Chapter Two

THEORETICAL FRAMEWORK AND EMPIRICAL MODELS

This chapter describes our theoretical framework and relates it to the empirical models we use in the data analysis. The theoretical framework provides a means to help explain why past deployment can affect a member's current decision to reenlist. The framework assumes that members have imperfect information about whether they will like or dislike deployment and that they face uncertainty about whether and for how long they will be deployed. The notion of learning from deployments about the frequency, duration, and utility of deployment is therefore a key to understanding why past deployment may affect reenlistment.

We assume that the member is interested in the expected utility of reenlisting for another term. Expected utility depends on deployment, but the features of deployment are not well known to the member. We first describe a mechanism for learning about deployment from actual deployment experience. Then, given the member's estimates of these features, we formulate a model of expected utility. The learning model illustrates how the member can learn from deployment experience, and the expected utility model illustrates how the features of deployment can affect expected utility. The learning and expected utility models are potentially estimable but not with available data. Thus, we rely on the models to clarify our understanding of the relationship between past deployment and reenlistment and to motivate the empirical work.

We describe the two kinds of models we estimate. These are a one-equation model of deployment and reenlistment and a two-equation model of promotion speed and reenlistment, both of which are dependent on deployment. Our working hypothesis is that the enlisted member does not influence the number and duration of deployments, but we discuss the alternative hypotheses that members can self-select deployment or that commanders handpick members for deployment.

LEARNING ABOUT EXPECTED UTILITY OF DEPLOYMENT

We assume that utility may be expressed as $u_d = f(y_d, \delta) + \varepsilon$, where y_d is income inclusive of deployment income, δ is a parameter affecting the level of utility, and ε is a random factor. (We make this function more explicit in the discussion of expected utility, below.) The member does not know the value of δ but knows that δ can take one of two values: $\overline{\delta}$ or $\underline{\delta}$. Utility when deployed is higher at $\delta = \overline{\delta}$ than at $\delta = \underline{\delta}$. The member learns through deployment experience about the probability that $\delta = \overline{\delta}$. The values of $\overline{\delta}$ and $\underline{\delta}$ can vary across members, reflecting heterogeneous tastes.

Deployment experience provides new information that allows beliefs to be updated. The member has a prior belief that the probability of $\delta = \overline{\delta}$ is π_o. The random factor ε has a zero mean and is identically and independently distributed through time with a single-peaked density $h(\varepsilon)$. Under the prior belief, expected utility when deployed is:

$$E u_d = \pi_o f\left(y_d, \overline{\delta}\right) + \left(1 - \pi_o\right) f\left(y_d, \underline{\delta}\right).$$

When a deployment occurs, the member realizes a level of utility $u_d = U_d$, and it is used to revise the prior π_o. Applying Bayes' Theorem, the posterior belief π_1 that $\delta = \overline{\delta}$ given $u_d = U_d$ is:

$$\begin{aligned}
\Pr\left(\delta = \overline{\delta} \mid u_d = U_d\right) &= \frac{\Pr\left(u_d = U_d \mid \delta = \overline{\delta}\right)\pi_o}{\Pr\left(u_d = U_d \mid \delta = \overline{\delta}\right)\pi_o + \Pr\left(u_d = U_d \mid \delta = \underline{\delta}\right)\left(1 - \pi_o\right)} \\
&= \frac{\Pr\left(f\left(y_d, \overline{\delta}\right) + \varepsilon = U_d\right)\pi_o}{\Pr\left(f\left(y_d, \overline{\delta}\right) + \varepsilon = U_d\right)\pi_o + \Pr\left(f\left(y_d, \underline{\delta}\right) + \varepsilon = U_d\right)\left(1 - \pi_o\right)}
\end{aligned}$$

$$= \frac{h\left(U_d - f\left(y_d,\bar{\delta}\right)\right)\pi_o}{h\left(U_d - f\left(y_d,\bar{\delta}\right)\right)\pi_o + h\left(U_d - f\left(y_d,\underline{\delta}\right)\right)\left(1-\pi_o\right)}.$$

