models. Reference [12, appendix B] reports negative coefficients on the SRB level in extend/leave log-odds models, at both the first-term and second-term decision points. These results indicate that P_E/P_L declines with increases in the SRB level, so that increased bonuses draw more reenlistments from those who otherwise would have extended than from those who otherwise would have left.

Figure 5 shows the difference between the conditional and multinomial logit models. Under the conditional logit model, a hypothetical SRB increase causes the probabilities of extending and leaving to both decrease by 20 percent. Under the multinomial logit model, more reenlistments are drawn from those who would have extended, so the extension probability decreases even more severely but the separation probability decreases less severely.

Figure 5. Reallocation of probabilities when SRB level increases

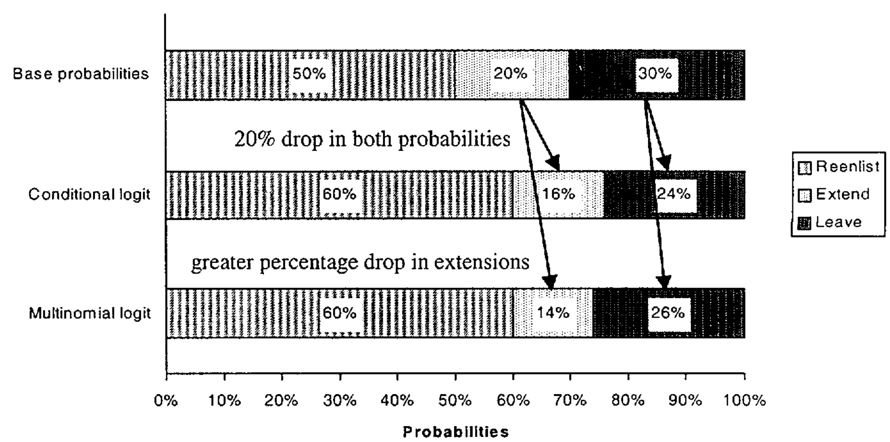

Hogan [45] cautions that the pay coefficients b_R, b_E, and b_L in the multinomial logit model are not the same as the partial effects and, further, may even differ in sign from the partial effects. For example, the partial effect of the SRB level on the reenlistment rate is given by:

$$\partial P_R/\partial \mathrm{SRB} = (b_R - b_L)\, P_R\,(1 - P_R) \;-\; (b_E - b_L)\, P_E\, P_R,$$
$$= P_R\,[b_R\,(1 - P_R) \;-\; b_E\,P_E \;-\; b_L\,P_L]. \qquad (9)$$

This expression will differ in sign from b_R if b_R, b_E, and b_L all have the same sign, but b_R has the smallest magnitude and P_R is close to 1.0. Moreover, the standard error of $\partial P_R/\partial \mathrm{SRB}$ is not immediately available from those of b_R, b_E, and b_L, but can be derived from the underlying variances and covariances.[17]

Interpretation of the multinomial logit model

Although the multinomial logit model breaks the "independence of irrelevant alternatives," it leads to other problems of interpretation. We argued earlier in favor of the ACOL approach, which combines all the elements of compensation into a single measure. The multinomial logit models, as estimated by Hosek and Peterson [12] and Lakhani and Gilroy [44], do not use a single measure of compensation. Instead, they segment compensation into two measures: the SRB level and an index of relative military pay.

Lakhani and Gilroy seem to believe that, if the ACOL approach were correct, segmenting compensation into multiple measures should produce equal elasticities on all of the measures. Conversely, if the elasticities prove unequal, the compensation measures should remain distinct rather than being combined into a single ACOL variable.

Lakhani and Gilroy report that the SRB elasticities across Army occupations are, if anything, negatively correlated with the relative-pay elasticities. They interpret this finding as evidence against the ACOL approach, concluding [44, p. 241]:

---

[17] The formula is given in Hosek and Peterson [12, appendix C].

> It is, therefore, somewhat presumptuous to assume that the effect of SRB is the same as that of relative pay, as is often done in the existing literature: Their dollar values are added to retirement to represent the cost of leaving in the Annualized Cost of Leaving (ACOL) model.

We will now argue that, based on economic theory, there is no reason to expect a positive correlation between the SRB and relative-pay elasticities. Further, we will argue that the ACOL approach can rationalize the observed negative correlation if there is, in turn, a positive correlation between SRB levels and civilian earnings opportunities. This will be the case if military occupations with superior civilian alternatives have chronically poor retention, and if SRBs are used to combat these retention problems. Thus, the observed negative correlation between the two elasticities, rather than being a paradox that vitiates the ACOL approach, is actually quite consistent with that approach.

To simplify the algebra, suppose that the difference between military pay (RMC) and civilian pay is constant over the individual's planning horizon. We denote the annual difference as $(M–C)$. The ACOL variable equals this quantity plus the annualized bonus. Again, to simplify the algebra, assume a lump-sum SRB. The annualized value of the SRB, over an s–year horizon, is given by:[18]

$$\text{SRB} \Big/ \sum_{j=1}^{s} (1+r)^{-j} = \text{SRB} / D. \tag{10}$$

Thus, ACOL is given by:

$$\text{ACOL} = (M - C) + \text{SRB} \;/\; D. \tag{11}$$

[18] For example, consider a lump-sum bonus, a 4–year planning horizon, and a 10-percent discount rate. Under these assumptions, the annualized bonus evaluates at 0.315 × SRB. If this amount were paid at the end of each year over the 4–year planning horizon, the *undiscounted* total payment would be 1.26 × SRB, but the *discounted* total payment would be exactly 1.00 × SRB.

Suppose we have an estimate of the elasticity of the retention probability (P) with respect to ACOL:

$$E = (\text{ACOL} \,/\, P)\,(\partial P \,/\, \partial \text{ACOL}). \tag{12}$$

We can derive the elasticity of the retention probability with respect to the military/civilian pay difference:

$$\begin{aligned}
&[(M - C) \,/\, P]\ [\partial P \,/\, \partial (M - C)] \\
&\qquad = [(M - C) \,/\, P]\ [\partial P \,/\, \partial \text{ACOL}]\ [\partial \text{ACOL} \,/\, \partial (M - C)] \\
&\qquad = E \times (M - C) \,/\, \text{ACOL};
\end{aligned} \tag{13}$$

and with respect to the SRB amount:

$$\begin{aligned}
&[\text{SRB} \,/\, P]\ [\partial P \,/\, \partial \text{SRB}] \\
&\qquad = [\text{SRB} \,/\, P]\ [\partial P \,/\, \partial \text{ACOL}]\ [\partial \text{ACOL} \,/\, \partial \text{SRB}] \\
&\qquad = (E \times \text{SRB}) \,/\, (D \times \text{ACOL}).
\end{aligned} \tag{14}$$

Because all of the other terms are common, the correlation (across occupations) between the latter two elasticities is essentially the correlation between $(M - C)$ and SRB. Thus, the ACOL framework is consistent with a negative correlation if SRBs are employed to compensate for salary shortfalls in selected military occupations.

Conclusions

We conclude that both the conditional logit model and the multinomial logit model have considerable merit. The former is more firmly grounded in economic theory, combining all the elements of compensation into a single measure of discounted pay. A wealth-maximizing individual would make his or her reenlistment decision based on this single measure, and no information is added by partitioning it into multiple components. Moreover, the rich sample variation in SRBs can be brought to bear in estimating the single pay coefficient.

On the other hand, the conditional logit model suffers from the independence of irrelevant alternatives. The multinomial logit model relaxes this restrictive assumption. However, as just demonstrated, the elasticities on the multiple pay components must be interpreted with caution. Moreover, some of the elasticities may be underestimated for lack of sample variation (e.g., the index of relative military pay).

We agree with Hogan [45, p. 258], who states:

> In the trichotomous logit model specified by Lakhani and Gilroy, both the SRB and relative wage variables are identical across choices, while the coefficients on the variables vary across the alternatives. Hosek and Peterson (1985) also specified the logit model in this way, whereas Goldberg (1984) constrained the coefficients to be the same and varied the level of the independent variable across choices. It is not clear to me which specification is preferable.

Finally, we will see later (in table 2) that the pay effects estimated from the two models are quite similar. Thus, a stark choice between the two models is not entirely necessary.

Reverse causation between bonuses and the reenlistment rate

Goldberg [11] and Hosek and Peterson [12] were concerned about reverse causation in the relationship between pay and the reenlistment rate. The goal of the analysis is to estimate the positive effect of pay, particularly reenlistment bonuses, in encouraging reenlistments. However, some enlisted occupations have suffered chronically poor retention because of arduous duty (e.g., Navy ratings with a high percentage of time at sea), slow promotions, or lucrative civilian opportunities. The enlisted occupations with chronically poor retention are generally awarded higher SRB levels. This pattern of reverse causation leads to a downward bias in the estimated effect of pay on the reenlistment rate.

Individual data

It is commonly believed that individual decision-makers are "price-takers" in the sense that, while their decisions may well be affected by SRB levels, their decisions do not, in turn, affect SRB levels. However, even individual data may be plagued by reverse causation that leads to biased estimation of the bonus effect.

Figure 6 illustrates the situation. We have drawn the supply and demand curves for reenlistments, both as a function of the SRB level. We have drawn the demand curve as a vertical line, to capture rigid personnel requirements that are insensitive to the price level. However, the analysis is virtually identical even if the demand curve exhibits some elasticity.

The military service attempts to set SRB levels to equate the supply of reenlistments to the desired level of demand. If the service errs on the side of too low an SRB level (e.g., SRB_1), too few people reenlist (point A), and a shortfall occurs (the distance AB). To alleviate the shortfall, SRB levels will have to be increased, either mid-year (if the problem is detected early enough and if funding is available) or during the following year. If the service sets too high an SRB level (e.g., SRB_2), too many people desire to reenlist, and a surplus occurs (the distance CD). In the latter instance, the service may suspend bonus payments partway through the fiscal year, only to resume payments at the start of the following fiscal year when new funding becomes available.

Figure 6. Supply and demand of reenlistments

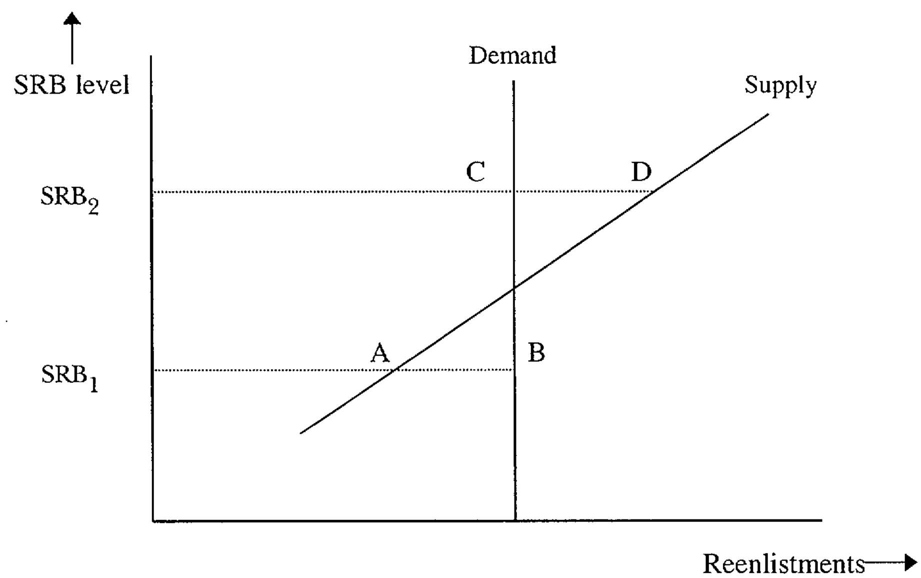

In either of these instances, a savvy person may wait until the start of the new fiscal year (executing a short extension, if necessary) in order to enjoy the restored (and possibly increased) SRB level. Thus, a degree of reverse causation exists in that the person's decision to wait until (or extend into) the new fiscal year may affect the SRB level that he or she is offered.

