The Heterogeneous Value of a Statistical Life: Evidence from U.S. Army Reenlistment Decisions*

Kyle Greenberg[†], Michael Greenstone[‡], Stephen P. Ryan[§], and Michael Yankovich[¶]

September 29, 2021

Abstract

We estimate the value of a statistical life (VSL), or the willingness to trade-off wealth and mortality risk, among 430,000 U.S. Army soldiers choosing whether to reenlist from 2002-2010. Using a discrete choice random utility approach and significant variation in retention bonuses and mortality risk, we recover *average* VSL estimates between \$500,000 and \$900,000, an order of magnitude smaller than U.S. civilian labor market estimates. We then document substantial heterogeneity by recovering indifference curves between wealth and mortality risk. The VSL increases with mortality risk within type, and soldiers in combat occupations have lower VSLs than those in noncombat occupations.

*We thank Daron Acemoglu, Josh Angrist, Orley Ashenfelter, David Autor, Eli Berman, Dan Black, Joseph Doyle, Jim Heckman, Christopher Knittel, Nicolai Kuminoff, Derek Neal, Heidi Williams, Tyler Williams, and participants in MIT's Labor Lunch, the NBER's Economics of National Security Working Group, and West Point's Economics Seminar Series for their helpful comments and suggestions. We thank Jinglin Yang, Claire Fan, Michael Galperin, Ananya Kotia, Jonathan Petkun, Dan Stuart, and Evelyn Yankovich for their exceptional research assistance. We also acknowledge generous and exceptional support from Luke Gallagher of the Army's Office of Economic and Manpower Analysis at West Point, New York. The views expressed in this paper are those of the authors and do not reflect the official policy or position of the Department of the Army, the Department of Defense, or the U.S. Government.

[†]U.S. Military Academy at West Point, kyle.greenberg@westpoint.edu.

[‡]University of Chicago and NBER, mgreenst@uchicago.edu.

[§]Washington University in St. Louis, CESifo, and NBER, stephen.p.ryan@wustl.edu.

[¶]U.S. Military Academy at West Point, michael.yankovich@westpoint.edu.

1 Introduction

A topic of extraordinary interest to researchers and policymakers is how to estimate individuals' willingness to pay for goods and services for which there isn't an explicit market (Greenstone, 2017). Rosen's (1974) seminal "hedonics" paper has provided the foundation for thousands of studies estimating the value of a wide range of "implicit" goods and services, including clean air (e.g., Chay and Greenstone (2005)), government policies (e.g., Gruber (1994)), school quality (e.g., Bayer et al. (2007)), alternative work arrangements (e.g., Mas and Pallais (2017)), and other job characteristics (e.g., Stern (2004); Lamadon et al. (2019); Folke and Rickne (2020)). The value of a statistical life (VSL), which is an individual's willingness to trade-off changes in wealth for changes in the probability of death, is a canonical example due to the first-order importance of mortality risk; it also features centrally in Thaler and Rosen (1976), who provide a framework for inferring preferences from the equilibrium relationship between wages and mortality risk. Further, the VSL is of tremendous practical value because it provides a measure of the benefits accruing from a wide range of policies that improve safety, such as public health mandates, environmental regulation, occupational safety standards, speed limits, and product safety requirements.

The resulting empirical VSL literature (see, for example, Ashenfelter and Greenstone (2004); Viscusi (2004, 2010); Evans and Taylor (2020)) has had enormous policy influence, but the interpretation and credibility of the estimates have been criticized by Black and Kniesner (2003), Ashenfelter (2006), and others for at least three reasons. First, it has failed to match the ambition of the Rosen (1974) and Thaler and Rosen (1976) framework by almost universally treating the VSL as a constant and thereby side-stepping the need to recover the bid functions that reflect heterogeneity in preferences across individuals and heterogeneity in the VSL within individuals at different mortality rates. This simplification, which is common in the broader hedonics literature, invalidates the strong casual observation that some people (e.g., Army enlistees) are more comfortable with safety risks than others, which in turn helps explain the sorting of individuals into occupations. Additionally, it rules out the standard assumption that there is declining marginal utility in the consumption of a good; in practice, this assumption means that the demand curve for safety is not downward sloping. Second, fatality risks are often correlated with other unobserved features of jobs and it has proven quite challenging to identify settings where it is possible to credibly disentangle the effect of mortality risk from these other factors.[1] Third, the mortality risk may be known imperfectly

[1] Lee and Taylor (2019) aim to solve this problem by instrumenting for mortality risk with randomly assigned federal safety inspections and estimate VSLs between $8.5 and $10.6 million. Rohlfs et al. (2015) use quasi-experimental evidence from changes to air-bag regulations to derive median VSL estimates between $10.6 and $13.0 million. All dollar values in this paper are converted to 2019 U.S. Dollars (USD).

to agents choosing jobs, complicating the interpretation of the empirical results.

This paper proposes and implements a novel approach to estimating the VSL that makes progress on all three of these challenges. Using detailed administrative data, we model the reenlistment decisions of the universe of 430,000 U.S. Army soldiers who completed their first term of service between 2002 and 2010. The military is a particularly salient setting for the application of the VSL, where decisions about personal and vehicular armor, the development of drones, the selection of operating procedures, soldier and veteran compensation, and a variety of other policies often come down to a choice between higher costs and lower fatality risks. The urgency of getting these choices right is especially great in the post-Vietnam era of an all-volunteer military that relies on reenlistments to fill senior ranks.

Our sample period offers a compelling setting to observe individuals' choices when faced with widely varying levels of compensation and risk. Regarding wealth, the Army uses lump-sum reenlistment bonuses to balance manpower needs with staffing levels. These bonuses vary widely by occupation and over time, are established through a centralized Army personnel process, change discontinuously, and are constant within an occupation at any point in time (i.e., not subject to negotiation). For example during the 2002-2010 period, the reenlistment bonus ranged between \$0 and \$18,600 (mean of \$4,100) for a 4-year reenlistment as a combat engineer, while it varied between \$0 and \$37,900 (mean of \$11,000) for human intelligence collectors. On the risk dimension, mortality rates varied tremendously both across and within specialties during the time period, because the sample begins in peacetime and extends through the height of the Afghanistan and Iraq wars, when the overall Army mortality rate more than quadrupled. There was also substantial variation in the change in mortality rates across occupations; for example, the four-year mortality rate for infantrymen increased by 600 percent between 2002 and 2007 (from 3.5 deaths per 1000 soldiers to 24.6 deaths per 1000 soldiers), while it increased by 145 percent (from 2.0 deaths per 1000 to 4.9 deaths per 1000) for wheeled vehicle mechanics over the same time period.

The foundation of our analysis is an empirical model of the decision to reenlist. Specifically, we use random utility discrete choice models where the bonus and fatality risk are the key covariates and a rich set of occupation- and soldier-specific controls are available. The (negative) ratio of the marginal utilities on the mortality rate and the reenlistment bonus is the VSL; this is the amount of money required to hold reenlistment rates constant for a marginal increase in mortality.

There are three primary findings. First, higher bonuses increase the reenlistment rate and higher mortality rates reduce it. In the preferred specification, a \$1,000 increase in the bonus for a 4-year term raises the probability of reenlistment by 0.46 percentage points. To put this in context, the overall reenlistment rate in this period was 45.3 percent, these soldiers earned about \$30,000 annually, and average bonuses increased by \$11,000 between 2003 and 2006. In the other direction, an increase in the expected 4-year mortality rate of 1 per 1000 decreases the reenlistment rate by roughly 0.26 percentage points. As a basis of comparison, we estimate that the overall expected 4-year mortality rate increased by 7.2 soldiers per 1000 from the low in 2002 to the peak in 2007. Overall, we interpret these results as evidence that reenlistment decisions are highly elastic to bonuses but less elastic to mortality rates.

Second, the *average* VSL estimated across the population of soldiers and observed range of mortality rates is generally between \$500,000 and \$900,000 across a wide set of specifications. These estimates are precise and at least at an order of magnitude smaller than the typical VSL estimate from overall U.S. labor market studies (Viscusi and Aldy, 2003; Kniesner et al., 2012; Banzhaf, 2021).[2] It is striking that the estimates are little affected by the inclusion of the rich set of controls available in these data, which is generally not the case in the hedonic studies of the civilian labor market (Black and Kniesner, 2003). Further, the results are qualitatively similar whether the reenlistment decision is modeled with a binary logit model with no unobserved heterogeneity, a moment forest binary logit model that uses a novel machine learning algorithm of Nekipelov et al. (2021) to relax parametric assumptions about the covariates, a binary random coefficients model that allows for unobserved heterogeneity in the bonus and mortality risk coefficients, or multinomial logit models that provide a richer characterization of the range of reenlistment possibilities.

Third, we fulfill Rosen's (1974) vision to recover indifference (or bid) curves between wealth and non-market goods (mortality risk in this setting) and find substantial heterogeneity in the marginal willingness to pay for safety within and across types of soldiers. Specifically, we model the bonus and mortality risk variables flexibly (i.e., with B-splines) to allow the marginal response of those variables to depend on their levels. We then determine indifference curves, based on the pairs of bonuses and mortality rates that return equal utility; the VSL is the slope along an indifference curve. Among men in combat occupations, the VSL varies by more than an order of magnitude within the observed range of mortality risk: the VSL is \$180,000 at a 4-year expected mortality rate of 1 per 1000 (1st percentile of their distribution of mortality rates), \$370,000 at 8 per 1000 (45th percentile) and \$2.02 million at 24 per 1000 (98th percentile). Further, we find evidence of substantial taste-based sorting, even within the Army, such that individuals with higher risk tolerance sort into the riskiest occupations: the VSL is more than 7 times greater for men in noncombat occupations versus men in combat

[2] Among draft-age men from the Vietnam Era, Rohlfs (2012) estimates marginal responses to the Vietnam Draft on the decision to attend college and the decision to voluntarily enlist to produce upper bound VSL estimates of \$8-14 million.

occupations (i.e., \$2.808 million compared to \$370,000) at an expected 4-year mortality rate of 8 per 1000.

The recovery of indifference curves is of great practical value because the welfare consequences of non-marginal changes and policy counterfactuals are not possible without them. For example, the loss in soldiers' surplus from the increase in the mortality rate due to the Afghanistan/Iraq wars would be understated by 40 percent if we used the average VSL derived from the entire Army sample rather than the separate indifference curves for men in combat occupations, men in noncombat occupations, and women.

Finally, this setting has several institutional features that help address the literature's challenges. First, we model reenlistment decisions with an individual-level discrete choice random utility model, allowing us to recover a rich and flexible distribution of VSLs. Specifically, we estimate specifications where the marginal utilities of wealth and risk vary with their levels, so we can reconstruct the theoretical indifference curves trading off risk and financial compensation as originally envisioned by Rosen (1974) and Thaler and Rosen (1976). This is a marked departure from the empirical VSL literature which to date has mostly (perhaps entirely) focused on estimating the market locus between wages and mortality risk; the market locus's limitation is that its economic interpretation is unclear for all but marginal and uniform changes in mortality rates across the population. In contrast, the recovery of indifference curves is the great promise of the hedonics literature because it allows for the complete welfare analyses that the empirical VSL and other hedonic literatures have been unable to conduct.[3]

Second, we believe the potential for omitted variables bias due to unobserved job characteristics is less of a concern in our setting than it is in most empirical VSL settings. The richness of the data allows us to adjust for many covariates, including year of initial enlistment, fixed effects for 145 occupations, initial entry term length (e.g., 4 years), the probability of deployment to a combat zone, age, gender, race, education, state of legal residence, Armed Forces Qualification Test (AFQT) score, and the local unemployment rate in the soldier's home county.[4] In our most demanding specifications, the models include the full interaction of fixed effects for every combination of year of initial enlistment, occupation, and initial entry term length. In this specification, the coefficients on the bonus and mortality rate are identified from variation in these variables among cohorts of soldiers who entered the Army

[3]We are not aware of any other papers that recover indifference curves within a hedonic context.

[4]The AFQT score represents the percentile-rank (1-99) of a soldier's arithmetic and verbal reasoning skills relative to a nationally representative sample of 18-23-year-olds (Department of Defense, 2004). Borgschulte and Martorell (2018) find that service member reenlistment decisions are responsive to home-state unemployment rates and reenlistment bonuses.

in the same year, in the same occupation, and for the same term of enlistment. For example, one of these cohort cells compared the reenlistment decisions of soldiers in 2006 who signed up in 2002 for four-year terms as combat infantryman (i.e., occupational specialty 11B). Within this cell, the identifying variation in mortality risk comes from the rolling 12 preceding month's average occupation-specific mortality hazard and in bonuses from the irregular changes in SRBs. Furthermore, the similarity of our results derived from the moment forest framework, which makes no assumptions about the functional form of the available covariates, lends additional credibility to our estimates.

Third, with respect to concerns that fatality rates are unknown, there is a single employer making it easier to share information, mortality is a part of daily life in the Army, and fatal casualties are published regularly online and in newspapers. Occupational fatalities are not only common, but also highly variable in our setting: between 2003 and 2010, the standard deviation of annual Army fatality rates in our sample was 14 times larger than the standard deviation of annual on-the-job fatality rates in U.S. construction and extraction occupations (i.e., the major civilian occupational group with the highest annual variation in fatality rates during the same period).[5] Additionally, information on the risk of service was widely available: 61 percent of soldiers in our sample deployed to a combat zone; recent evidence suggests 90 percent of soldiers who deployed prior to 2010 experienced some form of exposure to combat; and intelligence briefings, unit exercises, and required training provided soldiers with ample information on the risks associated with combat (Hosek and Martorell, 2009). Indeed, our informal discussions with soldiers confirm that they are aware of fatality rates and how they vary across occupations and within occupations over time, consistent with reports from soldiers in the popular press (see, for example, Aitken (2021)). Finally, to address concerns that our preferred measure of expected mortality risk does not adequately capture how soldiers form expectations of risk, we show that our results are robust to several alternative formulations of the mortality hazard that potentially account for different ways that soldiers might infer risk, including alternative hazards derived from a regression tree machine learning algorithm that does not stipulate how soldiers infer mortality.

[5]Counts of deaths by major occupational group come from the Occupational Injuries and Illnesses and Fatal Injury Profiles in the Injuries, Illnesses, and Fatalities (IIF) Databases maintained by the US Department of Labor, Bureau of Labor Statistics (see https://www.bls.gov/iif/, accessed on September 9, 2021). We divide these counts by estimates of the number of total workers by major occupational group from the Employment Projections tables published on February 28, 2012 by the US Department of Labor, Bureau of Labor Statistics (see https://www.bls.gov/opub/ted/2012/ted_20120228_data.htm, accessed on 10 September 2021). If we calculate Army on-the-job fatality rates from deaths directly attributable to combat alone, then the standard deviation in annual Army fatality rates is still 13 times larger than the standard deviation in fatality rates within construction and extraction occupations, but some of the additional variation in Army fatalities may be indirectly attributable to combat.

We proceed as follows: Section 2 draws on the theory of Thaler and Rosen (1976) to emphasize how recovering indifference curves for wealth and risk reveals important heterogeneity in the VSL; Section 3 details the setting for our empirical analysis; Section 4 provides details about the data and presents summary statistics; Section 5 describes our estimating framework; Section 6 presents and discusses results from standard logit approaches and our moment forest analysis framework; Section 7 reports on our recovery of the indifference curves envisioned by Thaler and Rosen and how they can be used to infer the welfare consequences of non-marginal changes in mortality risk; and Section 8 concludes.

2 The Economics of the VSL

Thaler and Rosen's (1976) canonical paper provides the economic underpinnings of the observed relationship between wages and mortality risk that has been noted since Adam Smith's observation that people must be induced to take risky jobs through a set of compensating differences in wage rates. Drawing from Rosen (1974), they derive a model of worker and firm optimizing behavior that results in an equilibrium relationship between wages and mortality risk. This section briefly reviews their contribution and uses it to interpret the extensive empirical VSL literature.

In their model, a job is comprised of a vector of characteristics, $C = (c_1, c_2, \ldots, c_n)$, which might include vacation days, strenuousness, exposure to extreme heat or cold, and mortality risk. The wage rate of job i can be written as: $P_i = P(c_{i1}, c_{i2}, \ldots, c_{in})$. The partial derivative of $P(\cdot)$ with respect to the j-th characteristic, $\partial P/\partial c_j$, is referred to as the marginal implicit price. On the worker side, Thaler and Rosen derive indifference, or bid, curves for workers that reveal the wage rate necessary to hold utility constant at different levels of job characteristics.[6] Heterogeneity in tastes for job characteristics and non-labor market income cause these indifference curves to vary across individuals. The other side of the market is comprised of employers that are heterogeneous due to differences in their costs of providing alternative job characteristics. In equilibrium, workers and employers assortively match, generating a market locus between wages and a given characteristic. For example, workers with a high tolerance for risk sort into firms that find it expensive to provide safe work environments.

Our setting differs from the canonical hedonic framework in that there is only a single employer. However, within the Army there are many occupations with different job characteristics. We assume that the Army's ability to provide alternative job characteristics

[6] We use the terms "indifference curve" and "bid curve" interchangeably throughout the paper.

requires investments that occur over a longer time horizon, and thus we treat them as fixed and restrict our attention to workers.

Since our interest is in estimating the VSL, we focus specifically on mortality risk and wealth. Figure 1 illustrates the endogenous sorting that occurs across occupations by workers with different tastes for risk. The heavy line in Figure 1 plots the market locus between wages and mortality risk or the market equalizing-difference wage function. It also plots indifference curves for three types of workers, denoted as types #1, #2, and #3. For worker type #2, two different indifference curves are depicted; it is apparent that utility level b is greater than a because the wage rate is higher at a given mortality risk. Each indifference curve reveals the standard declining marginal rate of substitution between mortality risk and the wage rate.

The three types choose jobs where their marginal willingness to pay for safety is equal to the market determined marginal implicit price, which occur at the pairs c_j^1, c_j^2, and c_j^3, respectively. Given market prices, the utility of these workers would be lower at jobs with higher or lower levels of mortality risk. Further, a comparison of the indifference curves across types reveals that workers in risky jobs have a comparative advantage in bearing mortality risk due to their relatively low distaste for mortality risk.

Along the market locus, there are a set of tangencies with worker indifference curves, so the marginal price of mortality risk is equal to an individual's marginal willingness to pay for safety (i.e., the negative of mortality risk). Therefore, the gradient of the market locus with respect to mortality risk gives the equilibrium differential that compensates workers for accepting greater mortality risk. Put another way, jobs with high levels of mortality risk must have higher wages to attract workers, and the locus reveals the price that allocates workers across jobs.

This figure illustrates several important issues, all of which have at their core that there is not a single VSL. First, the slope of the market locus reveals the average of the VSLs (i.e., the marginal willingness to pay for safety) across an unknown number of worker types facing different levels of mortality risk. This is because the market locus is comprised of a collection of single points from an unknown number of worker types on each worker type's respective indifference curve at different risk levels. It is noteworthy that the enormous empirical VSL literature has almost exclusively focused on estimating the market locus, although its economic interpretation is much more limited than the estimation of the demand primitives or indifference curves that Thaler and Rosen originally envisioned.

Second, the market locus is not informative about an individual's or type's willingness to pay for a non-marginal change in mortality risk. For example, the market locus suggests that worker #1 would be willing to accept an increase in mortality risk from c_j^1 to c_j^2 for an increase in wages equal to $w_2 - w_1$. However, their indifference curve reveals that the true required increase in compensation is $w_4 - w_1$. The broader point is that the locus cannot reveal willingness to pay for non-marginal changes in mortality risk for any class of workers, except in the extreme case where there is no heterogeneity in tastes for mortality risk. This limits its value, because it means that the locus cannot be used to make inferences on the welfare consequences of alternative or counterfactual policies.

Third, the focus on estimating the market locus can lead to the false impression that there is limited heterogeneity in the population. For example, Figure 1 depicts the market locus as a straight line, suggesting that each of the worker types has the same VSL. However, a much more meaningful measure of heterogeneity in risk preferences comes from comparing the VSL across worker types at a fixed level of mortality risk. The figure depicts the slope of all three worker types' indifference curves at mortality risk level c_j^2. At this (or any other risk level), it is apparent that there is great heterogeneity in risk preferences with worker type #1 having the highest VSL and worker type #3 having the lowest. It is therefore not surprising that worker type #1 has sorted to the safest job and worker type #3 to the riskiest one.

The full hedonic method includes a two-step approach to recover the demand primitives and resulting indifference curves that allow for a complete welfare analysis (Rosen, 1974).[7] However, Brown and Rosen (1982), Bartik (1987), and Epple (1987) describe the strong assumptions necessary to identify these structural parameters. From a practical perspective, there is a consensus that researchers have been unable to identify a single wage or other hedonic setting where these assumptions hold (Deacon et al., 1998; Chay and Greenstone, 2005).[8]

The remainder of this paper describes and implements our approach to developing new estimates of the VSL derived from a discrete choice random utility model of soldiers' reenlistment decisions that exploits substantial variation in both wealth (through reenlistment bonuses) and mortality risk. This is an alternative approach from the typical hedonic one where wages are regressed against mortality rates (see Mrozek and Taylor 2002; Viscusi and Aldy 2003; Cropper et al. 2011; O'Brien 2018 for reviews).[9] An important appeal of this

[7]In the context of the VSL, the first step of this method is to regress wages on all job characteristics, allowing for non-linear effects. To uncover the bid function, the second step is to regress the estimated relationship between wages and mortality risk on different values of risk while controlling for all covariates that might influence demand for safety.

[8]Ekeland et al. (2004) outline the assumptions necessary to identify the demand (and supply) functions in an additive version of the hedonic model with data from a single market. Heckman et al. (2010) examine identification and estimation of nonadditive hedonic models. Heckman et al. (2003) examine the performance of estimation techniques for both types of models.

[9]León and Miguel (2017) and O'Brien (2018) also use discrete choice methods to infer the VSL, but they

approach is that we are able to recover the demand primitives, resulting indifference curves, and marginal willingness to pay for mortality risk functions that were the goals of Thaler and Rosen's original model for, as far as we are aware, the first time in the VSL literature and possibly the broader hedonics literature. As we have noted, this is so important because it allows for assessing the welfare consequences of *non-marginal* changes in mortality risk.

3 Institutional Setting

3.1 Background on the Army's Use of Labor

The setting for our analysis is the U.S. Army between 2002 and 2010 when the number of people serving on active duty in the U.S. Army averaged between 475,000 and 550,000 soldiers. Each year, approximately 65,000 non-prior service young people volunteer to serve in the Army for a period of between two and six years. When an individual enlists in the U.S. Army, s/he chooses his/her term of service and selects an occupation. Once enlisted and subsequently trained, the Army chooses a soldier's unit of assignment based on the needs of the Army (Lyle, 2006; Lleras-Muney, 2010).[10]

The organization of work is generally structured to encompass hierarchical job ladders in which soldiers serving at higher levels of responsibility and skill must come from lower levels of the internal labor market. There is no lateral entry. Thus, the Army relies on a substantial proportion of each new cohort to reenlist for additional periods of obligated service in order to meet its requirements for higher-ranked, experienced workers. The Army has relatively fixed demand for soldiers in the short run, and such demand is defined by requirements at the level of occupation and rank.

With no lateral entry and the need for higher-ranked individuals, the Army must balance enlistment and reenlistment rates with attrition. In the long run, the Army can modify tables of organization and equipment so as to substitute equipment for people. The military can also change strategy or doctrine, essentially altering the way that it conducts operations, so as to require different amounts of labor. However, these types of changes extend over

do not infer indifference curves. There is also an extensive literature that uses individuals' stated preferences to infer the VSL (Cropper et al., 2011).

[10]Though some first term enlisted soldiers may have some scope over their choice of duty location conditional on occupation, they arguably have no scope over the choice of unit of assignment or of deployment timing. Once a soldier completes the initial entry training associated with his occupation of choice, the needs of the Army dictate which unit the soldier serves with and when the soldier deploys. As such, variation of within-occupation risk exposure is generally orthogonal to the soldier's preferences for deployment or risk, conditional on the soldier's choice of occupation.

the course of several months and in many cases years. Adjustments to compensation are a natural tool to obtain the necessary number of soldiers in each occupation. However, Congress controls annual salaries and they institute changes slowly and at the cohort level, not the occupation level, so these changes generally do not help fill relative shortages. Thus, the Army's only compensation tool to address short-run, and even medium-run, staffing needs are enlistment and reenlistment bonuses. The next section discusses the reenlistment process and the central role bonuses play in them.