The expression $h(U_d - f(y_d,\bar{\delta}))$ is the likelihood that the random term takes the particular value $\varepsilon = U_d - f(y_d,\bar{\delta})$ given $\delta = \bar{\delta}$. If, for example, the density $h(\varepsilon)$ is bell-shaped around zero and the value $\varepsilon = U_d - f(y_d,\bar{\delta})$ is approximately equal to zero, then this likelihood is high. By comparison, $h(U_d - f(y_d,\underline{\delta}))$ is the likelihood that the random term takes the particular value $\varepsilon = U_d - f(y_d,\underline{\delta})$ given $\delta = \underline{\delta}$. If, as mentioned, $\varepsilon = U_d - f(y_d,\bar{\delta})$ is approximately equal to zero and therefore has a high likelihood, then $\varepsilon = U_d - f(y_d,\underline{\delta})$ is likely to be farther from zero and have a lower likelihood. The higher likelihood of $\varepsilon = U_d - f(y_d,\bar{\delta})$ versus $\varepsilon = U_d - f(y_d,\underline{\delta})$ means that $\bar{\delta}$ fits the realized utility U_d better than $\underline{\delta}$. As a result, the posterior belief that $\delta = \bar{\delta}$ is higher than the prior belief—that is, π_1 is greater than π_o.[1]

An increase in the probability that deployment is a high-utility experience increases the expected value of δ and therefore expected utility. Because the likelihood of reenlistment depends on expected utility, the likelihood of reenlistment also increases. Furthermore, multiple deployments provide multiple opportunities to revise beliefs about deployment. If each deployment proved to be a positive experience, for example, the probability of reenlistment would rise with the number of deployments.[2] Similarly, the member can update his or her belief about the probability and duration of deployment.

The same framework can be applied to different types of deployment, such as those that involve hostile duty. Hostile deployments have higher danger, which could mean that $\bar{\delta}$ and $\underline{\delta}$ are both lower than they are for nonhostile deployments. By treating hostile and

[1]The analysis can be extended to allow the parameter to take a continuum of values, but this does not add insight.

[2]We identified the separation from family and friends as a generic aspect of deployment. The utility loss from this separation might change as deployments increase. For instance, a military spouse or close friend might get used to handling things on his or her own and become less distressed with each deployment.

nonhostile deployments separately, we allow the data to determine whether they have different relationships to reenlistment.

EXPECTED UTILITY

The member's willingness to reenlist depends in part on deployment, but future deployment is uncertain. Given this uncertainty, the member considers the expected utility of reenlisting. We assume that the member has subjective estimates of the frequency and duration of deployment, knows about deployment-related pays, and has a sense of the fixed and variable costs of deployment, e.g., arranging to have someone look after personal belongings and perhaps the cost of additional child care as the spouse copes with the member's absence. The member has preferences over the amount of time deployed versus time at home station, the variance of the number of deployments, and the variance of the duration of deployment.

We develop an expression for the expected utility of the term. We show that the expected utility can increase and then decrease as the expected deployments increase. Also, it can be positively or negatively related to the expected length of a deployment.[3] As mentioned, the connection between the expected utility of the term and the learning model is that, through past experience, the member learns about his or her preferences for deployment and about the mean and variance of deployments and deployment duration. The expected utility model provides a framework to put this learning to use.

Number of Deployments

The number of deployments during a three- to four-year term can be reasonably well described by a Poisson distribution. Given the actual

---

[3]The derivation of expected utility is conditional on the member's subjective estimates of the variance of deployments and deployment length and those preferences regarding time deployed and the variances of deployments and deployment length. We could extend the derivation of expected utility to take the expectation over these estimates by using the posterior distribution of their values, but no further insight would be gained.

distribution of deployments, we assume there are four possible outcomes: zero, one, two, or three deployments. The probability of four or more deployments is small enough to be negligible. We can see this by considering several Poisson distributions that reflect the observed distribution of deployments. The Poisson distribution has a single parameter, λ. For a given λ, the probability of n deployments (n = 0, 1, 2, ...) is $e^{-\lambda}\lambda^n/n!$. The mean and variance equal λ, and the probability of having one or more deployments rises with λ. Table 2.1 shows the probability of n deployments for values of λ that approximately correspond to the Army and Air Force ($\lambda = 0.5$) and to the Navy and Marine Corps ($\lambda = 1.1$).

Length of Deployment

We approximate the distribution of deployment length by a continuous distribution defined over the range of zero to 2μ. The mean duration is μ and the variance of duration is $\mu^2/3$. This is not as good an approximation as the Poisson is for the number of deployments, but it is good enough for our purpose of showing how learning can be applied to the expected utility calculation. The probability of a deployment of length s equals $1/(2\mu)$. We assume deployment lengths are independent of the number of deployments and also independent of each other. Therefore, the probability of n deployments of lengths $s_1, s_2, s_3, \ldots, s_n$ is $(e^{-\lambda}\lambda^n/n!)(1/(2\mu)^n)$. The total time deployed is $d(n, s_1, s_2, s_3, \ldots, s_n) = \sum_{i=1}^{n} s_i$.