A careful analysis of this situation would examine the pattern of reenlistments and extensions, accounting for seasonality over the course of the fiscal year and, in particular, mid-year bonus freezes and adjustments. Econometric techniques for disequilibrium models could be applied, although such models require an auxiliary equation to determine whether a particular observation is drawn from the supply curve (e.g., point A) or the demand curve (e.g., point C).[19] This approach has never, to our knowledge, been applied to reenlistment models, but seems worthy of serious consideration.

Panel data

Reverse causation always presents an estimation problem when using grouped data because the collective reenlistment decisions of the group will feed back (albeit possibly with a lag) into the SRB levels that they are offered. However, the downward bias can be alleviated by applying a fixed-effect estimator. Under this approach, each enlisted occupation is assigned a dummy variable intended to capture permanent deviations between that occupation's reenlistment rate and the overall sample average. Computationally, it is not actually necessary to include the multiple dummy variables in the regression equation. Instead, an exactly equivalent approach is to measure each observation (both left-hand and right-hand variables) as a deviation from the sample average for that occupation across all of the time periods.[20]

Both Goldberg [11] and Hosek and Peterson [12] applied a fixed-effect estimator to grouped data when estimating the two log-odds equations for reenlist versus leave and extend versus leave. Hosek and Peterson (in their table 5) report that the SRB effect on the second-term probability of reenlistment is actually *negative* when estimated without the fixed effects. Incorporation of fixed effects restores the expected positive coefficient and considerably increases

---

[19] Disequilibrium estimation is discussed in Maddala [46, chapter 10]. These techniques have been successfully applied to distinguish supply-constrained from demand-constrained observations in enlisted recruiting models; see Daula and Smith [47] and Dertouzos [48].

[20] See Baltagi [49, pp. 9-13] or Hsiao [50, pp. 25-32]. Goldberg [11, p. 96] noticed that differencing around the occupational averages introduces both serial correlation and heteroskedasticity. However, Baltagi [51; 49, p. 23] has shown that applying generalized least squares (GLS), in an effort to circumvent these statistical problems, is equivalent to applying ordinary least squares (OLS) in this situation.

the magnitude of the (already) positive coefficient on the first-term probability of reenlistment.

Yet another alternative would be to explicitly model the SRB levels in an auxiliary regression equation. The SRB equation and the retention equations could then be jointly estimated by two-stage least squares. Although this approach does not appear to have been attempted, it, too, seems worthy of consideration.

Joint models of attrition and retention

We have already discussed the possibility that conditions at the first-term reenlistment point (e.g., SRB levels) may affect subsequent second-term reenlistment rates. Similarly, conditions at the accession point (e.g., the civilian unemployment rate, accession bonus levels) may affect subsequent first-term reenlistment rates. More generally, the probability of surviving to the first-term reenlistment point may be correlated with the outcome of that reenlistment decision. When several years of data are pooled, the various accession cohorts may differ in both the conditions that prevailed at their respective accession points and the resulting survival rates. The reenlistment model is designed to pick up the effects of changes in SRBs and other variables on a fixed population. However, the reenlistment model may be confounded if these variables are correlated with changes in the population composition itself.

One way to model attrition is as a binary outcome: the person either survives to the first-term reenlistment point, or does not. A variety of functional forms, such as logit and probit, may be used for this purpose. The logit and probit functions monotonically map a linear combination of regressors (in principle, taking on any real value, either positive or negative) into an attrition probability that is restricted to the unit interval. In fact, any cumulative density function (CDF) defined over the entire real line has the same property and, thus, potentially qualifies as a binary attrition model.

An alternative approach is to model attrition as a continuous-time process, and to attempt to predict the exact number of months of service at which a person attrites (if indeed he or she attrites at all within the sample period). For example, Baldwin and Daula [52] modeled Army first-term attrition using a Weibull distribution. Depending on the estimated shape parameter, the Weibull distribution implies that the hazard rate (i.e., the instantaneous probability of attrition) is (a) constant, (b) always increasing, or (c) always decreasing. In particular, the Weibull distribution does *not* allow the hazard rate to behave non-monotonically (i.e., first increase, then decrease; or first decrease, then increase).

Binary attrition models

The proportional hazards model is considerably more flexible than the Weibull distribution. The hazard rate depends on time (t) and a set of regressors (X) in the following manner:

$$h(t, X) = g(t) \times \exp[-(\alpha + X\beta)]. \tag{15}$$

In this formulation, $g(t)$ is a step function that may behave non-monotonically if so indicated by the data. Note also the sign convention: because of the minus sign inside the exponential, a positive coefficient β_i implies that an increase in the corresponding variable X_i serves to *reduce* the hazard rate, and thus *increase* the survival probability.

It is also interesting to consider the binary attrition model that results if the underlying hazard function follows the proportional hazards model. Suppose we have two month-of-service markers, $0 \leq t_a < t_b$. Given that a person is still on active duty at time t_a, the probability that he or she will remain on active duty at the later time t_b is given by:[21]

$$P(t_a, t_b) = \exp\left[-e^{-(\alpha + X\beta)} \times \int_{t_a}^{t_b} g(s)\,ds\right]. \tag{16}$$

Again, given our sign convention, a positive coefficient β_i implies that the corresponding variable X_i is positively related to the survival probability. Moreover, if we absorb the integrated hazard into the intercept α, the double-exponential form is actually the CDF for

[21] See Prentice and Gloeckler [53] or Kalbfleisch and Prentice [54, pp. 36-37, 98-99].

Gumbel's Type I extreme-value distribution[22] with argument $z = (\alpha + X\beta)$. This observation is consistent with our earlier assertion that any CDF qualifies as a binary attrition model.

This result implies that neither logit nor probit is the correct model for binary attrition under the proportional hazards assumption. Interestingly, however, Warner and Solon [14] showed that the logit model may be recovered if the intercept α is itself randomly distributed across individuals according to an exponential distribution[23] with mean 1.0.

More precisely, suppose that the integrated hazard in equation 16, $\int_{t_a}^{t_b} g(s)ds$, has a gamma distribution across individuals with mean 1.0 and variance θ^{-1}. The assumption of unit mean is innocuous because any non-unit portion can be absorbed into the intercept, α. Then the probability of surviving from t_a to t_b becomes:

$$P(t_a,t_b) = \left[1 + \theta^{-1} e^{-(\alpha+X\beta)}\right]^{-\theta}. \tag{17}$$

This probability reduces to a logit form if the variance θ^{-1} equals 1.0. However, a gamma distribution with both mean and variance of 1.0 is just a unit exponential distribution.

---

[22] See Johnson and Kotz [55, chapter 21] or Mann, Schafer and Singpurwalla [56, pp. 108-111]. This is the same distribution that McFadden [38] used to derive the conditional logit model. The context, however, was quite different. McFadden assumed that the taste factor associated with any particular outcome is distributed across individuals according to the extreme-value distribution; further, the distributions across outcomes are statistically independent. Maximization of utility then leads to a logistic probability of choosing any particular outcome. In the situation discussed in the text, an underlying proportional-hazards model implies an extreme-value (*not* logistic) survival probability over a discrete time interval. The latter result follows solely from the proportional hazards assumption, and bears no relationship to utility maximization.

[23] This result actually goes back to Dubey [57].

Warner and Solon estimated a model of this form to predict survival to the first-term reenlistment decision point (though *not* the exact months-of-service if the individual attrites), jointly with a probit model to predict first-term reenlistments among those surviving to the reenlistment decision point. They allowed for correlation between the disturbances in the two models, but found that the correlation was not statistically significant.

Continuous-time models of attrition and reenlistment

Yet another alternative would be to estimate a single, continuous-time model to predict the exact months-of-service at which a person:

- Attrites before the reenlistment decision point; or
- Survives to the reenlistment decision point, but decides to separate at that point; or
- Survives to the reenlistment decision point, and decides to reenlist at that point, but separates at some subsequent point [either before or at the individual's updated expiration of term-of-service (ETS)].

When estimating this type of model, one danger is that the proportional hazards assumption constrains the regressor effects to be the same in the attrition phase as in the reenlistment phase. For example, suppose that a particular dummy variable (perhaps identifying a particular demographic group) has a coefficient of 0.15. This coefficient indicates that members of the highlighted group are 14 percent less likely than members of the base group to separate at any point in time (conditional on survival to that point).[24] However, using separate attrition and reenlistment models, Warner and Solon found that high school graduates (HSGs) are more likely than nongraduates to survive to the reenlistment decision point, but less likely than nongraduates to

---

[24] A coefficient of 0.15 serves to scale down the hazard function by the factor $e^{-0.15} = .86$, a 14-percent reduction. It is convenient—but imprecise—to associate, for example, a coefficient of 0.15 with a 15-percent reduction.

reenlist. Clearly, no single coefficient on HSG status could capture both effects.

Follmann, Goldberg, and May [13, 58] addressed this problem by developing a continuous-time model with different regressor effects along different intervals of the time axis. In their example, they modeled the probability of an unemployed worker finding a job as a function of the duration of his or her unemployment spell. A disproportionate number of unemployed workers find jobs during the week that their unemployment insurance (UI) benefits expire, typically the 26th or 39th week of the unemployment spell. This phenomenon could apparently be accommodated using a proportional hazards model with a nonmonotonic step function, $g(t)$. However, the conventional proportional-hazards model constrains the regressor effects to be the same throughout the entire time axis.

Instead, Follmann, Goldberg, and May modeled the probability of finding a job as an extreme-value regression (equation 16) during the week that UI benefits expire. They modeled the probability of finding a job using a conventional proportional-hazards model over the remainder of the time axis. Importantly, they allowed for different regressor effects over the two time intervals. Indeed, they found that college graduates residing in counties with low unemployment rates were most likely to find jobs during the week their UI benefits expire; but that older, white workers were more likely to find jobs at all other times.

