ABSTRACT

Title of dissertation: SIMULATING AND OPTIMIZING: MILITARY MANPOWER MODELING AND MOUNTAIN RANGE OPTIONS

Andrew O. Hall, Ph.D. Candidate, 2009

Dissertation directed by: Professor Michael Fu

In this dissertation we employ two different optimization methodologies, dynamic programming and linear programming, and stochastic simulation. The first two essays are drawn from military manpower modeling and the last is an application in finance.

First, we investigate two different models to explore the military manpower system. The first model describes the optimal retirement behavior for an Army officer from any point in their career. We address the optimal retirement policies for Army officers, incorporating the current retirement system, pay tables, and Army promotion opportunities. We find that the optimal policy for taste-neutral Lieutenant Colonels is to retire at 20 years. We demonstrate the value and importance of promotion signals regarding the promotion distribution to Colonel. Signaling an increased promotion opportunity from 50% to 75% for the most competitive officers switches their optimal policy at twenty years to continuing to serve and competing for promotion to Colonel.

The second essay explores the attainability and sustainability of Army force

profiles. We propose a new network structure that incorporates both rank and years in grade to combine cohort, rank, and specialty modeling without falling into the common pitfalls of small cell size and uncontrollable end effects. This is the first implementation of specialty modeling in a manpower model for U.S. Army officers. Previous specialty models of the U.S. Army manpower system have isolated accession planning for Second Lieutenants and the Career Field Designation process for Majors, but this is the first integration of rank and specialty modeling over the entire officer's career and development of an optimal force profile.

The last application is drawn from financial engineering and explores several exotic derivatives that are collectively known Mountain Range options, employing Monte Carlo simulation to price these options and developing gradient estimates to study the sensitivities to underlying parameters, known as "the Greeks". We find that IPA and LR/SF methods are efficient methods of gradient estimation for Mountain Range products at a considerably reduced computation cost compared with the commonly used finite difference methods.

SIMULATING AND OPTIMIZING: MILITARY MANPOWER MODELING AND MOUNTAIN RANGE OPTIONS

by

Andrew O. Hall

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2009

Advisory Committee:
Professor Michael Fu/Advisor
Professor Zhi-Long Chen
Professor Bruce L. Golden
Professor Itir Z. Karaesmen
Professor John P. Rust

Acknowledgments

Thanks to my professors and my committee who have helped me in my journey, the students of the Robert H. Smith School of Business, and the Mathematical Finance RIT. Thanks to Professor Karaesmen for sharing your love of Operations Management and to Professor Rust for sharing your depth of knowledge on Dynamic Programming and Computational Economics. Thanks to Professor Golden for selecting the great group of students that entered with me and taking us on a continuing adventure in education and research. Special thanks to Scott Nestler for encouraging me to join him in College Park. Thanks to Paul Hogan and John Warner for their helpful suggestions on my Army Officer Retirement Model at the Western Economic Association International Conference, 2008. Thanks to Robert Dell for his helpful suggestions on my Officer Manpower Model at the INFORMS Computing Society Conference, 2009. Thanks to COL Michael Phillips for selecting me for the opportunity to further my studies and my teaching position at USMA, and thanks to Tom Weimer for always being there. Thanks to Marilyn and Stan for your love and taking care of us all at Quarters 4760. Thanks to Marty, Jen and Jasper, Kristi, Adam and Matta, and Mom for being such a great family.

Special thanks to Professor Michael Fu for his leadership and guidance.

Thanks to my wife Mary Lou, and children, Catherine and Oscar, for putting up with me during my studies and for their continual love and support.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

A	action space
α	taste for military lifestyle parameter
β	discount rate
c	military compensation for housing and subsistence
η	age at commissioning
K	strike
L	limit
m	military pay
$\hat{m}$	average of high three years of military pay
π	a policy
π^*	an optimal policy
Π	the space of all Markovian policies
P	transition matrix
q	quit
$Q_t^*(x,a)$	Q-function value giving the expected reward for taking action a in state x
s	stay
r	retirement percentage
R	reward function
S	stock
t	time
T	final time
τ	life expectancy
V^π	value function for policy π
V^*	optimal value function
$V_t^\pi(x)$	reward-to-go value function for policy π from state x at time t
$V^*(x)$	optimal value for starting at state x
$V_t^*(x)$	optimal reward-to-go from state x at time t
ξ	civilian raises
X	state space
$\mathbb{Z}$	the set of integers
$1\{\cdot\}$	indicator function of the set $\{\cdot\}$
x^+	$\max(x, 0)$

ACOL	Annualized Cost of Leaving model
AAMMP	Active Army Military Manpower Program
BAH	Basic Allowance for Housing
BAS	Basic Allowance for Subsistence
EG	Enlisted Grade Manpower Model
AVF	All Volunteer Force
CFD	Career Field Designation
COL	Colonel
COMPLIP	Computation of Manpower Programs using Linear Programming
CPT	Captain
DOD	Department of Defense
DOPMA	Defense Officer Personnel Management Act
DRM	Dynamic Retention Model
ELIM	Enlisted Loss Inventory Module
FY	Fiscal Year
GBM	Geometric Brownian Motion
GWOT	Global War on Terror
G1	Deputy Chief of Staff for Personnel
G3	Deputy Chief of Staff for Operations
IPA	Infinitesimal Perturbation Analysis
LT	Lieutenant
LTC	Lieutenant Colonel
LR/SC	Likelihood Ratio / Score Function
MAJ	Major
MDP	Markov Decision Process
NDAA	National Defense Authorization Act
OCS	Officer Candidate School
PA	Perturbation Analysis
PMAD	Program Manning Authorization Document
POM	Program Objective Memorandum
PRS	Strength Forecasting Directorate of Plans and Resources Division, Army G1
P&L	Profit and Loss
RMC	Regular Military Compensation
ROTC	Reserve Officer Training Corps
SPA	Smoothed Perturbation Analysis
SSA	Subsystem Aggregate
VG	Variance Gamma
1LT	First Lieutenant
2LT	Second Lieutenant

Chapter 1

Introduction

Management Science models problems from many diverse application areas and also uses a diverse set of tools including those from optimization and simulation. In this dissertation we employed two different optimization methodologies, dynamic programming and linear programming, and gradient estimation from stochastic simulation. Two of the following three essays are drawn from the application area of military manpower modeling and the last from financial engineering.

Management Science and Operations Management are inherently inter-disciplinary endeavors. Increasing research in the areas of interface between Operations Management and Organizations was mentioned in Chopra et al. (2004) as part of the way ahead for operations management in the 50th Anniversary Article for Management Science. Human resource management systems were in turn recognized in Boudreau (2004) as an important area of organizations for the application of management science models. Here we will investigate two different models to further explore human resource management within the military manpower system. The first model describes optimal retirement behavior for an Army officer at any point in their career, and the second explores the attainability and sustainability of Army force profiles.

Simulation research has increasingly turned to Financial Engineering for applications as financial products have become increasingly complex, especially with the development of complex exotic derivatives and structured products that are traded over-the-counter. Many of these products cannot be priced by analytical models, and traders must rely on Monte Carlo simulation or other numerical methods both for pricing as well as for hedging. We investigate several exotic derivatives, collectively known Mountain Range options, using Monte Carlo simulation to price and develop gradient estimates to study the sensitivities to underlaying parameters, known as "the Greeks".

1.1 Military Manpower Models

Manpower models have been employed to study human resources in the military since the conception of military operations research (Abrams 1957). Manpower models were created to model the type of personnel system used by the military, but have not been restricted in application strictly to the military modeling domain. Military manpower models have primarily been employed where the personnel system can be modeled as a closed system with several distinct stages or ranks, and where a predominant feature is a bar to lateral entry. In this manner, manpower models must account for the need to grow the experienced personnel that are needed within the system. The academic workforce has also been routinely modeled as a multi-class manpower system, as entry is traditionally at the assistant professor level, followed by promotions to associate professor, and finally to professor.

Manpower models have also been used to analyze assigning workers to shifts and other problems that address matching personnel to jobs and job scheduling. Another early application of manpower models addressed problems within healthcare. Nurse scheduling was one of the early applications of linear programming to job scheduling models (Warner and Prawda 1972). The assignment of interns to hospitals has been another successful use of operations management (Roth and Peranson 1999). Recent work by Fry et al. (2006) extends manpower modeling to study annual staffing levels for firefighters.

One method of analyzing manpower models breaks the model into three components: recruiting, retention and compensation. The United States Army uses eight life cycle functions to describe human resource management: personnel structure, acquisition, distribution, development, deployment, compensation, sustainment, and transition (Cashbaugh et al. 2007).

These eight functions can be grouped into the three components that are included in much of the academic research, and we explore the inter-disciplinary research that has been completed in regards to each of these three areas. The issues of recruiting have been addressed in the marketing domain as well as in organizations and behavioral psychology. The area of retention has been addressed in a majority of manpower models, as the current workforce is the only source of experience within the military manpower system. Labor economics has studied compensation as well as the interface of compensation, recruiting, and retention.

Three types of mathematical models have been used in the preponderance of operations management, operations research, and management science manpower models. The early models were predominantly solved using dynamic programming.

Dynamic programming papers often cited manpower models as one of the sources of their application and motivation for the developed mathematical methods (Flynn 1975b). A second type of model is the transition-rate/Markovian manpower model. These probabilistic models are treated both in the operations research literature as well as the probabilistic literature. The last major type of model is linear and goal programming. In mathematical programming the manpower system is traditionally modeled as a network, and policy decisions are modeled using constraints on the multiple competing components of the objective function.

A series of texts provide an introduction to manpower models from different disciplines perspective. Wilson (1969) contains an early collection of papers and summarizes mathematical models and military manpower systems. Bartholomew et al. (1991) describes statistical methods for manpower planning; focusing on regression, renewal and Markov models. Vajda (1978) provides a mathematical modeling description of manpower planning. Edwards (1983) acknowledges the need for additional multi-disciplinary research in manpower planning, focusing first on technical economic models and then on behavioral models. The TIMS Studies in Management Science series produced a 1978 special issue on manpower planning and organizational design (Charnes et al. 1978). This special issue provides a good introduction to the management science approach applied to manpower planning.

1.1.1 Early Work in Military Manpower Models

There are a host of questions that arise when discussing the management decisions and policies that surround manpower planning. Military manpower planning covers the spectrum of human resource management policy in which employees enter at an initial state and flow though the system until they are eventually a loss to the system. The military manpower system is one where there is extremely limited lateral entry, and the system is constrained to hire entry-level workers, and retain experienced workers, and designed to provide the public good of defense.

Manpower models seek to answer questions of the number of personnel needed and what skills are required for operational capabilities. This optimization is traditionally viewed as constrained to defense budgets. Thirty years have passed since the United States transitioned from a last conscript Army to the All-Volunteer Force (AVF) that is currently waging the Global War on Terror (GWOT). The period of transition to the AVF was a rich time for operations research, management science and operations management research in personnel systems for the Department of Defense (DOD). The military roots of operations research, the immensity of defense budgets, complexity of military systems, and the consequences of failure all encourage research directed at solving questions concerning how to acquire, retain and compensate the AVF.

The transitions from peacetime overage to wartime shortages and conscription motivated early research even before the questions arose concerning the AVF. Dailey (1958) examined Navy re-enlistment rates finding that the re-enlistment rates for the Navy were amazingly consistent in total numbers, but varied considerably with regard to the percentage of sailors that would reenlist each year. They proposed a model with an inverse relationship between the number of sailors inducted in any year and the reenlistment percentage. A number of reasons for the inverse relationship were posed, but the model allowed a normalization of the annual percentages and provided clear comparison across time periods. The model set standards for evaluating the goodness of a retention year, and provided a predictive model of re-enlistment for retaining sailors past their first enlistment.