Table 2.1

Deployment Frequency Based on Poisson Distribution

Deployments	$\lambda = 0.5$	$\lambda = 1.1$
0	0.61	0.33
1	0.30	0.37
2	0.08	0.20
3	0.01	0.07
4	0.002	0.02
5	0.0002	0.004

Deployment-Related Pay and Cost

The member receives a base income of m_o dollars for the term. When deployed, the member receives deployment pay of w' dollars per unit time and incurs a cost of c dollars per unit time, for a net deployment pay rate of $w = w' - c$ The member incurs a fixed cost k for each deployment. Total income is $m = m_o + wd - nk$. (We comment further on net deployment pay below.)

Utility Function

We assume utility depends on purchased goods *x*, time not deployed *h*, time deployed *d*, the variance of deployments λ, and the variance of deployment length, which is proportional to μ^2. For the Poisson and Uniform distributions, which have one parameter, a higher variance implies a higher mean. The parameters λ and μ are dictated by the needs of the service, although the member must estimate their values.

If the member were free to choose time deployed as well as purchased goods, he or she would select the values that maximize utility:

$$Max\, L = U(x, T - d, d) + \varphi(m_0 + wd - nk - x).$$

This leads to the first-order conditions:

$$\begin{aligned} U_x &= \varphi, \\ U_{T-d} &= U_d + \varphi w, \\ m_0 + wd - nk - x &= 0. \end{aligned}$$

The first condition states that the marginal utility of purchased goods equals the marginal utility of income (φ) multiplied by the price of purchased goods, which is assumed to equal one. In the second condition, the marginal utility of time at home station *(T – d)* is equated to the marginal utility of time deployed *(d)* plus the marginal utility of the net deployment-related pay. If net deployment pay is zero, the member would prefer an amount of time deployed such that its marginal utility equaled the marginal utility of time at home station. If net deployment pay is positive, the member prefers more time deployed. Even though additional time deployed can be assumed to have a lower marginal utility, the additional pay offsets the decrease in marginal utility. The third condition states that the member exhausts the budget constraint.

The first-order conditions implicitly define the member's demands for purchased goods and time deployed as functions of base income, net deployment pay, fixed cost of deployment, price of purchased goods (which we set equal to 1), and parameters of the utility function. We can use the indirect utility function to describe how utility is affected when the member cannot choose the time deployed but must accept what the service assigns. The member's utility is highest if the assigned time deployed equals the amount the member would have chosen according to the first-order conditions. Higher or lower levels of time deployed reduce utility relative to that optimum.

Because members have chosen to be in military service, it is reasonable to focus on interior solutions where the preferred time deployed is positive but does not use all available time. Moreover, because time deployed trades off against time not deployed, utility rises as time deployed increases from zero to its optimal value, and then declines as time deployed crowds out time at home station. We use a quadratic function to approximate this relationship.

As shown above, the number and duration of deployments depend on the distribution parameters λ and μ. The member's income and total time deployed depend on the number and length of deployments. Therefore, one possibility is to write utility, given the occurrence of n deployments of lengths $s_1, s_2, s_3, \ldots, s_n$, as

$$U(m, n, s_1, s_2, s_3, \ldots, s_n) = \log m + ad - bd^2 - c\lambda - f\mu^2.$$

Utility increases with a and decreases with b ($a, b > 0$). If there were no deployment pay or cost so income did not depend on deployment, the optimal amount of time deployed would be $a/2b$. The member's expected utility, developed below, is a weighted sum of the probability of the outcome of n deployments of given lengths multiplied by the utility of that outcome. Because the utility of any outcome with positive deployment ($n > 0$) is positively related to a and negatively related to b, it follows that expected utility is also positively related to a and negatively related to b. The learning model described how the member learned about his or her preferences for deployment (i.e., about the values of a, b, c, and f in the utility function) and about the mean and variance of the number and duration of deployment (i.e., about λ and μ).