3.2 The Reenlistment Process

As a matter of level-setting, roughly 50,000 soldiers annually were eligible to reenlist between 2002 and 2010. The annual reenlistment rate during these years varied between 36 and 55 percent. About a year before the end of a soldier's initial contractually obligated term of service, she enters the reenlistment eligibility window that is depicted in Figure A1. During this window, soldiers confront the choice of reenlisting in the active force or finishing their term of service and exiting active duty. All soldiers participate in a mandatory career counseling program during which they are exposed to options for reenlisting for a follow-on term in the active military as well as to information about outside options including civilian employment and educational opportunities. This counseling takes place quarterly starting two years prior to the end of the soldier's initial contract, as prescribed in The Army Retention Program, Army Regulation 601-280.

The reenlistment options are characterized by a triplet of occupation, home post location (i.e., their permanent military base) and term-length, as seen in the bottom part of Figure A1. With respect to the first, soldiers can choose to remain in their current occupation (i.e., military occupation specialty or MOS) or retrain and switch to another occupation if their current one is over-strength and their preferred one is under-strength. Roughly 90 percent of soldiers who reenlist choose to remain in their current occupation. Additionally, soldiers can allow the Army to choose their military base, they can elect to stay at their current post location, or they can choose to serve at a post that is either overseas or inside the continental United States if the Army has a requirement for soldiers in the post where they wish to serve. About one-third of reenlistees choose to remain at their current home post location. None of the reenlistment options restrict a soldier's potential for selection to deploy to a combat zone as units from all locations deployed to Iraq and Afghanistan during the time period of our analysis, and the needs of the Army, which are not known to reenlisting soldiers, dictate the precise timing and location of unit combat deployments (Lyle, 2006). The third dimension of choice is on the term-length, which can vary between 2 and 6 years, with 4 years being the most common (see Table A1) and 2 years not being available to soldiers who select a new occupation or who choose a new post.

Through the Army's Selective Reenlistment Bonus (SRB) program, the Army's Human Resources Command (HRC) can offer bonuses to increase reenlistment in occupations "that do not have adequate retention levels to staff the force" (Army Regulation 601-280). Whether a soldier receives a bonus offer and the size of the bonus offer depends on, the occupation they would choose, how much time the soldier has served, their rank, and how many years the soldier is willing to extend his service on active duty. During the period of our study, nearly all soldiers who reenlisted into occupations eligible for bonuses received a lump-sum bonus payment at the time they signed their contract.[11] Bonus offers do not depend on any of the soldiers' personal characteristics or record within a category, so there is no discretion available to Commanders or others to alter a bonus for a particular soldier. Bonuses are subject to federal income tax except when the soldier reenlists in a combat zone.

The Army's HRC uses a computer program to determine reenlistment bonuses for each occupation according to the Army's demand for experienced soldiers and the Army's expected supply of experienced soldiers. Demand for soldiers is a function of the Army's authorized strength in each occupation and available funding. The Army calculates the expected supply of experienced soldiers as a function of the number of first-term soldiers currently in each occupation, historical attrition and reenlistment rates, and anticipated soldier training requirements. Occupations with higher expected shortages typically receive higher bonus offers than occupations with smaller expected shortages.

Finally, the Army announces changes to bonus offers through Military Personnel (MILPER) messages at frequent but unexpected intervals that do not follow an established schedule. For example, between 2002 and 2010, the Army released over 90 MILPER messages related to changes in bonuses, with some years experiencing substantially more announcements than others (e.g., 16 in 2004 vs. 5 in 2010).

[11] The Army began paying reenlistment bonuses as a single lump-sum in 2004. Prior to 2004, the Army paid some bonuses as a single lump-sum but paid others through the following schedule: a lump-sum payment worth 50 percent of the bonus at the time of reenlistment, then the remainder of the bonus paid over equal, annual installments through the duration of the reenlistment contract. We cannot observe which bonuses were not paid as a single lump-sum, but our results are not sensitive to assuming pre-2004 bonus offers have a slightly lower present value than bonus offers in 2004 or later.

4 Data and Descriptive Statistics

Our sample consists of 429,375 non-prior service, first-term Army enlisted soldiers who became eligible for reenlistment between 2002 and 2010.[12] We conduct our analysis with a database of military administrative data that links the complete menu of choices available to an Army soldier at their reenlistment decision with their demographic characteristics, variables that describe their experiences during service, their choice-specific reenlistment bonus and mortality rate, and their actual reenlistment choice. The remainder of this section describes how we map soldiers to bonus offers and measures of mortality risk for each reenlistment option before concluding with sample summary statistics.

4.1 Construction of Bonus Offers

We used publicly-available MILPER messages to construct a database of time, location, rank, and occupation-specific reenlistment bonus offers that provides the menu of bonuses soldiers faced at any point in time. After constructing an exhaustive database of bonus offers for each occupation from 2002 through 2010, we matched the individuals in our sample to the bonus offer database using a crosswalk based on the soldier's occupation, years-of-service, rank, and years of (potential) reenlistment.

To avoid concerns that soldiers strategically game the reenlistment process by holding out for a higher bonus offer, we assign soldiers to the bonus they would have received if they selected to reenlist 12 months prior to their initial Expiration of Term of Service (ETS) date since this corresponds to the month that most soldiers enter their reenlistment window and is the modal month of reenlistment (see Figure A2). Going forward, we refer to 12 months prior to the initial ETS date as the month a soldier enters her reenlistment window, but we acknowledge that sometimes the Army permits soldiers to reenlist more than 12 months prior to their initial ETS date—such decisions can apply to all soldiers or to soldiers in particular occupations or locations. Other times the Army will extend soldiers' ETS dates, which results in some soldiers reenlisting after their initial ETS date. In the analysis that follows, we find that VSL estimates exhibit little sensitivity with regard to different rules for the date that a soldier decides whether to reenlist (e.g., 9 months prior to the initial ETS date, 12 months prior, 15 months prior, etc.).

For reasons of parsimony and computational speed, we reduce the soldiers' choice set in a couple of ways. First, we treat the choice of new occupation as the average across all potential

[12] We exclude soldiers who served for less than one year and soldiers who did not reach the rank of Private First Class (PFC) by the last year of their initial enlistment. Neither group is eligible for reenlistment.

occupations that a soldier could choose. Second, we treat the choice of a new overseas home post location as a single post, averaging bonus offers across all possible overseas locations. Third, we treat the choice of a new home post location inside the United States as the average bonus offer across domestic stations. Sections 2.1 and 2.2 of the data appendix provide more details on how we develop the database of bonus offers, including how we treat bonuses that apply to soldiers in specific locations and bonuses that apply to soldiers with specific skills.[13]

Three institutional features of bonuses are particularly important in our setting. First, the Army establishes bonuses through a centralized personnel system that prevents any negotiation of bonus offers for specific soldiers. Second, both bonus levels and changes in bonus amounts were much larger for some occupations than others, as seen in Figure A3. For example, bonuses for infantry rose from around \$10,000 in 2002 to \$16,000 in 2006 and then declined to \$5,000 in 2010, whereas bonuses for combat medics rose from close to \$0 in 2002 to \$13,000 in 2006 before declining to \$3,000 in 2010. Third, soldiers are well informed about changes to bonuses as every Army unit has a Retention Non-Commissioned Officer who is responsible for informing soldiers of the latest bonus-related MILPER messages.

4.2 Construction of Fatality Rates

The other determinant of reenlistment decisions that we focus on is the expected mortality rate, or hazard, that soldiers associate with their various reenlistment choices. This requires us to develop reliable measures of their expectations about future mortality risk for each of their potential choices. We assume the ex-ante expected hazard is a function of a soldier's occupation and reenlistment timing. Consider a soldier facing a decision to reenlist into occupation k who is exactly 12 months prior to her initial ETS date in month t, where 12 months corresponds to the set of bonus offers we map to each soldier as described in Section 4.1. Our preferred approach is to estimate this soldier's expectations of mortality risk for that reenlistment alternative as the average annual mortality rate, per 1000 soldiers, of soldiers in occupation k over month t and the previous 11 months:

$$p_{kt} = \frac{\sum_{n=t-11}^{t} fatalities_{kn}}{\left(\sum_{n=t-11}^{t} strength_{kn}\right)/12} \cdot 1000, \tag{1}$$

[13]The Army will occasionally offer bonuses to soldiers with particular skills (e.g., parachutist) and to soldiers who are stationed at a particular base or unit. Further, some bonus offers are specific to soldiers who are deployed to a combat zone. Any bonuses based on skills, current location, or deployment status are still governed by the Army's HRC. See Asch et al. (2010) and the data appendix to this paper for additional details on the SRB Program.

where $strength_{kn}$ is the number of soldiers assigned to Army occupation k in month n.[14]

A practical challenge is that reenlistments range from 2 to 6 years, so the expected mortality risk associated with a reenlistment option depends on the contractual length of service. In contrast, the bonus is paid all at once for the entire term of service. For ease of interpretation it is important that the reenlistment estimation equations include variables for the total bonus and the total expected mortality risk. We therefore developed a measure of the present value of the full expected mortality risk that reflects greater risk for longer periods of service and also discounts future risk relative to current risk. The discounting is motivated by the fact that across a wide range of contexts individuals exhibit behavior consistent with valuing the future less than the present.

The literature has little to say on how soldiers or others value current versus future expected mortality risk; sufficient treatment of this topic is beyond the scope of this paper. However, Simon et al. (2015) estimate that enlisted soldiers who served after 2000 have discount rates of 7.2 percent while they are in the service. Guided by this finding, we develop an accumulated expected mortality risk that discounts future risk at a rate of 7.2 percent per year:

$$(p_{ktjl}) = \sum_{s=0}^{l-1} \frac{(p_{ktj})}{1.072^s} \tag{2}$$

where p_{ktjl} is the accumulated expected mortality risk for occupation k, in month t, for choice alternative (i.e., reenlistment option) j, and reenlistment term-length l, and p_{ktj} is the single-year measure of expected mortality risk for choice alternative j as estimated from equation (1).[15] We use this measure of accumulated mortality risk in the summary statistics and results that follow. Section 2.3 of the data appendix offers more details on how we construct this preferred measure of accumulated mortality risk to represent soldiers' expectations of future mortality. This includes a discussion on how we use the Army's broader occupational classification, the career management field (CMF), to construct mortality rates for Military Occupational Specialties (MOSs) with few soldiers.

We considered several alternative measures of a soldier's expected mortality hazard to account for different methods of inferring risk. These include alternatives that vary the population from which soldiers might infer risk (e.g., from other soldiers in the same narrow

[14]We constructed death rates from Army loss files. Each time a soldier departs active service, an entry is created in the loss file which records a full set of administrative characteristics about the type of discharge, including whether the soldier died. To validate the accuracy of mortality through the loss file, we compared it to a separate database containing all combat-related deaths that have occurred since September 2001.

[15]We calculate the expected mortality risk for the choice of a new occupation as the average mortality risk across all potential occupations that a soldier could choose.

occupational grouping (MOS), from other soldiers in the same broad occupational grouping (CMF), or from soldiers in all occupations), alternatives that place more weight on the months closer to a soldier's reenlistment decision date, alternatives that consider mortality rates in the 24 months prior to a soldier's reenlistment decision, and alternatives derived from mean-reverting time series measures. We also explored several alternatives derived from a regression tree machine learning method that constructs a predicted mortality hazard from data on monthly mortality rates at the MOS level in the 12 or 24 months leading up to a soldier's reenlistment decision.

Table A2 documents that all formulations of the expected mortality hazard exhibit a strong correlation with occupational mortality rates in the 12 months following a soldier's reenlistment decision, with correlation coefficients generally ranging from 0.65 to 0.73. One notable exception is the Army-wide mortality hazard (i.e., the mortality rate of all soldiers in the Army with no attempt to group soldiers into similar occupations), which has a correlation coefficient with future mortality of 0.206. The machine learning alternatives account for four of the five hazard measures that are most strongly correlated with future mortality, but a key drawback of the machine learning measures is that our data do not permit us to train the sample on mortality for soldiers with 2010 decision dates. It is also not clear that soldiers would form their expectations of future mortality in the same way that a machine learning algorithm would.

Ultimately, our preferred measure is the mortality rate of soldiers in the same occupation in the 12 months leading up to a soldier's reenlistment decision. This measure is a straightforward, parsimonious measure that captures how we think soldiers form their beliefs about mortality at the time of their reenlistment decision. It is also a very good predictor of future MOS-level mortality—it is more strongly correlated with future mortality than three of the eight machine learning formulations and all but two of the non-machine learning formulations (see Table A2). Regardless, as we discuss below, the paper's findings are qualitatively similar across a wide variety of choices for the expected mortality risk variable, including the machine learning formulations.

There are several appealing aspects of this setting for assessing how mortality risk affects individuals' employment choices. First, Figure A3 shows that, like bonus offers, there is substantial variation in mortality hazards both across and within military occupations over time. For example, the overall Army mortality rate rose from 2.06 per 1000 in 2002 to 9.26 per 1000 in 2007 and then declined to 5.36 per 1000 in 2010; and these movements were much larger for some occupations like infantry, than for others like supply specialists.[16] Second,

[16] Average annual mortality rates for infantry soldiers rose from 3 per 1000 soldiers in 2002 to nearly 25

in contrast to much of the previous VSL literature, this is a setting where it is plausible to assume that people have reliable information on the mortality risk associated with different employment choices. This is because there was widespread publication of military deaths by both internal military information dissemination networks and external public news media. Additionally, our informal discussions with soldiers (two of us are active duty soldiers) confirm that they are aware of fatality rates and how they vary across occupations and over time. As such, the scope for measurement error is relatively small in the military setting, unlike hedonic wage approaches which rely on aggregated occupation and industry classification data from civilian labor markets. Third, it seems appropriate to assign expected mortality rates at the occupational level in our setting because a soldier's personal exposure to risk is less a function of her individual actions and more a function of her unit's mission, and both the unit she is assigned to and that unit's mission is determined by the "needs of the Army" conditional on her occupational choice. We think it is reasonable to assume that Army needs are orthogonal to any soldier's individual decision to reenlist.

4.3 Summary Statistics

Table 1 summarizes key explanatory variables and observable characteristics of soldiers in our sample. Column (1) displays summary statistics for the whole sample. Panel A reports on a series of renlistment statistics. Overall, 45.3 percent of eligible soldiers chose to reenlist. The average bonus associated with a four-year reenlistment in the same MOS was $7,200. The average four-year, discounted mortality rate associated with a four-year reenlistment was 5.43 per 1000 soldiers, implying an annual (non-discounted) mortality rate of 1.50 per 1000 soldiers. The average annual probability of deployment, which we define as the percentage of soldiers in an individual's MOS who were deployed for at least one month in the year before the individual enters her reenlistment window, was close to 21 percent.

Panel B reports demographic characteristics of the sample. Roughly 17 percent of soldiers in the full sample are women, 17 percent are Black, 13 percent are Hispanic, and the majority are high-school graduates who did not attend college. The average AFQT score is 60, or 10 percentiles higher than the national AFQT median of 50. Further, the average age upon entering the reenlistment window is 24 years. Finally, the average unemployment rate in their home county in the year of their reenlistment window was 6.5 percent.

We also examine subsamples by grouping individuals into one of three gender-by-occupation classifications: men in noncombat occupations, men in combat occupations, and women

---

per 1000 in 2007, then declined to 11 per 1000 in 2010, while average annual mortality rates for supply specialists remained between 1 and 3.5 per 1000 soldiers from 2002 through 2010.

in all occupations (mostly noncombat).[17] Column (2) reports summary statistics for men in noncombat occupations while columns (3) and (4) report statistics for men in combat occupations and women in all occupations, respectively. On average, women reenlist at lower rates than men, and men in combat occupations at lower rates than men in noncombat occupations. Of interest, mortality rates among men in combat occupations are nearly three times as high as mortality rates among men in noncombat occupations and women. Men in combat occupations also have slightly higher bonus offers and deployment probabilities than other soldiers in our sample.

Table 2 reports key compensation and reenlistment variables on an annual basis. Even though Army soldiers faced increased mortality risk from 2002 through 2007 in the aftermath of the invasion of Iraq in 2003, with the average mortality rate increasing from 2.06 per 1000 in 2002 to 9.26 per 1000 in 2007, reenlistment rates grew by more than 30 percent during the same period. Reenlistment bonuses appear to have played an important role in achieving the required manning goals as the average value of bonus offers across all soldiers (i.e., those who did and did not receive a bonus) rose from \$1,772 in 2003 to \$12,952 in 2006 and total expenditures on bonuses also increased sharply. It is noteworthy that regular military compensation during this period did not noticeably increase until 2009, after the mortality rate began to decline (see column (3)). This suggests the Army relied on bonuses to manage manpower as the political system, which plays a more direct role in determining regular military compensation, did not respond in real time to the increase in risk.

5 Estimating Framework: Discrete Choice

5.1 Empirical Approach

We use a random utility model of discrete choice to model a soldier's decision whether to select among a range of reenlistment options or to exit active military service. At the time of the reenlistment decision, each soldier makes a single utility-maximizing choice from a set of J mutually-exclusive reenlistment alternatives plus the option of leaving the Army. The utility of soldier i choosing reenlistment alternative j is given by:

$$u_{ij} = U(b_{ij}, h_{ij}, x_{ij}; \delta_i, \gamma_i, \beta_i) + \epsilon_{ij}, \tag{3}$$

[17]Noncombat occupations include MOSs related to administrative, intelligence, logistical, medical, and communication functions. Combat occupations include air defense artillery, armor, aviation, combat engineer, field artillery, infantry, and special forces. Roughly 93 percent of women in our sample are in noncombat occupations. Seven percent of women served in air defense artillery and aviation occupations that were exempt from the female combat exclusion policies in effect during the time of our analysis.

where b_{ij} is the choice-specific bonus, h_{ij} is the choice-specific expected mortality hazard rate, x_{ij} is a vector of individual- and choice-specific observables (such as the deployment probability, state of residence fixed effects, local area unemployment rates, and soldier-specific demographics such as their AFQT scores, age, race, gender, and education level), and ϵ_{ij} is an idiosyncratic shock which is assumed to be distributed Type I Extreme Value. The function is parameterized by (individual-specific) δ, γ, and β. The value of leaving the Army is normalized to have a utility of zero.

The resulting probability of soldier i selecting reenlistment choice j is given by the familiar logit form:

$$Pr\,(i \text{ chooses } j) = \frac{U(b_{ij}, h_{ij}, x_{ij}; \delta_i, \gamma_i, \beta_i)}{1 + \sum_k \exp\left(U(b_{ik}, h_{ik}, x_{ik}; \delta_i, \gamma_i, \beta_i)\right)}. \tag{4}$$

The VSL is defined as the amount of money that a soldier is willing to accept for a marginal increase in the expected mortality rate such that utility remains constant. In our discrete choice setting, this naturally translates into the (negative) ratio of the marginal utility of the mortality rate over the marginal utility of the bonus in Equation (4):

$$VSL(h_{ij}, b_{ij}; \delta_i, \gamma_i) = -\frac{\left(\dfrac{\partial U(b_{ij}, h_{ij}, x_{ij}; \delta_i, \gamma_i, \beta_i)}{\partial h_{ij}}\right)}{\left(\dfrac{\partial U(b_{ij}, h_{ij}, x_{ij}; \delta_i, \gamma_i, \beta_i)}{\partial b_{ij}}\right)}. \tag{5}$$

In general, the VSL is a function and depends on the level of the both the hazard and the bonus. The dependence of the utility parameters allows for heterogeneity at the individual level (which may or may not be observable to the econometrician).

We estimate a variety of specifications of the discrete choice model. The simplest specifications are binary choice (stay in the Army or leave) with additively separable linear utility functions, implying that the VSL in Equation 5 is a constant. In many respects, the assumption of a homogeneous VSL is consonant with the previous empirical literature, however it is quite inconsistent with Rosen and Thaler's original conception of the VSL that was defined by heterogeneity in preferences for safety across workers. For example, there could be important differences in risk preference between individuals who initially enlist to serve in the infantry and those who initial enlist to serve as a line cook. One way that we will allow for this heterogeneity is to estimate separate models for the full sample, men in noncombat occupations, men in combat occupations, and women in all occupations and calculate a separate VSL for each of these groups.

To account for other sources of observable heterogeneity, we include a rich set of occupation and soldier characteristics as fixed effects. Among the potential sets of fixed effects, the analysis focuses attention on two. The first includes cohort (i.e., year of entry), occupation (i.e., 145 different MOS categories) and initial-entry term-length (e.g., 4 years). The specifications that control for these fixed effects adjust for all unobserved level differences in utility associated with cohort (e.g., changes in the sense of patriotism) that are common across enlistees in a year (i.e., entry year), occupation (e.g., any social rewards for volunteering for particular jobs), and the chosen initial term (e.g., perhaps those choosing 2 years may have different expectations about their future or qualify for different benefits post-service). Thus, the identifying variation from these specifications come from comparisons within an occupation over time, after accounting for time-varying factors common to those signing up for an initial term in the same year and factors common to people who initially sign up for the same number of years. Therefore, these models utilize differences in the within occupation variation in the expected mortality hazard rate and bonuses over time.

The second includes fully interacted occupation-by-cohort-by-term fixed effects. These specifications are very demanding of the data because the coefficients of interest are only identified from differences in the initial enlistment dates within an entry year, and the entry year is several years (generally 3) prior to the reenlistment decision, within these narrow cells. Thus, for example, the identification comes from comparing the reenlistment decisions of soldiers in 2006 who signed up in 2002 for four-year terms as combat infantryman (i.e., occupational specialty 11B). Within this cell, the identifying variation in mortality risk comes from the rolling 12 preceding month's average occupation-specific mortality hazard and in bonuses from the irregular changes in SRBs.

We explore heterogeneity in the VSL in a few other ways. First, we use the moment forest machine learning method described in Nekipelov et al. (2021) to estimate a simple binary logit specification with three parameters (a constant, the bonus, and the expected mortality), but permitting the moment forest to vary estimates of these parameters across all other observable characteristics. Moreover, the moment forest approach serves as a robustness check on our logit estimates because it estimates parameters for subsets of the population that are determined by a machine learning process and not researcher intervention. Second, we explore specifications with normally-distributed random coefficients on the retention bonus and fatality rate to control for unobserved heterogeneity in the marginal utilities across soldiers. The models with unobserved heterogeneity generate a distribution of VSLs which are calculated by integrating over the distribution of unobserved heterogeneity and the joint empirical distribution of bonuses and mortality rates, respectively.

Finally, it seems plausible and indeed likely that there is declining marginal utility in income and safety. For this reason, we consider models where we replace the value of those variables with cubic B-spline basis functions. B-splines allow the marginal utilities to vary smoothly with the level of the variable, are easy to implement, and can be constrained to impose monotonicity or convexity. Relative to regular polynomials, B-Splines have the advantage of being numerically stable and locally adaptive.

We use the results from the B-spline approach to develop indifference curves that are comprised of bundles of bonuses and mortality risk that return the same level of utility. The VSL is the slope along these indifference curves. Because the B-spline approach allows for varying marginal utilities, we can assess whether the VSL varies with the level of risk (i.e., whether the MWTP function is constant). This connects the empirical literature to the original intent of the Thaler and Rosen model and allows for counterfactual analysis with more realistic assumptions about the shape of indifference curves. We recover these indifference curves and MWTP functions for the full sample, as well as the subsamples of men in noncombat occupations, men in combat occupations, and women in all occupations.

5.2 Potential Threats to Identification

The cleanest method for estimating the effect of bonus offers on reenlistment is to isolate the demand-side factors that influence bonus offers, a point discussed in Borgschulte and Martorell (2018). Unfortunately, the Army's demand for soldiers is governed by personnel structure documents that often take several months or even years to change. As a result, short-term fluctuations in the supply of enlisted soldiers drive much of the variation in reenlistment bonuses. This raises the possibility that unobserved determinants of reenlistment may be correlated with bonus offers. There are a couple of reasons to believe why this potential bias is not large. Most importantly, commanders do not have the ability to adjust bonuses for specific soldiers, thus negating the possibility of targeting soldiers who are on the margin of reenlistment, who are strong performers, or who possess unique skills. Even though commanders cannot influence bonuses for specific soldiers, the Army might target bonuses to occupations that have experienced recent drops in reenlistment rates. Such a scenario would imply higher bonus offers to soldiers who are less likely to reenlist, which would bias our estimates of the reenlistment bonus downwards, implying an upward bias in our VSL estimates. However, the frequent and often dramatic changes in bonus levels (see Figure A3) suggest that the Army's Selective Reenlistment Bonus program is effective because bonuses can change rapidly at unannounced intervals, not because a single bonus offer is particularly good at producing a precise reenlistment rate.