Fisher and Morton (1967a) and Fisher and Morton (1967b) investigated reenlistment rates for electronics personnel in the U.S. Navy. They developed a model of operational usefulness that could compare two different possible navies based on human resource decisions. Their model looked at the value of additional incentives over the period of service to keep experienced technical experts in whom a considerable training cost had been invested. They developed a Cobb-Douglas production function to provide an ordinal ranking of different possible configurations, with four distinct groups of sailors providing different amounts to the production based upon experience, but at a corresponding increasing cost. They found that their model was not sensitive to choice of a particular operational usefulness function, as their results obtain for any monotonic function of the individual units operational usefulness. They found that myopic policies can be suboptimal, as the system falls behind on experienced sailors if only enough sailor recruits are assessed to meet current year demand and keep the operational usefulness above the required minimum. They determined that enlistment and re-enlistment decisions must look forward to first term survival rates and the needs of the future force.

McCall and Wallace (1969) addressed the issue of re-enlistments in the Air Force. They study the retention of electronic specialists within the Air Force and addressed the need for monetary compensation for skills attained within the organization to avoid high retraining costs associated with low retention rates. They found that personnel who perceived that they had a marketable skill were increasingly sensitive to differences between military and civilian pay. They conducted an empirical analysis to develop a logit model relating the probability of re-enlisting and the difference between military and civilian pay. They also concluded that in the absence of a draft, the enlistment decision could be modeled similarly and indicate the importance of competitive initial pay levels for military personnel.

Flynn (1975a) created a dynamic programming model to study the issue of retaining productive units. The model was created to find an optimal retention policy in units where productivity of an element is determined by age with a predetermined retirement age for each element. His model used a linear production function similar to previous work and incorporated a continuously increasing function of compensation to balance the increased productivity with age. They find conditions under which it is optimal to bring in just enough units at the entry level to satisfy linear production constraints. Interestingly, the condition that allows this myopic solution is when costs are monotonically decreasing. However, they do find policies that are uniformly good steady state policies, and always minimize the average cost per period. They look for policies that converge in N or fewer steps, and have a low cost of staying at steady state. A good target state is defined such that it is minimum and can be reached from any state in the state space. In this manner, if a good steady state policy exists, any steady state solution with a good target is a good policy.

Jaquette and Nelson (1976) developed a mathematical model to determine steady stage wage rates and the force distribution by length of service. The model adjusts the accession and retention policies to optimize the steady state system. Their model is a nonlinear program formulation and compares both linear and Cobb-Douglas production functions. Gradient search methods are used to find solutions to their model, but they were not able to prove global optimality. They make some assumptions that have become standard in the literature: re-enlistment points at four year intervals, twenty year career and no lateral entry. They find that lower accessions, higher first term re-enlistment rates and slightly lower career reenlistment rates are optimal policy to maximize production subject to existing budget constraints.

The foundations of manpower modeling within operations research and management science are clearly laid out in a series of textbooks. Grinold and Marshall (1977) covers an introduction to manpower models and Vajda (1978) addresses the mathematics of deterministic manpower models. Bartholomew (1982) addresses sochastic social process, of which manpower modeling is a featured example, and Bartholomew et al. (1991) further extends statistical methods for manpower models.

In addition the manpower models presented here, another thread of research focuses on the personnel problem: assigning individuals to tasks. The personnel assignment problem can be modeled as a transportation problem, to assign personnel in each class to job categories. Tien and Kamiyama (1982) presents a survey of methods for this variant of manpower models. Warner and Prawda (1972) presents solutions to this variant of the assignment problem, studying the nurse scheduling problem, which has become a standard application for this class of problem in the literature. Green et al. (2001) provides a more recent examination of staffing issues in hospital emergency departments.

1.1.2 Mathematical Model Classes

Gass (1991) provides a primer on military manpower models in operations research. Although written over fifteen years ago, this work captures the types of models currently used in production by the military, and the Army in particular. He describes a modeling construct where individuals are aggregated by a set of attributes and grouped into classes. Each combination of a set of attributes at each time period represents a state. Each state will contain a number of individuals with common attributes. His primer describes Markov transition-rate models, network flow models, and goal-programming models.

Transition rate models require transition probabilities as an underlying assumption. Given an initial distribution, Markov models answer the question of what attributes the force will have at each phase in the planning period. When modeling with a transition rate model, each individual follows the same Markov process and individuals are independent. Markov manpower models in discrete time are described and solved in Feichtinger and Mehlmann (1976). They find existence results when searching for an asymptotic solution with fixed recruitment, transition probabilities and initial conditions. Mehlmann (1977) extends their results to continuous Markovian manpower models. Davies (1982) considers a model with fixed promotions controlled by management. In this model the number of employees in each stage can be modeled as independent binomial random variables. This independence assumption facilitates calculation of one and two step transition probabilities.

Manpower models can also be modeled as a network flow. Network models ensure conservation of flow and have the advantage of integer solutions. The models are especially powerful when modeling systems with a limited number of characteristics at each time period.

Goal programming is a branch of mathematical programming that uses policy constraints as the right hand side of the programming constraints. The objective function is created as weighted deviations from the policy goals. The goal programming solutions will always be a compromise, as not all the policy objectives can be feasibly met. Price and Piskor (1972) provide a goal programming formation for the officer manpower system in the Canadian Forces. Price (1978) showed how to re-formulate goal programming problems as capacitated network flow problems. Gass et al. (1988) describes a goal programming formulation of the Army Manpower Long-range Planning System.

Mathematical programming can expand upon network structures to add multi-year constraints and goals to the modeling formulation. The different years have constraints on the number of separations and promotions as well as grade, total strength and operating strength targets. The objective function penalizes deviations from the goals, but relies on the underlying network structure to simplify the optimization. Current Army enlisted manpower network models have a 96-month planning horizon, and aggregate on months of service, grade, term of service, and gender (Cashbaugh et al. 2007).

Dynamic programming models are very popular in the operations research literature as well as the economics literature, reflecting the interdependence of the two disciplines. Many of the early manpower models were dynamic programming models, and manpower modeling applications provided an interesting area of application for expanding research in dynamic programming. Some of the early works described earlier (Fisher and Morton 1967b,a, Flynn 1975b,a, 1979), used dynamic programming models. The objective of deterministic dynamic programming is to develop a steady-state policy, within the defined state space, decision space, cost functions and transition functions. Linear and Cobb-Douglas production functions have both been prevalent in the literature.

Gotz and McCall (1983) used Markov decision processes to analyze Air Force officer retention. In this paper they created a model that could be used to study the policy changes in the two main drivers of officer retention: promotion and retirement policies. Officers have a different promotion system than enlisted soldiers and do not reenlist on a four-year horizon as is assumed in many of the previous models. They do not address nonpecuniary factors, but do find that Air Force officers make retirement decisions in a optimal sequential fashion. Their model shows that retirement pay is the most important factor for officer retention between the 10 and 20 year marks.

1.1.3 Research Opportunities

Manpower modeling is an important area of interdisciplinary research. We have detailed some of the interesting articles in Operations Management, Economics, and Management and Organizations. Many streams of research still await integration into decision-making models, for both military manpower planners and strategic business human resource planners.

When examining the literature on manpower modeling and retention, several clear patters are present. First, manpower modeling provides a rich field of application for mathematical modeling. The early work in dynamic programming looked to the manpower problem for clear links to application, and generated papers both in the application areas as well as operations research and applied mathematical journals.

The second observation is that changes in public policy require these previous results to be revisited. Much of the early work was motivated by a need to move away from conscription. The Reagan years buildups and policy changes necessitated better models to justify increasing budgets. The downsizing of the Clinton era provided another rich set of data to analyze the military members personnel discount rate and the effectiveness of force-shaping tools. Within the policy arena, Rostker (2006) provides a historical account of the development of the AVF and includes descriptions of the studies conducted to support the policy decision surrounding its creation. The newly published *Handbook of Defense Economics* has an excellent chapter on the "New Economics of Manpower in the Post-Cold War Era" (Asch et al. 2007, Hartley and Sandler 2007).

The effects of several policy decisions have not yet been fully analyzed. The full effect of the "High Three" retirement system with increasing compensation though 40 years has not yet been explored. Another question that will need to be addressed is the structure and compensation of the officer corps. The behavior of the force is almost surely in flux. Within the increased operational tempo, there is a perceived decrease in competition for promotion, and military-civilian compensation difference. We address these questions in the second and third chapter.

In light of these timely and critically important questions, the field of manpower modeling offers many opportunities for research. There already exists a rich body of research, but there are many areas that need to be expanded and analyzed. The field offers many practical problems to which analytic and empirical modeling will produce interesting and policy influencing research. These research challenges require the unique balance of qualitative and quantitative analysis that is aligned within the inherently multi-disciplinary purview of Operations Management.

1.2 Simulation and Financial Engineering

The theory of asset pricing and mathematical finance has come a long way since Louis Bachelier's *Theory of Speculation* in 1900. Many of the advances have come with the advances in computing, but the seminal works on options pricing (Merton 1973, Black and Scholes 1973) and martingale methods (Harrison and Kreps 1979), highlight the key developments of mathematical finance, replication and martingale pricing. Replication and martingale pricing allow for trading strategies to perfectly hedge a derivative in a complete market.

Vanilla options have moved from being traded over the counter, to having established markets and a well defined theory of pricing. Options themselves are now used along with stocks to hedge more complex derivatives and to allow risk managers to diversify their portfolios and to balance exposure to risk. Mathematical methods using numerical solutions to partial differential equations and Monte Carlo simulation have emerged as the most common methods used for derivative pricing. Developing the theory to price new assets is a fundamental goal of mathematical finance research.

1.2.1 Gradient Estimation

In addition to developing theory to price complex assets, financial engineering is equally interested in risk mitigation though replication and hedging. However, in order to hedge positions, the sensitivities to market and model parameters must be computed. Prices of vanilla puts and calls can be observed in the market, but market sensitives must always be calculated. The need to calculate these sensitivities has increased the use of simulation in financial engineering and provides a rich research area for gradient estimation and variance reduction.

Fu and Hu (1995) first applied gradient estimation techniques to option pricing focusing on IPA estimation for both European and American options. Broadie and Glasserman (1996) applied both IPA and LR/SF methods to European and Asian options. Within financial engineering, Glasserman (2004) has become the standard reference. For broader application, Fu (2006) provides an detailed survey of gradient estimation within stochastic simulation and explains how the direct techniques are much more efficient than “brute-force” finite difference methods. Fu (2008) provides a more gentle introduction with a focus on application to option pricing.

Several different techniques have been employed for gradient estimation in financial applications. The first class are indirect methods that rely on finite difference methods, where the value of a parameter is perturbed and re-simulated. Forward and central finite difference methods are quite easy to implement but will be biased, in addition to requiring additional simulations.

In addition to indirect methods of gradient estimation, direct methods can be used that tend to provide unbiased estimators and are computationally more efficient. These methods are based upon calculating the derivative of the simulated random variable, and vary depending on the dependency upon the parameter of interest. The resulting methods can be grouped into sample or path-wise methods and measure or distributional methods. Perturbation Analysis (PA) encompasses path-wise methods and includes infinitesimal perturbation analysis (IPA) and smoothed perturbation analysis (SPA) and other techniques that allow transformation using the chain rule into two components: a sample path derivative and a derivative of the random variable. SPA uses additional information about the stochastic process through conditioning to increase the applicability of PA techniques. When the parameters of interest appear in the distribution, likelihood ratio/score function (LR/SF) or weak derivative methods are effective as these two methods focus on the distributional parameters. In these two methods, the derivative of a measure must be calculated, as opposed to the derivative of the random variable as in the PA methods.