Although similar in spirit, there are some important differences between the approaches of Follmann et al. and Warner and Solon. Follmann et al. define the "spike event" as the period during which a disproportionate number of transitions occur. In their example, the spike event was the week in which UI benefits expire. If applied to attrition and reenlistment in the military, the spike event would be the period immediately preceding ETS. They modeled the probability of a spike event using extreme-value regression, although other models, such as logit or probit could have been used instead. They modeled events away from the spike using a continuous-time, proportional-hazards model. The advantage of using extreme-value regression is that the coefficient vector β in equation 16 is commensurable with the corresponding coefficient vector in the proportional-hazards model (equation 15) away from the spike. Although Follmann et al. allowed the two coefficient vectors to differ, it remains meaningful to statistically test their equality. A similar result would hold if generalized logit regression (equation 17) were used in place of extreme-value regression at the spike. However, a probit coefficient vector at the spike would be scaled differently from the coefficients in the proportional-hazards model, and the two could not easily be tested for equality.

Warner and Solon modeled their spike event—reenlistment during the period immediately preceding ETS—using probit regression. They modeled attrition away from the spike as a binary outcome using, variously, probit, logit, or generalized logit regression (i.e., equation 17). However, they did *not* estimate the entire, continuous-time hazard function, $g(t)$. Instead, they estimated the annual survival rates for the baseline demographic group (i.e., the group with regressors $X = 0$). The annual survival rates are the values of expression 17 when evaluated at successive annual intervals; i.e., $(t_a, t_b) = \{(0,1),(1,2),(2,3),(3,4),\}$.[25] Thus, Warner and Solon claim too much when they state [14, p. 263], "The main advantage of the proportional hazard results is that they trace out the temporal pattern of attrition." By estimating only the annual survival rates, rather than the entire hazard function $g(t)$, their resolution is limited to an annual view of the attrition process. The month-to-month hazard function could have been estimated using, for example, the Kaplan-Meier nonparametric procedure.[26]

The approaches of Follmann et al. and Warner and Solon both have merit, and it would be interesting to compare their performance on a common data set. Off-the-shelf statistical software could be used if the

---

[25] The annual hazard rates are obtained from the parameters α_1 through α_4 in table 6.7 of [12], using the transformation $\left[1+\theta^{-1}\exp(-\alpha_t)\right]^{-\theta}$. The survival rates are obtained by cumulating the annual hazard rates.

[26] See Kalbfleisch and Prentice [54, pp. 10-16].

reenlistment decision were modeled using logit or probit regression.[27] However, extreme-value regression would require specialized software, and incorporation of a permanent-transitory error structure would be even more difficult.[28]

---

[27] The earlier paper by Follmann et al. [13] actually made a start in this direction, examining attrition of non-prior-service Marine Corps reservists. At the time that paper was written, the authors were not yet aware of the special commensurability between extreme-value regression and the proportional-hazards model. Thus, they used logit regression rather than extreme-value regression at the spike (the end of the 4-year contract); and a Weibull model rather than a proportional-hazards model away from the spike. Nonetheless, they found that distinct subsets of regressors had significant effects on the hazard function at the spike versus away from the spike.

[28] Extreme-value regression could be estimated using LIMDEP. See chapter 21.7, "User Defined Index Models for Binary Choice," in the on-line manual for LIMDEP version 7.0 (www.limdep.com).

THIS PAGE INTENTIONALLY LEFT BLANK

Elasticity computation

Care must be exercised in computing the elasticity of reenlistment with respect to military pay. The elasticity is sensitive to both the definition of "reenlistment" and the definition of "military pay."

Definition of reenlistment

The definition of "reenlistment" is complicated by the following factors:[29]

- Early attrition
- Early reenlistment
- Reenlistment eligibility
- Extensions
- Reenlistments to retrain in a different military occupation.

Presumably, personnel who attrite before their ETS have a larger net preference for civilian life. Failure to control for early attrition could confound the reenlistment model, if attrition rates are correlated with variables included in the model. However, Warner and Solon's [14] results imply that the correlation between attrition and reenlistment may not be a major source of bias.

Some personnel reenlist early (e.g., more than 6 months before their ETS date), perhaps in preparation for an overseas assignment. Their inclusion in a reenlistment model may lead to a slight bias if the pay variable is expressed in ACOL form. We have noted that ACOL values tend to increase over a person's career. Those who reenlist early, therefore, would tend to have smaller ACOL values, leading to a

29 The next few paragraphs borrow heavily from the excellent *discussion in* Smith, Sylwester, and Villa [59, pp. 56-61].

downward bias in the ACOL coefficient. Smith, Sylwester, and Villa [59] solve this problem by evaluating ACOL at the original ETS date rather than at the early decision point.

Some personnel are declared ineligible to reenlist. Among the many reasons are poor test scores, failure to meet medical or physical fitness standards, disciplinary problems, or missed promotions. One temptation is to exclude these people from any model of voluntary reenlistment. There are two counterarguments, however. First, some people are declared ineligible after expressing a disinclination to reenlist (e.g., turning down a required overseas assignment). Excluding such people will overstate the reenlistment rate, and may also confound the model estimates if eligibility is correlated with variables in the model. Second, for the purposes of force planning, models that include ineligibles are probably more valuable because the total population at ETS (including ineligibles) is the base to which predicted reenlistment rates are generally applied.

Extensions have already been discussed at length. One difficulty with some of the older studies is that they do not document their treatment of extensions. For example, some studies drop extensions entirely and model the dichotomy between immediate reenlistments and separations. Other studies combine extensions with reenlistments and model total retention. Still other studies track extensions to determine whether they ultimately reenlist, thus retrospectively classifying extensions as either reenlistments or separations. It is difficult to compare the elasticities from studies that differ in their treatment of extensions.

Finally, individuals who reenlist to retrain in a different military occupation are not eligible to receive an SRB. Not all studies (particularly those using grouped data) correctly identify this group. Thus, the military pay variables are measured with error, and the pay coefficients may be downward biased.

Definition of military pay

We now turn to the definition of "military pay." Warner and Goldberg [18, p. 33] describe the preferred method to compute a pay elasticity:

> The pay elasticity is calculated as the percentage increase in the reenlistment rate brought about by a one-level SRB increase, divided by the percentage increase in [annualized military pay] implied by the SRB increase.[30]

Along similar lines, Smith et al. [59, p. 86] construct a "typical military pay raise" as a simultaneous 1-percent increase in basic pay, basic allowance for subsistence (BAS), and basic allowance for quarters (BAQ). These methods are preferred because the increase in military pay is measured as a percentage relative to the base value of military pay, and the latter is always strictly positive.

When the retention model is estimated from individual data, it is preferable to compute the increase in the reenlistment probability for each person in the sample and then aggregate, rather than to work directly with the sample averages. For example, a 1-percent increase in military pay might lead to a 1.0-percent increase in the reenlistment probability for one person, a 1.5-percent increase for a second person, and so on. The simple average of these probability increases should be used to form the numerator of the elasticity. Using, instead, the percentage increase in the reenlistment probability for the "average" person is less accurate and—because the model is nonlinear—could lead to a numerically different answer.

Some other studies compute the pay elasticity with a denominator equal to the percentage change in the military/civilian pay *difference*. For example, Daula and Moffitt [60] measure the percentage change in reenlistment generated by a (presumably sustained) increase in the military/civilian pay difference: $\Delta (M-C) / (M-C)$. In a later paper, Daula and Moffitt [24] make the same computation for some of their estimates. However, they also report elasticities using the Warner/Goldberg

---

30 Under the installment arrangement, the annual SRB payment equals an SRB "level" or "multiple" (an integer or half-integer in the range 0 to 6) × monthly basic pay at the date of reenlistment. Under the lump-sum arrangement, the single payment equals the product of the SRB level, monthly basic pay, and the number of years of reenlistment. In either case, a one-level SRB increase implies an increase in undiscounted dollars equal to monthly basic pay × number of years of reenlistment. However, this increase is enjoyed only by reenlisting, not by merely extending.

method, increasing military pay alone: $\Delta M / M$. They even report elasticities using the percentage change in ACOL, ΔACOL / ACOL, and the percentage change in the *total* (not annualized) discounted pay difference over a predetermined (not necessarily "optimal") horizon (their so-called Total Cost of Leaving or TCOL).

Elasticities estimated by these latter methods are unstable, and should not be used for policy evaluation. A pay difference near zero implies a zero elasticity; even a slightly negative pay difference implies a negative elasticity, even when retention and military pay move in the same direction. To see these points, consider the elasticity when the denominator is measured relative to the military/civilian pay difference:

$$[\Delta P / P] / [\Delta (M - C) / (M - C)]$$
$$= [(M - C) \times (\Delta P)] / [P \times \Delta (M - C)]. \tag{18}$$

Suppose that the base military/civilian pay difference happens to equal zero. Then the percentage change appearing in the denominator on the left-hand side is infinite, implying that the elasticity itself (the entire right-hand side) equals zero. As another example, if the base military/civilian pay difference is negative, the elasticity will also be negative, even if increased pay has a positive effect in encouraging more reenlistments (i.e., even if the changes ΔP and $\Delta (M - C)$ have the same algebraic sign). In neither case is the computed elasticity useful for policy evaluation. Our situation is unique because, unlike most policy evaluations, the base value of our independent variable does not always take the same algebraic sign, leading to instability in the conventionally computed elasticity.

Finally, recall that Hosek and Peterson [12] used two compensation measures: the SRB level and an index of relative military pay. They report the percentage-point increase in the reenlistment rate associated with a one-level SRB increase. Although not an elasticity, the interpretation of this quantity is straightforward.

Hosek and Peterson also report the percentage increase in the *retention* (i.e., reenlistment plus extension) rate associated with a 1-percentage-point increase in the military/civilian wage index.

Unless the base value of the wage index is 1.0, a *1-percentage-point* increase in the wage index does not equate to a *1-percent* increase. In fact, the mean value of the wage index was 0.94 in Hosek and Peterson's sample. Thus, to obtain the elasticity of retention with respect to relative military pay, their reported percentage increase in the retention rate must be multiplied by the factor 0.94.[31]

[31] This correction is sometimes neglected, for example, by Warner and Asch [61, table 5].

THIS PAGE INTENTIONALLY LEFT BLANK

Elasticity estimates

Pay elasticities

Table 2 summarizes the pay elasticities from various studies. The left-hand panel shows the elasticity of the reenlistment rate (or, in some cases, the total retention rate) with respect to some measure of military pay. As we have discussed, the most stable estimates use military pay alone, rather than the military/civilian pay difference, in the denominator of the elasticity. We will limit our discussion to the former estimates.

Goldberg and Warner [30] report first-term pay elasticities in the range of 1.1 to 2.7, whereas Warner and Goldberg [18] report a similar range of 1.1 to 3.4. Considering the other studies that use military pay alone in the denominator, most of the elasticities cluster within the even narrower range of 1.2 to 2.2. These other studies include Cooke, Marcus, and Quester [16], Daula and Moffitt [24], Shiells and McMahon [17], Smith et al. [59], and Warner and Solon [14]. All of the studies cited in this paragraph use individual data, except for Goldberg and Warner [30], who use grouped data.