Another limitation is that just like with the bonuses, experimental variation in fatality risk is unavailable. This raises the possibility that the resulting estimate of the effect of mortality risk on reenlistment confounds mortality risk with other factors. Perhaps, the most likely confounder is unobserved expectations about the amount of combat a soldier will face upon reenlistment. Whether more unobserved expected combat increases reenlistment is ambiguous because it involves higher non-fatal casualty rates which is a negative but also the possibility of seeing more action or excitement in combat which is a positive for at least some soldiers.

We experimented with including a variable for the expected injury rate, measured as the percentage of soldiers in an MOS who were wounded in action over the past year, but found that it is not possible to separately identify the effects of the expected mortality and injury rates (nor are we confident soldiers can distinguish between them).[18] The inability to statistically distinguish between the effects of on-the-job injury and mortality risk is a common problem in the VSL literature and our read of this literature is that it is uncommon to include both variables in the same estimating equation. Our preferred solution to avoid confounding mortality risk with the role that expectations for action play in reenlistment decisions is to control for occupational deployment rates. Our results below include specifications that do and do not adjust for this variable, but our preferred specification includes it.

Ultimately the paper's claims for causality rest on conditioning on the rich set of covariates described above. Our data permits us to estimate very rich models, including some specifications with fully interacted occupation-by-cohort-by-term fixed effects, which means we are relying on within-year variation in bonus offers and mortality rates among soldiers who enlisted into the same occupation at similar times. Our specifications also condition on occupation-specific mortality rates, which help control for risk, and include additional controls for occupation-specific deployment rates, several individual characteristics such as AFQT score, education level, race, sex, and home state of record, and home-county unemployment rates. Nevertheless, we cannot rule out the possibility that unobserved determinants of reenlistment covary with offered bonuses or mortality rates.

6 Empirical Estimates of the VSL

Our analysis begins with a relatively parsimonious set of models to recover the VSL before progressing to richer specifications. Section 6.1 starts with the simplest binary choice model, where the dependent variable is an indicator for whether the soldier chose to reenlist. We consider several different specifications that vary the number of fixed effects that enter the model, as well as a series of other controls. Section 6.2 then considers several specifications

[18]The expected mortality and injury rates are highly correlated, having a correlation coefficient of 0.93.

which allow the parameters to vary by sub-population, exploits the recent moment forest technique of Nekipelov et al. (2021) to estimate arbitrary group-level heterogeneity in the VSL, and reports results from a random coefficients binary logit approach. Since the binary choice framework aggregates all reenlistment possibilities into a single choice, Section 6.3 turns to a series of multinomial choice models that allow for a richer choice set along both reenlistment options and reenlistment term-lengths.

6.1 Full Sample Evidence on the VSL: Evidence from Binary Choice Logit Models

Panel A of Table 3 summarizes results from the estimation of the binary choice model on the full sample of 429,375 soldiers considering reenlistment. The table reports results from five different specifications for reenlistment, all of which include the maximum bonus offer associated with a four-year reenlistment (the most common reenlistment term-length) in the same occupation and the expected mortality hazard rate associated with a four-year reenlistment in the same occupation. The table reports their logit coefficients and their marginal effects evaluated at the mean of the covariates. Importantly, the table also includes the estimated VSL (i.e., the negative of the ratio of the mortality and bonus logit coefficients) reported in millions of 2019 dollars and its 95 percent confidence interval.

The specifications differ with regard to the set of included covariates and are arranged so that as one moves from left to right the specifications become increasingly rich, which is reflected in the log likelihood statistic reported at the bottom of panel A. Column (1) includes fixed effects for military occupational specialty (MOS), entry year (cohort), and initial-entry term-length. Column (2) adds the deployment probability control described above. Column (3) adds controls for individual demographic characteristics, including gender, race, education, the home state of the soldier at the time of their initial enlistment, age, and AFQT score. Column (4) adds to column (3) by including the county level unemployment rate in a soldier's home of record in the month that the soldier enters the reenlistment window to account for the influence of local unemployment rates on reenlistment decisions, as revealed in Borgschulte and Martorell (2018). Column (5) replaces the fixed effects from columns (1) - (4) with fully interacted occupation-by-entry year-by-term fixed effects, while still including all other controls from column (4). Standard errors are clustered on each combination of MOS and reenlistment-decision month, the level of variation in our mortality measure.

The first row of Panel A reports logit coefficient estimates and standard errors for the bonus offer, and the second row reports the average marginal effect and its standard error. Across all specifications in the full sample, the coefficient on the bonus offer has the expected sign, would easily be judged statistically significant by conventional criteria, and is remarkably stable across the various specifications that use very different sources of variation for identification and adjust for different controls. The estimates suggest that increasing the bonus offer for a four-year reenlistment by $1000 increases the probability of reenlistment by about 0.48 percentage points.[19] Thus, holding all else constant, the estimates suggest that the increase in the average bonus offer from $1,800 in 2003 to its peak of $13,000 in 2006 increased reenlistment rates by 5.2 to 5.4 percentage points. An alternative way to put this in context is to note that the annual salary for a first-term soldier in this period was roughly $30,000, so a 1-time payment of roughly 35 percent of annual salary increases four-year reenlistment rates by about 5 percentage points. This response is more than 10 percent of the overall reenlistment rate of 45 percent in our sample. Overall, it is apparent that reenlistment decisions are highly responsive to cash bonuses.

The next set of rows in Panel A report logit coefficient estimates and average marginal effects on the mortality hazard rate, as well as their standard errors. This coefficient also has the expected sign across all specifications: increasing the four-year mortality rate by 1 death per 1000 soldiers decreases the probability of reenlistment by roughly 0.24 to 0.28 percentage points. The stability of the coefficient is striking, especially to the inclusion of the deployment probability variable and the different sets of fixed effects. There is no absolute standard here, but this appears to be a relatively small effect. To see this, consider the increase in mortality rates from 2.06 per 1000 soldiers (over a four-year reenlistment) in 2002 to 9.26 per 1000 soldiers at its peak in 2007. Holding all else constant, including bonus offers, the estimates suggest that the full wartime increase in mortality risk reduced reenlistment rates by only 1.7 to 2.0 percentage points.

The estimate of the implied VSL reported in Panel A of Table 3 is the (negative) ratio of the logit coefficient on the expected mortality hazard divided by the coefficient on the bonus offer. Its corresponding 95 percent confidence interval is presented below (in brackets) and is calculated using the delta method. Point estimates of the VSL range between about $520,000 and $600,000. The large sample size is evident in the 95 percent confidence intervals,

[19] This estimate is in the range of several other studies that find that reenlistment bonuses increase reenlistment among U.S. service members. Borgschulte and Martorell (2018) and Patterson et al. (2020) estimate marginal responses to bonuses that are slightly smaller than our estimates while Hosek and Miller (2011) estimate slightly larger marginal responses. Our point estimate is closest to Borgschulte and Martorell (2018) who find that a 10 percent increase in earnings through reenlistment bonuses (roughly a $12,000 bonus) increases reenlistments by 3.7 to 5.0 percentage points. Modest differences in point estimates across studies are likely due to differences in cohorts analyzed and other sample differences. For example, Borgschulte and Martorell (2018) exclude women and Patterson et al. (2020) include soldiers on their second and third enlistment terms.

which are tight especially in the context of the VSL literature (see, for example, Black and Kniesner (2003) and Kniesner et al. (2012)). Our estimates in column (4) indicate that the VSL for the population of young people who volunteer for active service in the U.S. Army has a 95 percent confidence interval covering \$343,000 to \$696,000. Although our richest specification in column (5) has considerably less power, even here we find a comparable 95 percent confidence interval covering \$231,000 to \$849,000.[20] Finally, we note that the column (2) specification will be our preferred one and the one we carry along throughout the analysis because Army computing limitations prevented us from including additional covariates in the estimation of some of the more flexible models (i.e., the multinomial logit and nested logit models as well as the B-spline analysis) that we discuss below. The similarity between the estimates in columns (2) and (5) support this choice.

Further exploration suggests that the binary logit estimates and resulting estimates of the VSL are robust to alternative modeling decisions. Table A4, which reports estimates when we consider bonus offers and expected mortality hazards faced by soldiers at different time-frames prior to their original ETS date, details results that are similar to our baseline estimates. Similarly, the results reported in Table A5 show that 18 alternative formulations of the expected mortality hazard produce VSL estimates that are very similar to the baseline VSL estimates derived from our preferred mortality measure. Additionally, Table A6 reports estimates from specifications with individual control variables determined after soldiers enlist but prior to the final 360 days of their initial enlistment, coinciding with our preferred choice for a soldier's decision date. Controlling for a soldier's rank, individual combat deployment history, deployment status at the soldier's decision date, and individual exposure to casualties produces VSL estimates similar to those reported in Table 3.

Two aspects of our findings thus far deserve emphasis. First, our estimates of the VSL among early in their career Army soldiers are generally far lower than those typically reported in the related literature for the full US labor market. For example, the median VSL implied by hedonic labor market studies based on US data was about \$10 million in 2019 dollars (Viscusi and Aldy, 2003). Second, our findings produce tightly-estimated results that change little across specifications, which gives some reassurance that the combination of our institutional setting and econometric model provides a reliable approach to estimating the behavioral tradeoffs that soldiers make when reenlisting. The stability of these estimates and their relatively tight confidence intervals stand in contrast to the instability of the estimates in

[20] Although not reported in Table 3, we find that higher local unemployment rates induce more soldiers to reenlist, consistent with the findings of Borgschulte and Martorell (2018). Table A3 reports average marginal effects on all of the controls, except fixed effects, for the regressions reported in Panel A, Table 3.

response to changes in specification in much of the hedonic labor market literature.[21]

6.2 Binary Choice Models with Heterogeneity

Heterogeneity Based on Occupation. It is natural to assume that individuals who choose to enlist in the Army have different, and presumably lower, VSLs than those who do not enlist in the Army because they find fatality risk less distasteful. Here, we assess whether there is heterogeneity in the VSL *among* soldiers based on their occupation. Specifically, we estimate separate binary logit models of the probability of reenlistment using subsamples based on broad occupation-by-gender groupings, where the measured occupation is their occupation when they first entered the Army.[22] We explore three subsamples: men in noncombat occupations (N = 189,270), men in combat occupations (N = 168,943), and women in all occupations (N = 71,162).

Panels B, C, and D of Table 3 summarize the estimated VSLs and documents some heterogeneity across these broad aggregates of occupation-types.[23] Most notably, men in combat occupations have the smallest estimated VSLs across all of the specifications, while women have the highest estimated VSLs across all specifications. In the richest specification (i.e., column (5)), the point estimates (95 percent confidence intervals) are $670,000 ($-9,100, $1.34 million) for men in noncombat occupations (Panel B); $449,000 ($-7,200, $969,000) for men in combat occupations (Panel C), and $1.4 million ($526,000, $2.30 million) for women (Panel D). It is possible to reject the null of equal VSLs for men in noncombat occupations and women at the 15 percent level and between men in combat occupations and women at the 8 percent level (see Table A7). The null of equal VSLs among the two male categories would not be rejected at conventional significance levels. Table A7 suggests that differences in the VSL are driven by a lack of responsiveness to mortality risk among men in combat occupations. For the richest specification (i.e., column 5), increasing the four-year mortality rate by 1 death per 1000 soldiers only decreases reenlistment by 0.13 percentage points among men in combat occupations, whereas the same increase in mortality reduces reenlistment by 0.36 percentage points for men in noncombat occupations and 1.15 percentage points for women. However, men in combat occupations are also significantly less responsive to bonuses than men in noncombat occupations and women, thus attenuating differences in the implied VSL.

[21] For example, Black and Kniesner (2003) show that measures of job risk exhibit non-classical measurement error, posing a challenge for the proper specification and biasing results of cross-sectional hedonic equations.

[22] It is highly unusual for a soldier to change her occupation during her initial enlistment before reenlisting.

[23] Table A7 reports bonus and mortality hazard average marginal effects by subsample and p-values from tests of the null hypothesis of equal VSLs across subsamples.

Although the difference in implied VSLs between the two male categories lacks statistical significance, the results are directionally consistent with the idea that less risk-averse individuals sort into riskier occupations. However, this is not an apples to apples comparison because the mortality rates differ so dramatically between soldiers in noncombat occupations and soldiers in combat occupations: across the full sample covering 2002-2010, the mean four-year mortality rate for women is 2.96 per 1000 with a range from 0 - 12.7, for men in noncombat occupations, the mean is 3.25 and the range is 0 - 12.7, and for men in combat occupations the mean is 8.92 with a range from 0.8 - 27.4. In Section 7, we directly explore how the VSL varies with the level of the expected mortality hazard for the full sample and these three subsamples.

Observable Heterogeneity Based on Moment Forest Analysis. The results from Table 3 suggest there could be meaningful heterogeneity in the VSL. At one extreme, the various full-sample binary logit models allow for a large amount of observable heterogeneity in the intercept, based on soldiers' cohort, MOS, and initial term-length, which determines the baseline propensity to reenlist. They do not, however, allow for differences in the marginal utility of the hazard and the bonus along those same observables. On the other end of the spectrum, running the logit models separately on subsamples allows for full interaction between the sub-population and the parameters, but also raises two problems. First, the partitioning of soldiers into subgroups faces a variance-bias tradeoff, as one needs to estimate an increasingly large number of parameters on a shrinking subset of observations. Second, the researcher must also decide which subgroups should be chosen for the subset estimation.

As an alternative, we turn to the moment forest method of Nekipelov et al. (2021), which uses a variation of a classification tree, a technique from the machine learning literature, to estimate and assign parameters to subsets of the population without requiring researcher intervention. They prove their method is guaranteed to recover the true assignment of parameters to sub-populations at a faster than parametric rate. This allows us to generalize the three subgroup models we considered in the previous section to consider all possible subsets of the population. Further, it allows us to calculate and report on individual-specific implied VSLs which we describe here.[24]

Specifically, we estimate a moment forest using a simple binary logit with three parameters (constant, bonus, and hazard) as the moment. In the analysis that follows, we permit the moment forest to split on all controls used in the column (5) specification of Table 3, except for the bonus, hazard, and constant. Thus, we allow these three parameters to vary with all observable variables, providing a more flexible approach to capture heterogeneity.[25]

[24]For a more detailed explanation of how the moment forest works, see Section 1 in the Online Appendix.

[25]The moment forest produces a random sequence of moment trees, each of which groups soldiers into

We report the moment forest results in several ways. The building blocks for all of these estimates are the distributions of the bonus and mortality hazard parameters that are reported in Figures A4 and A5 for the full sample and the three occupation subgroups. The estimated implied individual-specific VSLs are then calculated by taking the negative of the ratio of each individual's mortality hazard estimate over her bonus estimate.

Figure 2 displays the distributions of the estimated individual-specific VSLs and makes apparent that there is meaningful heterogeneity. In the full sample, we estimate a median VSL of $980,000 with an inter-quartile range of ($450,000, $1,410,000). Further, 10.1 percent of the population has an implied VSL above $2 million, while 11.1 percent have a negative implied VSL, due to 0.6 percent of the sample having negative responses to bonuses and 10.5 percent having positive responses to expected mortality. Perhaps the most striking feature of the three occupation-based distributions is that all women have positive implied VSLs, because they all have the expected sign for the responses to the bonus (i.e., positive) and mortality hazard (i.e., negative) which is not true for all men. Women also have a much stronger response to the mortality hazard then the two groups of men, with a median response of -0.30, compared to median responses of -0.16 and -0.15 for men in noncombat and combat occupations, respectively. It is therefore not surprising that women have a higher median VSL.

To further analyze the results from the moment tree analysis, Figure A6 reports the median of the individual-specific VSL estimates and their inter-quartile ranges with box plots for different demographic, education, and AFQT groups. The distributions for each group are just subsets of the implied individual-specific VSLs reported in Figure 2a. It is evident that the median VSL for Black and Hispanic soldiers is larger than White soldiers' median, while the median VSL of high-school dropouts is more than $300,000 lower than the median VSL of soldiers with high school diplomas. Interestingly, the median implied VSL for soldiers in the lowest AFQT tercile is roughly $300,000 larger than the median VSL among soldiers in the highest AFQT tercile. VSLs do not appear to vary substantially with the age at initial entry into the Army or the unemployment rate in soldiers' county of residence. As a reminder, these estimated differences in median VSLs do not account for differences in these subgroups' expected mortality hazard.

To conclude our moment forest analysis, we assess whether the appealing flexibility of the moment forest framework affects the conclusions from the more traditional binary logit approach reported in Table 3. Specifically, column (1) of Table 4 reports the average estimated

subsets sharing the same parameter vector. Since each tree in the forest is built on re-sampled data, the moment forest may produce a different estimate of the VSL for each each unique set of soldier characteristics (the smallest possible grouping of soldiers) in the data.

VSL across the full distribution of responses for the full sample in Panel A and then for the three occupation subsamples in Panels B - D (Table A8 reports average estimated bonus and hazard logit coefficients from the moment forest framework). The average implied VSL is calculated as the negative of the ratio of the average hazard estimate to the average bonus estimate (rather than the average of the individual-specific ratios, to match what is reported in Table 3). Its 95 percent confidence interval is obtained by executing 50 bootstraps of the moment forest analysis, then using the 50 bootstraps to construct standard errors for the average bonus estimate and the average mortality estimate.[26] We then apply the delta method to the bonus and mortality standard errors to calculate a 95 percent confidence interval for the implied VSL. Over the full sample, the estimated VSL is $874,000, which is roughly $300,000 larger than the VSL estimate from the standard binary logit framework in column (5) of Table 3, although the 95 percent confidence intervals overlap. For each of the three subsamples, the average estimated VSL derived from the moment forest analysis is similar to the VSL derived from the standard binary logit approach reported in Table 3, again with overlapping 95 percent confidence intervals. Women continue to have the highest estimated VSL. Overall, the moment tree approach does not appear to appreciably alter conclusions about the average VSL from the standard binary logit analyses, lending some additional credibility to the choice of controls in our preferred logit specifications.

Unobservable Heterogeneity. Finally, as an alternative to allowing for heterogeneity based on occupation and other observables while remaining within the binary logit approach, column (2) of Table 4 reports estimates from a random coefficients model that allows for unobservable heterogeneity in soldiers' responses to the bonus and expected military rate by assuming these parameters are distributed normally (see Table A9 for the full results from the random coefficients model). The implied VSLs are reported separately for the full sample and the three subsamples using the preferred column (2) specification from Table 3, which includes fixed effects for MOS, cohort, and initial term-length, and the deployment probability control. For the full sample, the estimated VSL from the random coefficients model is roughly $200,000 larger than the corresponding estimate from the standard binary logit specification, but the 95 percent confidence interval from the random coefficient estimate overlaps with the standard binary logit estimate. Furthermore, the random coefficient estimates for each subsample are similar in magnitude, and not statistically different from, the standard binary logit estimates reported in column (2) of Table 3.

[26]Each bootstrap produces a bonus and a mortality estimate for each individual. We calculate the mean bonus estimate for each bootstrap, then calculate the standard error of the bonus estimate as the standard deviation of the 50 bonus means. We calculate the standard error of the morality rate in the same manner.

6.3 Multinomial Logit Models

Estimates from the binary choice framework are potentially biased through the aggregation of the full slate of reenlistment alternatives into a stay or go decision. In order to estimate the model in the binary choice framework, we assumed that the choice of reenlisting in the same occupation (MOS) for a four-year term was a sufficient representation of all the reenlistment alternatives. However, there are more dimensions to the decision to reenlist for an additional term of service, including the possibilities of selecting a new MOS, picking the home-base location of one's duty assignment, and choosing the term-length.

Columns (3) - (6) of Table 4 explore alternative characterizations of the choice set. In columns (3) and (4), we extend our model to a multinomial choice framework where the decision space includes the outside option (exiting the Army) and five inside options: 1) reenlist for the same MOS to be stationed at a location of the Army's choice; 2) reenlist for the same MOS to stay at one's current duty location; 3) reenlist to train for a new MOS; 4) reenlist for the same MOS to be stationed at an overseas location of choice; or 5) reenlist for the same MOS to be stationed at a location of choice within the continental U.S. This expanded choice set allows us to examine whether soldiers are willing to trade-off wealth and risk at different rates across alternatives with varying degrees of non-pecuniary benefits. We standardize each of these choices by applying the bonus and expected mortality rate associated with a four-year reenlistment. The full sample is used for estimation in column (3), while in column (4) it is restricted to high density MOSs (defined as having an average strength of at least 5,000 soldiers during the sample period) that reduces the sample from 429,375 to 277,877. The restriction is necessary due to the computational burden associated with the estimation of nested logit and richer multinomial logit models (discussed below).

The implied VSL estimate reported in in column (3), Panel A, of Table 4 is about \$700,000, which is roughly \$125,000 higher than the estimated VSL from the standard binary logit framework in the full sample, but this difference is not statistically significant.[27] VSL estimates derived from multinomial logit estimates for the three subsamples, reported in Panels B through D, are also qualitatively unchanged from the corresponding binary logit estimates.

A shortcoming of the J=5 multinomial choice framework is that it assumes there is no correlation between different alternatives or that there is independence of irrelevant alternatives (IIA). This assumption is unlikely to be valid in this setting where reenlistment options 1,

[27] All specifications in Table 4 include interactions between each reenlistment option and all control variables that do not vary with reenlistment alternatives (i.e., all right-hand-side variables except bonus offers, mortality hazards, and deployment probabilities). We also produced estimates with reenlistment option fixed effects that are not interacted with all right-hand side variables and found similar results. These results are available from the authors upon request.

2, 4, and 5 involve staying in the same occupation, whereas reenlistment option 3 involves the choice of training for a new occupation. It seems natural to expect alternatives related to staying in the same MOS to be highly correlated with one another, particularly if choice of home-base duty location is orthogonal to the probability of deployment. Column (5) of Table 4 therefore reports estimates from nested logit regressions where we permit errors to be correlated between reenlistment options that allow soldiers to remain in their original occupation. The similarity in VSL estimates between the multinomial and nested logit specifications suggest that any violations of the independence of irrelevant alternatives (IIA) assumption likely have minimal impact on our estimates.

Finally, column (6) reports on the estimation of multinomial logit models where the choice set is enriched to account for the 22 observed combinations of reenlistment options and reenlistment lengths. The consequence is that each option is defined by a triplet of occupation, home post location, and term-length.[28] The VSL estimate of \$605,000, as reported in column (6) of Table 4, is qualitatively unchanged from the estimated VSL in column (3) derived from the J=5 choice set. The confidence intervals of the VSL estimates in the three occupation subsamples also overlap the column (3) ones. In summary, the multinomial and nested logit results from Table 4 suggest that aggregating reenlistment alternatives into a single decision to reenlist does not appear to appreciably alter our estimates of the VSL.

7 Empirical Estimates of Indifference Curves

7.1 Recovering Indifference Curves Through B-Splines

The ultimate aim of hedonic theory is to to recover individuals' demand primitives and indifference (or bid) curves that are necessary for welfare analysis. However, the empirical VSL literature has settled for estimating the average VSL over heterogeneous populations that have limited value for welfare analysis. Indeed, the standard approach that effectively assumes a "constant" VSL is at odds with economic theory that assumes there is declining marginal utility from income and safety. As we have pointed out about this paper's setting, there were substantial differences in mortality rates across groups. This raises the possibility that efforts to infer differences in risk preferences across groups confound the risk preferences with differences in the levels of risk.

[28] Recall that reenlistment term-lengths are in annual increments of 2 through 6 years, but 2 year reenlistments are only available for soldiers who remain in the same occupation and who either allow the Army to choose their post or who elect to remain at their current post (see Figure A1).