Expected Utility Function

We form the expected utility function from its parts:

- The expected utility, given zero deployments:

$$EU_0 = m_0 - c\lambda - f\mu^2.$$

- The expected utility, given one deployment:

$$EU_1 = \int_0^{2\mu} (\log(m_0 + ws - k) + as - bs^2 c\lambda - f\mu^2) \frac{1}{2\mu} ds.$$

- The expected utility, given two deployments:

$$EU_2 = \int_0^{2\mu} \int_0^{2\mu} (\log(m_0 + w(s_1 + s_2) - 2k) + a(s_1 + s_2) - b(s_1 + s_2)^2 - c\lambda - f\mu^2) \frac{1}{(2\mu)^2} ds_1 ds_2.$$

- The expected utility, given three deployments:

$$EU_3 = \int_0^{2\mu} \int_0^{2\mu} \int_0^{2\mu} (\log(m_0 + w(s_1 + s_2 + s_3) - 3k) + a(s_1 + s_2 + s_3) - b(s_1 + s_2 + s_3)^2 - c\lambda - f\mu^2) \frac{1}{(2\mu)^3} ds_1 ds_2 ds_3.$$

Keeping the foregoing expressions in mind, the expected utility can be written compactly as:

$$EU = e^{\lambda} EU_0 + e^{\lambda} \lambda EU_1 + \frac{e^{\lambda}\lambda^2}{2} EU_2 + \frac{e^{\lambda}\lambda^3}{6} EU_3.$$

When the integrals for EU1, EU2, and EU3 are evaluated, we obtain an explicit form for expected utility. We have completed the integration and can use the explicit form to show how expected utility varies with its parameters: a, b, c, f, λ, and μ^2.[4]

This expected utility function is flexible enough to capture a variety of relationships between deployment and expected utility. We have already mentioned that expected utility increases with a and decreases with b, and that apart from net deployment pay, the member's preferred time deployed is $a/2b$. Because sailors join the Navy knowing they can expect a rotation of six months at sea and twelve months at home port, they probably have a higher value of a relative to b than do soldiers or airmen, who may not expect or prefer to be away as much. With respect to the learning model, we have argued that incoming members hold naive preferences about time deployed. Members learn more about deployment by being deployed, and, based on their experience, they may revise the prior values of their preferences. For instance, if a is revised upward because of a deployment, the level of expected utility rises and the member is more likely to reenlist. If a is revised upward with each deployment, the probability of reenlistment should rise with the number of deployments. This is not a necessary relationship, but it is a possibility that can be readily handled within the learning model and the expected utility model. Because our empirical work often shows an increase in reenlistment with the number of deployments, this is a relevant possibility to keep in mind.

We can also show that for reasonable parameters, expected utility is likely to increase with λ, up to a point, even though it has a negative direct effect on utility. (Again, λ equals the variance and the mean of the number of deployments.) Expected utility increases because an increase in λ increases expected time deployed, which initially has a high marginal utility.

Furthermore, we can show the relationship between expected utility and the variance of deployment duration. It seems reasonable that, controlling for the number of deployments, members with more actual time deployed are likely to revise upward their estimated mean and variance of deployment duration. (Again, the mean duration is μ and the variance of duration is $\mu^2/3$.) The theoretical

[4]The results of the integration are available from the authors on request.

model shows that a higher variance (or mean) can lead to either an increase or decrease in expected utility, depending on parameter values. In our empirical work, we find that time deployed has a positive effect on reenlistment for first-term Army members and a negative effect on reenlistment for first-term members in the other services. Yet, these seemingly contradictory effects are not inconsistent with the expected utility model.

We use two sets of parameter values to illustrate these points. Parameter set 1 has \$75,000 base income during the term, net deployment pay of \$75 per month, and a fixed cost of \$200 per deployment. The member prefers about 7.5 months deployed (a = 0.15, b = 0.01) and is indifferent to the variances of deployments and deployment duration (c = 0, f = 0). Parameter set 2 has a base income of \$75,000, net deployment pay of \$200 per month, a fixed cost of \$200 per deployment, a preference for about 15 months deployed (a = 0.30, b = 0.01), and an aversion to variance in the number and duration of deployment (c = 0.02, f = 0.04). Parameter set 1 roughly corresponds to a preference for time deployed that we might expect among soldiers and airmen, whereas parameter set 2 seems more descriptive of sailors and marines.

Figure 2.1 shows the relationship between expected utility and the variance of the number of deployments (λ) for parameter sets 1 and 2.[5] We show the relationship at several different values of the variance of deployment duration ($\nu \equiv \mu^2/3$). The similarity in the curves implies that the relationship is not sensitive to duration variance. With parameter set 1, expected utility rises with λ over the range of λ that seems relevant to soldiers and airmen. That is, from Table 2.1 we know that when $\lambda = 0.5$, about 40 percent of members have one or more deployments, and three-fourths of those members have a single deployment. Thus, in this range, learning that led a member to increase the subjective value of λ would be associated with higher expected utility, hence a higher probability of reenlistment.