1.3 Three Essays: Simulation and Optimization

In the following three chapters, we address two very different application areas, and three different modeling techniques: Markov decision processes, simulation, and mathematical programming. The first essay provides a new model of the retirement decision made by an Army officer, obtaining optimal policies for the current compensation and retirement system, as well as providing a framework within which to evaluate proposed changes. The second essays again looks at the military manpower system, but instead of evaluating the system from the perspective of the officer, the model provides opportunities to optimally control the system, creating the required pool of manpower. This facilitates exploration of several potential force profiles, by controlling accessions, promotions, transfers and reductions, in addition to easing two of the fundamental planning assumptions: bar to lateral entry and the up-or-out career pattern. The third essay studies a group of over-the-counter exotic options, Mountain Range options. Relatively little has been written about this family of exotic options, providing an opportunity for new work studying the payoffs and Greeks for these options.

Chapter 2

Optimal Army Officer Retirement

2.1 Introduction

A large incentive to continued military service has historically been the generous retirement opportunities. This delayed compensation for service, even considering the potential service obligation for officers on the retired rolls, provides a large monetary incentive for officers to serve until retirement. However, the cost of military retirement makes changes to the retirement system a continued point of debate. In addition to cost, the fairness of a system that cliff vests at twenty years has been debated by Congress on numerous occassions.

Here we explore the changes to optimal retirement behavior for Army officers. The Army is our largest service, currently authorized 532,400 soldiers, around 70,000 of which are officers. The Army must train its own officers, and has recently enacted the first bonus for Army Officers to attempt to increase retention past the critical ten years of service point to ensure the required strength of field grade officers.

Army Commissioned Officers are divided into three groupings: company grade, field grade and General officers. Company grade officers specialize in direct leadership at the small unit level and comprise the ranks of 2nd Lieutenant (2LT), 1st Lieutenant (1LT) and Captain (CPT). The field grade ranks specialize in leadership of larger organizations and leading and coordinating staff functions. The field grade ranks in the Army are Major (MAJ), Lieutenant Colonel (LTC) and Colonel (COL). General officers are recognized both by their title and by the number of stars they wear on their collar. General officers are the senior executives of the Army and comprise the ranks of Brigadier General (one star), Major General (two star), Lieutenant General (three star) and General (four star).

In a zero-sum budget game, resources that are shifted from those who serve a full career and retire to those who leave prior to the current vesting point will decrease compensation to those serving full careers. Such changes will not necessarily change the decision to serve full careers, but changes to retirement and compensation can be delicate issues. Changes to retirement and compensation must be tailored to ensure that the changes induce the desired changes in behavior.

In this chapter we will address the optimal retirement policies for Army officers, incorporating the current retirement system, pay tables, and Army promotion opportunities. The first two factors are common to each of the services, while the promotion opportunities, although regulated by Congress, differ between the services. Although the model presented here is tailored to the Army, the methodology can be easily adapted for use by the other services.

2.1.1 Military Retirement

Many characteristics of the global environment and the military compensation system have changed significantly in the last 25 years, yet military retirement as a form of deferred compensation remains a powerful incentive for soldiers to stay in the

Army. The current retirement system provides an immediate annuity after 20 years of service, based upon the average of the soldier's highest three years of military compensation, a system known as the "High Three". Soldiers who entered service prior to 1986 are entitled to an annuity equal to 50% of their High Three while Soldiers who entered service after 1986 are offered the choice between the existing system or a lump sum payment at 15 years of service followed by 40% retirement annuity. The annuities are adjusted for inflation, and increase by a percentage for each additional year of service.

In addition to the changing retirement system, there are several other changes which alter military compensation directly and correspondingly the retirement annuity. The current military compensation system is based upon a pay table, with increasing pay for increasing rank and years of service. The pay tables have been significantly reworked, attempting to close a perceived pay gap between civilian pay and military pay. Also, longevity raises have been included for some ranks out to 40 years of service. With the additional of a new pay raise for the rank of COL at 30 years, an officer would have to serve 33 years to fully realize the last pay raise in their retirement annuity. The last pay raise for LTC is at 22 years, and correspondingly is fully realized in retirement pay at 25 years. Similarly, the last pay raise for MAJ is at 18 years, which is fully realized at 21 years of service.

The promotion opportunities have also changed for Army officers. Under the Defense Officer Personnel Management Act (DOPMA) of 1980, Congress established guidelines for promotion to MAJ, LTC and COL at 80%, 70%, and 50%, respectively (Rostker et al. 1993). At present, there is very little probability that an officer will not be allowed to serve as long as legally permissible. There is currently a shortage of field grade officers, especially acute in the rank of MAJ, the rank of most officers with ten to fifteen years of service. As a result, recent promotions rates for CPT, MAJ and LTC are at near 100%, with the first promotion that is truly competitive is a 50% promotion rate to COL (Henning 2006). The promotion opportunity for most officers to COL is after 21 years of service. Current promotion rates are unlikely to continue after the current shortage year groups created by the 1990s military downsizing pass though the promotion windows. Decreasing promotion rates will, however, only decrease the expected value of a military career.

Although a valuable part of the entire military compensation package, military retirement is very expensive, and the current policy of cliff vesting at twenty years is arguably inconsistent with current labor policies in the civilian sector. The Defense Advisory Committee on Military Compensation highlighted three reasons to adjust the retirement system: inequity, inefficiency and inflexibility (DACMC 2006). The perceived inequities are that most service members are never vested, a higher percentage of officers become vested than enlisted soldiers, and the percentage of enlisted members of each service that become vested vary significantly between the services. The current system is considered inefficient because annual accrual for retirement is 27% of basic pay, even though most service members are never vested. The retirement system is considered inflexible because of the strong attraction for all that serve past ten years to remain to retirement, and a service culture that is adverse to forcing service members to leave between ten and twenty years of service.

However, before addressing additional proposed changes to the current system, we have to address the behavior currently induced by Army compensation. Assuming that officers make their decisions based on optimizing their financial status, the current system will incentivize behavior based upon the retirement system, promotion opportunities, and pay tables. The pay tables are updated annually and provide the base compensation by grade and years of service which is common to all the services. In addition to monthly basic pay, Regular Military Compensation (RMC) includes two non-taxed allowances: Basic Allowance for Housing (BAH) and Basic Allowance for Subsistence (BAS).

2.2 Modeling Officer Decisions

The career decisions of an officer may be influenced by many factors, but clearly economics will be a major consideration. With all the pressures of serving as an officer, there must be some compensation that outweighs the other career options at different stages of the officer's career. The officer does not know how many times they will be promoted when they are commissioned, but the uncertainty begins to be resolved over time. In our model, the officer knows their current rank, years in grade and years of service, and is then assumed to choose optimally between the compensation of continuing to serve versus retiring from the military and starting a second career. With the transition to the All Volunteer Force (AVF), the compensation offered officers is expected to be competitive in the labor market, both during time of peace and conflict.

Here we present a normative model of the retirement decision by an individual

Army officer. Our model is built upon the theory of Markov Decision Processes (MDP), using the framework of an optimal stopping problem to describe the optimal policy regarding when an officer should retire from the Army.

2.2.1 Retirement Modeling

Modeling retirement decisions as an MDP has been employed both in research focusing on military retirement, and on retirement problems in general. Most of the literature in this field has created econometric discrete choice models and used panel data to estimate model parameters. An exception is Gotz and McCall (1983), who developed a normative model of Air Force officer retention and retirement decisions called the Dynamic Retention Model (DRM). Their model was the first to incorporate rank structure into the decision making framework for military retirement modeling. Gotz and McCall (1983) assumed that the Air Force officer knew their promotion probability distribution and that officers were homogenous in that they did not incorporate any non-pecuniary aspects of military compensation when they estimated optimal retirement points.

The competing model paradigm, explained in Warner and Goldberg (1984), focused on the modeling of enlisted personnel. The Annualized Cost of Leaving model (ACOL) presented by Warner and Goldberg (1984) is used in their examination of the effects of sea-duty on Navy enlisted personnel. They present a model with a single military state and examine the compensation level that would make the sailor indifferent to leaving the Navy. They estimate taste factors, which are assumed to be jointly normally distributed, to explain an individual's relative preference for military and civilian life. Black et al. (1990) adapt the model to study the quit decisions of federal employees.

Gotz and McCall (1984) was one of the first papers to conduct structural estimation of a MDP or discrete choice model. Gotz and McCall (1984) estimated differences between "rated" officers, i.e. pilots and navigators, and non-rated officers from various commissioning sources by incorporating a taste random variable, assumed to be drawn from a extreme value distribution, to characterize each set of officers, estimating the parameters of the DRM by maximum likelihood using Air Force officer data from 1973 - 1977. Daula and Moffitt (1995) created a dynamic choice model with discrete choice variables based upon the work in Gotz and McCall (1984) to model enlisted soldier retention decisions. Their stochastic dynamic programming model allows them to explore the differences in military and civilian pay levels and the value of retirement for infantry soldiers, a combat specialty, who enlisted between 1974 and 1984. They found that the perceived difference in military and civilian pay levels had a significant effect on re-enlistment rates.

Arguden (1988) addressed questions concerning changes in the military retirement system and the unintended effects on retention and work force mix with empirical results and forecasts created using both the ACOL and the DRM calibrated to Air Force retention rates from 1971-1981. The model predicts that the average service per enlistment will fall under the reduced retirement program from 6.9 years to 6.5 years, resulting in a 22 percent decrease in airmen in the 8-20 year cohorts. They suggest that the overall attractiveness of the military with reduced compensation would alter the quality of personnel that are willing to stay. Having the benefit of hindsight, this paper is prophetic of the issues the services would face in the years after implementing the policy of a reduced retirement. Congress has since reversed the policy after observing the retention shortfalls of the 1990s. It was changed as a part of the National Defense Authorization Act of 2000 to re-instate the provisions of the 1980 retirement policy (Asch et al. 2002).

Asch and Warner (2001) develop an individual decision model for large hierarchical organizations such as the military, focusing on several unique military compensation features, including a much flatter compensation structure than most large hierarchical organizations, an unusual retirement system, and an up-or-out promotion system. Their individual decision model strives to explain why the military deviates from standard practice in these three areas. This individual decision model provides insight into both the level of effort and retention decisions of individuals. Their results indicate that the current combination of an up-or-out promotion system and generous retirement make the compensation package incentive compatible and internally consistent without an extreme skew in compensation between senior and junior personnel.

Related research has addressed social security, medicare and civilian retirement decisions using dynamic programming. Rust (1987) presented a dynamic programming model to describe the decision to partially or fully retire and expanded the work to explore the effects of Social Security and Medicare on retirement decisions in Rust and Phelan (1997). Where the previous model focused on a population for whom social security was their only pension, Berkovec and Stern (1991) present another discrete choice model allowing the workers to chose to continue working, change jobs, partially retire or fully retire, focusing on wage and pension benefits, without including Social Security. Burkhauser et al. (2004) used a similar model to explore the timing of application for Social Security Disability Insurance.

Our model follows the approach of Gotz and McCall (1983) in that we present a normative model of officer retention and retirement decisions. However, changes in the military system in the ensuing 25 years since their paper, including the "High Three" retirement system, increasingly competitive 40 year pay tables for military base pay, the elimination of Regular versus Reserve Active Duty commissions, and changing career timelines and promotion opportunities, lead to a very different underlying Markov chain to control the rank state in our model. Furthermore, our model allows us to analyze the sensitivity of current officers' decisions to personal discount rates, promotion opportunities and taste for military service.

2.2.2 Modeling Uncertainty

Optimal stopping problems are characterized by a decision maker who observes the current state of the system and decides whether to stop or continue. A key aspect of most optimal stopping problems is the Markovian nature of the decision to stop based upon the value of the current state. The system has either a reward or cost for each additional period, and once the decision has been made to continue with the system, the reward for the next period is received, and the system transitions according to a probability distribution to the next state. At this next state, the decision must be revisited.