This section seeks to recover these more fundamental measures of preferences through a variation of the logit model that uses cubic B-splines to allow for nonlinearities in the responses to the bonus and mortality rate.[29] The estimation of these models produces the parameters of the utility functions and, with those in hand, it is straightforward to develop indifference curves defined as the set of all bundles of bonuses and mortality hazard rates that give the same utility level. Because we allow these two variables to influence utility nonlinearly, the slope along the indifference curve reveals the marginal rate of substitution between the money and mortality risk (i.e., the VSL) and how it varies with the mortality rate.

We begin by using the column (2) specification from Table 3 and estimating B-splines using a stochastic Laplace Type Estimator (LTE) optimization process from Chernozhukov and Hong (2003).[30] The estimation results are used to construct the blue curve in Figure 3 (left plot) that is an indifference curve calculated as the level of the bonus at each mortality rate that produces the same level of utility as found at the average bonus and hazard.[31] The light blue shaded area indicates the 95 percent confidence interval along the support of the mortality hazard and is estimated using the method described in Chernozhukov and Hong (2003).[32] Finally, the empirical distributions of the expected mortality hazard and bonus offers are in the smaller figures and are aligned with the horizontal and vertical axes to indicate the areas where there is support in the data, versus those that rely more heavily on extrapolation based on functional form. As a point of reference, 77 percent of the sample has a mortality rate below 8 and 84 percent has one below 9.5.

The raw data broadly conform to standard economic theory with some exceptions. Specifically, the required bonus is generally increasing in the mortality hazard over the range of observed mortality rates. The estimated VSLs at several mortality rates are reported in

[29]Unidimensional B-splines are locally-adaptive polynomial expansions of a variable constructed by aggregating a series of simple functions of increasing complexity. The simplest b-splines are piecewise constants defined on a set of disjoint intervals, while higher-dimensional approximations add lines, quadratic terms, cubic terms, etc. B-splines are analagous to the common power series polynomials (e.g. $x, x^2, x^3, \dots$) except that they are defined over a local neighborhood and are more numerically stable. The local nature of B-splines allows for the functional approximation to be influenced disproportionately by nearby points in the domain and not governed by outliers, which is a particular problem with globally-defined polynomials like the power series.

[30]The LTE method is a stochastic optimization procedure that has two attractive properties: it has good performance against non-smooth and non-convex objective functions and it traces out the covariance matrix of the estimated parameters at the same time it finds minima to the objective function.

[31]Details available from the authors upon request.

[32]We constructed the indifference curve's 95 percent confidence interval from the quasi-posterior distribution of moments from 500 LTE draws, where each draw produced a separate indifference curve. The confidence interval reflects the 2.5th and 97.5th percentiles of solved bonus amounts at each expected mortality rate. The indifference curve from every LTE moment includes a point at the average bonus and average mortality rate of the sample, so the confidence interval shrinks to zero at this point. However, the slope of each indifference at this point will still differ across LTE moments, producing a range of implied VSLs.

Table A10 and up to a mortality rate of 8 per 1000 it generally ranges between \$500,000 and \$1,500,000. The departures from monotonicity are concentrated in regions where there is little support in the data and in those regions the confidence intervals indicate that the data are consistent with a very wide range of VSLs.[33]

We turn to economic theory to make further progress. Specifically, the dashed brown curve in Figure 3 (left plot) plots a version of the indifference curve after imposing monotonicity on the estimated B-splines such that higher bonuses are associated with higher utility and higher mortality rates are associated with lower utility. It is apparent that these restrictions fix many of the "irregular" parts of the indifference curve; they also greatly reduce the confidence intervals for the curve and the VSLs (Table A10). Nevertheless, a few regions on the indifference curve violate standard assumptions about diminishing marginal utility for income and reductions in mortality risk, especially where the data is sparse. For this reason, the remainder of this section will additionally impose convexity in the B-splines such that higher bonuses increase utility at a weakly decreasing rate and higher mortality rates reduce utility at a weakly increasing rate. These restrictions are both intuitive and produce indifference curves that are nicely behaved, even in regions with relatively little data.

The black line in Figure 3 (right plot) is the estimated indifference curve for the full sample after imposing both monotonicity and convexity. The red lines are tangent to the indifference curve at the indicated hazard rates and we interpret their slope (reported in the figure) as the implied VSL. The most striking finding is the shape of the bid curve and its similarity to the bid curves from Thaler and Rosen's original theoretical framework (reproduced in Figure 1) that predicted that the VSL increases with expected mortality. Specifically, Figure 3 and the third row of Table A10 reveal that soldiers with a relatively low mortality risk of 1 death per 1000 soldiers (between the 2nd and 3rd-percentile of the distribution of the mortality rate among the full sample) have estimated VSLs near \$482,000, while soldiers with a mortality risk of 5 deaths per 1000 soldiers (65th-percentile of the mortality rate) have an estimated VSL near \$672,000, and soldiers with a mortality risk of 16 deaths per 1000 (96th-percentile of the mortality rate) have an estimated VSL of \$1.44M. We believe this is the first-of-a-kind empirical construction of bid curves depicting the wealth versus risk of death trade-off.

A potentially even more informative set of results is reported in Figure 4, which displays the bid-curves and their confidence intervals for the three subgroups on the same axes and is summarized in Table 5. The substantial heterogeneity across subsamples is apparent

[33] This highlights a drawback of using locally-adaptive approximations like B-splines, where sparse data can result in locally non-monotone curves that result in poorly-behaved indifference curves. On the other hand, locally-adaptive approximations are not driven by outliers in the tails, which is especially important when the density of the data is non-uniform over the support.

from the shape and steepness of the estimated bid curves. Reassuringly, the pattern of this heterogeneity has broad parallels to the heterogeneity in the subsample-specific VSL estimates reported in Table 3. For example, the estimated bid curves for women and for men in noncombat occupations are substantially steeper than the bid curves for men in combat occupations, revealing these occupation subgroups' relative dislike of mortality risk. It is noteworthy that men in noncombat occupations facing an expected mortality rate of 1 death per 1000 soldiers have an estimated VSL of $717,000, and that this more than doubles to $1.667M when they face a mortality rate of 5 deaths per 1000. Among men in combat occupations, the estimated VSL does not reach $1.667M until the mortality rate exceeds 23 deaths per 1000 soldiers, which is the 97th percentile of their mortality distribution.

Importantly, Figure 4 and Table 5 reveal how comparing average estimated VSLs between subsamples confounds differences in preferences for safety and wealth with differences in expected mortality rates. Although our estimates up to this point suggest that men in combat occupations have lower average VSLs then women and, to a lesser extent, men in noncombat occupations, these differences are significantly larger and statistically meaningful when we compare soldiers across groups while holding the expected mortality rate constant. For example, when the four-year mortality rate is 8 per 1000 soldiers, the implied VSL for men in noncombat occupations is 7 times larger than the implied VSL for men in combat occupations ($2.808 million for noncombat men compared to $374,000 for combat men). The confidence intervals reported in Table 5 confirm that when the occupation subsamples' implied VSLs are compared at the same expected mortality rates, the VSL estimates for men in combat occupations are statistically distinguishable from the VSLs for men in noncombat occupations and women at all but the smallest levels of mortality, making clear that the difference in risk preferences across these groups is economically and statistically significant.

7.2 Welfare Analysis of Non-Marginal Changes in Fatality Risk

A primary benefit of estimating dis-aggregated bid curves is the ability to compute welfare implications of non-marginal changes in mortality risk that also vary across subgroups. We illustrate these benefits in two ways. First, we calculate the welfare implications of mortality increases of the order experienced by U.S. Army soldiers in recent conflicts. We then evaluate the efficiency of a safety investment that the U.S. Army started in 2005.

In the first example, we estimate the money required to keep reenlistment rates constant as mortality risk increases from occupational mortality rates in 2002 to the highest annual occupational mortality rates we observe in our data.[34] We interpret mortality risk in 2002 as a baseline period that most closely resembles peacetime eras. At the other end, we interpret maximum annual mortality rates at the occupation level as a measure of risk during the most lethal periods of the wars in Iraq and Afghanistan. Although this gives an upper bound on the welfare losses associated with increased risk during Iraq and Afghanistan, it could underestimate the welfare losses from other recent conflicts that were more lethal, such as the wars in Korea and Vietnam (Defense Industry Daily Staff, 2005).

We then map each occupation's 2002 mortality rate and peak mortality rate to the estimated bid-curves in Figure 4 to identify the bonuses required for each occupation to keep reenlistments constant. For example, mortality rates for infantry soldiers increased from 3.5 deaths per 1000 in 2002 to 24.6 deaths per 1000 in 2007. The indifference curve for men in combat occupations suggests that the average infantry soldier requires a lump-sum bonus of \$12,900 to remain equally well-off when mortality increases from 3.5 to 24.6. In other words, the typical infantry soldier suffered a welfare loss of \$12,900. Applying this framework to all occupations, we estimate that the increase in mortality risk due to the Iraq and Afghanistan wars reduced welfare by \$8,500 for the average soldier. Since our mortality measure is a four-year sum with future years discounted at 7.2 percent, \$8,500 is equivalent to a \$2,355 annual reduction in welfare, or 8 percent of annual military pay (see column (3) of Table 2).

This example also helps illustrate the importance of estimating indifference curves to assess the welfare consequences of non-marginal changes in mortality risk or other implicit goods and services. If we used our average full sample, linear binary logit VSL estimate of \$575,000 (panel A, column (2), Table 3), we would estimate that the increased risk from the wars in Iraq and Afghanistan only reduced welfare for the average soldier by \$5,100. This is 40 percent less than the welfare loss calculated from our estimates of separate indifference curves for men in noncombat occupations, men in combat occupations, and women.[35]

We now turn to a second example to further emphasize the policy implications of understanding how the VSL varies by both occupation and the level of risk within occupations. In response to a Department of Defense study citing shortcomings in standard body armor issued at the start of the wars in Iraq and Afghanistan, the U.S. Army began issuing enhanced body armor to soldiers in all occupations in 2005. The enhanced body armor cost \$760 per unit, came with a total projected system cost of \$200M, and was estimated to

[34] Recall that our measure of risk is a 12 month rolling average, but summed over four years with future years discounted at 7.2 percent to keep our risk measure consistent with our bonus measure (see Section 4.2). Roughly 48 percent of soldiers in our sample are in occupations where the highest mortality rate occurs in 2007, consistent with column (6) of Table 2.

[35] Of course, we acknowledge that exact welfare calculations would require us to estimate separate indifference curves for every occupation. Unfortunately, power limitations make this infeasible.

reduce mortality by roughly 15 percent.[36] Ignoring heterogeneity, our full sample binary logit VSL estimate of $575,000 would imply that the $200M enhanced body armor program was welfare-improving if it saved 350 lives.[37] This is a reasonable possibility considering that nearly 4,000 service-members died in Iraq between 2003 and 2007.[38]

However, this standard cost-benefit analysis derived from a single VSL could miss crucial welfare implications when comparing changes in mortality risk across occupations and within occupations but at different mortality rates. For example, the mortality rate for truck drivers in 2005 was 7.8 per 1000, implying a VSL for male truck drivers of roughly $2.73M when interpolating between estimated VSLs for men in noncombat occupations, as reported in columns (3) and (4) of Table 5. Meanwhile, average mortality rates for infantry soldiers were 11.3 per 1000 in the same year, implying a VSL of $428,000 (see columns (4) and (5) of Table 5). If the enhanced body armor reduced mortality by 15 percent across all occupations, then the program would clearly be welfare improving for male truck drivers, who would be willing to pay roughly $3,194 for the enhanced body armor $((\frac{7.8}{1000})(0.15)(2,730,000) = 3,194)$. On the other hand, infantry soldiers would only be willing to pay roughly $725 for the enhanced body armor $((\frac{11.3}{1000})(0.15)(428,000) = 725)$, just under the system's unit cost. In other words, the welfare maximizing policy in 2005 would have been to issue enhanced body armor to truck drivers and maintain infantry staffing levels through higher bonuses.[39]

8 Conclusion

Nearly a half century ago, Rosen (1974) laid out an approach to recover indifference curves between wealth and non-market goods; Thaler and Rosen (1976) extended this approach to infer the VSL. In the subsequent decades, the empirical compensating differentials and VSL literatures have almost exclusively assumed a constant marginal willingness to pay for non-market goods and in so doing assumed away the need to estimate the demand primitives necessary for welfare analysis in the presence of heterogeneity (Greenstone, 2017). The result has been to turn a rich model of the world into a mechanical empirical summary that fails to

[36] Defense Industry Daily Staff (2005) provides total and per units costs, in 2006 USD, of $160M and $600, respectively. Moss (2006) explains the marginal effectiveness of the enhanced body armor.

[37] If we assume the occupations of soldiers receiving body armor mirrored the share of occupations in our sample, then the subsample binary logit estimates from Table 3 would permit us to estimate that the $200M body armor program would be welfare improving if they saved 87 men in noncombat occupations, 180 men in combat occupations, and 26 women.

[38] See Operation Iraqi Freedom (OIF) Casualty Summary By Month and Service (accessed April 2021).

[39] It is noteworthy, however, that the infantry mortality rate in 2007, at 24.6 per 1000, was high enough to justify a $760 enhanced body armor system as welfare improving at that particular time period.

accommodate the widespread casual observation that people's willingness to pay for safety and other non-market goods vary tremendously across types and contexts.

This paper returns the empirical literature to its theoretical origins in the case of the VSL by using a random utility discrete choice model coupled with rich data on the first-time reenlistment decisions of U.S. Army soldiers. This combination allows us to recover indifference, or bid, functions. Our study's setting also offers potential to make important progress on two other problems common to the VSL literature: that mortality risk is correlated with unobserved job features and might not be well known (Ashenfelter, 2006).

There are two principal empirical findings. First, we estimate that the *average* VSL estimated across the population of soldiers and observed range of mortality rates is generally between \$500,000 and \$900,000 across a wide set of specifications. These estimates are precise and are at least at an order of magnitude smaller than the typical VSL estimate from overall U.S. labor market studies (Viscusi and Aldy, 2003; Kniesner et al., 2012). Second, we find that the assumption of a single and universal VSL is rejected in this setting. For example, the estimated VSL varies substantially within the observed range of mortality risk among men in combat occupations. Additionally, there is clear evidence of sorting across occupations based on risk preferences such that there is a seven-fold difference in the VSL across occupation categories when we compare those categories at the same level of risk.

The broader message is that accounting for heterogeneity in the marginal willingness to pay for non-market goods is critical. In the case of the VSL, there is not a single VSL among soldiers, which also raises the possibility of similar heterogeneity in the broader population. Beyond the VSL literature, a greater focus on capturing heterogeneity with discrete choice random utility models or other approaches will improve understanding of how consumers value a wide range of non-market goods and services, including clean air, public transportation, school quality, and other job characteristics to name a few. Such understanding could be the basis for policy reforms that improve social welfare.

References

Aitken, Peter, "9/11 had the biggest effect on military recruiting since Pearl Harbor," 2021.

Asch, Beth J, Paul Heaton, James Hosek, Francisco Martorell, Curtis Simon, and John T Warner, "Cash Incentives and Military Enlistment, Attrition, and Reenlistment," Technical Report, RAND National Defense Research Institute 2010.

Ashenfelter, Orley, "Measuring the Value of a statistical Life: Problems and Prospects.," *The Economic Journal*, 2006, *116* (510), C10.

_ **and Michael Greenstone**, "Using Mandated Speed Limits to Measure the Value of a Statistical Life.," *Journal of Political Economy*, 2004, *112* (1), S226 – S267.

Banzhaf, H Spencer, "The Value of Statistical Life: A Meta-analysis of Meta-analyses," Technical Report, National Bureau of Economic Research 2021.

Bartik, Timothy J., "The Estimation of Demand Parameters in Hedonic Price Models," *Journal of Political Economy*, 1987, *95* (1), 81–88.

Bayer, Patrick, Fernando Ferreira, and Robert McMillan, "A unified framework for measuring preferences for schools and neighborhoods," *Journal of political economy*, 2007, *115* (4), 588–638.

Black, Dan A. and Thomas J. Kniesner, "On the Measurement of Job Risk in Hedonic Wage Models.," *Journal of Risk and Uncertainty*, 2003, *27* (3), 205 – 220.

Borgschulte, Mark and Paco Martorell, "Paying to Avoid Recession: Using Reenlistment to Estimate the Cost of Unemployment," *American Economic Journal: Applied Economics*, 2018, *10* (3), 101–127.

Brown, James N. and Harvey S. Rosen, "On the Estimation of Structural Hedonic Price Models," *Econometrica*, 1982, *50* (3), 765–768.

Chay, Kenneth and Michael Greenstone, "Does Air Quality Matter? Evidence from the Housing Market," *Journal of Political Economy*, 2005, *113* (2), 376–424.

Chernozhukov, Victor and Han Hong, "An MCMC approach to classical estimation," *Journal of Econometrics*, 2003, *115* (2), 293–346.

Cropper, Maureen, James K. Hammitt, and Lisa A. Robinson, "Valuing Mortality Risk Reductions: Progress and Challenges.," *Annual Review of Resource Economics*, 2011, *3* (1), 313 – 336.

Deacon, Robert T., David S. Brookshire, Anthony C. Fisher, Allen V. Kneese, Charles D. Kolstad, David Scrogin, V. Kerry Smith, Michael Ward, and James Wilen, "Research Trends and Opportunities in Environmental and Natural Resource Economics," *Environmental and Resource Economics*, 1998, *11* (3-4), 383–97.

Defense Industry Daily Staff, "Up to $160M for Enhanced Interceptor Vest Protection," *Defense Industry Daily*, 2005.

Department of Defense, "Population Representation in the Military Services: Fiscal Year 2002," 2004.

Ekeland, Ivar, James J Heckman, and Lars Nesheim, "Identification and Estimation of Hedonic Models," *Journal of Political Economy*, 2004, *112* (1 pt 2).

Epple, Dennis, "Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products," *Journal of Political Economy*, 1987, *95* (1), 59–80.

Evans, Mary F and Laura O Taylor, "Using revealed preference methods to estimate the value of reduced mortality risk: Best practice recommendations for the hedonic wage model," *Review of Environmental Economics and Policy*, 2020, *14* (2), 282–301.

Folke, Olle and Johanna Karin Rickne, "Sexual Harassment and Gender Inequality in the Labor Market," 2020.

Greenstone, Michael, "The continuing impact of Sherwin Rosen's "Hedonic prices and implicit markets: product differentiation in pure competition"," *Journal of Political Economy*, 2017, *125* (6), 1891–1902.

Gruber, Jonathan, "The incidence of mandated maternity benefits," *The American economic review*, 1994, pp. 622–641.

Heckman, James J, Rosa L Matzkin, and Lars Nesheim, "Simulation and estimation of hedonic models," *IZA Discusion Paper*, 2003, (843).

_ , _ , and _ , "Nonparametric identification and estimation of nonadditive hedonic models," *Econometrica*, 2010, *78* (5), 1569–1591.

Hosek, James and Paco Martorell, *How have deployments during the war on terrorism affected reenlistment?*, Rand Corporation, 2009.

Hosek, James R. and Trey Miller, *Effects of Bonuses on Active Component Reenlistment Versus Prior Service Enlistment in the Selected Reserve.*, Rand, 2011.

Kniesner, Thomas J, W Kip Viscusi, Christopher Woock, and James P Ziliak, "The value of a statistical life: Evidence from panel data," *Review of Economics and Statistics*, 2012, *94* (1), 74–87.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler, "Imperfect competition, compensating differentials and rent sharing in the US labor market," Technical Report, National Bureau of Economic Research 2019.

Lee, Jonathan M and Laura O Taylor, "Randomized safety inspections and risk exposure on the job: quasi-experimental estimates of the value of a statistical life," *American Economic Journal: Economic Policy*, 2019, *11* (4), 350–74.

León, Gianmarco and Edward Miguel, "Risky transportation choices and the value of a statistical life," *American Economic Journal: Applied Economics*, 2017, *9* (1), 202–28.

Lleras-Muney, Adriana, "The Needs of the Army: Using Compulsory Relocation in the Military to Estimate the Effect of Air Pollutants on Children's Health.," *The Journal of Human Resources*, 2010, *45* (3), 549.

Lyle, David S., "Using Military Deployments and Job Assignments to Estimate the Effect of Parental Absences and Household Relocations on Children's Academic Achievement.," *Journal of Labor Economics*, 2006, *24* (2), 319 – 350.

Mas, Alexandre and Amanda Pallais, "Valuing alternative work arrangements," *American Economic Review*, 2017, *107* (12), 3722–59.

Moss, Michael, "Pentagon Study Links Fatalities to Body Armor," *The New York Times*, 2006.

Mrozek, Janusz R. and Laura O. Taylor, "What Determines the Value of Life? A Meta-Analysis.," *Journal of Policy Analysis and Management*, 2002, *21* (2), 253.

Nekipelov, Denis, Paul Novosad, and Stephen P Ryan, "Moment Forests," 2021.

O'Brien, James, "Age, autos, and the value of a statistical life," *Journal of Risk and Uncertainty*, 2018, *57* (1), 51–79.

Patterson, Christina, Jonathan Petkun, and William Skimmyhorn, "I Want You! (But Not You): Selection in Military Retention," 2020.

Rohlfs, Chris, "The Economic Cost of Conscription and an Upper Bound on the Value of a Statistical Life: Hedonic Estimates from Two Margins of Response to the Vietnam Draft," *Journal of Benefit-Cost Analysis*, August 2012, *3* (3).

_ **, Ryan Sullivan, and Thomas Kniesner**, "New estimates of the value of a statistical life using air bag regulations as a quasi-experiment," *American Economic Journal: Economic Policy*, 2015, *7* (1), 331–59.

Rosen, Sherwin, "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *Journal of Political Economy*, 1974, *82* (1), 34–55.

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

Stern, Scott, "Do scientists pay to be scientists?," *Management science*, 2004, *50* (6), 835–853.

Thaler, Richard and Sherwin Rosen, "The Value of Saving a Life: Evidence from the Labor Market," in "Household Production and Consumption," National Bureau of Economic Research, Inc, 1976, pp. 265–302.

United States Army, "Army Regulation 601-280: Army Retention Program," 2019.

Viscusi, W. Kip, "The Value of Life: Estimates with Risks by Occupation and Industry," *Economic Inquiry*, 2004, *42* (1), 29–48.

_ , “The Heterogeneity of the Value of Statistical Life: Introduction and Overview,” *Journal of Risk and Uncertainty*, 2010, *40* (1), 1–13.

_ **and Joseph E. Aldy**, “The Value of a Statistical Life: A Critical Review of Market Estimates throughout the World.,” *Journal of Risk and Uncertainty*, 2003, *27* (1), 5 – 76.

Figures

Figure 1: Hedonic Wage-Safety Theoretical Framework

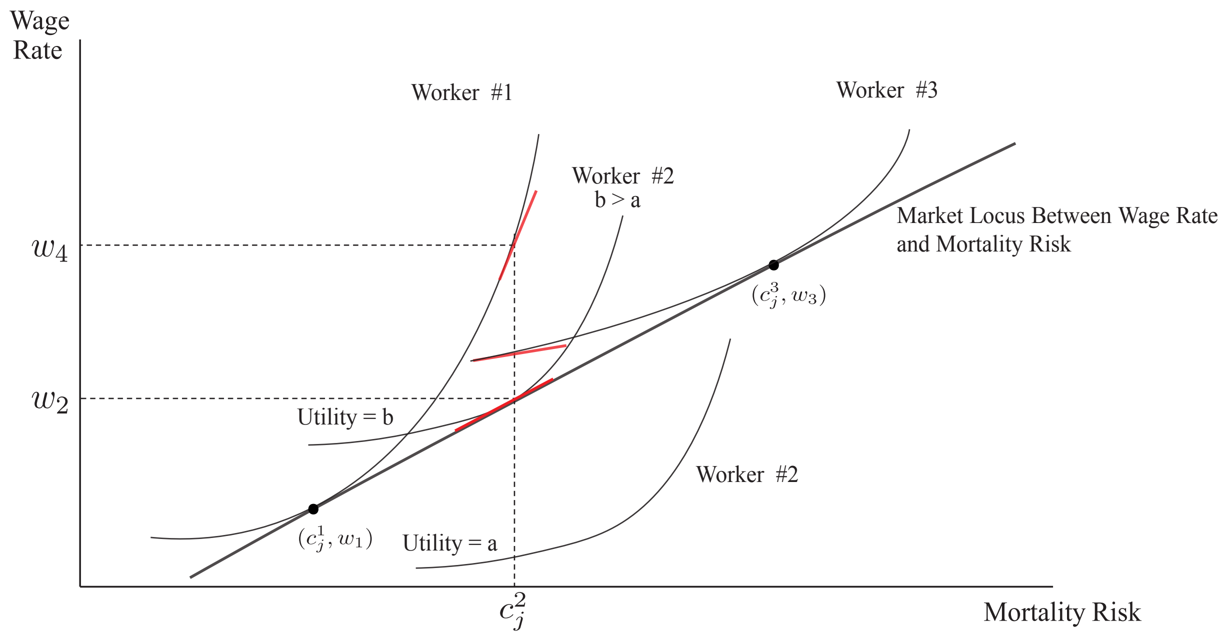

Notes: This figure depicts a theoretical example of the Hedonic Wage-Safety framework described in Section 2. The figure plots families of indifference curves for three types of workers, denoted as types #1, #2, and #3. For worker type #2, two different indifference curves are depicted. The heavy line in black plots the equilibrium market locus between wages and mortality risk or the market equalizing-difference wage function. The three types choose jobs where their marginal willingness to pay for safety is equal to the market determined marginal implicit price, which occur at mortality risk levels c_j^1, c_j^2, and c_j^3, respectively. The figure includes a red tangency line for each type. The slope of each tangent line is the VSL for a worker type when the mortality risk equals c_j^2.