Three studies do not conform to this pattern, instead reporting considerably lower first-term pay elasticities. Mackin et al. [27] estimated pay elasticities by Navy occupational group. Even their most responsive occupational group had an elasticity of only 1.5. Two of the studies used the ACOL–2 approach, but correctly computed the pay elasticities with respect to military pay alone, rather than with respect to the ACOL pay difference. Mackin [62] reports elasticities by service, ranging from 0.5 to 1.4. Black, Hogan, and Sylwester [31] report an elasticity of 0.8 to 0.9 for Navy enlisted personnel. As we discussed earlier, the ACOL-2 approach was designed, in part, to correct an

Table 2. Pay elasticities from various studies

Study	Pay variable	Sample restrictions	Pay elasticity: First-term	Pay elasticity: Second-term	SRB effect on reenlistment (*not* total retention) rate (pp's = percentage points): First-term Installment	First-term Lump-sum	Second-term Installment	Second-term Lump-sum
Black, Hogan, and Sylwester [31]	ACOL–2; but elasticity of reenlistment with respect to military pay	Navy enlisted	0.8 – 0.9	0.3				
Cooke, Marcus, and Quester [16]	Military/civilian pay index; SRB	Navy enlisted	1.6		2.5 pp's			
Daula and Moffitt [60]	Military/civilian pay difference	Army infantry	1.2					
Daula and Moffitt [24]	Military/civilian pay difference	Army infantry	0.5	0.5				
	Military pay alone	Army infantry	2.2	1.3				
	ACOL–2	Army infantry	0.8	0.6				
Goldberg and Warner [30]	Total retention; military pay alone (RMC)	Navy enlisted, by occupational group	1.1 – 2.7	0.9 – 3.8	1.5 – 3.0 pp's	2.0 – 3.9 pp's	1.6 – 9.1 pp's	2.1 – 11.4 pp's
Gotz and McCall [9]	Military/civilian pay difference	USAF pilots, YOS 7-8		0.8				
		USAF non-rated officers, YOS 6-7		1.4				
Hosek and Peterson [12]	Military/civilian pay index; SRB	Enlisted males, four services	3.6	1.6	1.8 pp's	2.5 pp's	2.3 pp's	2.2 pp's

Table 2. Pay elasticities from various studies (continued)

Study	Pay variable	Sample restrictions	Pay elasticity: First-term	Pay elasticity: Second-term	SRB effect on reenlistment (*not* total retention) rate (pp's = percentage points): First-term Installment	First-term Lump-sum	Second-term Installment	Second-term Lump-sum
Mackin [62]	ACOL–2; but elasticity of reenlistment with respect to military pay	Army enlisted	1.2	1.0				
		Navy enlisted	1.0	0.8				
		USAF enlisted	0.5	0.4				
		USMC enlisted	1.4	1.1				
Mackin et al. [27], conditional logit model	Reenlistment; military pay alone	Navy enlisted, by occupational group	0.2 – 1.5	0.5 – 1.1	0.4 – 2.8 pp's		1.9 – 2.6 pp's	
	Total retention; military pay alone		0.2 – 0.9	0.5 – 1.0				
Shiells and McMahon [17]	Military/civilian pay index; SRB	Navy enlisted	1.9		1.5 pp's			
Smith et al. [59]	Military pay alone	Army infantry	1.3	0.9	2.2 pp's		1.7 pp's	
		Army maintenance	1.8	1.1				
		Army administration	1.9	1.8				
Warner and Goldberg [18]	Military pay alone (SRB)	Navy enlisted, by occupational group	1.1 – 3.4		1.8 – 5.5 pp's			
Warner and Solon [14]	ACOL; but elasticity of reenlistment with respect to military pay	Army infantry	1.2					

Notes: pp's = percentage points, USAF = United States Air Force, USMC = United States Marine Corps, YOS = years of service.

upward bias in simple ACOL models when applied at the second-term (or later) decision points. Thus, we would expect the ACOL–2 approach to yield lower pay elasticities for these later decisions. However, it is surprising that the approach yields such lower pay elasticities at the first-term decision point.

Figure 7 plots the various first-term pay elasticities. In constructing this figure, we arrayed the elasticities from left to right, alphabetically by the last name of the first author, except that we grouped the ACOL–2 studies to the extreme right. For studies that estimated separate elasticities by occupational group, we used a vertical bar to depict the range of elasticities, and a circle to indicate the midpoint of the range. We see that there is considerable dispersion in the elasticity estimates, but the ACOL–2 estimates tend to be concentrated at the lower end.

Figure 7. Elasticity of first-term retention with respect to pay

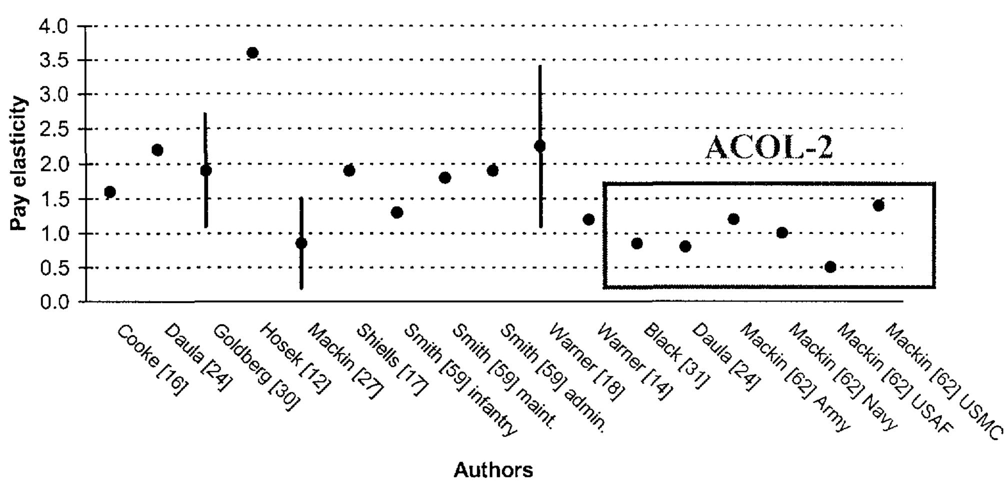

We now turn to the second-term pay elasticities. Goldberg and Warner [30] report pay elasticities in the range of 0.9 to 3.8. Daula and Moffitt [24] and Smith et al. [59] report pay elasticities that cluster within the considerably narrower range of 0.9 to 1.8. Mackin et al. [27] report somewhat lower pay elasticities, in the range 0.5 to 1.1. Using the ACOL–2 approach, Mackin [62] reports elasticities by service ranging from 0.4 to 1.1. Finally, Black, Hogan, and Sylwester [31] report an elasticity of only 0.3 for Navy enlisted personnel.

In figure 8, we plot the second-term pay elasticities. Not surprising in this case, the ACOL–2 estimates again tend to be concentrated at the lower end.

Figure 8. Elasticity of second-term retention with respect to pay

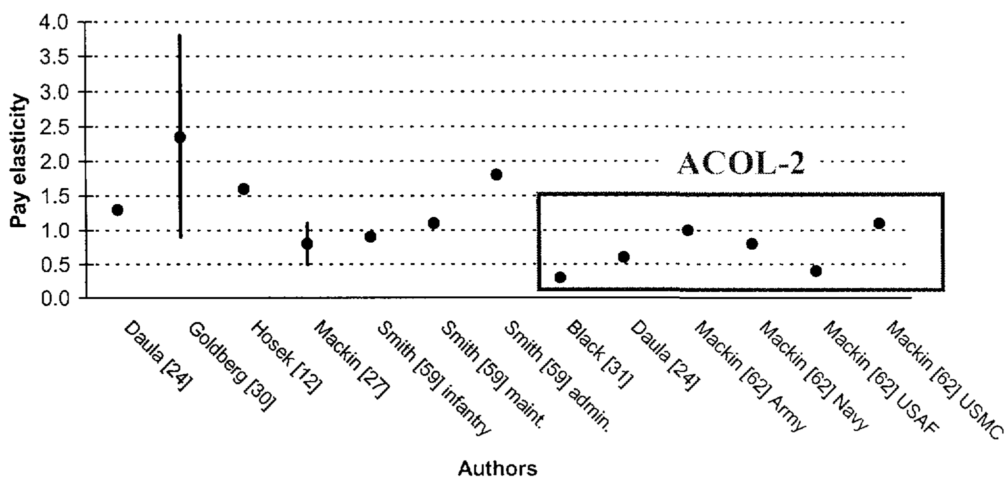

SRB effects

Table 2 also summarizes the effect of a one-level SRB increase on the reenlistment (*not* total retention) rate. The most interesting comparison is between the two studies that used grouped data: Goldberg and Warner [30], who estimated a conditional logit model, and Hosek and Peterson [12], who estimated a multinomial logit model. According to Goldberg and Warner, a one-level increase in the first-term SRB, if paid in annual installments, serves to increase the reenlistment rate by 1.5 to 3.0 percentage points. Hosek and Peterson's point estimate of 1.8 percentage points falls within this range. Similarly, Goldberg and Warner report second-term installment SRB effects in the range of 1.6 to 9.1 percentage points; excluding a high outlier narrows the range to 1.6 to 5.2. Hosek and Peterson's point estimate of 2.3 percentage points falls within that range. The other studies tend to confirm the ranges established by Goldberg and Warner. Three other studies [16, 17, and 59] report first-term SRB effects in the range of 1.5 to 3.0 percentage points. Both Mackin et al. [27] and Warner and Goldberg [18] estimated SRB effects by Navy occupational group. In both cases, the occupational groups were rather dispersed, with a few producing SRB effects outside the range of 1.5 to 3.0. The first-term SRB effects are plotted in figure 9.

Figure 9. Effect of one-level SRB increase on first-term reenlistment rate

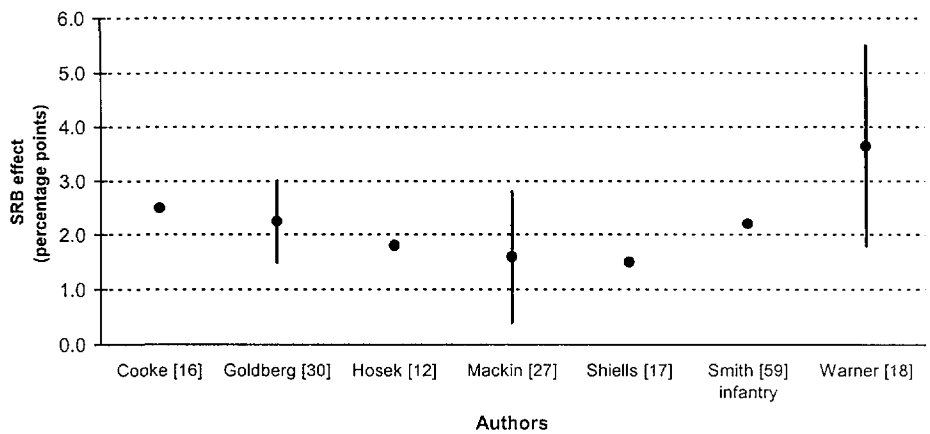

The second-term SRB effects are plotted in figure 10. All of the second-term SRB effects fall within the range of 1.6 to 5.2 percentage points established by Goldberg and Warner. In fact, the other estimates cluster within the considerably narrower range of 1.7 to 2.6 percentage points.