This optimal stopping problem will be modeled as an MDP with finite state space and fixed time horizon. The fixed time horizon is a reasonable modeling assumption, with current limits on service career lengths. The model will be analyzed for optimal stopping policies for each of the grades, looking at sensitivities to civilian pay expectation as well as sensitivities to personal discount rate. The MDP used to model retirement behavior can be effectively solved by dynamical programming techniques. Many references exist explaining modeling and solutions techniques for this class of problem (Bertsekas 1995, Miranda and Fackler 2002, Adda and Cooper 2003, Puterman 2005).

A policy is a decision rule that prescribes the action that a decision maker should make upon reaching any particular state in a certain period. The policy may be a deterministic or a randomized policy and may be time dependent or stationary. An optimal policy is a policy that results in the maximum (minimum) value of the stated objective function. The theory for solution of discrete state MDP is based upon the principle of optimality due to Bellman (Bertsekas 1995).

Given the MDP state space X and the action space A, define the state $x_t \in X$, and action $a_t \in A_t(x), t = 1, 2, \ldots T$, with β as the personal discount factor, Π the space of all policies, $A(x) \subset A$, denoting the feasible action set for a given state x, and $R(x, a)$ the reward function for a given state x and action a. Solving this optimal stopping program amounts to finding a policy $\pi = \{\pi_t, t = 1, 2, \ldots\} \in \Pi$, where $\pi_t : X \to A$, that maximizes the expected total discounted reward

$$V^{\pi}(x) = E\left[\sum_{k=1}^{T} \beta^{k-1} R(x_k, \pi_k(x_k))|x_1 = x\right].$$

We define the corresponding reward-to-go for a policy π at time t as

$$V_t^{\pi}(x) = E\left[\sum_{k=t}^{T} \beta^{k-t} R(x_k, \pi_k(x_k))|x_t = x\right].$$

The optimality equations for the optimal stopping problem are

$$V^*(x) = \max_{\pi \in \Pi} V^{\pi}(x) \tag{2.1}$$

and

$$V_t^*(x) = \max_{\pi \in \Pi} V_t^{\pi}(x) \tag{2.2}$$

where V^* is the optimal value function. The solution to (2.1) or (2.2) is normally referred to as stochastic dynamic programming. Dynamic programming is equivalent to the solution method of backward induction and is founded in the principle of optimality, which provides this traditional solution method for a discrete MDP with a finite horizon. The solution of the MDP will consist of an optimal value function V^* and the associated optimal policy, π^*, that provides the prescription for the decision to be made in each state.

2.2.3 Retirement Model

Building upon this mathematical structure, we next describe the MDP used to model Army officer retirement decisions. We explain the states of the MDP, the rewards or cash flows at each state of the Markov chain, and the optimality conditions or value function used to model this decision process.

2.2.3.1 States

The model will address officers in the ranks of 2LT though COL. Each state in the model will represent a combination of grade, years-of-service and time-in-grade. Each of the ranks is modeled by several states to address differences in compensation, as well as promotion opportunity. The rank state space in this model has 29 different rank and years in grade combinations, each indexed by years of service. The model in Gotz and McCall (1983) modeled the ranks of CPT though COL with 13 states to address promotion timing, and a single state to address involuntary separations. The state space was enlarged to model Air Force regular and reserve commissions, resulting in 27 total states. However, when the High Three retirement structure in the new objective function is combined with the states described in Gotz and McCall (1983), the model is no longer Markovian, motivating the creation of a new model.

Currently, all Army officers receive the same commission, and there are three opportunities for early promotions, and as such states for COL will have 18-40 years of service and 0-3 years in grade. Although a COL at 30 years who was promoted at year 22 will have 8 years in grade, there is no compensation difference after three years in grade for retirement, and so the additional states do not add fidelity to the model unless Brigadier General promotions are modeled. LTCs will possibly have 14-40 years of service and 0-6 years in grade. MAJs will have 9-40 years in service and 0-6 years in grade. CPTs can have 3-40 years of service, and 0-8 years in grade. 1LTs will have 1-40 years in service and 0-3 years in grade. 2LTs will have 0-40 years in service and 0-3 years in grade.

Table 2.1 displays the states in the network model. These states are numbered in sequence from 1 to 29, with the traditional career path shown in the fourth column. Most Army officers are promoted in accordance with this career timeline, although there are three opportunities for early promotions, to MAJ, LTC and COL. With the three opportunities for early promotions, there are eight career paths to COL that reflect the combinations of early and due-course promotions. Due-course is the term used by the Army for the standard career model. This timeline is reflected the fourth column of Table 2.1.

The years-of-service are captured as the time variable in the model. The years-in-grade are critical to identify the promotion probabilities, as well as calculating the rank held for the final three years of service. The 29 distinct pay/rank combination states, each indexed by years of service results in 1160 states.

Figure 2.1 represents the transition diagram for the Markov chain modeling an officer's career. The transition probabilities are slightly stylized, but reflect institutional beliefs within the Army (Haight and Lewis 2008). Currently promotions to 1LT are approved locally, resulting in the near 100% selection rates. However, those

State	Rank	Years in Grade	Army Yeargroup Timeline
1	2LT	1	1
2	2LT	2	2
3	2LT	3+	~
4	1LT	1	3
5	1LT	2	4
6	1LT	3+	~
7	CPT	1	5
8	CPT	2	6
9	CPT	3	7
10	CPT	4	8
11	CPT	5	9
12	CPT	6	10
13	CPT	7+	~
14	MAJ	1	11
15	MAJ	2	12
16	MAJ	3	13
17	MAJ	4	14
18	MAJ	5	15
19	MAJ	6	16
20	MAJ	7+	~
21	LTC	1	17
22	LTC	2	18
23	LTC	3	19
24	LTC	4	20
25	LTC	5	21
26	LTC	6+	~
27	COL	1	22
28	COL	2	23
29	COL	3+	24

Table 2.1: Rank States in MC

that the Commander deems not ready for promotion, are given a 50% probability of qualifying each year. The promotions rates to CPT, MAJ, and LTC are adjusted to 99% to reflect the current perception that all officers are being promoted. The promotion to COL is modeled at 50%, the current legal floor under DOPMA (Rostker et al. 1993). The below the zone or early promotions are targeted at 5%-10% of

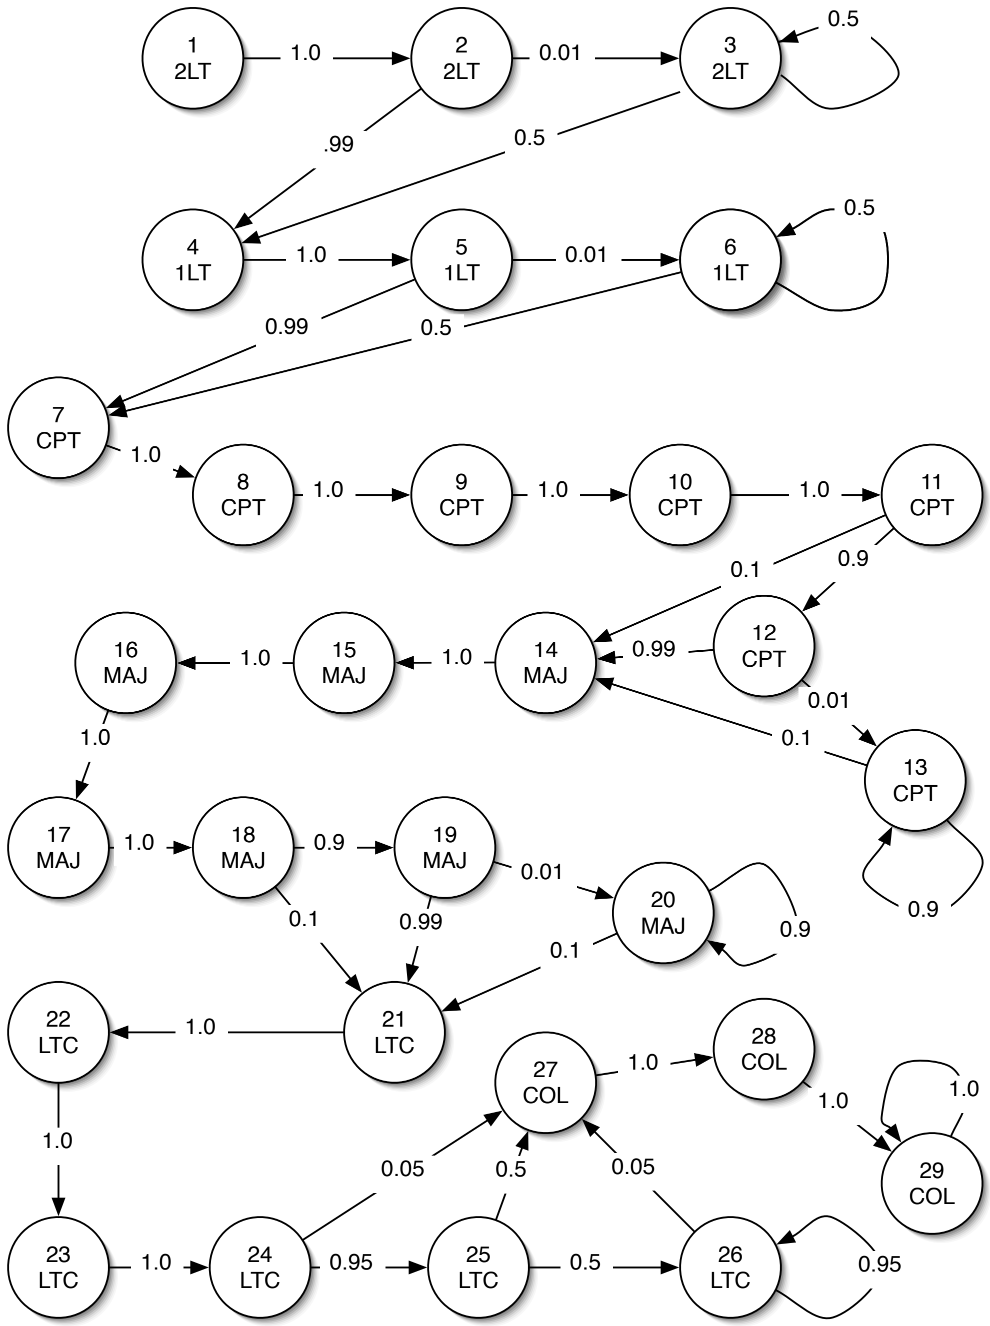

Figure 2.1: Markov Chain for Rank State Space

each year group. As the Army is currently short MAJ and LTC, and over-strength COL, we modeled the first two early promotions, from state 11 to 14 and 18 to 21, at 10% and modeled the early promotion from 24 to 27 at 5%. Sensitivity to transition probabilities was explored in numerical experiments.

2.2.3.2 Transitions

Each year is modeled as a period, and the officer progresses though the Markov chain described in Figure 2.1. The officer knows his/her current state, which includes information on rank, years at current rank, as well as years-in-service. Let $P_{x,y,t}$ denote the probability of transitioning from the current state x in period t to state y in the next period, $t+1$. The transition matrix $P \equiv [P_{x,y,t}]$ reflects the Markov chain described by Figure 2.1 with promotion opportunities greater than the floors established by Congress.

The transition matrix does not allow for demotions, and as such has all zero elements below the diagonal. States are naturally grouped by rank with the transition matrix describing increasing seniority within the current rank, as well as promotions to the next rank. The extra states within each rank are necessary to capture the information necessary for retirement and promotion calculations. Pay is not dependent on the path taken to any state, only on the current rank and time in service, information which is contained in each rank state and the model's time variable.