Figure 2: Moment Forest VSL Estimate Distributions, Full Sample and by Subsample

(a) Full Sample

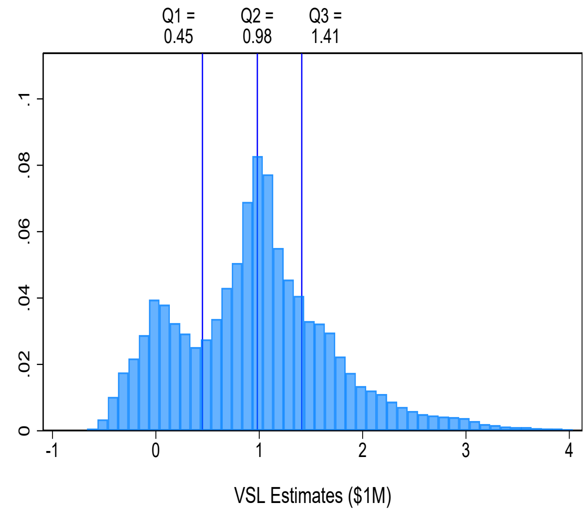

(b) By Subsample

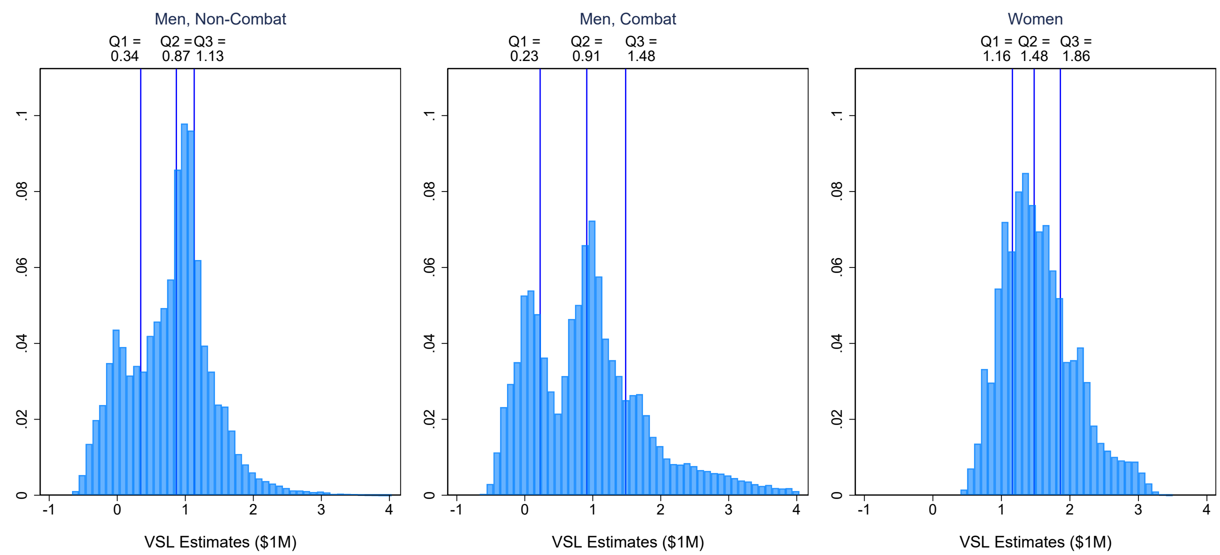

Notes: The histograms display the distribution of implied VSL estimates derived from our moment forest analysis for the full sample and subsamples. This analysis permits potential splitting on MOS, initial term-length, entry-cohort, deployment probability, gender, race, education, AFQT, age, state, and local unemployment rates. Q1, Q2, and Q3 are the 25th percentile, median, and 75th percentile of estimates. To reduce computing time, we convert continuous variables into discrete variables by grouping soldiers with similar covariate values. Specifically, we group soldiers into deployment probability categories with bandwidths of 5 percentage points (i.e., one group with deployment probabilities less than 5 percent, another group with deployment probabilities between 5 and 10 percent, and so on), we group soldiers into AFQT categories with bandwidths of 5 percentiles, we group soldiers into entry-age categories with bandwidths of 2 years, and we group soldiers into unemployment rate categories with bandwidths of one percentage point. VSL estimates below -$1 Million (0.6 percent of the full sample and 1.5 percent of the Men, Combat subsample) or above $4 Million (1.9 percent of the full sample and 4.9 percent of the Men, Combat subsample) are not depicted.

Figure 3: Bid Curve from B-Spline Estimates of Reenlistment Response to Bonus Offers and Mortality Hazards, Full Sample

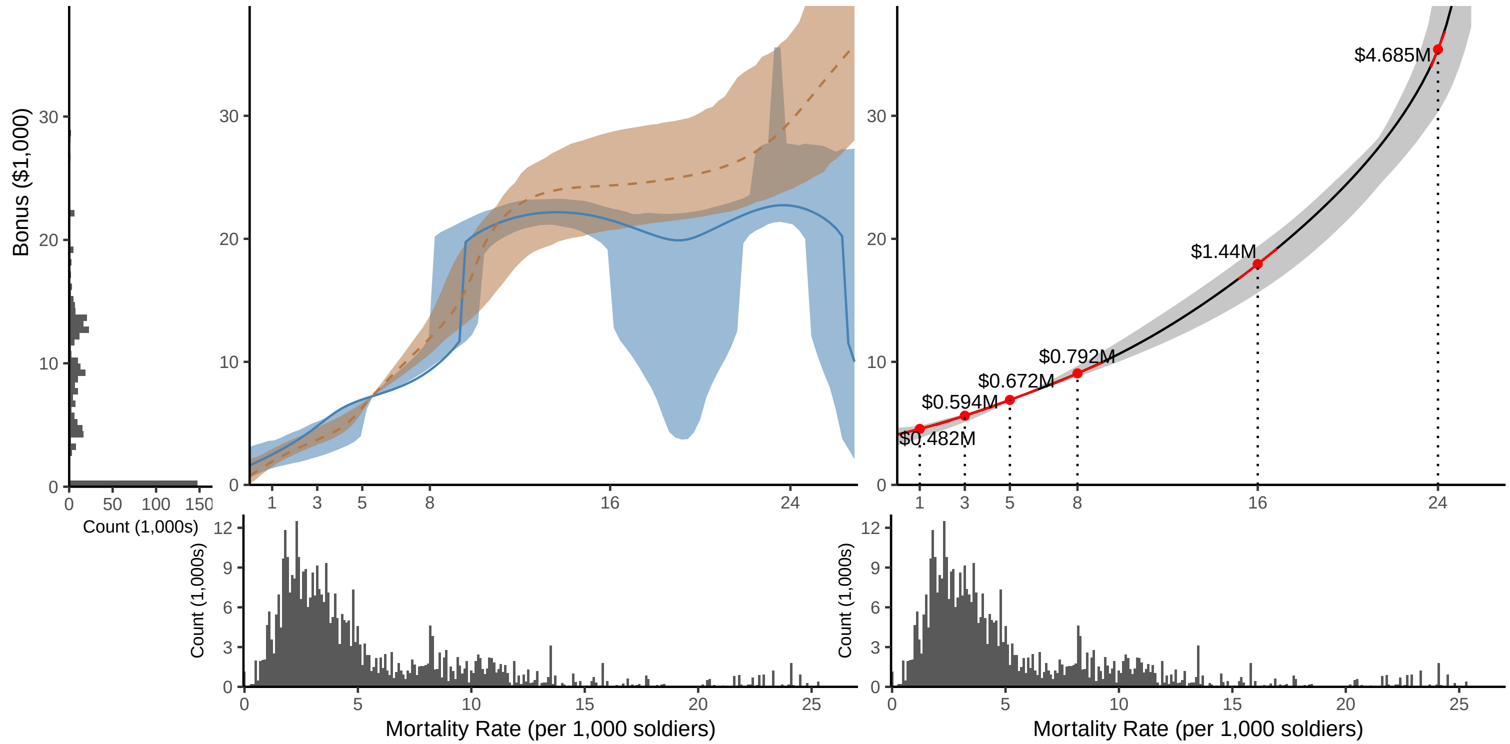

Notes: This figure depicts the indifference, or bid, curves derived from B-spline estimates of the column (2) specification from Table 3. The blue curve (left plot) is derived from B-spline estimates with no restrictions. The brown curve (left plot) imposes monotonicity. The black curve (right plot) imposes both monotonicity and convexity. The text labels on the right plot denote the implied VSL at each point, which is equal to the slope of the red tangent line at each point. Table A10 reports implied VSL estimates and confidence intervals at different mortality rates for all curves in this figure. We use estimation results from the stochastic Laplace Type Estimator (LTE) optimization process described in Chernozhukov and Hong (2003) to calculate the level of the bonus at each mortality rate, or hazard, that produces the same level of utility as found at the average bonus and hazard. The shaded area around each curve is a 95 percent confidence interval constructed from the quasi-posterior distribution of moments from 500 LTE draws, where each draw produced a separate curve. The confidence interval reflects the 2.5th and 97.5th percentiles of solved bonus amounts at each mortality rate. The bid curve from every LTE moment includes a point at the average bonus ($7,200) and average hazard (5.4 deaths per 1000) of the sample, so the bid curve confidence interval shrinks to zero at this point. However, the slope of each curve at this point will still differ across LTE draws, producing a range of implied VSLs.

Figure 4: Bid Curves Derived from B-spline Estimates by Subsample

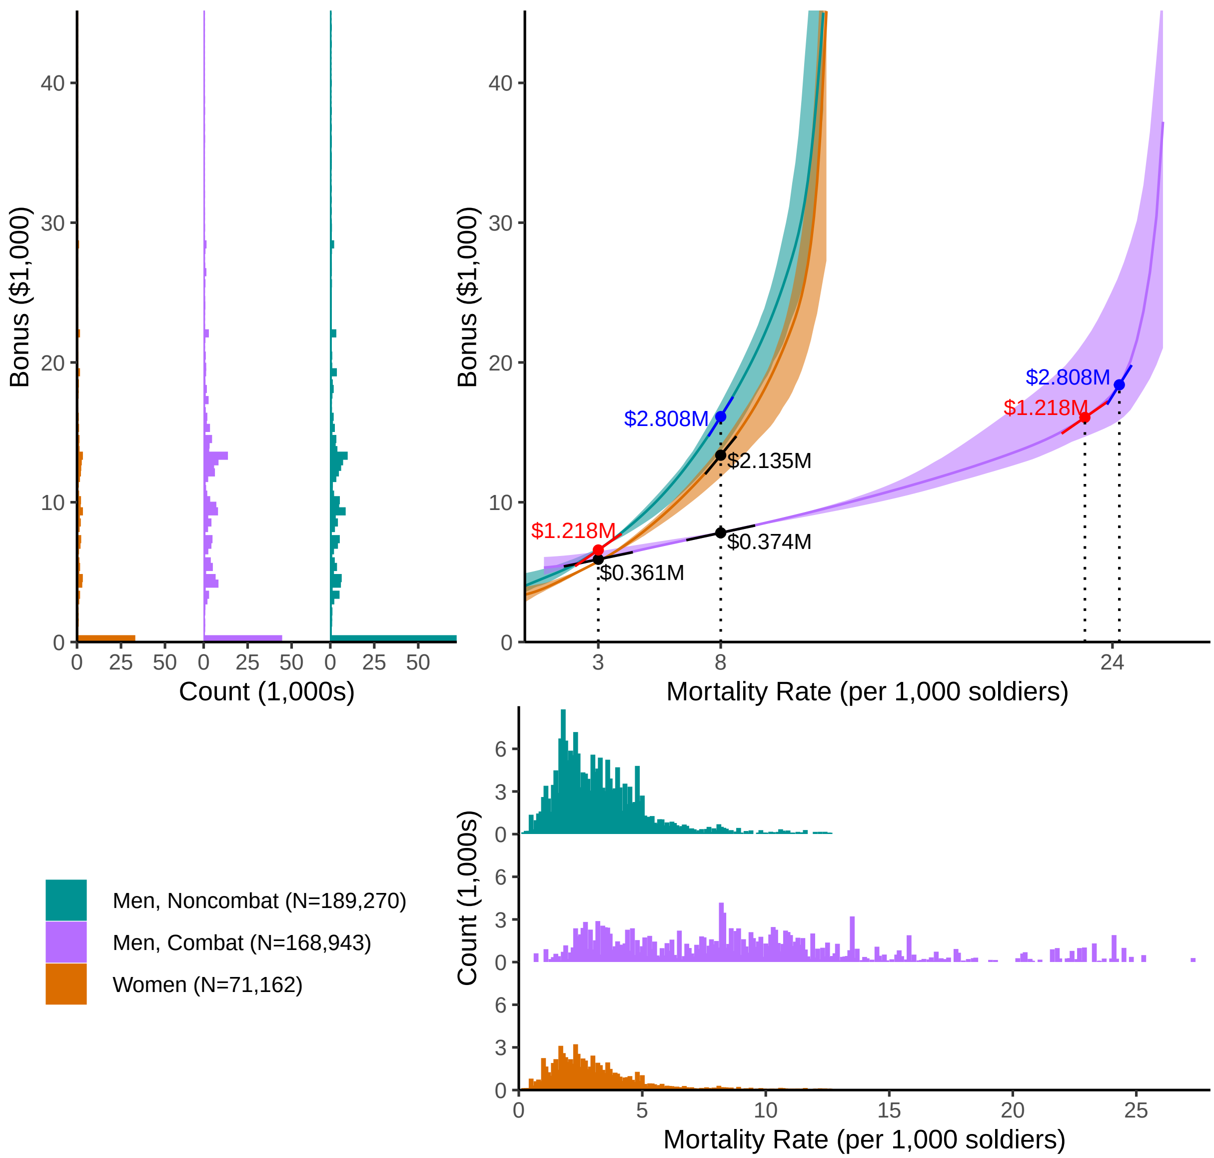

Notes: This figure depicts bid-curves, by subsample, derived from B-splines estimated using the column (2) specification from Table 3 when imposing monotonicity and convexity conditions. Shaded bands around each curve indicate the 95 percent confidence intervals, constructed from the 2.5th and 97.5th percentiles of solved bonus amounts at each expected mortality. The notes for Figure 3 describe how we construct the bid-curves and their confidence intervals. Text labels denote the implied VSL at each point.

Tables

Table 1: Descriptive Statistics

	Full Sample (1)	Men, Non-Combat (2)	Men, Combat (3)	Women (4)
Panel A: Reenlistment Statistics				
Proportion Reenlisted	0.453	0.488	0.431	0.409
Bonus Offer ($ 2019)	7,198	6,890	8,164	5,727
Mortality Hazard (per 1000 soldiers)	5.43	3.25	8.92	2.96
Annual Deployment Probability	0.209	0.193	0.242	0.175
Panel B: Demographic Characteristics				
Female	0.166	0	0	1
Black	0.172	0.201	0.080	0.317
Hispanic	0.126	0.126	0.118	0.143
High School GED / Dropout	0.138	0.128	0.178	0.071
High School Graduate	0.733	0.733	0.721	0.762
Some College	0.081	0.085	0.067	0.104
College Graduate	0.042	0.048	0.029	0.058
AFQT Score	59.50	60.37	60.04	55.87
Age at reenlistment decision	24.05	24.33	23.66	24.22
Unemployment Rate	6.51	6.57	6.47	6.48
Observations	429,375	189,270	168,943	71,162

Notes: This table reports group means for first-term, non-prior service enlisted soldiers with reenlistment decisions between 2002 and 2010. Bonus offers, measured in 2019 USD, are the maximum bonus associated with a four year reenlistment in the same military occupational speciality (MOS). A soldier's mortality hazard is a multiple of the annual mortality rate (per 1000 individuals) of soldiers in the same MOS averaged over the month she enters her reenlistment window and the 11 previous months. Since the modal reenlistment term length is four years, we convert the annual mortality rate into a four-year sum where mortality in future years is discounted at 7.2 percent. Soldiers in MOSs with an average strength of less than 5000 individuals are assigned to their CMF's (broader occupation) mortality hazard. The deployment probability is the percentage of soldiers in the same MOS who were deployed for at least one month in the year immediately before a soldier enters her reenlistment window. The unemployment rate is the annual unemployment rate in a soldier's home county in the year that she must decide to reenlist. 0.5 percent of the sample is missing education data.

Table 2: Bonus and Reenlistment Trends

Decision Year	Total Bonus Dollars, All Reenlistments ($ M)	Regular Annual Military Pay	Number of Soldiers Eligible for Reenlistment	Proportion of Sample Who Reenlisted	Average Mortality Hazard (per 1000)	5th - 95th Percentile of Mortality Hazard	Mean Bonus Offer	5th - 95th Percentile of Bonus Offers	Average Bonus Value for Bonus Recipients
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
2002	181	28,582	43,374	0.388	2.06	0.70 - 3.50	6,464	0 - 21,437	6,968
2003	142	29,081	51,178	0.364	2.92	1.05 - 6.10	1,772	0 - 10,071	9,793
2004	193	29,412	48,956	0.370	5.04	1.13 - 11.40	8,784	0 - 13,511	13,135
2005	653	29,080	46,298	0.396	5.53	1.66 - 11.47	9,870	0 - 19,359	15,326
2006	932	29,376	49,691	0.474	6.31	1.82 - 13.70	12,952	0 - 28,452	17,175
2007	697	29,217	47,567	0.506	9.26	1.67 - 24.15	9,992	0 - 16,033	13,998
2008	862	28,818	48,658	0.543	7.29	1.87 - 21.62	8,084	0 - 14,073	11,973
2009	629	30,331	45,488	0.545	4.92	1.95 - 9.41	4,290	0 - 11,880	8,952
2010	258	31,007	48,165	0.491	5.36	1.72 - 10.83	2,602	0 - 8,339	6,134
All Years	506	29,376	429,375	0.453	5.43	1.23 - 14.36	7,198	0 - 19,359	12,192

Notes: This table reports reenlistment, pay, mortality, and bonus trends for first-term soldiers with reenlistment decisions in the year indicated in the first column. All dollar values are 2019 USD. Column (2) reports the sum of reenlistment bonuses the Army paid to all soldiers (not just first-term soldiers); these numbers are from Army Financial Management: Assistant Secretary of the Army for Financial Management and Comptroller, accessed April 2021. Annual pay, column (3), is the military base pay for a Specialist (E4) with four years of service. The bonus offer and mortality hazard variables are constructed as described in Section 4.2 and the notes for Table 1. The All Years row displays the sum across years in column (3), the average weighted by number of soldiers eligible for reenlistment in columns (2), (4), (5), (6), (8), and (10), and the 5th and 95th percentile of the total distribution across all years in columns (7) and (9).

Table 3: Binary Logit Estimates, Full Sample and Subsamples

	(1)	(2)	(3)	(4)	(5)
Panel A: Full Sample (N=429,375)					
Bonus Offer Logit Coefficient	0.0193***	0.0192***	0.0205***	0.0210***	0.0210***
	(0.0010)	(0.0010)	(0.0010)	(0.0010)	(0.0012)
Effect of $1000 Bonus Increase	0.0046***	0.0046***	0.0048***	0.0049***	0.0048***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0003)
Mortality Hazard Logit Coefficient	-0.0100***	-0.0111***	-0.0121***	-0.0109***	-0.0113***
	(0.0018)	(0.0019)	(0.0018)	(0.0019)	(0.0034)
Effect of Mortality Increase (1 per 1000)	-0.0024***	-0.0026***	-0.0028***	-0.0025***	-0.0026***
	(0.0004)	(0.0005)	(0.0004)	(0.0004)	(0.0008)
Estimated VSL ($ M)	0.520	0.575	0.592	0.520	0.540
VSL CI	[0.334, 0.706]	[0.379, 0.771]	[0.412, 0.772]	[0.343, 0.696]	[0.231, 0.849]
Log Likelihood	-287,673	-287,668	-282,794	-282,771	-276,921
Panel B: Men in Non-Combat Occupations (N=189,270)					
Estimated VSL ($ M)	0.853	1.012	0.738	0.629	0.670
VSL CI	[0.283, 1.423]	[0.422, 1.602]	[0.187, 1.289]	[0.079, 1.179]	[-0.091, 1.430]
Panel C: Men in Combat Occupations (N=168,943)					
Estimated VSL ($ M)	0.530	0.438	0.451	0.362	0.449
VSL CI	[0.218, 0.841]	[0.101, 0.776]	[0.168, 0.734]	[0.079, 0.644]	[-0.072, 0.969]
Panel D: Women (N=71,162)					
Estimated VSL ($ M)	1.261	1.274	1.104	1.038	1.415
VSL CI	[0.622, 1.899]	[0.637, 1.912]	[0.478, 1.729]	[0.418, 1.658]	[0.526, 2.304]
Cohort FE, MOS FE, Term-length FE	X	X	X	X	
Deployment Probability		X	X	X	X
Individual Controls			X	X	X
County Unemployment Rate				X	X
Cohort x MOS x Term FE					X

Notes: This table reports implied VSL estimates in millions of 2019 USD derived from binary logit estimates of equation 4. Panel A also reports logit coefficients and average marginal effects for the full sample. The outcome is an indicator for whether a soldier reenlisted. Section 4 and the notes for Table 1 describe the construction of the bonus offer, mortality hazard, deployment probability, and unemployment rate variables. Individual controls include gender, race indicators (White, Black, Hispanic, Other), education indicators (dropout, high school graduate, some college, college, missing education (0.5 percent of sample)), state fixed effects, and linear terms for AFQT score and age. Standard errors clustered on each combination of MOS and decision month are reported in parentheses and VSL 95 percent confidence intervals, reported in brackets, are derived via the delta method. ***,**, and * denote significance at the 1%, 5%, and 10% level.