Combining the first-term and second-term results, it appears as a rough rule of thumb that a one-level SRB increase leads to an increase in the reenlistment rate of about 2 percentage points.

Figure 10. Effect of one-level SRB increase on second-term reenlistment rate

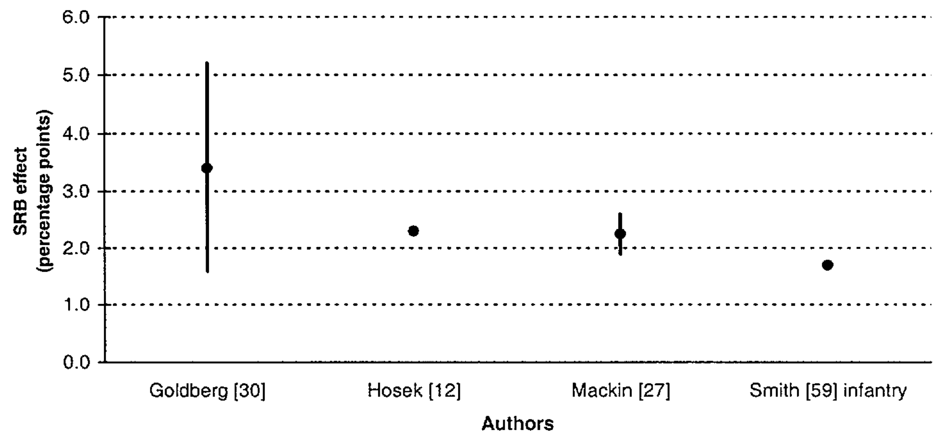

Relative stability of SRB effects

We observe a pattern in which the SRB effects are much more stable across studies than are the pay elasticities. To understand this pattern, recall that we usually estimate binary retention models using logit or probit functional forms, both of which are S-shaped curves. Both functions are very nearly linear over moderate ranges. The SRB effect is the slope of the retention function with respect to a particular type of pay increase, and the slope is essentially a constant throughout the range over which the retention function is nearly linear. On the other hand, the pay *elasticity* is not constant along a nearly linear retention function. Put differently, the elasticity is a measure of curvature for an iso-elastic approximation to the retention function. The best-fitting iso-elastic function is extremely sensitive to the point of evaluation (i.e., the sample average retention rate). Thus, elasticity estimates from a logit or probit function tend to be unstable.

These arguments are illustrated in figures 11 and 12. A logit function and a linear function are virtually indistinguishable over a range of reenlistment rates. However, an iso-elastic function has greater curvature and bends noticeably away from the logit function, even over a relatively modest 12-percentage-point range of reenlistment rates.

Figure 11. Logit function is closer to linear than to iso-elastic

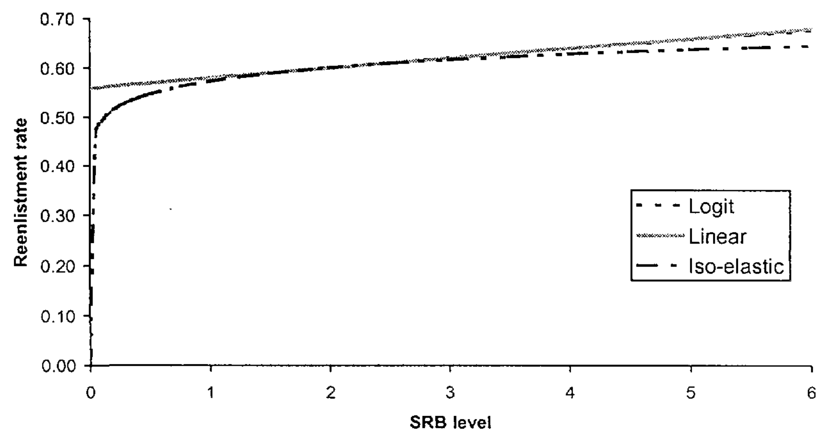

Figure 12 increases the magnification within a plausible range of reenlistment rates. The logit function is essentially linear, with a slope of 0.02 reflecting the 2-percentage-point increase in the reenlistment rate with each unit increase in the SRB level. Also shown are three iso-elastic curves with very different elasticities. A small perturbation in the point of evaluation can lead to a nearly threefold variation in the pay elasticity.

Figure 12. Point estimates of elasticities are unstable

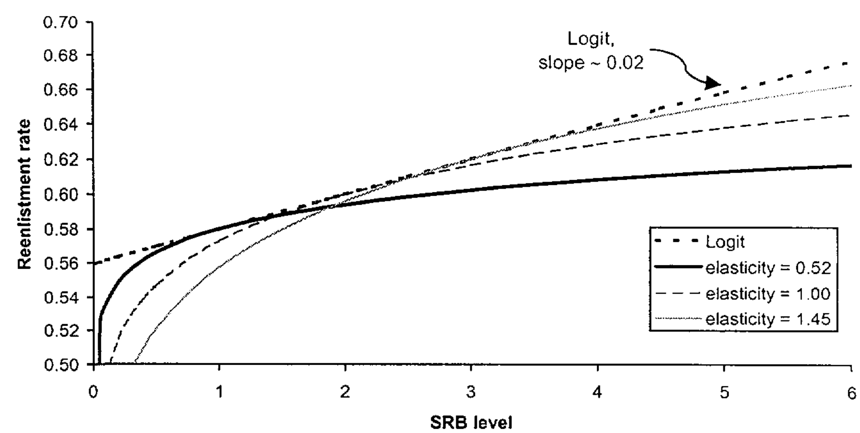

The studies summarized in table 2 vary considerably in the sample average retention rate. The theoretical threefold variation in pay elasticities is supported by figure 7: most of the first-term elasticities fall within the range of 1.2 to 2.2, but a few fall below 1.0, and others exceed 3.0. Thus, while the SRB effects are relatively stable, the pay elasticities are rather unstable. If either, it is the former that is closer in character to a "natural constant."

THIS PAGE INTENTIONALLY LEFT BLANK

Estimation of discount rates

Hosek and Peterson [12, p. 1] state that, "The chief purpose of [their] study is to determine whether lump-sum reenlistment bonuses are more cost-effective than installment bonuses." The primary factor in this determination is a comparison of discount rates between military members and the federal government. Interestingly, Hosek and Peterson did *not* use their findings to explicitly infer a military member's discount rate. Nonetheless, they argued that lump-sum reenlistment bonuses are cost-effective as long as the federal government's real discount rate lies anywhere in the range of 4 to 10 percent. Their report appeared during a period when OMB was still mandating use of a 10-percent real discount rate on government investment projects. Thus, they concluded that lump-sum bonuses are the preferred method of payment.[32]

Hosek and Peterson exploited a natural experiment that occurred in April 1979, when the method of SRB payment switched from annual installments to lump-sum payments. They estimated both a dichotomous model of staying (i.e., either reenlisting or extending) versus leaving, and a trichotomous model of reenlisting, extending, or leaving. They found that, in the former case, installment bonuses were only 82.7 percent as effective as lump-sum bonuses in encouraging first-term enlisted personnel to stay. In the latter case, installment bonuses were 72.8 percent as effective as lump-sum bonuses in encouraging first-term personnel to reenlist. We can use these findings to roughly estimate a military member's real discount rate.

According to economic theory, individuals make their decisions by comparing the discounted present values of the various alternatives.

---

[32] In 1992, OMB revised its guidance and tied the real discount rate to inflation-adjusted market rates on Treasury bonds. Those rates have generally been in the range of 3 percent to 4 percent. The rationale for OMB's revised guidance was provided after the fact by Goldberg [63].

If installment bonuses are only 72.8 percent as effective as lump-sum bonuses, the present value of the former must be only 72.8 percent the present value of the latter. Assume a 4-year reenlistment horizon and, following Hosek and Peterson, a 95-percent annual survival rate within the second term of service. Considering a notional \$1,000 bonus, we have the following present-value equation:

$$0.728 \times \$1{,}000 = \$250 \times \{1 + 0.95/(1+r) + [0.95/(1+r)]^2 + [0.95/(1+r)]^3\}. \quad (19)$$

Note that we have not deflated the \$250 annual installments by a price index, so that the installment stream is expressed in *nominal* (i.e., current) dollars. Therefore, the solution to this equation provides an estimate of the *nominal* discount rate.[33] The solution is easily computed as 20.1 percent. Repeating the exercise using a relative effectiveness of 82.7 percent (based on the dichotomous model) yields a nominal discount rate of 8.7 percent.

To convert to real discount rates, we use the following relationship between the nominal discount rate (r), the real discount rate (d), and the rate of inflation (f):

$$(1 + r) = (1 + d) \times (1 + f). \quad (20)$$

[33] Hosek and Peterson's analysis appears to contain an error. Their equation on p. 42 of [12] is essentially the same as our equation 19. They express the annual installment bonus in nominal terms, as 25 percent of the corresponding lump-sum bonus. To properly discount the stream of installment payments, they should be using a *nominal* discount rate. However, they state on p. 43, "In keeping with our having adjusted the bonus amounts in the empirical work for inflation, the interest rate [that solves the equation] should be interpreted as the 'real' rate—that is, the inflation-adjusted rate." This statement is a *non sequitur*; having normalized the various years' bonus amounts in the regression analysis does not relieve the requirement to either discount a nominal payment stream with a nominal discount rate, or a real payment stream with a real discount rate. Hosek and Peterson attribute to real discounting *all* of the military member's preference for a lump-sum bonus. In fact, some of that preference should instead be attributed to the automatic inflation protection provided by an immediate, lump-sum payment.

Over the sample period of FY 1976 through FY 1981, the geometric average rate of increase in the Consumer Price Index (CPI) was 9.2 percent. Applying the above formula, we estimate real discount rates of 9.9 percent from the trichotomous model and –0.5 percent from the dichotomous model. Cylke et al. [15] argue that consideration of progressive income taxes tends to increase the estimated discount rate. We have not performed the detailed analysis of Hosek and Peterson's results, including tax effects. However, their results imply that military members' real discount rates are surely positive, and may well exceed 10 percent.

Other discount-rate estimates

Table 3 summarizes the real discount rates estimated from various studies. Cylke et al. [15] followed a procedure similar to Hosek and Peterson, comparing the effectiveness of SRBs when paid in annual installments (pre-April 1979) versus a single lump-sum (post-April 1979). Daula and Moffitt, by contrast, used the method of maximum likelihood to estimate the discount rate as one parameter in the dynamic program. Warner and Pleeter [64] compared military members' choices between installment and lump-sum severance pay when the two were offered concurrently; we will discuss their study in the next section.

Table 3. Estimates of real discount rates

Study	Sample restrictions	Term of service	Real discount rate
Cylke et al. [15]	Navy enlisted	1st term only	17%
Daula and Moffitt [60]	Army infantry	1st term only	4.0% – 5.5%
	Army infantry	1st and 2nd terms	10.5%
Daula and Moffitt [24]	Army infantry	1st and 2nd terms	10% – 14%
Hosek and Peterson [12]	enlisted males, four services	1st term only	9.9%
Warner and Pleeter [64]	Army, Navy, Air Force officers	YOS 7 through 15	6% – 26%
	Army, Navy, Air Force enlisted	YOS 7 through 15	26% – 37%

Note: Real discount rate from Hosek and Peterson [12] is inferred in the current paper.