[5]The left- and right-hand panels in Figure 2.1 use different scales. Because utility functions are unique only up to a monotone transformation, the scale is arbitrary. The main purpose of the figure is to show how expected utility varies with the variance of deployments. Similar comments apply to Figure 2.2, which shows expected utility with respect to the variance of deployment length.

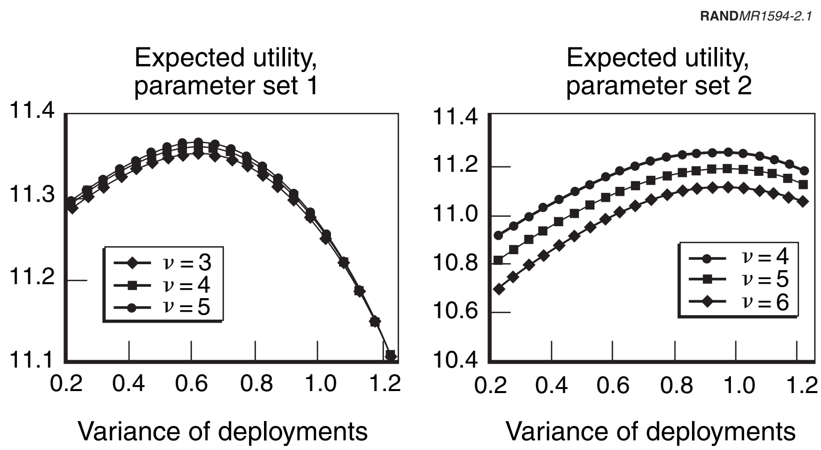

Figure 2.1—Relationship Between Expected Utility and the Variance of Deployments for Parameter Sets 1 and 2

With parameter set 2, expected utility is relatively flat in the range from λ = 0.85 to λ = 1.05, a range consistent with the number and variance of deployments sailors and marines might have expected when they signed up. However, for lower values of λ, expected utility declines. This suggests that if deployment was much lower than initially expected and preferred, λ would be revised down and expected utility would be reduced. A high amount of deployment, leading to a large upward revision of λ, would also be associated with lower expected utility.

Figure 2.2 shows the relationship between expected utility and the variance of deployment duration for each parameter set. With parameter set 1, expected utility rises with duration variance because the increase in variance implies an increase in mean duration; for these parameter values, expected utility rises with duration. With parameter set 2, expected utility declines with duration variance. In this case, a decline occurs because of the negative effect of the variance on expected utility, and also because the rise in variance implies an increase in mean duration, which leads to more time deployed than the member prefers.

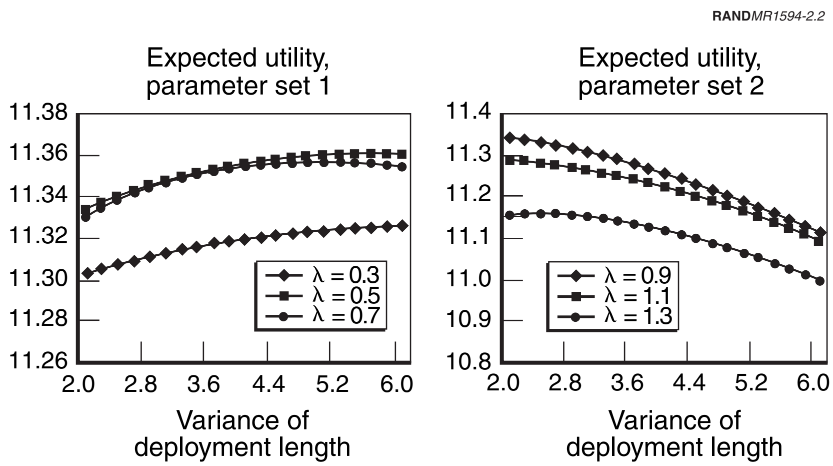

Figure 2.2—Relationship Between Expected Utility and the Variance of Deployment Duration for Parameter Sets 1 and 2

Net Deployment Pay

Because utility when deployed depends on net deployment pay, policy can affect utility by the use of deployment pays and by steps to reduce a member's fixed and variable cost of deployment.

Deployment pays—such as Family Separation Allowance (FSA), Hostile Fire/Imminent Danger Pay (referred to in this report as HFP), Combat Zone Tax Exclusion, Career Sea Pay, and (as of February 1999) Hardship Duty Pay—compensate for separation from dependents, unusual danger, arduous duty, and inhospitable circumstances. By increasing income during deployment, deployment pays reduce the extent to which these adverse aspects of deployment decrease a member's utility. Because deployment pays are set by policy, they should tend to be higher for more-demanding or riskier deployments. Because they are set ahead of time, they may not be well targeted for a particular deployment; however, it is easy for a member to factor them into expected utility.[6]

---

[6]We could not estimate the effect of deployment pays on reenlistment because their levels changed little in our data period.