Table 4: Moment Forest, Random Coefficient, Multinomial Logit, and Nested Logit Results

	All MOSs			High Density MOSs		
	Moment Forest (1)	Random Coefficients (2)	MNL J=5 (3)	MNL J=5 (4)	Nested Logit J=5 (5)	MNL J=22 (6)
Panel A: Full Sample						
VSL ($ M)	0.874	0.769	0.702	0.717	0.653	0.605
VSL CI	[0.697, 1.051]	[0.571, 0.967]	[0.514, 0.889]	[0.513, 0.921]	[0.452, 0.854]	[0.478, 0.732]
Observations		429,375			277,877	
Panel B: Men in Non-Combat Occupations						
VSL ($ M)	0.685	0.758	0.858	1.023	1.012	0.703
VSL CI	[0.525, 0.844]	[0.394, 1.122]	[0.241, 1.474]	[0.413, 1.633]	[0.408, 1.615]	[0.234, 1.171]
Observations		189,270			104,051	
Panel C: Men in Combat Occupations						
VSL ($ M)	0.789	0.599	0.595	0.559	0.458	0.301
VSL CI	[0.634, 0.944]	[0.379, 0.818]	[0.237, 0.952]	[0.041, 1.077]	[-0.061, 0.977]	[0.086, 0.517]
Observations		168,943			131,022	
Panel D: Women						
VSL ($ M)	1.503	1.530	1.149	1.318	1.266	1.140
VSL CI	[1.047, 1.959]	[1.102, 1.958]	[0.490, 1.809]	[0.537, 2.099]	[0.544, 1.989]	[0.444, 1.836]
Observations		71,162			42,802	

Notes: Column (1) reports average bonus and average mortality hazard estimates across the distribution of individual-specific bonus and mortality estimates from the moment forest analysis described in Section 6.2. Column (2) reports implied VSLs in millions of 2019 USD derived from binary logit random coefficient estimates of equation 4 that allows for unobservable heterogeneity in soldiers' responses to the bonus and mortality rate by assuming these parameters are distributed normally. Columns (3) - (6) report implied VSLs derived from multinomial and nested logit regression estimates of equation 4. The choice alternatives for the J=5 multinomial logit specifications correspond to options 1-5 in Figure A1. The nested logit regression in column (5) permits errors to be correlated between reenlistment options 1, 2, 4, and 5, options where the soldier remains in the same MOS. The alternatives for the J=22 specification include each possible combination of reenlistment option and reenlistment term length (two through six years for options 1-2, three through six years for options 3-5). All random coefficient, multinomial logit, and nested logit estimates include the deployment probability control, entry cohort fixed effects, initial term-length fixed effects, and MOS fixed effects, corresponding to the column (2) specification from Table 3. All multinomial and nested logit specifications also include interactions between each reenlistment alternative and each control that does not vary with the reenlistment alternatives. VSL 95 percent confidence intervals, reported in brackets, are derived by applying the delta method to bonus and mortality standard error estimates. We construct standard errors for the moment forest framework by executing 50 bootstraps of the moment forest analysis, then calculating the standard deviation of the average bonus (or mortality) estimates derived from each of the 50 bootstraps. We estimate bootstrapped standard errors, estimated from 600 draws with replacement, for random coefficient models. We cluster standard errors on MOS for multinomial and nested logit estimates.

Table 5: Implied VSLs from Bid Curves at Different Mortality Levels by Subsample

	Mortality Hazard Rate (per 1000 soldiers)					
	1	3	5	8	16	24
Men, Noncombat	0.717	1.218	1.667	2.808		
	[0.425, 0.921]	[0.781, 1.397]	[1.214, 1.865]	[2.144, 3.281]		
Men, Combat	0.177	0.361	0.387	0.374	0.506	2.022
	[0.042, 0.512]	[0.246, 0.438]	[0.274, 0.395]	[0.290, 0.374]	[0.414, 0.918]	[0.983, 3.901]
Women	0.789	0.940	1.423	2.135		
	[0.429, 0.980]	[0.833, 1.094]	[1.138, 1.592]	[1.551, 2.462]		

Notes: This table reports implied VSL estimates in millions of 2019 USD for each subsample at six mortality hazard rates. VSL estimates are the slopes of the B-spline-derived bid curves in Figure 4. The range of mortality rates for noncombat men and women do not extend to 16 and 24 deaths per 1000 soldiers. 95 percent confidence intervals, constructed from the 2.5th and 97.5th percentiles of solved bonus amounts at each expected mortality rate, as described in the notes for Figure 3, are displayed in brackets.

1 Online Appendix: Overview of the Moment Forest

The moment forest approach of Nekipelov et al. (2021) uses a two-step procedure to estimate moment-based models with observable heterogeneity. The goal of their procedure is to estimate models of the general form:

$$E[Y - m(x; \theta(Z))] = 0, \tag{6}$$

where Y are outcomes, m is a set of moments, x are observable covariates, and $\theta(Z)$ is a set of parameters that are indexed by observables Z. Their approach generalizes the standard moment model, where θ is a universal vector of parameters that applies to all observations in the sample. The motivation is that parameters often vary by observable characteristics, e.g. demand parameters vary with demographics. They propose an ensemble estimator, a moment forest, that consists of a number of moment trees. In each tree, the estimator first partitions observations based on Z, and then recovers estimates of that θ within each partition. To achieve the partitioning, they leverage a modification of classification trees. The basic idea is to split the data into two parts and recursively partition the first sample on the basis of moment fit by cutting of data along a single dimension of X at each step. This process is repeated until some convergence criteria are met. In a second step, the structure of the partitioning from the first step is applied to a second sample of observations to recover estimates of θ in each partition. The moment forest's estimate of θ is then the average of θ estimated in each tree. They provide proofs of consistency and uniform convergence, and also show that the rate of convergence for the first step is, under weak regularity conditions, faster than parametric. This allows the econometrician to ignore the partitioning step when calculating standard errors in the second step.

In our setting, Y is an indicator of whether a soldier reenlisted, and the moment function is a logit discrete choice model with a simplified utility structure consisting of three components: a constant, a bonus, and a hazard rate. At the initial node of the tree, the estimator finds a model that fits the whole sample. It then searches over each of the X, which as illustrated by Figure A6 includes demographic variables, educational attainment, and test scores. It splits the sample into two partitions based on the single best improvement in the fit of the moment function at that node. It then repeats this split for each partition, continuing until convergence criteria are met. In the second step, we take the structure from the first step and estimate parameters within each of the final partitions using the second sample. Doing so generates the moment forest estimates reported in the paper.

2 Online Appendix: Data

The Army's Office of Economic and Manpower Analysis (OEMA) provided all military administrative data to Kyle Greenberg and Michael Yankovich as part of a restricted use agreement that specifies the data can only be stored, accessed, and analyzed using government information systems within government facilities. For access to the raw personnel data, one must make direct application to OEMA or the Defense Manpower Data Center. Data preparation and analysis programs are available from the authors on request.

The population for our setting begins with 589,211 first-term soldiers with no prior enlistments who would have been 12 months away from their initial End of Term of Service (ETS) date, the date a soldier's first enlistment contract ends, between 2002 and 2010 had they completed the their full enlistment contract. We exclude 92,415 soldiers who did not serve for at least one year in the Army as these soldiers are not eligible for reenlistment. We then exclude another 63,849 soldiers who had not reached the rank of Private First Class by the last year of their initial enlistment as these soldiers are also not eligible to reenlist. We exclude another 16 soldiers who achieved a rank of Sergeant First Class or higher during their initial enlistment as this is highly unusual and likely the result of a unique enlistment contract. We also drop 3,551 soldiers with an occupation of cryptologic linguist as we cannot identify the bonus they are eligible for (more on this below). Finally, we exclude 5 additional soldiers with occupations we cannot identify, leaving us with an analysis sample of 429,375 soldiers.

2.1 Details on the Selective Reenlistment Bonus Program

The Selective Retention Bonus (SRB) Program is a monetary incentive offered to qualified Soldiers who reenlist in the Regular Army for continued duty in certain military occupational specialties (MOSs).

Eligible occupations and locations are determined by the Army's Human Resources Command (HRC) to meet rank, skill, and mission requirements. Bonus policies change frequently at irregular intervals. If a soldier's current occupation or desired location is not eligible for an SRB, then the soldier might still be eligible for an SRB if he is currently serving in a full-strength or over-strength occupation and is willing to retrain for service in a shortage occupation.[40] Additionally, during the period 2004 through 2010, soldiers in all occupations

[40]The Army's Personnel Manning Authorization Document (PMAD) establishes authorized manpower strengths by occupation, rank, and unit. Over-strength occupations are those where the total number of soldiers in the occupation exceed the authorization established in the PMAD.

were generally eligible to reenlist and receive a bonus if reenlisting during a deployment.

Bonus offers with specific effective and expiry dates are periodically announced by the Army's HRC via Military Personnel (MILPER) messages. MILPER messages contain a complete menu of SRBs and reenlistment options available to each soldier at the time, thus providing policy announcements on the Army's reenlistment program several times each year.[41] Although these documents come in a variety of formats, they usually consist of tables that offer bonuses based on observable characteristics of soldiers. These include:

- *Military Occupational Specialty (MOS)*: a 3-digit code describing a soldier's occupation.
- *Rank*: these include Private 1st Class (E3), Corporal or Specialist (E4), Sergeant (E5), and Staff Sergeant (E6) (63 percent of first-term soldiers in our sample have a rank of Specialist when they enter their reenlistment window).
- *Skills*: codes that indicate additional skills acquired by soldiers, such as parachuting.
- *Locations*: countries or specific army bases to which the soldier applies for reenlistment, which may or may not be the soldier's present station.
- *Deployment Status*: Some bonus premiums apply to deployed soldiers.

About 1-2 years before the ETS date, the soldier enters the "reenlistment window," a time period during which they have the option to reenlist for another term of service. During this time, each soldier meets with a reenlistment officer, who explains reenlistment options and the bonuses that the soldier would receive depending on the reenlistment type they choose.

There are five reenlistment options which classify types of reenlistment.

1. *Regular Army Reenlistment:* the soldier has no control over where they are stationed in this option, and are not eligible for any location-specific bonuses. The soldier must remain in their current MOS.
2. *Current Station Stabilization:* this is the same as Option 1, but the soldier may elect to stay at their current station. Soldiers who choose this option are eligible for both location-specific and non-location-specific bonuses.
3. *Re-training:* the soldier has no control over where they are stationed in this option, and are not eligible for any location-specific bonuses. The soldier re-trains to a new MOS, and receives the bonus offer for the new MOS. Any skills the soldier has are carried over.

[41] For details, see `www.hrc.army.mil/Milper`.

4. *Outside Continental United States (OCONUS) reassignment:* the soldier chooses an overseas location, and is stationed there. The soldier must remain in their current MOS.

5. *Contiguous States and the District of Columbia (CONUS) reassignment:* the soldier chooses a location in the continental U.S., and is stationed there. The soldier must remain in their current MOS.

2.2 Details on the Construction of Bonus Offers

MILPER messages, released by the Army's HRC, are the data source for bonus information. We compile them into a time series of bonus offers by MOS, rank, skill, location and language.

The MILPER messages list an effective date and an expiration date for each bonus. Frequently, a MILPER message provides a bonus offer with an effective date that is in the future. On occasion there are multiple updates to a future bonus offer before the bonus takes effect. To avoid duplicate observations due to this feature of the data, we only keep the observation from the MILPER message with the latest issue date.

A MILPER message may be rescinded by subsequent messages. If this happens before a bonus's original expiry date, we update its expiry date. If a bonus is subsequently rescinded before its original effective date, we drop it from our data observation (because the bonus never came into effect).

During the period 2002-2010, the Army released at least 90 MILPER messages which announced changes to reenlistment and bonus policy. Each military personnel message specifies the occupation, rank, and level of bonus authorized. Some messages, such as the one shown in Appendix Figure A7, associate a bonus multiplier and a maximum bonus amount with each occupation and rank. Other messages associate a specific dollar amount with each occupation in rank. For example, Appendix Figure A7 shows that a Truck Driver (military occupation specialty (MOS 88M)) in the reenlistment eligibility window in July 2007 with rank of specialist (SPC) in zone A (i.e., has between 17 months and 6 years of active service) is eligible for a Selective Reenlistment Bonus (SRB) with multiplier 2. The bonus offer is the multiplier times the soldier's basic monthly pay, which increases with the amount of time a soldier has already served in the military, times the number of years of the reenlistment term that he selects. Since there is a cap on the bonus of \$10,000, the formula for calculating the bonus offer is given by:

$$B_{k=88M,t=July2007} = max\left\{Multiplier \cdot BasePay_{E4,4yos} \cdot ReupTerm,\ \$10,000\right\} \qquad (7)$$

For a truck driver with four years of service (monthly basic pay in 2007 of $1,978.50), the menu of reenlistment options available to the soldier is depicted in Table A11.[42]

As discussed above, SRBs for different occupations are often location or skill-specific. Location refers to a soldier's duty location. A skill refers to a certification achieved through specific military training. For example, a soldier who successfully complete Airborne training earns the skill identifier associated with being a "parachutist." If no location or skill is mentioned, however, we assume that a bonus offer applies in the broadest sense. For example, if a MILPER message states that MOS 89R has multiplier 3 for rank E4, we assume that this applies to all soldiers with MOS 89R and rank E4, regardless of skill or location, unless there is a location or skill-specific bonus offer for the same MOS active at the same time. If a skill is specified for a bonus offer, we assume all soldiers with that MOS-skill combination are eligible for that bonus offer, regardless of their present location. Similarly, if a location is specified for a bonus offer, we assume all soldiers with that MOS-location combination are eligible for that bonus offer, regardless of their skills.

While creating a database of potential SRBs, we had to make a few simplifying assumptions. First, our data does not contain information on the languages spoken by a soldier. Therefore, we cannot include language-specific bonus offers in our analysis. For this reason our analysis excludes 3,551 soldiers with the MOS of 98G (cryptologic linguists, later 35P), since the probability of soldiers in this MOS taking language-specific bonuses is high. We also remove language-specific bonuses for MOS 98C (signal intelligence analyst) and 97E (human intelligence collector, interrogator).[43]

Second, we do not have information on the specific divisions or brigades in which soldiers serve. However, a few bonuses are specific to these attributes (e.g., bonuses for the 75th Ranger Regiment or 82nd Airborne Division). If such a category constitutes an overwhelming majority of the soldiers at a given base (e.g., the 82nd Airborne for Ft. Bragg), we code bonuses for that category as being specific to the base. If a category is spread evenly over multiple bases, we remove these bonuses, since the best we can do on the soldier side is matching soldiers to locations, and we cannot guess that a soldier is in that unit from location alone.

[42] Soldiers who select new MOS training (reenlistment option 3) receive a bonus based on the new MOS they enter. We inferred bonus offers for this option by constructing an average of MOS-specific bonus offers among shortage occupations. Soldiers who select option 3 are only eligible to reenlist into shortage occupations.

[43] The Army redesignated the MOS code for cryptologic linguists (98G) and several other occupations between 2004 and 2007. We constructed an MOS-by-month crosswalk to match every MOS code to a time-consistent MOS code to ensure that every soldier matched to the appropriate occupation-specific bonus offers.

Finally, when constructing bonus offers for Option 3 (New Occupation), Option 4 (Choice of Overseas Post), and Option 5 (Choice of Continental U.S. Post), we treat the bonus associated with the choice of new occupation or new post as the average of all potential bonuses across occupations or posts a soldier could select with each particular option. We explored constructing bonus offers for these options as the maximum bonus among all possible choices, but found that the average bonus among possible choices was more predictive of actual bonuses received than the maximum bonus among possible choices.

To assess the measurement error introduced by the aggregation of bonuses across choice categories, we used the Army's RETAIN database to construct a matrix of actual reenlistment decisions, including actual bonus received, option selected, and term of additional service. After matching our constructed bonus offers to the actual bonuses that reenlisting soldiers received, we found that the correlation coefficient between bonus offers and bonuses received was 0.56. Among soldiers who reenlisted exactly 12 months prior to their initial ETS date, the correlation coefficient was 0.80. If our constructed bonus offers exhibit classical measurement error, then our estimates of the responsiveness to bonus offers will be biased downwards, implying that the true VSL is smaller than this paper's estimates of the VSL.

2.3 Details on the Construction of Expected Mortality Rates and Alternative Formulations

We assume that soldiers' expectation of future risk are based on recent mortality rates of other soldiers in their same military occupation code (MOS). This assumption seems reasonable for occupations with several soldiers who generally work with other soldiers in the same occupation within their military unit, such as infantrymen (11B) or combat engineers (12B). Yet this assumption might not be reasonable for soldiers in low-density occupations, such as microwave systems operator-maintainer (25P). Consistent with these assumptions, we find that soldiers in low-density MOSs exhibit little responsiveness to mortality rates of soldiers within their same MOS, but are much more responsive to mortality rates of soldiers within their same career management field (CMF), where a CMF represents a broader occupational category defined by the first two digits of a soldier's MOS (for example, CMF 11 includes MOS 11B, infantryman, and MOS 11C, indirect fire infantryman (mortarman)). Meanwhile, soldiers in high-density occupations exhibit a similar level of responsiveness to MOS mortality rates and CMF mortality rates (see Table A12).

Considering this, we define our preferred measure of a soldier's expected fatality rate as the average annual mortality rate among soldiers in the same occupation (MOS) over the month of a soldier's decision date and the 11 months prior, where the decision date is 12 months prior to the soldier's initial ETS date. However, for soldiers in low-density occupations, which we define as occupations that have an average, Army-wide annual strength of fewer than 5000 soldiers, we define a soldier's expected fatality rate as the average annual mortality rate of soldiers in the same CMF over the month of a soldier's decision date and the 11 months prior. Our results are not sensitive to alternative definitions of low-density occupations (see Table A5, discussed below).

We construct several alternative measures of expected mortality rates to account for different ways that soldiers might develop expectations of their mortality rate upon reenlistment. In addition to our preferred hazard rates, we also construct hazard rates based purely on MOS mortality rates, CMF mortality rates, MOS-CMF mortality rates that use different definitions of low-density MOSs, and Army-wide mortality rates that do not group mortality rates by occupational categories. Other alternative hazard formulations assume soldiers place more weight on mortality rates in months that are closer to the soldier's reenlistment decision date than in months that are further from the soldier's reenlistment decision date, while others assume soldiers infer mortality risk from the past 24 months rather than the past 12 months. We also also explore estimates derived from a mean-reverting time series mortality prediction. Finally, we consider hazard measures derived from regression tree machine learning algorithms that uses monthly mortality rates at the MOS-level to predict MOS-mortality rates in the 12 months following a soldier's reenlistment decision date, with variations on the number of months used to predict future mortality and variations on the minimum number of observations in each terminal node of the tree-growing algorithm. Table A2 reports estimates from a regression of soldier's future occupational mortality on these alternative hazard measures. We define a soldier's future occupational mortality as the average annual death rate of soldiers in the same occupation (MOS) in the 12 months after a the soldier's decision date. The notes for Table A2 provide additional details on how we construct each alternative hazard measure.

Figure A1: Enlistment to Reenlistment Process Diagram

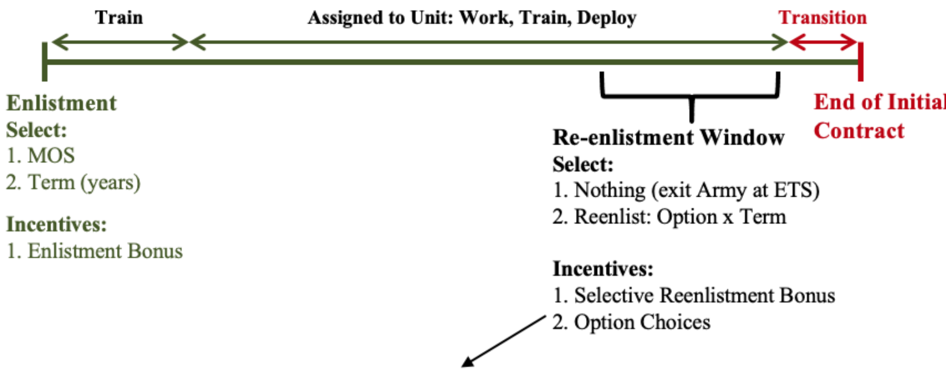

Reenlistment Alternatives

Option	Occupation	Home Post Location*	Term Length
1	No Change	Army Chooses	2, 3, 4, 5, or 6 years
2	No Change	Stay at Current Post	2, 3, 4, 5, or 6 years
3	New Occupation	Army Chooses	3, 4, 5, or 6 years
4	No Change	Pick Overseas Post	3, 4, 5, or 6 years
5	No Change	Pick Post inside U.S.	3, 4, 5, or 6 years

** Home post location refers to the permanent military base, post, or station to which the Soldier is formally assigned. This is not the same as a deployment location. Soldiers are assigned to units, which have their home-bases at permanent locations, and units get deployed to operational locations.*

Figure A2: Months between Initial Expiration of Term of Service (ETS) Date and Observed Reenlistment Date

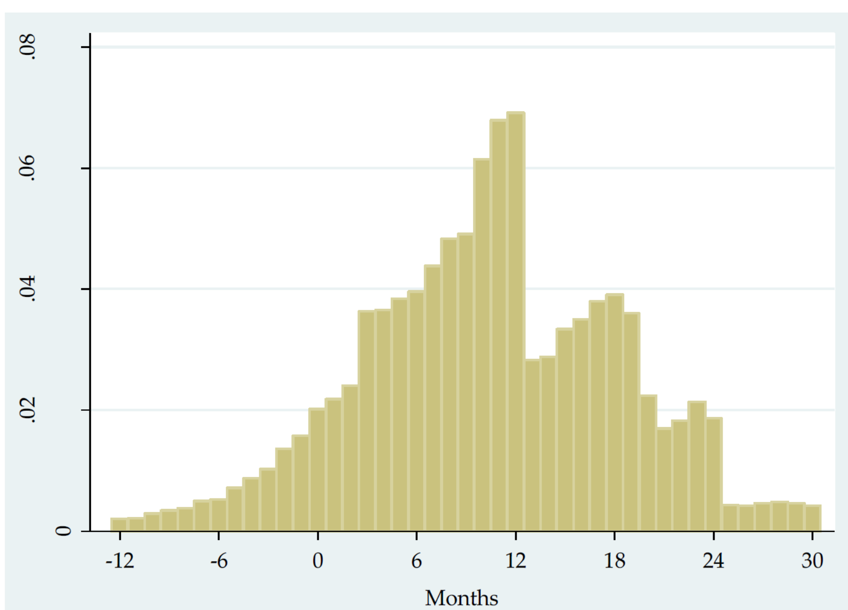

Notes: This figure plots the histogram of months between a soldier's initial ETS date and her observed reenlistment date among soldiers in the analysis sample.

Figure A3: Reenlistment Bonus Offers and Fatality Rates for Select Army Occupations, 2002-2010

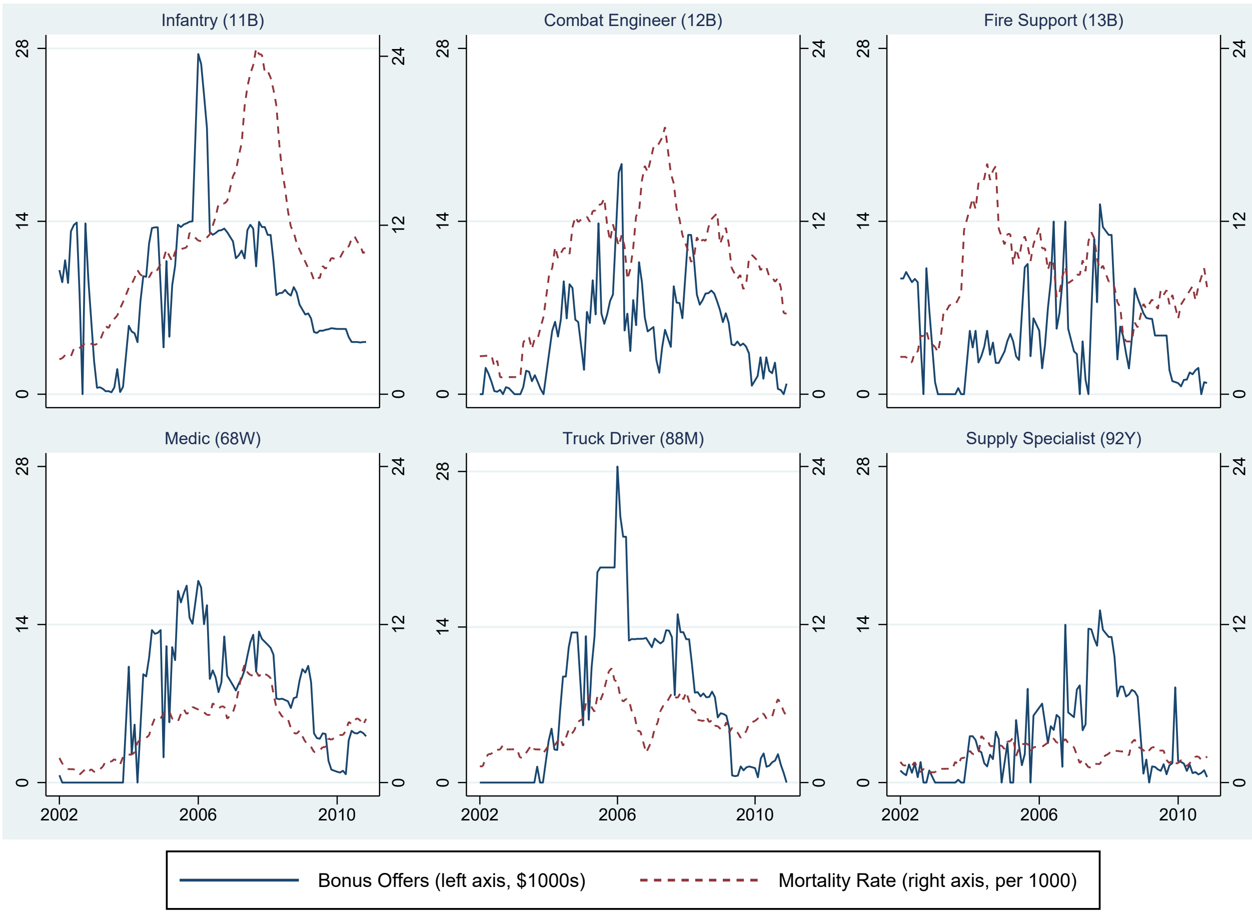

Notes: This figure reports average bonus offers and average mortality rates over time for selected occupations. The military occupational speciality (MOS) code for each occupation is in parentheses next to the name of the occupation. The bonus offers reflect lump-sum bonus amounts (in thousands of 2019 USD) available to soldiers with 3 years of service and the rank of Specialist who choose a 4-year, Regular Army Reenlistment in the same occupation (Option 1 per Figure A1). The mortality rate is the fatality rate of soldiers in the same occupation over the 12 previous months, but scaled to reflect the mortality rate per 1000 soldiers over the next four years where future years are discounted at 7.2 percent, as described in Section 4.2. For reference, a mortality rate of 12 implies an annual mortality rate of 3.32 death per 1000 person-years.