The "High Three" retirement system is also path independent, because pay is not dependent on years in each state, but only on the rank held and years in service. Since the first state of each rank can only be accessed from the immediately junior rank and the time spent at each rank ≥ 2, each state contains the necessary information to accurately calculate retirement pay. Therefore, the first state of each rank is immediately preceded by at least two states of the previous rank, which can be used for calculation of retirement pay, independent of the path the officer actually took though the Markov chain. Under a no demotions assumptions, the last three years of pay are the "High Three", which we define as

$$\hat{m}_{x,t} = \frac{m_{x-2,t-3} + m_{x-1,t-2} + m_{x,t-1}}{3}.$$

2.2.3.3 Reward

An officer will make a decision each period to continue on active duty and receive their next period pay and benefits or to leave active service. The value of staying in the military for an additional period for an officer in state x is recursively defined by

$$V_t(x) = m_{x,t} + c_{x,t} + \beta \sum_{y=1}^{29} P_{x,y,t} V_{t+1}(y),$$

where $m_{x,t}$ is base military compensation in state x in period t, and $c_{x,t}$ is additional compensation, housing and subsistence, only received while on active duty, also in state x at period t. Discounting the next period value function at the personal discount rate β, and summing over all possible transitions from the current state x to possible future states y captures the period t expected value of compensation of staying in the military for an additional period.

Although single period compensation is strictly increasing with time in service and grade, the expected total earnings in any given state incorporates future promotions and uncertainty that is resolved over time and as a result is only monotonic for COL.

The value of leaving military service is a combination of the pay that will be received from working in the civilian sector and the value of any military retirement that has been earned up to the epoch that the soldier leaves active duty. The percentage of $\hat{m}$ that will be received in retirement compensation from the military is represented by

$$r_t = 1\{t \geq 20\}(.5 + 0.025(t - 20))$$

where $1\{.\}$ is the set indicator function. This equates to a retirement consisting of 50% of base pay at 20 years, increasing an additional 2.5% for each additional year of service. This cash flow is modeled by

$$r_t \hat{m}_{x,t} \sum_{k=t}^{\tau-\eta} \beta^{k-t} + \sum_{k=t}^{T-1} \beta^{k-t} m_{x,t}(1 + \xi_{k-t}) \tag{2.3}$$

where ξ_n represents a civilian raise in year n of working in the civilian sector and $m_{x,t}$ represents the starting civilian pay in retirement, assumed to be equal to final pay in the Army. The expected number of years that the officer will receive retirement benefits is represented by $\tau - \eta$ in (2.3), where life expectancy is represented by

τ and the age at commissioning is represented by η. Past mandatory retirement at 40 years, months of civilian earnings available will be assumed to be equivalent across all states, but the distribution of income based on years in the civilian sector may very significantly. Time T represents full military retirement at forty years of service, i.e. $r_t = 1.0$, and $T - t$ captures the number of periods that the officer has remaining to work in the civilian sector.

The starting assumption for level of civilian compensation is that the officer is being fairly compensated by the Army. As such, the officer is expected to be able to command a similar salary in the civilian sector. Civilian pay is modeled as a monotonically increasing function, beginning at the pay level of the officer when they exit military service. The sensitivity to this initial level of compensation will be explored in numerical experiments.

2.2.3.4 Value Function

Each officer will make a decision each period to continue to serve or retire and seek civilian employment. The assumption is made that the officer will continue to pursue full-time employment for a 40 year work career. The optimal stopping policy will be found by evaluating the value of all possible decisions at each state and period and choosing optimally at each state. In order to exploit the principle of optimality from dynamic programming, the value of each state at final time, $T = 40$, must be calculated. This time reflects mandatory retirement, and assuming $\tau > \eta + T$, no decision can be made but to accept the cash flow associated with the final state

$$V_T(x) = r_T \hat{m}_{x,T} \sum_{k=T}^{\tau-\eta} \beta^{k-T} = \hat{m}_{x,T} \frac{1-\beta^{\tau-(\eta+T)+1}}{1-\beta}.$$

Creating the $T-1$ value function is a choice of serving one additional period, or retiring at 39 years of service and receiving one year of civilian pay and retirement compensation:

$$\begin{aligned} V_{T-1}(x) = \max \Bigg\{ & m_{x,T-1} + c_{x,T-1} + \beta \sum_{y=1}^{29} P_{x,y,T-1} V_T(y); \\ & r_{T-1} \hat{m}_{x,T-1} \sum_{k=T-1}^{\tau-\eta} \beta^{k-(T-1)} + m_{x,T-1} \Bigg\}. \end{aligned}$$

Solving for the value of each period by backwards induction, the value function for each period is represented as

$$\begin{aligned} V_t(x) = \max \Bigg\{ & m_{x,t} + c_{x,t} + \beta \sum_{y=1}^{29} P_{x,y,t} V_{t+1}(y); \\ & r_t \hat{m}_{x,t} \sum_{k=t}^{\tau-\eta} \beta^{k-t} + \sum_{k=t}^{T-1} \beta^{k-t} m_{x,t}(1+\xi_{k-t}) \Bigg\}. \end{aligned} \tag{2.4}$$

2.3 Numerical Results

We solved the retirement model using a dynamic programming formulation, allowing for some sensitivity analyses, in addition to obtaining the expected value of continuing or retiring in each state at each time period. The value function is the optimal solution for a given state in a given time period. This is computed using a backward induction algorithm, so it is dependent on making optimal decisions up to that point. The civilian opportunity at 20 years incorporates vesting in military retirement and is the first epoch where the civilian opportunities are greater than military compensation, providing an optimal point to retire for many states.

Represent the action of continuing to serve in the Army by s and the decision to resign or retire as q, so that $a \in \{s, q\}$, and define

$$\begin{aligned} Q_t^*(x,s) &= m_{x,t} + c_{x,t} + \beta \sum_{y=1}^{29} P_{x,y,t} V_{t+1}(y), \\ Q_t^*(x,q) &= r_t \hat{m}_{x,t} \sum_{k=t}^{\tau-\eta} \beta^{k-t} + \sum_{k=t}^{T-1} \beta^{k-t} m_{x,t}(1+\xi_{k-t}), \end{aligned} \tag{2.5}$$

then

$$V_t^*(x) = \max_{a \in \{s,q\}} Q_t^*(x,a).$$

The addition of a taste parameter α to allow for differing valuations of civilian wages to (2.5) results in the following modification for the optimal reward-to-go function defined by (2.2):

$$V_t^*(x) = \max \Big\{ Q_t^*(x,s); \alpha Q_t^*(x,q) \Big\}.$$

A taste-neutral officer will have a taste parameter $\alpha = 1$ and will value compensation gained from the military or civilian sector equivalently. When $\alpha > 1$, the officer would prefer to be working in the civilian sector, and correspondingly $\alpha < 1$

represents a preference for the military lifestyle. The assumed personal discount rate is 13.6%, i.e. $\beta = 0.88$, with starting civilian pay equal to last military pay. In our numerical experiments, we assume the officer is married and value basic allowance for housing at the "with dependents" rate.

The taste parameter α allows us to explore non-cash compensation and intangible benefits. An officer may have aversion to change, believing that switching careers would entail significant effort. In addition, each officer has some value they place on their total compensation, including tax advantages and prestige benefits. In this way, decreasing the value of α describes an individual that would require an incrementally higher civilian wage in retirement to feel equivalently compensated for the effort of changing careers and the value of non-monetary compensation and prestige.

Figure 2.2 provides a graphical representation of the optimal decision as a function of rank/years in grade and number of years in service, where the retirement boundary is given by $Q_t^*(x,q) = Q_t^*(x,s)$. As seen in Figure 2.2, most states first cross the retirement boundary at 20 years. In addition, the states in the upper left corner of the graph are practically unattainable since any officer with four years or less of service would still be obligated to the Army under an initial service obligation.

Without discounting, the continuous approximation of $Q_t^*(x,q)$ is strictly decreasing in time and increasing in state. However, the Markov chain includes semi-absorbing states, ones for which the probability of leaving is very low, which keeps $Q_t^*(x,s)$ from being increasing in state. Neglecting these states, $Q_t^*(x,s)$ is also increasing in state and decreasing in time. With discounting, $Q_t^*(x,s)$ and $Q_t^*(x,q)$

Figure 2.2: Retirement Boundary $Q_t^*(x, q) = Q_t^*(x, s)$

are both concave functions with a single optimal for each state over time. In this model, most of the states are transient, and therefore the change over time and state in combination is most critical to the behavior described here. As shown in (2.4), the value of continuing in the Army, $Q_t^*(x, s)$, incorporates the future value function, and incorporates optimal choice in each future state. The comparison is made to civilian opportunity, $Q_t^*(x, q)$, which in this model does not depend on future decisions.

An interesting point illustrated by Figure 2.2 is that the LTC states, excluding 25, display an optimal decision to retire at 20 years. State 25 reflects a 50%

promotion opportunity to COL, and with the increase in military pay and civilian opportunity, the optimal decision is to stay for an additional year. Interestingly, the difference in pay with discounting at 12% only has a one period effect, as an officer would optimally retire in state 24, 26, and 27. Although state 27 is not reachable until 22 years for a due-course officer, it is interesting to note that each COL state is a retirement state.

Another point to notice from Figure 2.2 is that the optimal decision for year 19 is to stay for all states, indicating that the optimal decision for all the 19th year states is to serve for one additional year. This relationship holds for all attainable rank states with $t < 20$. We conducted extensive sensitivity analysis, and found that sensitivity to personal discount rate and to civilian pay opportunity (or taste for military lifestyle) were the most interesting and will be detailed in the following sections.

2.3.1 Discount Rate

In the numerical experiments, initially the personal discount rate was set at 13.6%, but as research has shown military officers' discount rates could be as high at 20% Warner and Pleeter (2001). The sensitivity to personal discount rate for LTC at 20 years is depicted in Figure 2.3. The cost to leave is defined as $Q_t^*(x,s) - Q_t^*(x,q)$, so that a positive cost to leave is associate with the optimal policy to stay in the Army. As shown in the graph, the retirement boundary for LTC at twenty years is between an 11% and 12% personal discount rate. Officers having lower discount rates should optimally choose to continue serving at 20 years, whereas those with higher discount rates would choose to retire. As expected, state 25 has the highest personal discount rate crossover point, and state 26, the semi-absorbing state, has the lowest. A risk-neutral individual would have a personal discount rate that is close to the rate of borrowing. As detailed in Warner and Pleeter (2001), different individuals may face different borrowing rates affecting their value of inter-temporal substitution and that personal discount rates decreases with increasing education and income. Warner and Pleeter (2001) estimated personal discount rates for officers between 10 and 20 years of service between 10% and 19%.