In addition to compensating for the adverse aspects of deployment, deployment pays help offset a member's fixed and variable costs of deployment. A junior enlisted member with no dependents may have a low cost of deployment. However, if the member owns a car, has bills to pay (e.g., loan payments, telephone bills), or lives off base and has belongings to take care of (e.g., TV, disc player, dog), then arrangements must be made. These arrangements represent a fixed cost for each deployment, while handling the arrangements when deployed is a variable cost that continues for the duration of deployment. Married members can rely on their spouse to handle personal affairs, which suggests that the fixed and variable cost may be shifted to the spouse. The spouse may have to adjust his or her schedule in response to the member's absence (e.g., work fewer hours or less convenient hours; buy, rather than cook, more meals; use more baby-sitting; perform more home maintenance).

On base, family support services are available to help military spouses cope with the stress of separation and the added responsibility of running a household when the member is deployed. These services may be thought of as in-kind deployment pays. Family support services can, for example, put a military spouse in touch with counselors and provide suggestions regarding child-care providers. Family support groups create telephone trees to relay messages about the deployed unit members to their spouses and friends. Also, the services provide such communication links as e-mail and weekly telephone calls so that deployed members can stay in touch.

It follows from Bayesian updating that the net deployment pay affects the posterior belief π_1 that $\delta = \bar{\delta}$. For a given realized utility U_d, as net income increases, the value of π_1 *decreases*. This occurs because utility depends on net pay and the value of δ, and the same level of utility can be produced by a low income and high δ ($\delta = \bar{\delta}$) or a high income and low δ ($\delta = \underline{\delta}$). Therefore, for a given realized utility, the probability that $\delta = \bar{\delta}$ is lower when the level of income is higher.

Deployment and Promotion Speed

Deployment might also affect reenlistment through promotion speed. Because income, responsibility, and authority increase with rank, we assume that reaching the next rank faster increases expected utility and reenlistment. The expected utility model can be extended into a dynamic programming model of retention, but we do not make that transformation here.[7]

Deployment could affect promotion speed in several ways. Deployment could increase or decrease the amount of time available for reading and studying for promotion. It could also affect the member's willingness to exert effort toward promotion. A member might infer that future utility when deployed is higher at a higher rank. Also, the services may value deployment experience in making promotions to the extent that it results in decorations, awards, improved physical condition, greater skill and knowledge, or a higher rating of future potential.[8]

As with reenlistment, the relationship between deployment and promotion speed is an empirical matter. The effect may differ between nonhostile and hostile episodes of deployment. Hostile deployment may provide less off-duty time for the member and be more physically demanding, making it harder to prepare for promotion. However, hostile deployment might be more likely to be recognized by a decoration or award.

Summary

We have presented a learning mechanism that describes how a member might revise his or her prior beliefs about deployment and a model of expected utility that describes how a member can utilize that knowledge when deciding whether to reenlist. In particular, the member may learn about preferences for time deployed relative to time not deployed, preferences for the variance of deployments and the variance of deployment duration, and the variances (and means) themselves. We also showed how deployment pays and the fixed and variable costs of deployment enter into expected utility.

---

[7]Hosek and Totten (1998) put deployment, promotion, and reenlistment in the context of a dynamic programming model, building on the work of Gotz and McCall (1984) and Asch and Warner (1994).

[8]Williamson (1999) describes the services' enlisted promotion systems; every service takes into account the factors we mention above.

The learning and expected utility models provide a conceptual framework for connecting past deployments to a member's current reenlistment decision. The learning model allows prior beliefs to be revised up or down, and therefore does not imply any particular relationship between deployment and posterior beliefs. The expected utility model, as it has been specified, allows for a number of relationships that helped to motivate our empirical work and aided in interpreting the results:

- The expected utility model implies that an upward revision in the preferred time deployed (i.e., an upward revision in *a* and a downward revision in *b*) causes an increase in expected utility, hence in the probability of reenlistment. Therefore, if deployment typically led to an upward revision, that would be reflected by higher reenlistment.
- Depending on parameter values, an increase in the mean or variance of deployment may increase or decrease expected utility. The relationship between expected utility and the mean or variance of deployment is an inverted u-shape. For parameter values that seem relevant to members of the Army and Air Force, an increase in the mean or variance of deployment leads to an increase in expected utility. For parameter values that seem relevant to the Navy and Marine Corps, an increase in the mean or variance of deployment has little effect on expected utility. However, a sizeable reduction in the mean or variance of deployment is likely to reduce expected utility. A sizeable reduction might occur if, for example, a member entered the Navy or Marine Corps expecting a high rate of deployment by going to sea, but actually had no deployment.
- Depending on parameter values, an increase in the mean or variance of the length of a deployment might increase or decrease expected utility. Controlling for the number of deployments, more time deployed might cause an upward revision in the value of the mean or variance of deployment length. Because this could either increase or decrease expected utility, it is possible to observe a positive or a negative effect of time deployed on the probability of reenlistment.
- Expected utility is positively related to income. Income is higher with higher base pay and higher with rate of deployment pay but

is lower with higher fixed and variable cost of deployment. These relationships are potentially testable, but during our study period deployment pays were nearly constant. Also, we have no data on a member's cost of deployment. Therefore, the effect of deployment pay and cost is not observed directly but intertwined with the variables indicating deployment.

- Deployment might affect reenlistment by speeding up, or slowing down, the time to promotion. Faster promotion leads to higher pay, and service at a higher rank, with its greater authority and responsibility, might be more satisfying. Perhaps deployment is more satisfying when experienced at a higher rank; perhaps not. If deployment speeds up promotion, we would expect an increase in reenlistment.
- The preference for deployment may depend on the characteristics of deployment. For instance, deployment involving hostile duty might have as many or more positive aspects than deployment that does not involve hostile duty, but hostile deployment probably has more negative aspects (high stress, poor conditions, long hours, disease, combat risks).

EMPIRICAL MODELS

Our basic model of deployment and reenlistment is a probit regression. Let y_i be the member's propensity to reenlist and x_i represent the explanatory variables. In the probit model:

$$\begin{aligned} y_i &= \beta x_i + \upsilon_i \\ \upsilon_i &\sim N(0,1) \\ \Pr(\mathit{reenlist}) &= \Pr(y_i > 0) \\ &= \Pr(\beta x_i + \upsilon_i > 0) \\ &= \Phi(\beta x_i) \\ \Pr(\mathit{not\ reenlist}) &= 1 - \Phi(\beta x_i). \end{aligned}$$

The error term υ_i represents unobserved factors that influence the reenlistment decision. The error term is normally distributed with zero mean and unit variance, and $\Phi(\cdot)$ is the standardized normal distribution. In the data, each member's reenlistment decision and explanatory variables are observed. A likelihood function is created by multiplying together the probabilities of reenlistment for those who reenlist, and the probabilities of non-reenlistment for those who do not. The likelihood function is maximized with respect to the parameters β to obtain estimates of the parameters and their standard deviations.

The explanatory variables include indicator variables for the number and kind (hostile/nonhostile) of deployments over a three-year period ending three months prior to the date at which the member made a decision to reenlist or leave. There are indicator variables for one, two, or three or more nonhostile deployments, and one, two, or three or more hostile deployments. We also define interactions among the deployment indicator variables, and in exploratory specifications we enter the total months deployed in addition to the deployment indicators. Other explanatory variables include the member's Armed Forces Qualification Test (AFQT) category, education, occupational area, race/ethnicity, gender, dependency status, fiscal year in which the member's current term ends or a reenlistment decision is made, the unemployment rate at the start of the current term, and the current unemployment rate.

Our promotion/reenlistment model allows deployment to affect reenlistment directly and indirectly through its effect on promotion speed. Promotion speed is measured by t_i, the number of months to E-5 (the first noncommissioned officer rank).[9] The structure of the model is as follows:

$$t_i = \alpha z_i + \eta_i$$
$$y_i = \beta x_i + \gamma t_i + \upsilon_i$$
$$\begin{pmatrix} \eta_i \\ \upsilon_i \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^2 & \rho \\ \rho & 1 \end{pmatrix} \right).$$

Here, γ is an estimate of the effect of promotion time on reenlistment. If a longer time to promotion indicates a poorer fit with the military, we expect γ to be negative. The model allows for the pos-

[9]Unlike the other services, Navy promotions occur on a six-month cycle. Therefore, the time unit in the Navy is six months rather than one month.

sibility that unobserved factors affect promotion speed and reenlistment. Such factors may reflect the member's effort, ability, and commitment to military service. After controlling for the observed variables, if a shorter time to promotion is associated with a higher probability of reenlistment, the error correlation ρ will be negative.