Figure A4: Moment Forest Bonus Estimate Distributions, Full Sample and by Subsample

(a) Full Sample

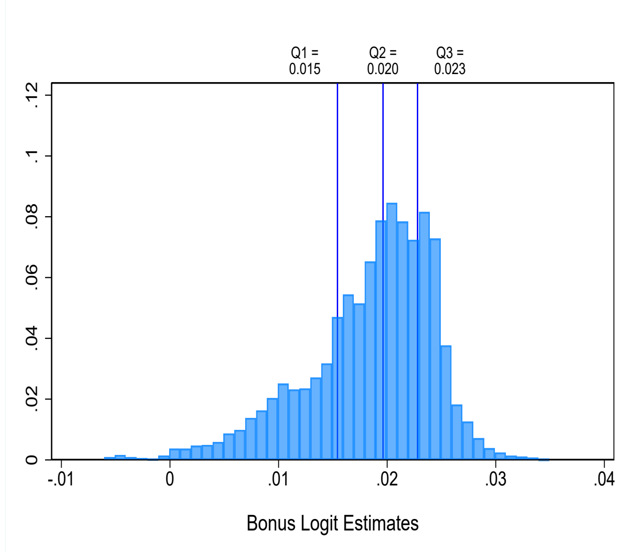

(b) By Subsample

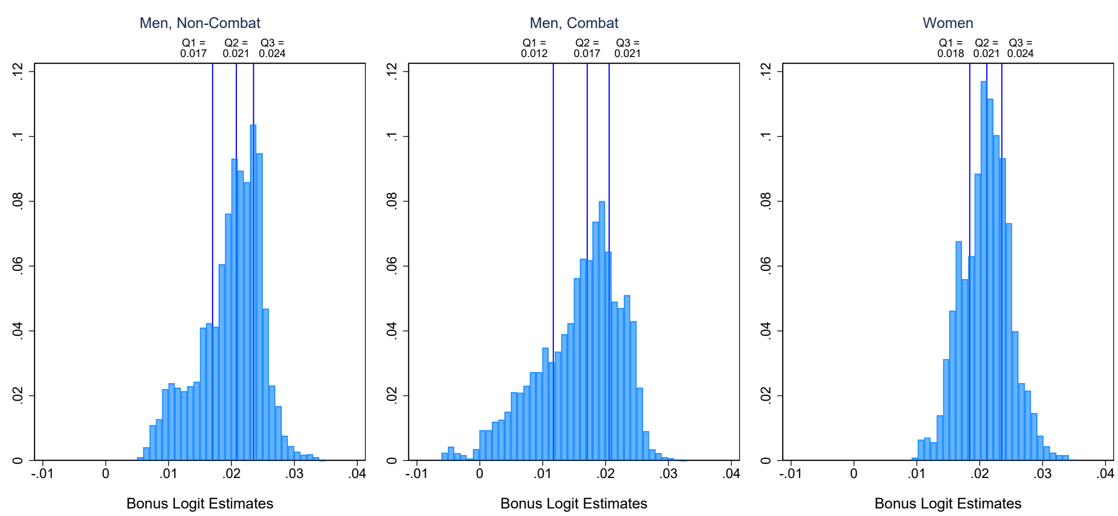

Notes: The histograms above display the distribution of logit bonus estimates derived from our moment forest analysis for the full sample, men in noncombat occupations, men in combat occupations, and women. Bonus offers are in 1000s of 2019 USD. Q1, Q2, and Q3 constitute the 25th percentile, median, and 75th percentile of estimates. This analysis permits potential splitting on MOS, initial term-length, entry-cohort, deployment probability, race, education, AFQT, age, state, and local unemployment rates. To reduce computing time, we convert continuous variables into discrete variables by grouping soldiers with similar covariate values. In particular, we group soldiers into deployment probability categories with bandwidths of 5 percentage points (i.e., one group with deployment probabilities less than 5 percent, another group with deployment probabilities between 5 and 10 percent, and so on), we group soldiers into AFQT categories with bandwidths of 5, we group soldiers into entry-age categories with bandwidths of 2 years, and we group soldiers into unemployment rate categories with bandwidths of one percentage point.

Figure A5: Moment Forest Hazard Estimate Distributions, Full Sample and by Subsample

(a) Full Sample

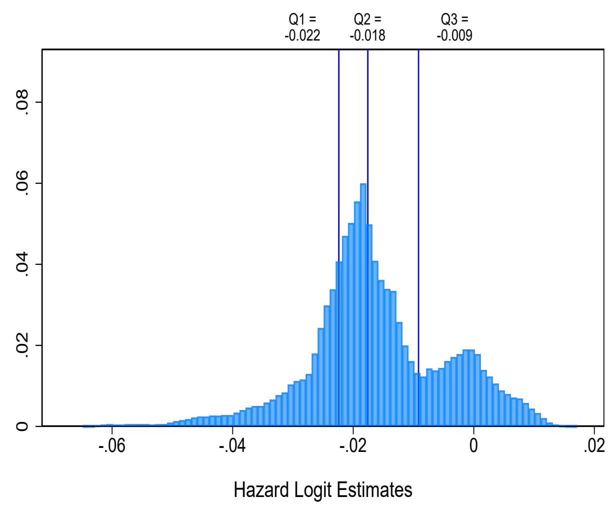

(b) By Subsample

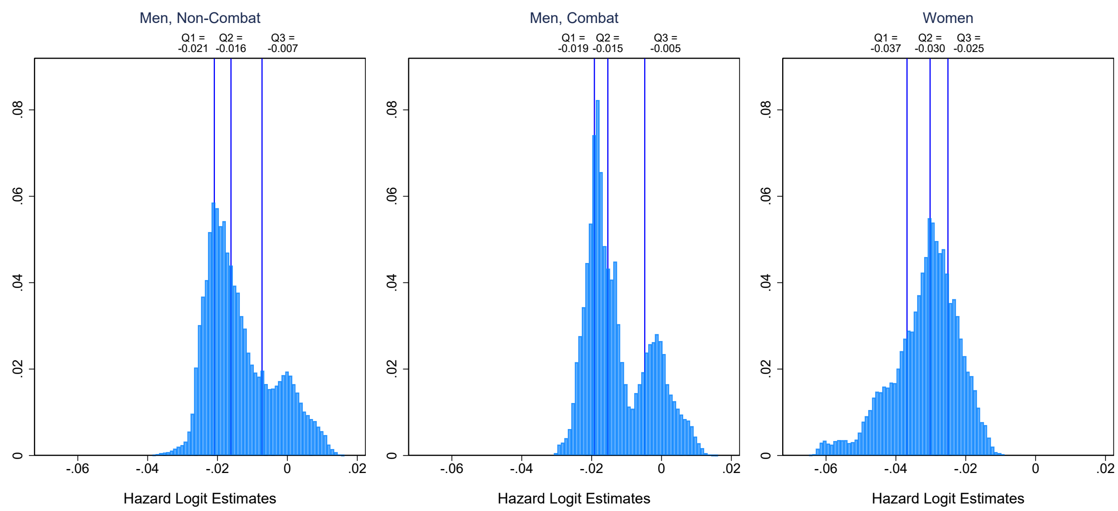

Notes: The histograms above display the distribution of logit mortality hazard estimates derived from our moment forest analysis for the full sample, men in noncombat occupations, men in combat occupations, and women. Mortality hazard rates are the four-year discounted sum of annual mortality rates per 1000 soldiers in the same occupation, as described in the notes for Table 1. Q1, Q2, and Q3 constitute the 25th percentile, median, and 75th percentile of estimates. See the notes for Figures 2 and A4 for additional details.

Figure A6: Moment Forest VSL Estimates for Alternative Subgroups

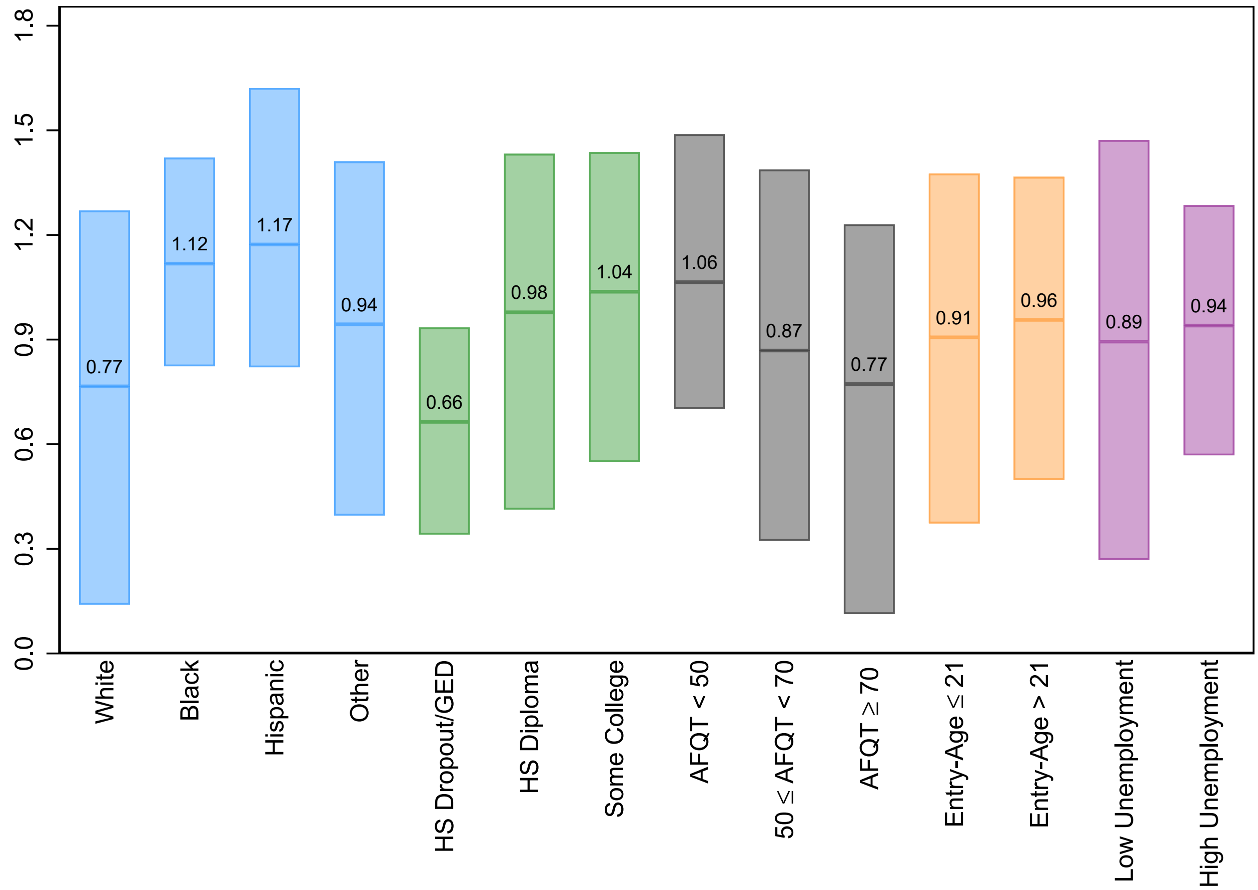

Notes: The box plots above display the 25th, 50th, and 75th percentiles of VSL estimates within each subgroup, as derived from the moment forest analysis. The moment forest analysis produces bonus and hazard estimates for each observation, which we use to construct percentiles within each subgroup. The text within each box reports the median VSL within the subgroup. AFQT scores are percentile-ranks (1-99) of a soldier's arithmetic and verbal reasoning skills relative to a nationally representative sample of 18-23-year-olds. 33 percent of the sample has an AFQT score below 50, 35 percent has an AFQT between 50 and 69, and 32 percent has an AFQT of 70 or higher. Entry-age refers to a soldier's age at initial enlistment into the Army. Soldiers in the low unemployment group are from counties with unemployment rates below the median unemployment rate within the full sample. Soldiers in the high unemployment group are from counties with unemployment rates above the median within the full sample.

Figure A7: Example of a Military Personnel Message Announcing Reenlistment Policy Updates

Milper Message Number 07-141
SELECTIVE REENLISTMENT BONUS (SRB) PROGRAM

Issued 04 JUNE 2007. This MILPER message will expire NLT 31 DECEMBER 2007.

This message announces changes to the Regular Army Active Component SRB Program. The effective date for additions to SRB multipliers or maximum payments is 05 JUNE 2007. The effective date of decreases and terminations to the SRB Program is 05 JULY 2007.

Zone A includes soldiers who have between 17 months and 6 years of service at time of reenlistment.

MOS/TITLE	SPC	SGT	SSG	CAP
11B Infantryman	1A	1.5A	1.5A	$15,000
13B Cannon Crew Member w/"P"	0	1A	0	$10,000
13F Fire Support Specialist	1.5A	2A	2A	$10,000
25U Signal Support Systems Spec.	1.5A	1.5A	0	$10,000
31B Military Police	1A	1.5A	1A	$10,000
68W Health Care Specialist	1A	1A	1A	$15,000
88M Motor Transport Operator	2A	1.5A	1A	$10,000
92Y Unit Supply Specialist	0.5A	0	0	$10,000
97E Interrogator	4A	4A	4A	$30,000

Table A1: Reenlistment Proportions by Option and Term, 2002-2010

			Reenlistment Term (Years)						
Option	Occupation	Home Post Location	0	2	3	4	5	6	Total
0	Exit Army		0.547						0.547
1	No Change	Army Chooses		0.034	0.018	0.030	0.013	0.023	0.117
2	No Change	Stay at Current Post		0.005	0.033	0.038	0.022	0.043	0.141
3	New Occupation	Army Chooses			0.006	0.021	0.006	0.007	0.040
4	No Change	Pick Overseas Post			0.003	0.033	0.008	0.008	0.051
5	No Change	Pick Post inside U.S.			0.036	0.035	0.014	0.017	0.102
Total			0.547	0.039	0.096	0.156	0.064	0.097	1.000

Notes: This table reports the proportion of the total sample that selected a particular combination of reenlistment option and reenlistment term-length. The sample consists of 429,375 first-term, non-prior service enlisted soldiers with reenlistment decisions between 2002 and 2010. Home post location refers to the permanent military base, post, or station to which the soldier is formally assigned. This is not the same as a deployment location. Soldiers are assigned to units, which have their home bases at permanent locations, and units get deployed to operational locations.

Table A2: Correlation Between Next Year's MOS Mortality Rate and Alternative Hazard Measures

	MOS/CMF (12mo) (Baseline) (1)	MOS Hazard (12mo) (2)	CMF Hazard (12mo) (3)	MOS/CMF (12mo, high dens MOS≥ 2500 soldiers) (4)	MOS/CMF (12mo, high dens MOS≥ 7500 soldiers) (5)	Army Hazard (12m) (6)	MOS/CMF (12mo) (Recency Weighting) (7)	MOS/CMF (24mo) (8)	MOS/CMF) (24mo) (Recency Weighting) (9)	MOS/CMF (12mo MRTS) (10)
Correlation Coef	0.681	0.659	0.676	0.680	0.674	0.206	0.694	0.651	0.691	0.678
Observations	381,208	381,193	381,208	381,208	381,208	381,208	381,208	381,208	381,208	381,208

	MOS/CMF (12mo MRTS with Iraq FE, Iraq Surge FE) (11)	MOS ML (12mo) (Stop-10) (12)	MOS ML (12mo) (Stop-20) (13)	MOS ML (12mo) (Stop-40) (14)	MOS ML (12mo) (Stop-80) (15)	MOS ML (24mo) (Stop-10) (16)	MOS ML (24mo) (Stop-20) (17)	MOS ML (24mo) (Stop-40) (18)	MOS ML (24mo) (Stop-80) (19)
Correlation Coef	0.677	0.728	0.728	0.721	0.686	0.664	0.664	0.695	0.671
Observations	381,208	381,156	381,156	381,156	381,156	381,063	381,063	381,063	381,063

Notes: This table reports the correlation coefficient between a soldier's occupational mortality rate in the 12 months after a soldier enters his reenlistment window on the hazard measures indicated in each column. A soldier's occupational mortality rate over the next 12 months is calculated as the mortality rate for soldiers in the same occupation (MOS). The correlations reported in this table exclude soldiers with decision years in 2010 because we do not have mortality data past that year. We adjust the mortality rate over the next 12 months to reflect a four-year discounted annual mortality, as we do with all other mortality rates (described in Section 4.2). For column 1, soldiers in high-density MOSs, occupations with an average strength of fewer than 5000 individuals, are assigned to their CMF (broader occupation) hazard rate defined over the 12 months leading up to their reenlistment decision date. Column 2 reports the correlation between next year's MOS-level mortality rate and the previous year's MOS-level mortality rate for all soldiers, and column 3 reports the correlation between next year's MOS-level mortality rate and the previous year's CMF-level mortality rate for all soldiers. For columns 4 and 5, we define high-density MOSs as soldiers in MOSs with an average strength exceeding 2500 and 7500 soldiers, respectively. For column 6, we map soldiers to the average Army-level mortality rate across all occupations. Column 7 reports the correlation when we replace the baseline hazard measure with an identical measure, but where we place more weight on occupational mortality rates in the months closer to a soldier's decision date. Specifically, column 7 calculates the hazard as $\frac{1}{78}\sum_{t=-11}^{t=0}(12+t)Hazard_t$, where $t=0$ is the month of a soldier's decision date, $t=-11$ is the 11th month prior to the month of a soldier's decision date, and $Hazard_t$ is the monthly occupational mortality rate in month t. Column 8 reports the correlation when we replace the baseline hazard measure with an identical measure, but where we average mortality rates over the 24 months leading up to a soldier's decision date (rather than the baseline of 12). Column 9 reports the correlation when we average mortality rates over the past 24 months, but placing more weight on months closer to a soldier's decision date as in column 7. Columns 10 and 11 replace the baseline hazard with a mean-reverting time-series hazard, with column 11 including a fixed effect for months after the March 2003 invasion of Iraq and another fixed effect for months between the January 2007 through May 2008 Iraq surge in the regression used to construct the mean-reverting time series hazard. Columns 12 through 19 replace the baseline hazard with a regression tree machine learning (ML) hazard prediction derived from monthly mortality rates at the MOS-by-month level. The ML hazard predictions in columns 12-15 use mortality data from the previous 12 months to predict MOS-level mortality rates over the next year. Columns 16-19 use mortality data from the previous 24 months to predict MOS-level mortality rates over the next year. For column 12, we constrain the ML hazard prediction such that the tree-growing algorithm results in terminal nodes with at least 10 observations. Columns 13-19 use the stopping criteria identified in the column heading.

Table A3: Binary Logit Estimates, Full Sample: Average Marginal Effects For All Variables

	(1)	(2)	(3)	(4)	(5)
Bonus Offer (1000s of $ 2019)	0.0046***	0.0046***	0.0048***	0.0049***	0.0048***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0003)
Mortality Hazard (per 1000 soldiers)	-0.0024***	-0.0026***	-0.0028***	-0.0025***	-0.0026***
	(0.0004)	(0.0005)	(0.0004)	(0.0004)	(0.0008)
Deployment Probability		0.0489**	0.0496**	0.0548**	0.2825***
		(0.0233)	(0.0229)	(0.0229)	(0.0356)
Female			-0.0946***	-0.0945***	-0.0980***
			(0.0026)	(0.0026)	(0.0026)
White			-0.0157***	-0.0157***	-0.0153***
			(0.0035)	(0.0035)	(0.0035)
Black			0.0858***	0.0853***	0.0836***
			(0.0039)	(0.0039)	(0.0039)
Hispanic			0.0245***	0.0237***	0.0233***
			(0.0039)	(0.0039)	(0.0039)
AFQT			-0.0019***	-0.0019***	-0.0020***
			(0.0001)	(0.0001)	(0.0001)
Age at Enlistment			0.0037***	0.0038***	0.0036***
			(0.0003)	(0.0003)	(0.0002)
High School GED / Dropout			0.1585***	0.1584***	0.1578***
			(0.0049)	(0.0049)	(0.0049)
High School Diploma			0.1253***	0.1251***	0.1251***
			(0.0045)	(0.0045)	(0.0044)
Some College			0.1076***	0.1075***	0.1079***
			(0.0049)	(0.0049)	(0.0049)
Education Missing			0.1439***	0.1437***	0.1445***
			(0.0116)	(0.0116)	(0.0119)
Home-County Unemployment Rate				0.0026***	0.0033***
				(0.0004)	(0.0004)
Cohort FE, MOS FE, Term-length FE	X	X	X	X	
Deployment Probability		X	X	X	X
Individual Controls			X	X	X
County Unemployment Rate				X	X
Cohort x MOS x Term FE					X

This table reports average marginal effects from estimates of equation 4 for all variables (excluding fixed effects) used to derive estimates in Table 3. All variables and estimating equations are identical to those of Table 3. The omitted race category is Other (5.6 percent of sample). The omitted education category is college graduate (4.2 percent of sample). Standard errors clustered on each combination of MOS and decision month are reported in parentheses. ***,**, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table A4: VSL Estimates with Alternative Reenlistment Decision Dates

	(1)	(2)	(3)	(4)	(5)
Panel A: Decision Date = 360 Days Prior to Initial ETS Date (Baseline Estimates from Table 3)					
Estimated VSL ($ M)	0.520	0.575	0.592	0.520	0.540
VSL CI	[0.334, 0.706]	[0.379, 0.771]	[0.412, 0.772]	[0.343, 0.696]	[0.231, 0.849]
Panel B: Decision Date = 180 Days Prior to Initial ETS Date					
Estimated VSL ($ M)	0.571	0.406	0.465	0.350	0.060
VSL CI	[0.379, 0.762]	[0.203, 0.608]	[0.279, 0.650]	[0.170, 0.530]	[-0.302, 0.421]
Panel C: Decision Date = 270 Days Prior to Initial ETS Date					
Estimated VSL ($ M)	0.531	0.518	0.550	0.480	0.379
VSL CI	[0.331, 0.731]	[0.308, 0.729]	[0.355, 0.744]	[0.287, 0.673]	[0.023, 0.735]
Panel D: Decision Date = 450 Days Prior to Initial ETS Date					
Estimated VSL ($ M)	0.442	0.513	0.532	0.505	0.693
VSL CI	[0.237, 0.647]	[0.300, 0.726]	[0.337, 0.727]	[0.309, 0.701]	[0.327, 1.059]
Panel E: Decision Date = 540 Days Prior to Initial ETS Date					
Estimated VSL ($ M)	0.363	0.448	0.466	0.459	0.096
VSL CI	[0.180, 0.546]	[0.257, 0.640]	[0.292, 0.641]	[0.282, 0.636]	[-0.205, 0.398]
Cohort FE, MOS FE, Term-length FE	X	X	X	X	
Deployment Probability		X	X	X	X
Individual Controls			X	X	X
County Unemployment Rate				X	X
Cohort x MOS x Term FE					X