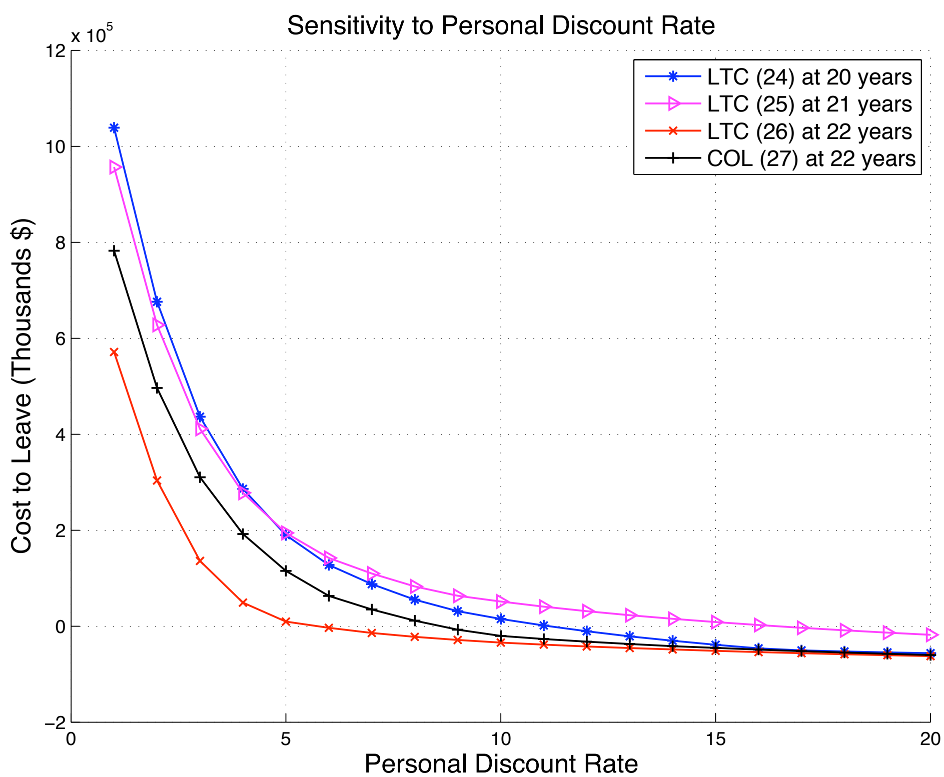

Figure 2.3: Personal Discount Rate Sensitivity

To gain some intuition for the difference in personal discount rate, we provide the following bounds on the value of retiring at any state. The personal discount rate adjusts the value of retirement, as does the number of years that the officer expects to spend in retirement. As the number of years in retirement is always finite, the value of retirement is also always finite. With discounting, even letting the number of years that the officer will spend in retirement tend to ∞, the value of retirement stream still has an upper bound which can easily be found by taking the limit of the geometric series. Taking the limit of the retirement portion of the $Q_t^*(x, q)$ function at three different personal discount provides some intuition concerning the difference in value of retirement. For personal discount rates of 5%, 10% and 20% the upper limits are $20r_t\hat{m}_{x,t}$, $10r_t\hat{m}_{x,t}$, and $5r_t\hat{m}_{x,t}$.

This shows that officers that choose to continue serving may, in addition to possibly having a higher taste for military service, also have a lower personal discount rate. It might be that the security of the military compensation and retirement system, combined with lack of compensation volatility entices those with lower personal discount rate to serve longer careers. This might suggest that officers serving past 20 years are serving for a combination of taste and monetary compensation that is tied to their low discounting of their military retirement.

Figure 2.3 also shows that the personal discount rate for COL that would result in an optimal choice of staying in the Army is between 8% and 9%. This would require a personal discount rate equivalent to the interest rate and much lower than estimated in the literature. Holding the personal discount rate at 12% would require discounting the civilian opportunity by 22-23%, as shown in Figure 2.5 and discussed in the next section.

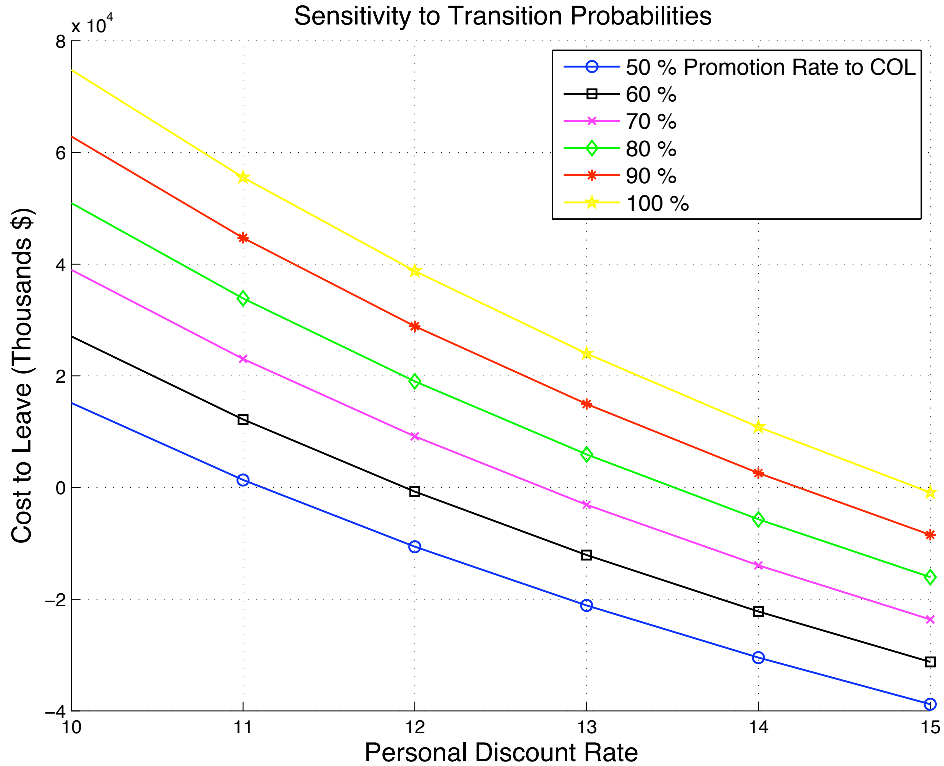

Figure 2.4: Transition Rate Sensitivity

Figure 2.4 displays the changing break-even personal discount rate as the promotion opportunity to COL varies from 50% to 100%. Although each of the promotions modeled by the transition probabilities effects the future value of continuing, only the promotion to COL is relevant to the exploration of LTC retirement decisions. As Figure 2.2 indicates optimal policy for $t < 20$ is to stay in the Army, the promotion rate to COL is the critical parameter. Interpret the changing probabilities of promotion to COL in Figure 2.4 as conditional probabilities based on some additional information revealed to the officer prior to $t = 20$. Assuming again a personal discount rate of 12%, all LTC that viewed their conditional probability of promotion to COL as greater that 60% would follow an optimal policy of staying in the Army at 20 years. This could be interpreted as attaining some other selection prior to $t = 20$, e.g. War College, Battalion Command, PhD Schooling, and adjusting the conditional probability of selection to COL appropriately.

2.3.2 Changing Civilian Opportunities

The assumption that a retiring officer can command the same salary in their second career as on active duty is a strong simplifying assumption. However, if the labor market is efficient, and officers possess transferable skills, the assumption that civilian salary would be approximately the same as the salary they are earning in the Army is reasonable (Stigler 1962). We assume that the mean of the distribution of potential salaries is their current salary. Also from the perspective of the officer, his/her current salary provides a benchmark to compare against potential future offers.

Some officers have the intention to "do something different" in retirement. These are taste differences, and the desire to start a new career may easily drive the retirement decision. Clearly doing something different might have very different compensation, but there are may field grade specialties that translate very closely to civilian occupations. Nuclear engineers, physicists, software engineers, operations research analysts and public affairs officers all closely align with post-military employment. As such, it would be a reasonable assumption that under an efficient labor market assumption, civilian compensation should mirror military compensation. However, it might be such that in expectation the compensation is

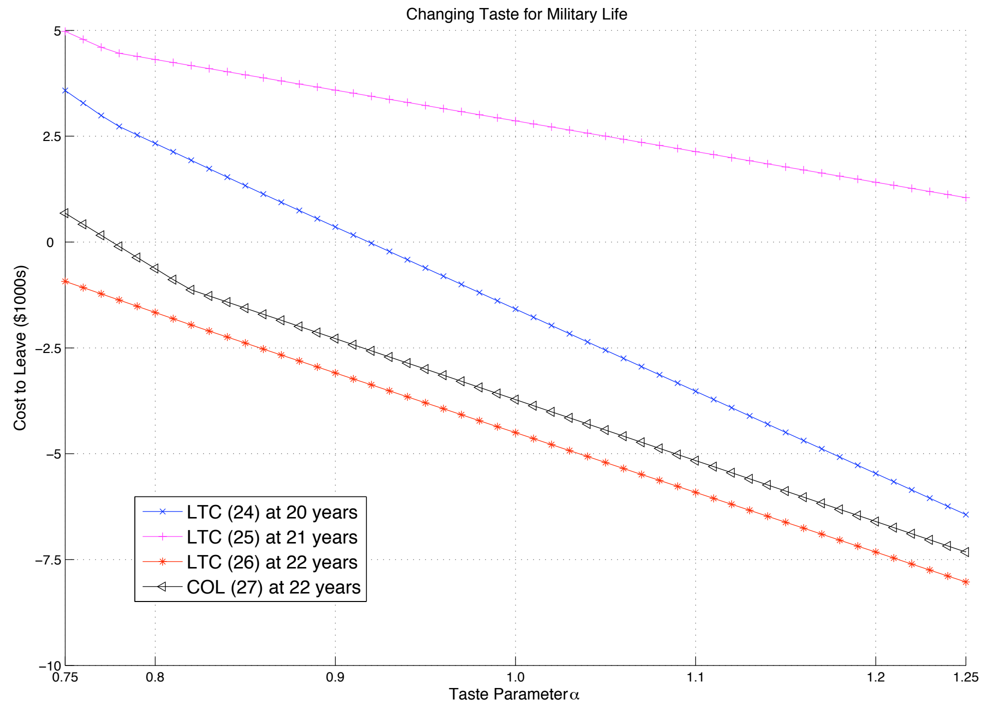

Figure 2.5: Changing Civilian Opportunities, personal discount rate = 12%

very comparable, but very different in distribution, especially considering the low variance and skew in the military pay system.

In this model, the expected civilian compensation is not dependent on future value functions, as seen in (2.5) and Figure 2.5. Varying α as a taste parameter changes the value of civilian compensation linearly. This can alternatively be interpreted as an increase (decrease) in taste for military lifestyle or decrease (increase) in available civilian compensation. In this model, increasing α represents an increase in available civilian compensation or decrease in taste for military lifestyle. Holding the personal discount factor at 12%, with $\alpha > 0.92$, the optimal decision for LTC is to retire at 20 years and for values of $\alpha < 0.92$ the decision is reversed. Figure 2.5

also shows that for COL at 22 years, the decision is to retire when $\alpha > 0.78$ and to continue serving when $\alpha < 0.78$. This result suggests that the current compensation package incentivizes officers that discount civilian opportunities by 25% to continue serving until their consideration for COL, and possibly even until mandatory retirement.

This model finds that the taste-neutral officer decision is very dependent on their personal discount rate, but the linearity of the taste parameters makes the model more dependent on the assumption of pay distribution rather than the α value. However, changing α values can be interpreted in a couple of different ways. Clearly, if civilian opportunity is less than the salary you are earning on active duty, it is optimal to continue serving. Also, the α can be viewed as a taste for military service. There are many unique features of serving, and this value can be viewed as a relative weighting of military compensation versus civilian opportunities. Similarly, an officer that had a large aversion to mid-career change, or placed a high value on the prestige of serving in the Army, would have a lower α and require a higher offered civilian wage to change their optimal policy from that of continuing to serve.

2.4 Summary and Policy Analysis

In Gotz and McCall (1983), an Air Force Lieutenant Colonel would optimally wait to see if they were promoted to Colonel, and then conditioned on not being selected, they would retire at 23 years. Here we found that a LTC not being considered for promotion would retire at 20 years. A significant issue is that the COL promotion opportunity for due course officers is at 22 years of service, and the current compensation system does not provide incentive for officers to serve the additional year after twenty to compete for this last promotion. Our results indicates that non-monetary factors play a significant role in the decision to compete for promotion to COL, and that COL compensation alone does not provide incentive for continued service.

Clearly the change to allow a 40 year career is only a meaningful innovation is if there is an incentive to serve that long in uniform. Of course, the existence of officers serving past the perceived fiscally optimal retirement point lends credence to any model that includes or weights non-pecuniary factors. Even without knowing exactly which non-pecuniary factors are the the most significant, modeling can price these factors by comparing model predication with empirical behavior of the force.

If the 1981 policy changes had the desired effects, optimal points would be *greater* under the new system. However, the policy change appears to have the *opposite* effect in the model. This shift agrees with anecdotal evidence from Army officers that carreer aspirations have shifted to retirement at LTC at 20 years. However, it is problematic if the selection point to COL is at 22 years, past the optimal retirement point for a LTC. Such a shift creates a system where the Army only incentivizes careers past twenty years for officers with aspirations to General officer rank and who strongly value the non-pecuniary benefits of military service and discount the advantages of civilian life.