Looking across all of the studies, the estimated discount rates range between 4 and 37 percent. However, excluding Daula and Moffitt's [60] low first-term estimates, as well as Warner and Pleeter's high estimates for enlisted personnel, the remaining estimates range between 6 and 26 percent. These rates all exceed the federal government's real discount rate of 3 to 4 percent, leaving little doubt that lump-sum bonuses are the preferred method of payment.

Warner and Pleeter study

Warner and Pleeter [64] exploited a natural experiment that occurred when DoD reduced military endstrength during FY 1992 and FY 1993. Recall that both Cylke et al. and Hosek and Peterson compared a time period during which SRBs were paid in annual installments to a time period during which SRBs were paid as a single lump-sum. By contrast, Warner and Pleeter examined a single time period during which both installment and lump-sum severance pay were offered concurrently. What makes their study unique is that military members were offered a *contemporaneous choice* between the two payment options.

Specifically, DoD offered severance packages to mid-career personnel (both officer and enlisted) in selected combinations of military occupation, paygrade, and years of service. The Voluntary Separation Incentive (VSI) provided annual payments equal to 2.5 percent of terminal basic pay, multiplied by terminal years of service. The annual payments would continue for a period of time equal to twice the terminal years of service, with no indexing for inflation. The Special Separation Benefit (SSB) provided a lump-sum payment equal to 15 percent of terminal basic pay, multiplied by terminal years of service.[34]

---

[34] Mehay and Hogan [65] report that, during FY 1992, less than 10 percent of the Navy enlisted force met the occupation/paygrade/YOS eligibility criteria. Among these individuals, 12 percent accepted some form of buyout. Mehay and Hogan did not explicitly analyze the choice between the two payment options. However, they report that among Navy enlisted personnel who accepted some form of buyout, 85 percent chose the SSB (lump-sum) payment option.

If individuals did not discount, the annuity option would be preferred as long as YOS > 3 (for then 2.5% × basic pay × YOS × 2 × YOS > 15% × basic pay × YOS). With discounting, the breakeven career length is somewhat longer. Put differently, for any given career length of YOS > 7 (the minimum for buyout eligibility), one can compare the discounted present values of the two payment options at various discount rates. In fact, when announcing the program, DoD published a pamphlet giving the comparison of present values at a 7-percent nominal discount rate, which was the typical yield on money market funds at the time. Using that discount rate, the present value of the annuity option was as much as twice the size of the lump-sum payment.

One can also calculate the breakeven discount rate—the rate at which a person must discount the annuity payments to yield a present value equal to the lump-sum payment. Warner and Pleeter computed before-tax breakeven discount rates ranging from 17.5 to 19.8 percent, varying solely as a function of YOS. They also computed after-tax breakeven discount rates ranging from 17.5 to 23.6 percent.

Despite these high breakeven rates, most people chose the lump-sum payment option, indicating that their personal discount rates were even higher. According to Warner and Pleeter:

> Among the officers with less than 10 years of service, more than half took the lump-sum. Among the E-5 enlisted personnel with less than 10 years, over 90 percent did so. Almost 75 percent of E-7 enlisted personnel with 15 years of service took the lump-sum. Even among the more senior officers, 30 percent or more took the lump-sum. Overall, about half of the officers chose the lump-sum while over 90 percent of the enlisted personnel did so.

Warner and Pleeter's estimated discount rates were shown earlier in table 3. They ranged from 6 to 26 percent for officers, and from 26 to 37 percent for enlisted personnel.

We use some diagrammatic tools from microeconomic theory to illuminate these calculations. Figure 13 is an indifference-curve diagram for a person choosing between the annuity and lump-sum payment options. The axes measure consumption of goods and services in the first and second time periods, respectively. Absent either severance pay or access to financial markets, the person would simply consume his or her income in each period. This income "endowment" is depicted as point E. Relaxing these assumptions gradually, suppose next that the person may either borrow or lend at the interest rate r. The resulting budget line passes through point E with negative slope of $1 + r$. As the figure is drawn, this individual would choose to lend money in the first period, reducing consumption in that period but increasing consumption in the second period when the investment comes due. Geometrically, the person moves along the budget line from point E to a consumption optimum at point B.

Figure 13. Comparison of annuity and lump-sum payment options

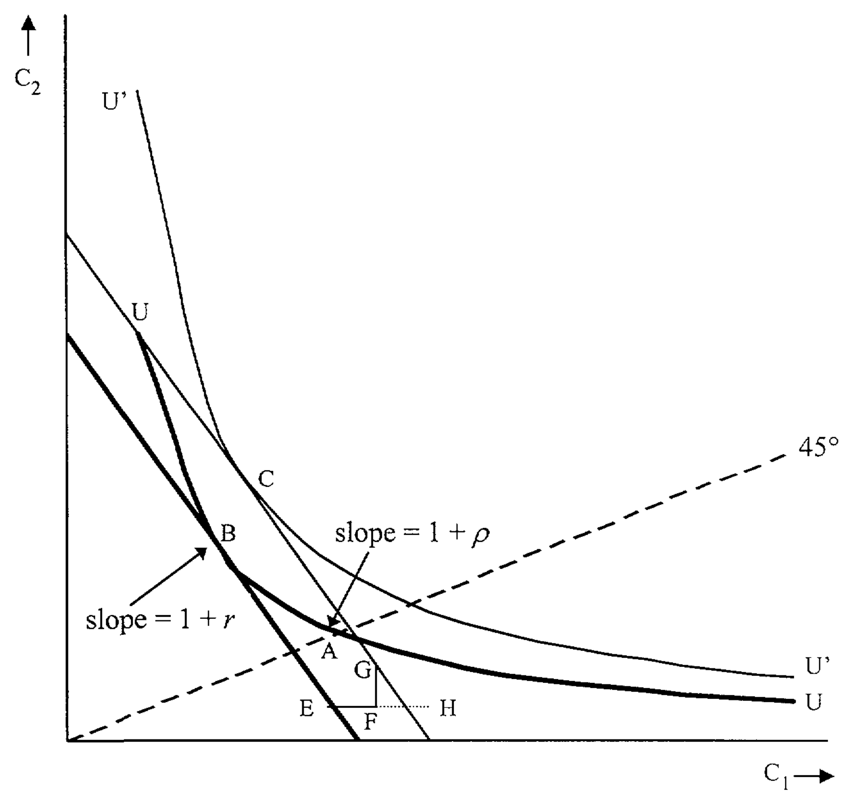

The slope of the person's indifference curve at point B equals $1+r$. Thus, the observed slope is solely a function of the interest rate at which the person may either borrow or lend. The observed slope is *not* a measure of underlying preference for current versus future consumption. The latter must be measured at some benchmark level of relative consumption that is independent of market opportunities. Conventionally, a measure of time preference (or "impatience") is derived from the slope of the person's indifference curve at the 45-degree line (the line along which current and future consumption are equal). The slope at the 45-degree line (e.g., at point A) is equated to $1 + \rho$, and ρ is defined as the consumer's "rate of time preference."[35]

To see this point in another way, consider the two individuals pictured in figure 14. These two have very different preferences for current versus future consumption. The "impatient" one is more concerned with current consumption, and has a steep indifference curve (drawn as a solid curve). By contrast, the "patient" person is more concerned with future consumption, and has a flat indifference curve (drawn as a dashed curve). These differences in time preference are evident by comparing the slopes of the solid and dashed indifference curves at the 45-degree line. Yet, if these two people borrow or lend at the same interest rates as each other, they will reach consumption optima at which each has an indifference curve with slope $1 + r$. Again, the observed slope measures market opportunities rather than underlying time preference.[36]

Finally, we compare the annuity and lump-sum severance packages. Returning to figure 13, the annuity option serves to increase income during both the first and second periods. Thus, the annuity option shifts the endowment point both horizontally (point F) and then vertically (point G). The new budget line passes through point G with

---

35 This definition is found in Epstein and Hynes [66], among other places.

36 When reviewing studies of "discount rates," one must carefully distinguish between those that measure market opportunities (r) and others that measure underlying time preference (ρ). All of the studies in our table 3 are measuring market opportunities. Other studies of market opportunities, outside the military sector, include Gately [67], Gilman [68], and Hausman [69]. Lawrance [70] is the best-known study of the consumer's rate of time preference.

slope 1 + *r*. The individual may now adjust his consumption path to reach a higher utility level at point C.

Figure 14. Two individuals with different time preferences but equal interest rates

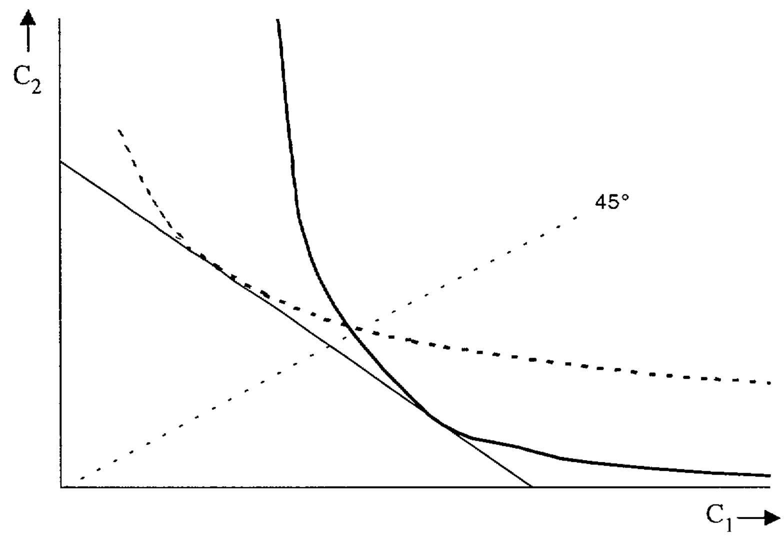

By contrast, the lump-sum option serves to increase income during only the first period. Thus, the lump-sum option shifts the endowment point horizontally beyond point F, perhaps to point H. As drawn in figure 13, point H lies to the northeast of the budget line passing through point G. Thus, the lump-sum option leads to an even higher budget line and, consequently, an even higher utility level (neither of which is explicitly shown in the figure, to avoid clutter).

The consumer's preference for point H (lump-sum option) over point G (annuity option) reveals that his or her budget line is *steeper* than the line segment GH. The slope of the budget line is, again, 1 + *r*. However, the slope of the line segment GH equals 1.0 plus the "breakeven rate" computed by Warner and Pleeter. To see why, note that points G and H would lie on the same budget line only if the two payment options yielded exactly the same discounted present value. The line segment GH is, in fact, a subset of the hypothetical budget line with slope equal to 1.0 plus the interest rate that equates the two present values—the breakeven rate.

As is clear from this analysis, the consumer prefers the lump-sum payment option only if his or her personal discount rate exceeds the breakeven rate. Thus, a consumer's choice of either the annuity or lump-sum payment option serves to bound his or her personal discount rate on one side or the other of the predetermined breakeven rate. It is this information that Warner and Pleeter exploit to estimate the distribution of personal discount rates in the military population.