Notes: This table reports implied VSL estimates derived from logit estimates of equation 4. Each panel reports VSL estimates when assigning soldiers to the bonus offer, mortality, and deployment probability corresponding to the date indicated in the panel heading. All variables are constructed in the same manner as described in the notes for Table 3. The equations used to produce regression estimates in columns (1) through (5) are identical to the estimation equations in the corresponding columns of Table 3. VSL 95 percent confidence intervals, reported in brackets, are derived by applying the delta method to standard errors clustered on each combination of MOS and decision month. ***,**, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table A5: VSL Estimates with Alternative Hazard Measures

	(1)	(2)	(3)	(4)	(5)
Panel A: MOS Hazard for MOSs with an Average Strength Exceeding 5000, CMF Hazard for Others (Baseline)					
Estimated VSL ($ M)	0.520	0.575	0.592	0.520	0.540
VSL CI	[0.334, 0.706]	[0.379, 0.771]	[0.412, 0.772]	[0.343, 0.696]	[0.231, 0.849]
Observations (Same In Each Panel)	429,375	429,375	429,375	429,375	429,375
Panel B: MOS Hazard					
Estimated VSL ($ M)	0.356	0.388	0.402	0.348	0.321
VSL CI	[0.204, 0.509]	[0.229, 0.547]	[0.255, 0.550]	[0.205, 0.492]	[0.088, 0.553]
Panel C: CMF Hazard					
Estimated VSL ($ M)	0.640	0.711	0.713	0.629	0.633
VSL CI	[0.434, 0.846]	[0.493, 0.930]	[0.514, 0.913]	[0.433, 0.824]	[0.287, 0.978]
Panel D: MOS Hazard for MOSs with an Average Strength Exceeding 2500, CMF Hazard for Others					
Estimated VSL ($ M)	0.510	0.568	0.589	0.520	0.484
VSL CI	[0.329, 0.691]	[0.376, 0.759]	[0.412, 0.765]	[0.347, 0.692]	[0.184, 0.784]
Panel E: MOS Hazard for MOSs with an Average Strength Exceeding 7500, CMF Hazard for Others					
Estimated VSL ($ M)	0.542	0.602	0.617	0.543	0.554
VSL CI	[0.356, 0.727]	[0.406, 0.799]	[0.435, 0.798]	[0.365, 0.721]	[0.243, 0.865]
Panel F: Army-Wide Hazard					
Estimated VSL ($ M)	0.823	0.951	0.719	0.244	0.929
VSL CI	[0.379, 1.267]	[0.477, 1.426]	[0.256, 1.182]	[-0.245, 0.733]	[0.227, 1.631]
Panel G: MOS-CMF Hazard Over Past 12 Months (Baseline), With Recency Weighting					
Estimated VSL ($ M)	0.423	0.461	0.487	0.425	0.338
VSL CI	[0.248, 0.597]	[0.277, 0.645]	[0.317, 0.657]	[0.260, 0.591]	[0.080, 0.596]
Panel H: MOS-CMF Hazard Over Past 24 Months					
Estimated VSL ($ M)	0.478	0.527	0.565	0.473	0.204
VSL CI	[0.240, 0.715]	[0.279, 0.776]	[0.334, 0.795]	[0.246, 0.699]	[-0.335, 0.744]
Panel I: MOS-CMF Hazard Over Past 24 Months, With Recency Weighting					
Estimated VSL ($ M)	0.432	0.472	0.496	0.433	0.393
VSL CI	[0.257, 0.608]	[0.288, 0.656]	[0.325, 0.667]	[0.266, 0.599]	[0.136, 0.650]
Panel J: MOS-CMF Mean-Reverting Time Series Hazard Measure					
Estimated VSL ($ M)	0.566	0.626	0.644	0.565	0.586
VSL CI	[0.364, 0.768]	[0.414, 0.839]	[0.448, 0.839]	[0.373, 0.756]	[0.251, 0.920]
Panel K: MOS-CMF Mean-Reverting Time Series Hazard Measure, With Iraq Invasion and Iraq Surge Fixed Effects					
Estimated VSL ($ M)	0.560	0.622	0.640	0.560	0.575
VSL CI	[0.359, 0.761]	[0.410, 0.834]	[0.445, 0.834]	[0.369, 0.751]	[0.243, 0.907]
Cohort FE, MOS FE, Term-length FE	X	X	X	X	
Deployment Probability Control		X	X	X	X
Individual Controls			X	X	X
County Unemployment Rate				X	X
MOS x Cohort x Term FE					X

Notes: Table continued on next page.

Table A5: VSL Estimates with Alternative Hazard Measures (Continued From Last Page)

	(1)	(2)	(3)	(4)	(5)
Panel L: Baseline Hazard, Restricted to Sample with Machine Learning Hazard Predictions					
Estimated VSL ($ M)	0.580	0.610	0.633	0.531	0.497
VSL CI	[0.372, 0.789]	[0.390, 0.829]	[0.431, 0.834]	[0.335, 0.727]	[0.173, 0.821]
Observations (Same In Each Panel)	381,156	381,156	381,156	381,156	381,156
Panel M: MOS Machine Learning Hazard Prediction, Last 12 Months (Nodes of size 10 or larger)					
Estimated VSL ($ M)	0.649	0.649	0.674	0.596	0.311
VSL CI	[0.384, 0.914]	[0.381, 0.917]	[0.425, 0.924]	[0.354, 0.837]	[0.037, 0.586]
Panel N: MOS Machine Learning Hazard Prediction, Last 12 Months (Nodes of size 20 or larger)					
Estimated VSL ($ M)	0.672	0.677	0.703	0.627	0.449
VSL CI	[0.408, 0.937]	[0.407, 0.948]	[0.448, 0.959]	[0.380, 0.875]	[0.195, 0.703]
Panel O: MOS Machine Learning Hazard Prediction, Last 12 Months (Nodes of size 40 or larger)					
Estimated VSL ($ M)	0.672	0.679	0.700	0.611	0.366
VSL CI	[0.372, 0.972]	[0.373, 0.984]	[0.411, 0.990]	[0.332, 0.891]	[0.063, 0.668]
Panel P: MOS Machine Learning Hazard Prediction, Last 12 Months (Nodes of size 80 or larger)					
Estimated VSL ($ M)	0.726	0.728	0.706	0.626	0.377
VSL CI	[0.383, 1.068]	[0.380, 1.076]	[0.371, 1.041]	[0.303, 0.949]	[0.048, 0.706]
Panel Q: MOS Machine Learning Hazard Prediction, Last 24 Months (Nodes of size 10 or larger)					
Estimated VSL ($ M)	0.833	0.833	0.835	0.757	0.384
VSL CI	[0.532, 1.133]	[0.531, 1.135]	[0.553, 1.117]	[0.486, 1.029]	[0.055, 0.713]
Panel R: MOS Machine Learning Hazard Prediction, Last 24 Months (Nodes of size 20 or larger)					
Estimated VSL ($ M)	0.833	0.833	0.835	0.757	0.384
VSL CI	[0.532, 1.133]	[0.531, 1.135]	[0.553, 1.117]	[0.486, 1.029]	[0.055, 0.713]
Panel S: MOS Machine Learning Hazard Prediction, Last 24 Months (Nodes of size 40 or larger)					
Estimated VSL ($ M)	0.751	0.757	0.777	0.688	0.281
VSL CI	[0.455, 1.047]	[0.458, 1.056]	[0.492, 1.062]	[0.412, 0.963]	[-0.023, 0.586]
Panel T: MOS Machine Learning Hazard Prediction, Last 24 Months (Nodes of size 80 or larger)					
Estimated VSL ($ M)	0.919	0.922	0.890	0.804	0.374
VSL CI	[0.565, 1.273]	[0.565, 1.280]	[0.554, 1.226]	[0.481, 1.128]	[0.028, 0.721]
Cohort FE, MOS FE, Term-length FE	X	X	X	X	
Deployment Probability Control		X	X	X	X
Individual Controls			X	X	X
County Unemployment Rate				X	X
MOS x Cohort x Term FE					X

Notes: This table reports implied VSL estimates derived from logit estimates of equation 4. Panel A reiterates the main VSL results from Table 4, obtained under our preferred estimates of expected mortality rates, while other panels report estimates from alternative definitions of the mortality hazard. The notes for Table A2 describe how we construct each alternative hazard measure. All variables are constructed in the same manner as described in the notes for Table 3. The equations used to produce regression estimates in columns (1) through (5) are identical to the estimation equations in the corresponding columns of Table 3. VSL 95 percent confidence intervals, reported in brackets, are derived by applying the delta method to standard errors clustered on each combination of MOS and decision month. ***,**, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table A6: Binary Logit Estimates with Individual Service Controls

	(1)	(2)	(3)	(4)	(5)
Panel A: Baseline Estimates from Table 3					
Estimated VSL ($ M)	0.520	0.575	0.592	0.520	0.540
VSL CI	[0.334, 0.706]	[0.379, 0.771]	[0.412, 0.772]	[0.343, 0.696]	[0.231, 0.849]
Panel B: With Controls For Rank					
Estimated VSL ($ M)	0.496	0.558	0.602	0.537	0.548
VSL CI	[0.301, 0.692]	[0.351, 0.764]	[0.406, 0.798]	[0.343, 0.730]	[0.185, 0.911]
Panel C: With Controls For Rank and Individual Deployment History					
Estimated VSL ($ M)	0.690	0.616	0.664	0.597	0.676
VSL CI	[0.446, 0.934]	[0.362, 0.870]	[0.427, 0.900]	[0.364, 0.830]	[0.164, 1.189]
Panel D: With Controls for Rank, Deployment History, and Current Deployment Status					
Estimated VSL ($ M)	0.668	0.601	0.641	0.575	0.657
VSL CI	[0.427, 0.908]	[0.351, 0.851]	[0.409, 0.872]	[0.347, 0.803]	[0.162, 1.152]
Panel E: With Controls for Rank, Deployment History, Current Deployment Status, and Exposure to Casualties					
Estimated VSL ($ M)	0.715	0.645	0.674	0.607	0.679
VSL CI	[0.472, 0.958]	[0.392, 0.898]	[0.440, 0.907]	[0.377, 0.837]	[0.181, 1.177]
Cohort FE, MOS FE, Term-length FE	X	X	X	X	
Deployment Probability		X	X	X	X
Individual Controls			X	X	X
County Unemployment Rate				X	X
Cohort x MOS x Term FE					X

Notes: This table reports implied VSL estimates derived from logit estimates of equation 4 where the outcome is an indicator for reenlisting. Panel A reiterates the main VSL results from Table 3. Panel B adds rank fixed effects to each specification, where rank is determined when a soldier enters the reenlistment window (360 days before initial ETS date). Panel C adds rank fixed effects and an indicator for a soldier having deployed to a combat zone prior to entering the reenlistment window. Panel D adds rank fixed effects, the individual deployment history indicator, and an additional indicator for being deployed when the soldier enters the reenlistment window. Panel D includes the additional controls from Panel D plus an indicator for having been exposed to casualties, which we define as being assigned to the same company where at least one other soldier was wounded or killed in action prior to entering the reenlistment window. The notes for Table 3 contain additional details. VSL 95 percent confidence intervals, reported in brackets, are derived by applying the delta method to standard errors clustered on each combination of MOS and decision month. ***,**, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table A7: Binary Logit Estimates for Subsamples

	(1)	(2)	(3)	(4)	(5)
Panel A: Men in Non-Combat Occupations (N=189,270)					
Bonus Offer	0.0054***	0.0053***	0.0056***	0.0057***	0.0053***
(1000s of $ 2019)	(0.0003)	(0.0004)	(0.0003)	(0.0003)	(0.0004)
Mortality Hazard	-0.0046***	-0.0054***	-0.0041**	-0.0036**	-0.0036*
(per 1000 soldiers)	(0.0016)	(0.0017)	(0.0016)	(0.0016)	(0.0021)
Estimated VSL ($ M)	0.853	1.012	0.738	0.629	0.670
VSL CI	[0.283, 1.423]	[0.422, 1.602]	[0.187, 1.289]	[0.079, 1.179]	[-0.091, 1.430]
Panel B: Men in Combat Occupations (N=168,943)					
Bonus Offer	0.0029***	0.0029***	0.0033***	0.0033***	0.0030***
(1000s of $ 2019)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0004)
Mortality Hazard	-0.0015***	-0.0013**	-0.0015***	-0.0012**	-0.0013*
(per 1000 soldiers)	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0008)
Estimated VSL ($ M)	0.530	0.438	0.451	0.362	0.449
VSL CI	[0.218, 0.841]	[0.101, 0.776]	[0.168, 0.734]	[0.079, 0.644]	[-0.072, 0.969]
Panel C: Women (N=71,162)					
Bonus Offer	0.0065***	0.0065***	0.0067***	0.0067***	0.0081***
(1000s of $ 2019)	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0005)
Mortality Hazard	-0.0083***	-0.0083***	-0.0074***	-0.0070***	-0.0115***
(per 1000 soldiers)	(0.0022)	(0.0021)	(0.0021)	(0.0022)	(0.0037)
Estimated VSL ($ M)	1.261	1.274	1.104	1.038	1.415
VSL CI	[0.622, 1.899]	[0.637, 1.912]	[0.478, 1.729]	[0.418, 1.658]	[0.526, 2.304]
Tests for Equality of Estimates (p-value): Men in Non-Combat Occs vs. Men in Combat Occs					
Bonus Offer	0.000	0.000	0.000	0.000	0.000
Hazard	0.074	0.018	0.121	0.169	0.305
VSL	0.339	0.110	0.374	0.406	0.630
Tests for Equality of Estimates (p-value): Men in Non-Combat Occs vs. Women					
Bonus Offer	0.003	0.002	0.005	0.007	0.000
Hazard	0.089	0.161	0.128	0.113	0.028
VSL	0.270	0.486	0.304	0.246	0.147
Tests for Equality of Estimates (p-value): Men in Combat Occs vs. Women					
Bonus Offer	0.000	0.000	0.000	0.000	0.000
Hazard	0.002	0.001	0.007	0.008	0.006
VSL	0.052	0.029	0.068	0.057	0.072
Cohort FE, MOS FE, Term-length FE	X	X	X	X	
Deployment Probability		X	X	X	X
Individual Controls			X	X	X
County Unemployment Rate				X	X
Cohort x MOS x Term FE					X

Notes: This table reports average marginal effects from binary logit estimates of equation 4 for key subsamples. All variables and estimating equations are identical to those used to produce estimates in columns (1) through (5) of Table 3. Implied VSL estimates are derived from logit coefficient estimates (not reported). Standard errors, clustered on each combination of MOS and decision month, are reported in parentheses. ***,**, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table A8: Moment Forest Results

	(1) Full Sample	(2) Men, Non-Combat	(3) Men, Combat	(4) Women
Bonus Offer	0.0185***	0.0199***	0.0158***	0.0210***
(1000s of $2019)	(0.0007)	(0.0007)	(0.0008)	(0.0010)
Mortality Hazard	-0.0161***	-0.0136***	-0.0125***	-0.0315***
(per 1000)	(0.0014)	(0.0014)	(0.0012)	(0.0044)
Estimated VSL ($ M)	0.874	0.685	0.789	1.503
VSL CI	[0.697, 1.051]	[0.525, 0.844]	[0.634, 0.944]	[1.047, 1.959]
Observations	429,375	189,270	168,943	71,162

Notes: This table reports moment forest logit estimates from a specification that permits splitting on all covariates used in column (5) of Table 3 (entry cohort, initial enlistment term-length, MOS, deployment probability, gender, race, education level, age, AFQT score, state, and local unemployment rates). To reduce computing time, we convert continuous variables into discrete variables by grouping soldiers with similar covariate values. In particular, we group soldiers into deployment probability categories with bandwidths of 5 percentage points (i.e. one group with deployment probabilities less than 5 percent, another group with deployment probabilities between 5 and 10 percent, and so on), we group soldiers into AFQT categories with bandwidths of 5, we group soldiers into entry-age categories with bandwidths of 2 years, and we group soldiers into unemployment rate categories with bandwidths of one percentage point. Bootstrap standard errors, estimated from 50 draws with replacement, are reported in parentheses. ***,**, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table A9: Random Coefficient Results

	Full Sample (1)	Men, Non-Combat (2)	Men, Combat (3)	Women (4)
Bonus Offer	0.0189***	0.0246***	0.0106***	0.0272***
(1000s of $ 2019)	(0.0005)	(0.0015)	(0.0016)	(0.0016)
Mortality Hazard	-0.0145***	-0.0187***	-0.0063***	-0.0417***
(per 1000 soldiers)	(0.0018)	(0.0052)	(0.0011)	(0.0049)
Sigma Bonus	0.0717***	0.0869***	0.0402***	0.0729***
	(0.0157)	(0.0260)	(0.0140)	(0.0216)
Sigma Hazard	0.0733***	0.0131***	0.0346***	0.0263***
	(0.021)	(0.004)	(0.011)	(0.008)
VSL ($ M)	0.769	0.758	0.599	1.530
VSL CI	[0.571, 0.967]	[0.394, 1.122]	[0.379, 0.818]	[1.102, 1.958]
Observations	429,375	189,270	168,943	71,162

Notes: This table reports estimates from a binary logit random coefficients model that allows for unobservable heterogeneity in soldiers' responses to the bonus and mortality rate by assuming these parameters are distributed normally. The outcome is an indicator for whether a soldier reenlisted. All regressions include entry cohort fixed effects, initial enlistment term-length fixed effects, MOS fixed effects, and the deployment probability control, corresponding to the column (2) specification from Table 3. Bootstrap standard errors, estimated from 600 draws with replacement, are reported in parentheses. VSL 95 percent confidence intervals, reported in brackets, are derived by applying the delta method to bonus and mortality rate standard errors. ***,**, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table A10: Implied VSLs from B-spline Estimates of the Reenlistment Response to Bonus Offers and Mortality Rates

	Mortality Hazard Rate (per 1,000 soldiers)					
	1	3	5	8	16	24
Full Sample	0.943	1.418	0.627	1.306	-0.534	-0.205
(w/o Convexity, w/o Monotonicity)	[-0.029, 1.547]	[0.392, 1.311]	[0.652, 9.257]	[-5.415, 32.254]	[-12.391, 0.955]	[-9.375, 3.417]
Full Sample	1.070	0.768	2.189	1.864	0.116	2.011
(w/o Convexity, with Monotonicity)	[0.460, 1.508]	[0.589, 0.964]	[1.228, 2.940]	[1.365, 3.300]	[0.054, 0.811]	[0.266, 4.667]
Full Sample	0.482	0.594	0.672	0.792	1.440	4.685
(with Convexity, with Monotonicity)	[0.253, 0.723]	[0.540, 0.722]	[0.590, 0.938]	[0.644, 1.021]	[1.100, 1.433]	[3.103, 14.200]

Notes: This table reports implied VSL estimates in millions of 2019 USD for the full sample at the mortality hazard rates indicated in the column headings. VSL estimates are the slopes of the B-spline-derived bid curves in Figure 3. The first row reports estimated VSLs and confidence intervals on the full sample when we impose neither monotonicity nor convexity. The next row reports estimated VSLs and confidence intervals when we impose monotonicity but not convexity. The last row reports estimated VSLs and confidence intervals when we impose both monotonicity and convexity. We use the method described in Chernozhukov and Hong (2003) to construct the confidence interval (see the notes for Figure 3).

Table A11: Menu of Reenlistment Options for Truck Driver in July 2007

			Reenlistment Term (Years)				
Option	Occupation	Home Post Location	2	3	4	5	6
1	No Change	Army Chooses	$0	$5,936	$7,914	$9,893	$10,000
2	No Change	Stay at Current Post	$0	$5,936	$7,914	$9,893	$10,000
3	New Occupation	Army Chooses	n/a	n/a	n/a	n/a	n/a
4	No Change	Pick Overseas Post	n/a	$5,936	$7,914	$9,893	$10,000
5	No Change	Pick Post inside U.S.	n/a	$5,936	$7,914	$9,893	$10,000

Notes: This table shows the example of the bonus offer slate for a Motor Transport Operator (military occupational specialty 88M) with the rank of Specialist (E4) and with four years of service when entering the reenlistment window in July 2007. Available choices of home post location in Options 4 and 5 depend on the needs of the Army.

Table A12: Logit Estimates in High vs. Low Density MOSs

	MOS-level inside hazard			CMF-level inside hazard		
	All MOSs (1)	High-Density MOSs (2)	Low-Density MOSs (3)	All MOSs (4)	High-Density MOSs (5)	Low-Density MOSs (6)
Panel A: Full Sample						
Bonus Offer	0.019***	0.020***	0.019***	0.019***	0.020***	0.019***
(1000s of $ 2019)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Mortality Hazard	-0.007***	-0.010***	-0.002	-0.014***	-0.013***	-0.028***
(per 1000 soldiers)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.006)
Estimated VSL ($ M)	0.39	0.52	0.10	0.71	0.65	1.41
VSL CI	[0.23, 0.55]	[0.31, 0.73]	[-0.13, 0.33]	[0.49, 0.93]	[0.42, 0.88]	[0.83, 1.99]
Observations	429,324	277,898	151,426	429,341	277,898	151,443
Panel B: Men in Non-Combat Occupations						
Bonus Offer	0.022***	0.030***	0.015***	0.022***	0.030***	0.016***
(1000s of $ 2019)	(0.001)	(0.003)	(0.001)	(0.001)	(0.003)	(0.001)
Mortality Hazard	-0.006**	-0.027***	-0.005	-0.026***	-0.028**	-0.028***
(per 1000 soldiers)	(0.003)	(0.010)	(0.003)	(0.008)	(0.014)	(0.008)
Estimated VSL ($ M)	0.30	0.88	0.31	1.18	0.92	1.76
VSL CI	[0.02, 0.57]	[0.30, 1.46]	[-0.10, 0.71]	[0.49, 1.88]	[0.06, 1.77]	[0.69, 2.83]
Observations	189,238	104,054	85,182	189,253	104,054	85,197
Panel C: Men in Combat Occupations						
Bonus Offer	0.012***	0.009***	0.022***	0.012***	0.009***	0.022***
(1000s of $ 2019)	(0.001)	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)
Mortality Hazard	-0.004**	-0.002	0.000	-0.008***	-0.005**	-0.025***
(per 1000 soldiers)	(0.002)	(0.002)	(0.004)	(0.002)	(0.003)	(0.009)
Estimated VSL ($ M)	0.33	0.25	-0.02	0.66	0.56	1.17
VSL CI	[0.02, 0.64]	[-0.26, 0.75]	[-0.36, 0.33]	[0.29, 1.03]	[0.02, 1.09]	[0.32, 2.02]
Observations	168,940	131,040	37,899	168,940	131,040	37,899
Panel D: Women						
Bonus Offer	0.028***	0.030***	0.027***	0.028***	0.030***	0.027***
(1000s of $ 2019)	(0.002)	(0.003)	(0.002)	(0.002)	(0.003)	(0.002)
Mortality Hazard	-0.004	-0.038***	0.005	-0.039***	-0.044**	-0.036***
(per 1000 soldiers)	(0.005)	(0.012)	(0.005)	(0.011)	(0.018)	(0.014)
Estimated VSL ($ M)	0.14	1.27	-0.19	1.40	1.50	1.33
VSL CI	[-0.18, 0.47]	[0.50, 2.05]	[-0.54, 0.17]	[0.61, 2.18]	[0.31, 2.68]	[0.33, 2.34]
Observations	71,141	42,804	28,337	71,143	42,804	28,339

Notes: This table reports binary logit coefficients from estimates of equation 4 for the samples identified in each panel heading. The outcome is an indicator for whether a soldier reenlisted. Low-density MOS are occupations with an average annual strength below 5,000 soldiers. All estimates are produced from regressions of the column (2) specification from Table 3, which include entry cohort fixed effects, term-length fixed effects, MOS fixed effects, and the deployment probability control. Standard errors clustered on each combination of MOS and decision month are reported in parentheses. ***,**, and * denote significance at the 1%, 5%, and 10% level, respectively.