Our model also demonstrates the value and importance of promotion signals regarding the promotion distribution to COL. Assuming that the Army has knowledge of the subset of officers most competitive for promotion, signaling to these officers can alter their optimal retirement policy. As shown in Figure 2.4, signaling an increased promotion opportunity from 50% to 75% for some officers would switch these officer's optimal policy at twenty years to continuing to serve in the Army and competing for promotion to COL.

2.5 Future Research

Several additional areas of exploration could expand the results of this chapter. An empirical analysis of the actual salaries of retired officers would increase the strength of the results. The civilian pay assumption that an officer retiring can continue to earn the same salary in retirement is reasonable if there are no significant pay disparities between civilian and military occupations. Empirical work to estimate the earnings paths of retired officers in different specialties would add richness to our results.

Also, empirical study of the current behavior of retiring officers will indicate any shifts in behavior. Although there are many factors influencing early retirement, the behavior of current LTCs is key to understanding the expected population of officers serving until they are considered for promotion to COL. In addition, the modeling of General officer opportunity might provide some insight into the retention behavior of those officers whose early promotions provide them a higher opportunity of promotion beyond COL.

Officers are not homogeneous, and adding differences between officer types could add richness to the model. The α coefficient examines the differences in how an officer values their current compensation, but the model could be expanded to model two different types of primary motivations: prestige and economic. In this type of a model, learning could occur as the individual updates transition/promotion probabilities over time and shifts primary motivation between that of prestige and economic motivations. This would allow signals, e.g. performance reports, school or command selections, to be incorporated into the individual's decision-making. Different utility functions, e.g. power or exponential, could also be used to examine possible different valuations between individuals.

The tax advantage of staying in the military and the security of the military career are two additional areas to explore. Many senior officers serve in the Pentagon, which provides tax-free housing compensation at a higher level than most Army posts. Also, while serving on active duty, officers do not have to pay state taxes based on their posting, but rather based upon their residence. This allows many to serve in the Virginia area and enjoy the additional benefits of the National Capital Region, without having to pay the taxes that would be required if working in the civilian sector. These additional modifiers to salary could also strengthen the results.

Although this discrete state space, fixed horizon problem is easily solved using dynamic programming, other solution methodologies might offer additional insight for sensitivity analysis. It has been established that MDPs can be alternatively modeled as linear programs and the resulting solution might offer additional opportunities for sensitivity analysis and for exploiting the dual formulation. Simulation methodologies might add additional insight in creating sample career paths, and to facilitate additional exploration of civilian pay distributions.

In this chapter we have found that the optimal policy for taste-neutral LTCs is to retire at 20 years. This suggests that research in areas of organizational theory to investigate what compensation a mid-career worker most highly values might provide additional insight into why and which LTCs stay. Such investigation could suggest policy and programatic changes that the Army could implement if there was a need to increase mid-career officer continuation rates.

Chapter 3

Army Officer Force Profiles and Markovian Manpower Models

3.1 Introduction

Illinois Senator Everett McKinley Dirksen is credited with commenting, "a billion here, and a billion there, and pretty soon you are talking about real money" (Senate 2009). The cost of officer manpower in 2008 was \$9.04 billion, accounting for 36% of the total \$47 billion manpower bill and 5% of the total Army funding of \$155 billion for fiscal year 2008 (OSD 2008).

The cost of the Army manpower program is the largest line in the President's budget, and a substantial part of that budget is devoted to the executives, managers and leaders of this force, the officer corps. The officer corps is currently roughly 75,000 soldiers in 10 ranks, and over 50 specialties. This manpower system is made increasingly complex by a restriction requiring most accessions at Second Lieutenant (2LT), the entry level commissioned officer rank. This restriction, termed a bar to lateral entry, is common to all the services and serves as a defining characteristic of military manpower systems.

It is relatively straightforward to document the specialty and ranks required within an organization. However, documenting time in grade or time in service requirements is more complex. In academia, a classroom can be filled by any professor in a specialty, with substitutability for rank and time at the university. In Army units, officer positions are coded for a particular branch and rank. The documented requirements do not detail a desired age distribution of the position and hence do not provide requirements for the age distribution of the workforce. This subtle nuance displays the difference in planning the work assignments for a given year versus planning for the career workforce of the entire organization.

Any organization would prefer some mix of junior and senior members of each rank and specialty. The employees need to have a structure in the organization that offers them a viable career. However, how the Army as a whole should aggregate such requirements and provide viable career paths for the soldiers and officers is a non-trivial manpower problem.

The manpower requirements for the Army are expressed in terms of a force profile. The focus on providing viable officer careers has traditionally focused on modeling year groups, i.e. time in service, in order to express the differentiation between alternate force profiles. Although the description of a force profile focuses on details of time in service or experience, the Army requirement documents are only detailed to rank and specialty. Here we look to combine modeling efforts based on time in service, with the requirements detailed at rank and specialty.

Here we look to find an optimal **commissioned officer force profile**. The officer corps are the most senior leaders of the uniformed Army. Clearly, an optimization of the force profile could take several forms. Best force for a given cost or cheapest cost for a required force would be two obvious alternative objective functions. Another alternative is to suggest that the current force is optimal, as it is clearly feasible, and then attempt to find a way to maintain the system at steady state in the neighborhood of the current state.

We will develop a model of the force profile and its evolution over time, controlled by the manpower planners of the Army Staff, but subject to the decisions of individual officer's, Congressional guidance and Department of Defense leadership. In times of change, where either the operational requirements or budget constraints might change drastically, a model can provide a prescription of how to get to the new required force, from the current force.

There are several methods of analysis that have an established history in manpower planning, and we will look to combine insights from several modeling paradigms and develop a linear programming model of the Army officer corps. We propose a new network structure that incorporates both rank and years in grade to combine cohort, rank, and specialty modeling without falling into the common pitfalls of small cell size and uncontrollable end effects.

This model follows in the tradition established by Gass et al. (1988) and implements modeling techniques and insights from the Army implementation of the Enlisted Grade and Enlisted Specialty model (Hall 2004, Cashbaugh et al. 2007). This chapter is the first implementation of specialty modeling in a manpower model for U.S. Army officers. Models of the U.S. Army manpower system have isolated accession planning for second lieutenants (Henry and Ravindran 2005) and the Career Field Designation process for majors (Shrimpton and Newman 2005), but this is the first integration of rank and specialty modeling over the entire career and development of an optimal force profile.

This chapter will highlight aspects of manpower modeling in operations research, and Section 3.2 discusses the literature that has contributed to the development of the different modeling approaches in this research area. Section 3.3 will explore characteristics of U.S. Army officer manpower planning before presenting a linear programming model to find the optimal force profile in Section 3.4. We conclude with results from our experiments in Section 3.5 and possibilities for extensions in Section 3.6.

3.2 Manpower Modeling

Manpower models need to answer questions for the planner. The structure of the force may be already set or may be the product of exploration. Models are often used for forecasting and cost estimation, as well as to investigate the feasibility of proposed structures. Manpower models are designed to help with some aspect of creating and sustaining a force of workers to accomplish the organization's objectives or mission. Much of the literature focuses in those industries and application areas where there is a human delivered element of the product or service, e.g. education, consulting and defense. The area of manpower planning has been an interdisciplinary area of research that is not unique to the military, but is a function of all large organizations.

Statistical models in manpower research fill a role that can be identified as belonging to one or more of four categories: descriptive, forecasting, control or design. Clearly there is some overlap, but this is the categorization framework presented in Bartholomew et al. (1991). Any given model can have strengths in one or more of these categories.

Descriptive models look at the distributions of the systems, with the intent of developing metrics to indicate the current state of the system. The number of personnel in different ranks, and the experience distribution of the workforce are common descriptive statistics used in manpower planning. Insight into the number of assignments that are available at different grades or the number of new hires needed at the different levels of the organization can be determined from a model that accurately describes the current state of the system.

Forecasting models focus on the changes that will take place in the system if current trends continue. A forecasting model that describes a system in steady state is not making a prediction about what trends will come to be, but instead provides prediction for the state of the system under that assumption that current trends continue. The forecast can help to guide decisions to be made, but traditionally forecasting models do not offer recourse within the model.

Control models reflect changes to the system based upon managerial action. The management is viewed as having some set of tools or levers, such as promotions or new hires, and the models are designed to illustrate potential impacts of different decisions. Often control and forecasting models are combined to allow the exploration of the effect of different policies given the assumption that trends continue.

Design models help the decision maker create new systems, and look to see what characteristics will be present in that system. Design models will borrow heavily from, or intersect with, concepts from, and used in, organizational design.

In manpower intensive and service industries, the workforce is the key product of operations, and design models can be crucial in the management of these operations. In the military, education and consulting, the workforce is the key element of the operation.

Two of the main questions addressed by these classes of models are attainability and sustainability. In a multi-class or rank system, e.g, military and academia, the number and experience in each grade can be of fundamental interest. In the military, this can be expressed as a force profile. The **officer force profile** is an expression of the number of officers by grade and time in service. One assumption is that productivity will increase with experience, suggesting that the organization would like to retain experience at all levels. Yet, serving for too long at any one stage in the organization can cause stagnation and motivational challenges. Moreover, the force profile must meet the services needs. The traditional graded military manpower system with a bar to lateral entry may place limits on the number of time periods required to move to a desired force profile.

Another important modeling characteristic is the closed or open nature of the system. Many social processes are modeled using a closed system, and the flow into or out of the system are the key metrics. Other models focus on an open system, and focus on the attrition or wastage, as personnel are lost from the system and workers must be promoted to fill vacancies or new workers must be hired to maintain the system. Models can focus either on modeling workers, jobs or a combination of the two. When modeling jobs, a model will often look to assign workers against vacancies. When focusing on modeling workers, there traditionally must be a time component which could be discrete or continuous.

3.2.1 A Textbook Introduction

The foundations of manpower modeling within operations research and management science are clearly laid out in a series of textbooks. Grinold and Marshall (1977) covers an introduction to manpower models and Vajda (1978) addresses the mathematics of deterministic manpower models. Bartholomew (1982) addresses sochastic social process, of which manpower modeling is a featured example, and Bartholomew et al. (1991) further extends statistical methods for manpower models.

Manpower planning covers a wide range of problems and applications, in addition to the military manpower planning problems that are our focus. Tien and Kamiyama (1982) provides a survey of algorithms for scheduling manpower, incorporating total requirements as well as vacation in creating shift and work schedules. Edwards (1983) provides a summary that focuses on how current models have been implemented, and issues that arise from trying to implement manpower planning models. They focus on civilian manpower planning in the United Kingdom, with a secondary application to other European countries. Green et al. (2001) provide an application of queueing theory to staffing issues in hospital emergency departments.

Gass (1991) provides a primer of military manpower modeling and breaks the research into three main groups: network models, transition rate models, and goal programming models. Common to all three groups is a modeling construct that describes individuals who are aggregated by a set of attributes and grouped into classes. Each combination of a set of attributes at each time period represents a state. Each state will contain a number of individuals with common attributes. Army manpower planners refer to the level of aggregation used in the model of a system as subsystem aggregates (SSAs). We will next explore some of the manpower research that has been completed in each of the different mathematical modeling classes.

3.2.2 Network Models

Network models are based upon an extension of the transportation problem to personnel assignments. Network models ensure conservation of flow and have the advantage of producing integer solutions. The models are especially powerful when modeling systems with a limited number of characteristics at each time period. The following are military examples of the personnel assignment problem that have been modeled using variations of network flow models.