THIS PAGE INTENTIONALLY LEFT BLANK

Effects of variables other than pay

Personal characteristics

Retention models have often included variables other than pay. Perennial favorites include the civilian unemployment rate, and personal characteristics, such as marital status, race, education, and mental group. One difficulty is that these are some of the same personal characteristics used to predict civilian pay in forming the military/civilian pay difference or pay ratio. Inclusion of these characteristics in the retention model introduces multicollinearity, which tends to depress the estimated coefficient on the relative pay variable. This problem was noted by Warner [71, pp. 222-223]:

> inclusion of individual attributes such as education, race, and Armed Forces Qualification Test (AFQT) score in the retention equation to control for non-pecuniary factors results in substantial changes in pay parameter estimates. Such changes may occur either because the model is properly specified only with these variables included or because of the multicollinearity introduced, since these factors also help determine the relative pay variable in the equation....A second source of sensitivity...may arise if variables that affect civilian wages are also entered directly in the retention function. I am not sure whether they should or should not be included. I will only comment that I have done it both ways, and I have found maximum-likelihood estimates of pay elasticities to be much more sensitive to inclusion or exclusion of these variables than estimates based on a grouped logit approach.

Whether or not to include personal characteristics remains an open question. Personal characteristics should certainly be included if there is an independent interest in their effects on retention. However, if the primary goal of a particular study is to estimate the pay effects, it may be preferable to exclude the personal characteristics because their inclusion tends to destabilize the pay coefficient.

Sea duty

Several studies have examined the effect of sea duty on reenlistment rates of sailors. Warner and Goldberg [18] modeled the first-term reenlistment rate as a function of the expected percentage of time spent on sea duty during the *second* term of service. They assigned each person the expected sea duty specific to his or her Navy rating. They estimated that a 10-percentage-point increase in prospective sea duty would reduce the reenlistment rate by a modest 1.6 percentage points.

The effect of sea duty was revisited by Shiells and McMahon [17]. In the numerical example to illustrate their findings, they increased the prospective sea/shore ratio from 2.6 : 1 (i.e., 2.6 years on sea duty for every year on shore duty) to 3.3 : 1. They estimated that the 25-percent increase in the sea/shore ratio would reduce the reenlistment rate by 1.9 percentage points. Note that Shiells and McMahon did not use the same metric for sea duty as did Warner and Goldberg. Using the latter authors' metric, the percentage of time spent on sea duty would increase from 72 percent (2.6/3.6) to 77 percent (3.3/4.3). Correspondingly, the reenlistment rate would fall by 0.8 percentage point. Comparing the two estimates, Shiells and McMahon's is over twice as large as Warner and Goldberg's. However, even the larger estimate implies that the hypothesized increase in sea duty could be offset by a one-level SRB increase.

Personnel tempo

More recently, a few studies have begun to look at the effect of personnel tempo (Perstempo) on retention. Cooke, Marcus, and Quester [16] examined three measures of Perstempo for male sailors:

- Length of deployment
- Turnaround ratio, defined as time between deployments divided by length of deployment
- Percentage of time under way while not deployed.

Reference [16] found few systematic patterns for sailors with 8 to 10 years of service. However, the authors found several significant effects among 4-year obligors at their first reenlistment point. For instance, they found that longer deployments adversely affected retention, especially among married sailors (roughly one-third of those making reenlistment decisions). Second, a lower turnaround ratio also adversely affected retention, but the effect was smaller and limited to married sailors. Third, the percentage of time under way while not deployed had an adverse effect on retention. The latter effect was most severe among both married sailors and sailors serving in relatively sea-intensive ratings (these two groups overlap).

Among their other results, Cooke, Marcus, and Quester [16] found that retention was lower among sailors deployed at their decision point, even controlling for a sailor's deployment history during the 3-year window leading up to the decision point. Finally, retention was lower among sailors serving on ships that had recently undergone a major maintenance activity (overhaul or restricted availability) that lasted 8 months or longer.

Hosek and Totten [19] developed some additional measures of Perstempo, and extended the analysis to all four military services. They not only examined long duty, but they appear to be the first to examine the effect of hostile duty on retention. They measured long duty as the incidence of the Family Separation Allowance, which is paid to personnel with dependents from whom they are separated for 30 or more consecutive days. They measured hostile duty as the incidence of Hostile Fire Pay, which is paid to personnel subject to hostile fire or explosion, or those on duty in areas deemed to be hostile.

Hosek and Totten found that some degree of long or hostile duty actually increases first-term retention, particularly among Army and Marine Corps personnel. Beyond a certain point, however, additional long duty reduces retention, especially if that duty is also hostile. Therefore, overall retention might improve if the burden of long and hostile duty were shared among large numbers of military personnel. However, Hosek and Totten caution that duty sharing must be balanced against operational factors, including unit cohesion:

> We found that the impact of added long or hostile duty differs for personnel who have had it from those who have not, and whether it is hostile or non-hostile. Thus, if the added hostile duty can be spread to troops who have not yet been deployed, then the effect on reenlistment is likely to be positive; if the added duty falls to those who have already been deployed, then the effect on reenlistment is likely to be negative. Of course, decisions about how to allocate additional assignments must include a variety of factors beyond effects on retention rates. A Service's capability to share long or hostile duty among units may be influenced by advantages gained from assigning units particular roles for a major theater war and assuring that these units stand at full readiness. For readiness reasons, it may not be advisable to spread such duty more broadly (pp. xvii-xviii);...because personnel are attached to units and develop specialized skills and knowledge about the unit's roles and missions, weaponry/equipment, and fellow unit members, simply swapping one person or unit for another is essentially infeasible. A more subtle means must be devised (p. 58).

Finally, consistent with Cooke , Marcus, and Quester [16], Hosek and Totten found much smaller effects of long or hostile duty for early-career personnel (i.e., those beyond the first term but with 10 or fewer years of service) than for first-term personnel.

Areas for future research

This survey has identified several fruitful areas for future research:

- Determine what joint distribution of taste factors, if any, would yield the sequential logit procedure used by Warner [42] to model the stay/leave decision along with the reenlist/extend decision among those who stay.
- Attempt full-information maximum-likelihood estimation of the nested logit model for the stay/leave and reenlist/extend decisions.
- Add cohort-composition effects to trichotomous logit (or probit) models for the reenlist/extend/leave decision.
- When using grouped data, explicitly model SRB levels jointly with reenlistment rates, to control for reverse causation.
- When using individual data, apply disequilibrium estimation to distinguish supply-constrained from demand-constrained observations.
- Continue the joint modeling of attrition and first-term reenlistment. Compare the methods of Follmann, Goldberg, and May [58] and Warner and Solon [14] on a common data set.
- Further investigate the apparently low first-term pay elasticities produced by the ACOL–2 approach.

The current paper has attempted to decompose the variation in pay elasticities in terms of differences in data handling (e.g., treatment of ineligibles and extensions), modeling technique, and elasticity computation. However, this decomposition is confounded by differences in service, occupational group, and time period among the many studies examined. A useful "controlled experiment" would be to apply the various estimation techniques to a common data set, thereby eliminating any confounding differences in sample composition. In light of the changing market for military labor, it seems imperative to conduct this experiment using the most recent available data.

THIS PAGE INTENTIONALLY LEFT BLANK

References

[1] S. Altman and A. Fechter. "The Supply of Military Personnel in the Absence of a Draft," *American Economic Review,* Vol. 57, 1967, pp. 19-31

[2] A. Fisher, "The Cost of the Draft and the Cost of Ending the Draft." *American Economic Review,* Vol. 59, 1969, pp. 239-254

[3] Lee Hansen and Burton Weisbrod. "Economics of the Military Draft." *Quarterly Journal of Economics,* Vol. 82, 1967, pp. 395-421

[4] J. Miller (ed.). *Why the Draft? The Case for the Volunteer Army.* Baltimore: Penguin, 1968

[5] Walter Oi. "The Economic Cost of the Draft." *American Economic Review,* Vol. 57, 1967, pp. 39-62

[6] John H. Enns. *Effect of the Variable Reenlistment Bonus on Reenlistment Rates: Empirical Results for FY 1971.* Santa Monica, CA: RAND Corporation, Jun 1975 (R-1502-ARPA)

[7] John H. Enns. *Reenlistment Bonuses and First-Term Retention.* Santa Monica, CA: RAND Corporation, Sep 1977 (R-1935-ARPA)

[8] Samuel Kleinman and William Shughart. *The Effects of Reenlistment Bonuses,* 1974 (CNA Research Contribution 269)

[9] Glenn A. Gotz and John J. McCall. *Estimating Military Personnel Retention Rates: Theory and Statistical Methods.* Santa Monica, CA: RAND Corporation, 1980 (R-2541-AF)

[10] Glenn A. Gotz and John J. McCall. *A Dynamic Retention Model for Air Force Officers: Theory and Estimates.* Santa Monica, CA: RAND Corporation, 1984 (R-3028-AF)

[11] Matthew S. Goldberg. "Compensation and Retention in the U.S. Navy," doctoral dissertation, University of Chicago, 1984

[12] James Hosek and Christine Peterson. *Reenlistment Bonuses and Retention Behavior.* Santa Monica, CA: RAND Corporation, 1985 (R-3199-MIL)

[13] Dean A. Follmann, Matthew S. Goldberg, and Laurie J. May. *Modeling Spikes in Hazard Rates,* 1987 (CNA Research Contribution 572)

[14] John T. Warner and Gary Solon. "First-Term Attrition and Reenlistment in the U.S. Army." In Curtis L. Gilroy, David K. Horne, and D. Alton Smith (eds.), *Military Compensation and Personnel Retention: Models and Evidence.* Alexandria, VA: U.S. Army Research Institute for the Behavioral and Social Sciences, 1991

[15] Steven Cylke et al. *Estimation of the Personal Discount Rate: Evidence From Military Reenlistment Decisions,* 1982 (CNA Professional Paper 356)

[16] Timothy W. Cooke, Alan J. Marcus, and Aline O. Quester. *Personnel Tempo of Operations and Navy Enlisted Retention,* 1992 (CNA Research Memorandum 91-50)

[17] Martha E. Shiells and Joyce S. McMahon. *Effects of Sea Duty and Advancement on First-Term Retention,* 1993 (CNA Research Memorandum 92-205)

[18] 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,* Vol. 66, 1984, pp. 26-35

[19] James Hosek and Mark Totten. *Does Perstempo Hurt Reenlistment?* Santa Monica, CA: RAND Corporation, 1998 (MR-990-OSD)

[20] John T. Warner. "Analysis of the Retention Effects of the Proposed Retirement System." In *Supplementary Papers of the*

President's Commission on Military Compensation. Washington, DC: U.S. Government Printing Office, Apr 1978

[21] John H. Enns, Gary R. Nelson, and John T. Warner. "Retention and Retirement: The Case of the U.S. Military." *Policy Sciences,* Vol. 17, 1984

[22] James H. Stock and David A. Wise. "Pensions, the Option Value of Work, and Retirement." *Econometrica,* Vol. 58, Sep 1990, pp. 1151-1180