Model estimation is complicated because many observations on promotion time are censored. Censoring arises when a member has not been promoted before leaving service or before the end of the data window. If t_i^c is the censoring date, the probability that promotion occurs after that date is:

$$\begin{aligned}\Pr(t_i \textit{ less than } t_i^c) &= \Pr(\alpha z_i + \eta_i > t_i^c) \\ &= \Pr(\eta_i > t_i^c - \alpha z_i) \\ &= 1 - \Phi(t_i^c - \alpha z_i).\end{aligned}$$

Because promotion time and reenlistment are assumed to follow a bivariate normal distribution, the model uses an expression for the joint probability of, say, reenlistment and censored time to promotion.[10] If a member does not reenlist, the promotion process is followed up to the time of departure. If a member reenlists, the promotion process is followed up to the time of promotion or the end of the data window.

To identify the effect of expected time to E-5 on reenlistment, the promotion equation includes some variables that are not in the reenlistment equation. These are indicators of whether the member was fast to the previous pay grade, E-4, and the quarter of the year in which the member entered service. The speed to E-4 is specified by indicators of whether the member's time to E-4 was in the 25th, 50th, or 75th percentile relative to those in his or her entry cohort who

---

[10]The probability of reenlistment and censored time to promotion is

$$\int_{-\beta x}^{\infty} \int_{t^c - \alpha z}^{\infty} \phi(0, \Sigma)\, d\eta\, du,$$

where ϕ is the normal density and Σ is the covariance matrix.

reached E-4.[11] Otherwise, the variables in the promotion equation include AFQT, education, occupational area, and fiscal year.

WHAT IF DEPLOYMENT IS SELECTIVE?

We think the assumption that deployment is exogenous to the member is appropriate for our empirical analysis of first- and second-term reenlistment. Junior enlisted members typically have little say in choosing their assignments and missions. However, Wardynski (2000) raised the possibility that a member or the member's commanding officer can affect whether the member deploys.

If a member influences deployment, the influence will be directed toward increasing the level of expected utility. Members who want more deployment will seek to increase their deployment, while those who prefer less deployment will seek the opposite. Therefore, if members can self-select, the probability of reenlistment should increase for those who want more deployment and those who do not. It is unknown whether the difference between their reenlistment probabilities would widen or narrow. We note this because the empirical analysis contrasts the reenlistment probability of deployers to nondeployers. Self-selection would not necessarily make deployers appear more likely to reenlist than nondeployers and therefore would not necessarily bias upward the effect of deployment on reenlistment.

The commanding officer presumably seeks to exclude from deployment those members with poor attitudes or poor performance. The commander's scope for culling the ranks depends on whether replacements can be found, if needed, to keep unit manning at the level required for the deployment. Some excluded members might have preferred not to deploy, so exclusion would increase their expected utility. Other excluded members might have preferred to deploy, so exclusion would reduce their expected income. If the commanding officer removed members with poor attitudes or poor performance who were unlikely to reenlist, the average reenlistment probability would increase among members who deploy and decrease among members who do not deploy. This may increase or

---

[11]We also estimated the model without the E-4 indicators, as discussed below.

reduce the estimated effect of deployment on the reenlistment probability. If the reenlistment probability of deployers is initially higher than that of nondeployers, then commander-selection increases this positive difference. If the reenlistment probability of deployers is initially lower than that of nondeployers, then commander-selection reduces this negative difference and possibly creates a positive difference. Thus, commander-selection may exaggerate or reduce the estimated effect of deployment on the reenlistment probability, depending on the initial values.

Anticipating the empirical results, we found that reenlistment tended to increase with nonhostile episodes of deployment. It is possible that commander selection biased upward an already-positive relationship between deployment and reenlistment. We also found that reenlistment was little affected by hostile deployments. If the true effect of hostile episodes on reenlistment were negative, commander selection might have changed the negative effect to a zero effect.

We should also ask whether commander selection of those who deploy affects our ability to identify the effect of expected time to E-5 promotion on reenlistment. This effect is identified by the two variables in the promotion equation that are not in the reenlistment equation, namely, time to E-4 and quarter of accession. Members with short times to E-4 promotion (controlling for AFQT, education, and occupational area) are high performers and may have high tastes for the military. If so, the expected time to E-5, which depends on time to E-4 and presumably, taste, might be correlated with the error term in the reenlistment equation. This could bias its coefficients. If a longer expected time to E-5 reduced reenlistment, the bias would probably make this negative relationship steeper.

Although these are possibilities, we have no firm evidence on the role played by commander selection or self-selection. If the role is minor and perhaps negligible, as we suspect, there should be little effect on our estimates.