[23] Robin L. Lumsdaine, James H. Stock, and David A. Wise. "Three Models of Retirement: Computational Complexity versus Predictive Ability." In David A. Wise (ed.), *Topics in the Economics of Aging.* Chicago: University of Chicago Press, 1992

[24] Thomas V. Daula and Robert Moffitt. "Estimating Dynamic Models of Quit Behavior: The Case of Military Reenlistment." *Journal of Labor Economics,* Vol. 13, No. 3, 1995, pp. 499-523

[25] Glenn A. Gotz. "Comment on the Dynamics of Job Separation: The Case of Federal Employees." *Journal of Applied Econometrics,* Vol. 5, 1990, pp. 263-268

[26] Matthew Black, Robert Moffitt, and John T. Warner. "Reply to Comment by Glenn Gotz." *Journal of Applied Econometrics,* Vol. 5, 1990, pp. 269-272

[27] Patrick C. Mackin et al. *U.S. Navy Enlisted Compensation and Retention Modeling.* Falls Church, VA: SAG Corporation, 1996

[28] David A. Wise. "Comment on Warner and Solon." In Curtis L. Gilroy, David K. Horne, and D. Alton Smith (eds.), *Military Compensation and Personnel Retention: Models and Evidence.* Alexandria, VA: U.S. Army Research Institute for the Behavioral and Social Sciences, 1991

[29] John T. Warner and Bruce Simon. *An Empirical Analysis of Pay and Navy Enlisted Retention in the AVF: Preliminary Results,* 1979 (CNA Memorandum 79-1878)

[30] Matthew S. Goldberg and John T. Warner. *Determinants of Navy Reenlistment and Extension Rates,* 1982 (CNA Research Contribution 476)

[31] Matthew Black, Paul F. Hogan, and Stephen D. Sylwester. *Dynamic Model of Military Reenlistment Behavior.* Arlington, VA: SRA Corporation, Jun 1987

[32] Matthew Black, Robert Moffitt, and John T. Warner. "The Dynamics of Job Separation: The Case of Federal Employees." *Journal of Applied Econometrics,* Vol. 5, 1990, pp. 245-262

[33] J. S. Butler and Robert Moffitt. "A Computationally Efficient Quadrature Procedure for the One-Factor Multinomial Probit Model." *Econometrica,* Vol. 50, No. 3, May 1982, pp. 761-764

[34] Richard L. Fernandez, Glenn A. Gotz, and Robert M. Bell. *The Dynamic Retention Model.* Santa Monica, CA: RAND Corporation, 1985 (N-2141-MIL)

[35] John T. Warner. *Military Compensation and Retention: An Analysis of Alternative Models and Simulation of a New Model,* 1981 (CNA Research Contribution 436)

[36] V. J. Hotz and R. Miller. "Conditional Choice Probabilities and the Estimation of Dynamic Models." *Review of Economic Studies,* Vol. 61, Apr 1994, pp. 265-289

[37] Paul F. Hogan and Matthew Black. "Reenlistment Models: A Methodological Review." In Curtis L. Gilroy, David K. Horne, and D. Alton Smith (eds.), *Military Compensation and Personnel Retention: Models and Evidence.* Alexandria, VA: U.S. Army Research Institute for the Behavioral and Social Sciences, 1991

[38] Daniel McFadden. "Conditional Logit Analysis of Qualitative Choice" In Paul Zarembka (ed.), *Frontiers of Econometrics.* New York: Academic Press, 1973

[39] Samuel Johnson and Norman Kotz. *Distributions in Statistics: Continuous Multivariate Distributions.* New York: Wiley, 1972

[40] Daniel McFadden. "Modelling the Choice of Residential Location." In A. Karlquist (ed.), *Spatial Interaction Theory and Residential Location.* Amsterdam: North Holland, 1978

[41] Daniel McFadden. "Quantitative Methods for Analyzing Travel Behavior by Individuals." In D. Hensher and P. Stopher (eds.), *Behavioral Travel Modelling.* London: Croom Helm, 1979

[42] John T. Warner. *Determinants of Reenlistment and Extension Rates in the U.S. Marine Corps,* 1982 (CNA Memorandum 82-1733)

[43] William H. Greene. *Econometric Analysis.* 4th ed. New York: Prentice Hall, 1999

[44] Hyder Lakhani and Curtis L. Gilroy. "Army Reenlistment and Extension Decisions by Occupation." In Curtis L. Gilroy (ed.), *Army Manpower Economics.* Boulder, CO: Westview Special Studies in Military Affairs, 1986

[45] Paul F. Hogan. "Comment on Lakhani and Gilroy." In Curtis L. Gilroy (ed.), *Army Manpower Economics.* Boulder, CO: Westview Special Studies in Military Affairs, 1986

[46] G.S. Maddala. *Limited-Dependent and Qualitative Variables in Econometrics.* Cambridge: Cambridge University Press, 1983

[47] Thomas V. Daula and D. Alton Smith. "Recruiting Goals, Enlistment Supply, and Enlistments in the U.S. Army." In Curtis L. Gilroy (ed.), *Army Manpower Economics.* Boulder, CO: Westview Special Studies in Military Affairs, 1986

[48] James N. Dertouzos. "Microeconomic Foundations of Recruiter Behavior: Implications for Aggregate Enlistment Models." In Curtis L. Gilroy (ed.), *Army Manpower Economics.* Boulder, CO: Westview Special Studies in Military Affairs, 1986

[49] Badi H. Baltagi. *Econometric Analysis of Panel Data.* New York: Wiley, 1995

[50] Cheng Hsiao. *Analysis of Panel Data.* Cambridge: Cambridge University Press, 1986

[51] Badi H. Baltagi. "Applications of a Necessary and Sufficient Condition for OLS To Be BLUE," *Statistics and Probability Letters,* Vol. 8, 1989, pp. 457-461

[52] Robert H. Baldwin and Thomas V. Daula. "Army Recruit Attrition and Force Manning Costs: Methodology and Analysis." In Ronald Ehrenberg (ed.), *Research in Labor Economics (Vol. 7).* Greenwich, CT: JAI Press, 1985, pp. 311-338

[53] Ross L. Prentice and L. A. Gloeckler. "Regression Analysis of Grouped Survival Data With Application to Breast Cancer Data." *Biometrics,* Vol. 34, 1978, pp. 57-67

[54] John D. Kalbfleisch and Ross L. Prentice. *The Statistical Analysis of Failure Time Data.* New York: Wiley, 1980

[55] Samuel Johnson and Norman Kotz. *Distributions in Statistics: Continuous Univariate Distributions–1.* New York: Wiley, 1970

[56] Nancy R. Mann, Ray E. Schafer, and Nozer D. Singpurwalla. *Methods for Statistical Analysis of Reliability and Life Data.* New York: Wiley, 1974

[57] S. D. Dubey. "A New Derivation of the Logistic Distribution." *Naval Research Logistics Quarterly,* Vol. 15, 1969, pp. 37-40

[58] Dean A. Follmann, Matthew S. Goldberg, and Laurie J. May. "Personal Characteristics, Unemployment Insurance, and the Duration of Unemployment." *Journal of Econometrics,* Sep 1990

[59] 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.* Alexandria, VA: U.S. Army Research Institute for the Behavioral and Social Sciences, 1991

[60] Thomas V. Daula and Robert Moffitt. "Estimating a Dynamic Programming Model of Army Reenlistment Behavior." In Curtis L. Gilroy, David K. Horne, and D. Alton Smith (eds.), *Military Compensation and Personnel Retention: Models and Evidence.* Alexandria, VA: U.S. Army Research Institute for the Behavioral and Social Sciences, 1991

[61] John T. Warner and Beth J. Asch. "The Economics of Military Manpower." In K. Hartley and T. Sandler (eds.), *Handbook of Defense Economics, Volume 1.* New York: Elsevier Science, 1995

[62] Patrick C. Mackin. *Reestimation of ACOL Coefficients for the CAPM Model.* Falls Church, VA: SAG Corporation, 1996

[63] Matthew S. Goldberg. "Discount Rates for Government Investment Projects: The Economic Logic Behind OMB Circular A-94." *The Engineering Economist,* Summer 1998

[64] John T. Warner and Saul Pleeter. "The Personal Discount Rate: Evidence from Military Downsizing Programs." *American Economic Review,* Vol. 91, No. 1, Mar 2001, pp. 33-53.

[65] Stephen L. Mehay and Paul F. Hogan. "The Effect of Separation Bonuses on Voluntary Quits: Evidence From the Military's Downsizing." *Southern Economic Journal,* Vol. 65, 1998, pp. 127-139

[66] Larry G. Epstein and J. Allan Hynes. "The Rate of Time Preference and Dynamic Economic Analysis." *Journal of Political Economy,* Vol. 91, Aug 1983, pp. 611-635

[67] Dermot Gately. "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables: Comment." *Bell Journal of Economics,* Vol. 11, 1980, pp. 373-374

[68] Harry J. Gilman. *Determinants of Implicit Discount Rates: An Empirical Examination of the Pattern of Voluntary Pension Contributions of Employees in Four Firms,* 1976 (CNA Memorandum 76-0984-30)

[69] Jerry A. Hausman. "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables." *Bell Journal of Economics*, Vol. 10, 1979, pp. 33-54

[70] Emily C. Lawrance. "Poverty and the Rate of Time Preference: Evidence from Panel Data." *Journal of Political Economy*, Vol. 99, February 1991, pp. 54-77

[71] John T. Warner. "Comment on Daula and Baldwin." In Curtis L. Gilroy (ed.), *Army Manpower Economics*. Boulder, CO: Westview Special Studies in Military Affairs, 1986

List of figures

THIS PAGE INTENTIONALLY LEFT BLANK

Distribution list

Research Memorandum D0004085.A2

ASN (M&RA)
  Attn: PDASN (M&RA)
    DASN (Manpower)
OUSD (P&R)
  Attn: Dr. Curtis Gilroy
    Dr. John Enns
OUSD (Compensation)
  Attn: Captain Kopang
    Dr. Saul Pleeter
BUPERS MILLINGTON TN
  Attn: PERS-1
    PERS-4
    PERS-6
USNA ANNAPOLIS MD
NAVPGSCOL MONTEREY CA
  Attn: Library
    Professor Stephen Mehay
NAVWARCOL NEWPORT RI
NPRST MILLINGTON TN
RAND SANTA MONICA, CA
  Attn: Dr. Beth Asch
RAND ARLINGTON VA
  Attn: Dr. John Ausink
RAND PITTSBURGH PA
  Attn: Dr. James Hosek
IDA ALEXANDRIA VA
  Attn: Dr. Glenn Gotz
THE LEWIN GROUP FALLS CHURCH VA
  Attn: Mr. Paul Hogan
SAG CORPORATION ANNANDALE VA
  Attn: Mr. Patrick Mackin
CLEMSON UNIVERSITY CLEMSON SC
  Attn: Professor John Warner
UNIVERSITY OF YORK UNITED KINGDOM
  Attn: Professor Keith Hartley

OPNAV

N8
N12
N13T