Klingman and Phillips (1984) addresses both topological and computational aspects of assigning enlisted Marine Corps personnel to positions. Their model is a network flow model to assign 10,000 Marines to billets, which groups billets into differing priority levels, and assigns fill percentages to each subset. Bausch et al. (1991) looks at wartime assignment of Marine Officers. Their model increases the pool of officers available to assign to wartime tasks by considering all active duty and reserve Marine officers to be available for reassignment to fill a wartime billet plan, even though loosening to this extent the constraint of the number of officers who are available for reassignment at any one period may not be feasible or reflect the needs of the Marine Corps. Krass et al. (1994) developed a model for Navy personnel assignment to ships and combat units. Their model incorporates Navy readiness metrics and seeks to optimize personnel readiness in the individual ships and across the Navy. Reeves and Reid (1999) also examines the readiness of a combat unit, in their case a reserve company. They focus on the schooling and training required for the company's soldiers, and maximize the percentage of soldiers with the required skills.

Holder (2005) looks at the problem of offering assignments to sailors while maintaining readiness for the Navy. In their model, a subset of the jobs for which a sailor is qualified are selected and offered to the sailor, when a sailor calls their assignment manager or detailer. The non-selected jobs are returned to the available list after the sailor makes a selection or after predetermined waiting period. Lewis et al. (2006) also looks at the Navy assignment problem, grouping sailors into teams that will be assigned to a command using interval bounds, and incorporating training and minimum class sizes for training that will be required enroute to the new assignment. Shrimpton and Newman (2005) uses a network model to describe the assignment of Army officers to functional areas. Additional specialties exist at the field grade ranks that are not prevalent at the company grade ranks, and this model decides how to allocate the existing inventory to the larger group of specialties. The network model minimizes the sum of metrics from an order-of-merit list while still meeting Army requirements

3.2.3 Markov Models

Transition rate models require transition probabilities as an underlying assumption. Given an initial distribution, Markov models will answer the question of what attributes the force will have at each phase in the planning period. When modeling with a transition rate model, each individual follows the same Markov process and individuals are assumed to behave independently. Whereas the network models described earlier focused on military personnel assignments, the following research focused on describing manpower in academia.

Examples of Markov manpower models in discrete time are described and solved in Feichtinger and Mehlmann (1976) and Feichtinger (1976), where the authors examine stable populations and age distributions in open systems. They link the Gani-type education manpower models with Leslie linear population dynamics models. Mehlmann (1977) extends their results to continuous Markovian manpower models. Davies (1976, 1982) considers the effect of recruitment policies and especially the importance of smoothing hiring actions. In this model, recruitment policies are derived using fixed promotions which are controlled by management, and the stochasticity in the model involves employee continuation decisions. In this model, employees in each stage are modeled as independent binomial random variables, and this independence assumption facilitates calculation of one and two-step probabilities and recruitment policies.

Grinold and Stanford (1974) creates a mathematical programming model to investigate fractional flow models. They examine both finite and infinite horizon, fixed and free final state constraints, steady-state distributions with staff distribution constraints, and temporarily transient systems moving towards steady-state. Their models are created to design possible structures, rather than to prescribe operational promotion policies. Grinold (1976) presents a manpower model for Naval aviators that is a hybrid between a linear-quadratic optimal control problem and Markov decision problem. The demand for Naval aviators is varied based upon six states of the world that correspond to various peace and wartime states. Dynamic programing solutions are presented to find steady-state solutions for the objectives of minimizing the weighted squared error between supply and demand and smoothing the flow of manpower.

Feuer and Schinnar (1984) link personnel and vacancy flows in a graded personnel system, focusing on promotion opportunities and outside hiring within a large university community. Their results examined the effect on junior faculty promotions when additional senior faculty are hired. Stanford (1980, 1985) develops a mobility process based upon the cumulative length of service of members of a given rank. The loss distribution and promotion distributions are taken as inputs and their model develops time in service estimates for each rank. Different promotion policies with regard to seniority are evaluated and used to derive expected length of service distributions.

Yadavalli and Natarajan (2001) examines a single grade system to study the cost to the organization of having vacancies and recruiting. Nilakantan and Raghavendra (2005, 2008) look at proportionality policies, bounding the recruitment to each grade to a proportion of the population. They address the controllability of such a system in term of attainability of a given structure and maintainability of such structure, focusing on the numbers of internal promotions and external hires.

3.2.4 Goal Programming

Price and Piskor (1972) first applied goal programming to military manpower planning, interpreting the constraints as goals as opposed to hard constraints. They created a control model for a three-year planning horizon to forecast strengths by specialty and provide promotions prescriptions for the Canadian Army officer corps. Zanakis and Maret (1980, 1981) describes a Markov chain model for forecasting available inventory, and then a one-period preemptive linear goal program to allocate workload in an engineering firm.

Holz and Wroth (1980) describes the development of ELIM-COMPLIP, Enlisted Loss Inventory Module Computation of Manpower Programs Using Linear Programming. The linear programming solution of the manpower program was the initial solution used to complete the transition from a conscript manpower system to the All-Voluneteer Force (AVF). The initial COMPLIP model was created to provide an alternative to the computerization of the manual system, and the incremental improvement that resulted in the ELIM-COMPLIP model were operationally employed by Army manpower planners until replaced by the Enlisted Grade (EG) model early in the 21st century (Hall 2004).

Collins et al. (1983) designs a goal programming formulation to balance the education level, mental category, i.e. CAT I -IV, race and sex of recruiting goals in order to meet manpower requirements. The model develops historical rates and factors, simulates losses over a one-year horizon, optimizes the quality mix for accessions to meet manpower requirements, and then calculates the cost of the resulting manpower. Gass et al. (1988) describes a goal programming model that prescribes accessions, promotions, separations to minimize the different from grade goals and total force goals. The model makes use of a Markov model to estimate projected manpower, and then optimizes to match requirements over a ten and twenty year planning horizon.

Bres et al. (1980) allocates Naval officer accession from each of the commissioning sources to different warfare communities. The objective function is to minimize deviation from strength goals, subject to policy constraints, including bounds to minimize inter-year deviations. Henry and Ravindran (2005) compares preemptive and non-preemptive models to assign Army officers to initial branches and branch detail programs.

3.2.5 Stochastic Programming

Martel and Al-Nuaimi (1973) models a oil industry manpower problem as a two-stage stochastic program with recourse. The application is made to industries where skilled labor and union rules make hiring and firing expensive and where it is not possible to create extra inventory when the manpower is not required. The recourse decision in this model is only to hire additional manpower, temporary workers or pay overtime when it is determined that the base staffing level is inadequate.

Martel and Price (1981) combines the results of goal programming and stochastic programming applied to human resource planning for the Canadian armed forces.

3.3 Modeling the Officer Manpower System

The Active Army Military Manpower Program (AAMMP) describes the manpower requirements of the active force of the Army. The AAMMP forms the basis for the President's budget and the Program Objective Memorandum (POM). The AAMMP answers the fundamental Army manpower questions of how many commissioned and warrant officers, enlisted soldiers and non-commissioned officers will be required, and an estimated number of soldier man-years and a forecasted end-strength for the next two fiscal years of the President's budget planning horizon and the eight years of the POM planning horizon. The Army program encompasses all aspects of Army funding. Interestingly, it was early work in building military programs that gave their name to mathematical and linear programming.

The Deputy Chief of Staff of the Army for Personnel, i.e. Army G1, has the responsibility for developing the military manpower portion of the Army program. This responsibility is carried out by four of the divisions of the Plans and Resources directorate of the Army G1. The divisions of Plans, Compensation, and Strength Forecasting are all responsible for portions of the inputs, which are assembled by the Resources division.

Congress has the constitutional authority to raise an Army and as such publishes the National Defense Authorization Act (NDAA) annually. The NDAA 2008 established Army active duty end-strength, defined as the number of soldiers on active duty on 30 September of a given year, at 525,400 soldiers and Army funding at $155.7 billion. In additional to end-strength restrictions, Congress provides quarterly upper and lower strength bounds that serve to control or regulate the man-years that will be executed each year.

In addition to the fiscal guidance contained in the NDAA, the Defense Officer Personal Management Act of 1980 (DOPMA) contains guidance on the officer portion of the services, and the Army in particular. One such piece of guidance is a chart of allowances for field grade officers, i.e. majors, lieutenant colonels and colonels, based upon the total number of officers on active duty. The requirements keep a balance of approximately 5 %, 11% and 24% of the total Army commissioned officers at the ranks of colonel, lieutenant colonel and major, respectively. The officer corps as a whole has traditional been limited to 15% of the total end-strength or force structure.

When focusing on officer programming, the overall Army strength limits will not constrain the officer system, while the DOPMA requirements will be constraining, and must be incorporated. However, most of the system constraints will be goals or regulations imposed by the Army to control manpower. The current methodology used by the Army G1 is to first determine the number of officers that will be contained within the program by statistical analysis, and then to run a suite of optimization and simulation models to determine the number, grade and specialty of the enlisted force (Cashbaugh et al. 2007). In this respect, our model provides additional analytic rigor by employing optimization for the officer force by grade and specialty. The second order effects of improved officer manpower forecasts are realized as input provided to enlisted forecasting models.

3.3.1 Characteristics of the Army Manpower Program

The ten ranks of the commissioned officer corps are divided into 4 General officer ranks, i.e. Brigadier General, Major General, Lieutenant General, General, 3 field grade ranks, i.e. Major, Lieutenant Colonel, Colonel, and 3 company grade ranks, i.e. Second Lieutenant, First Lieutenant and Captain. The 4 General officer ranks are tightly controlled and limited to 307 officers, a limit imposed by Congress. Very little forecasting is need to estimate the General officer strength, as all General officers, by virtue of promotion to General officer, share the same specialty. However, the field grade officers, also bounded by Congress, are a large pool of manpower which requires modeling to forecast their strength by grade and specialty. The rank pyramid is structured so that there are approximately 16,000 MAJ, 9,600 LTC and 3,800 COL. The requirements for LT and CPT are 12,000 and 23,000, respectively, with no distinction made between the 2LT and 1LT.

The Army uses branches to group officers into specialties. Further, the war-fighting specialties are grouped into the Army Competitive Category (ACC), and the Medical, Chaplains and Judge Advocate General, i.e. lawyers and judges, are each grouped into their own subgroup. The ACC is divided into 43 specialties, 15 basic commissioning branches, 3 special operations branches and 25 functional areas specialties. Not all specialties are required at lieutenant rank, and many of the functional area specialties are predominantly required at field grade ranks.

The challenge to adding specialty to this model is to avoid the "curse of dimensionally" while allowing for specialty modeling without an order of magnitude increase in the state or decision space. A method to increase the fidelity without enumerating the entire states space is to create a more detailed network structure to describe rank states that contain additional time in service or system characteristics. In this way, we aggregate officers which can be thought of as substitutes, either by rank or specialty or time in service, to reduce the dimensionality of the state space.

The Army personnel databases contain many fields we might suggest for incorporating into a network model, e.g. rank, time in service, time in grade, assignment history, education level, to efficiently model officer manpower we must chose some subset of the available data. A model at the level of monthly projection, over six grades, with fifty specialties, with 480 possible months of service, would result in 144,000 cells per period and assuming an eight-year projection horizon would result in 1,152,000 cells. Modeling 60,000 officers with 144,000 potential cells would result in a very high percentage of empty cells. Combining states which might be viewed as substitutes over an annual time horizon can allow for a higher level of aggregation without losing model fidelity.

To explain the structure of the force profile a little further, we will look to a descriptive model based upon the Army's yearly cohort structure. The three Army commissioning sources are the Reserve Officers Training Corps (ROTC), the United States Military Academy at West Point (USMA), and Officer Candidate School (OCS) produce between 4,000 and 6,000 officers annually. Officers are all assigned a year group, based upon the fiscal year in which they are commissioned. The USMA officers have been in a training pipeline for 4 years, the ROTC officers from 2-4 years, and the OCS officers for just under a year. These lead times all lend predictability to officer accessions and credibility to a model based upon annual accessions goals and cohorts.

3.3.2 Cohort Modeling

The Army pyramid structure and up-or-out policy requires fewer officers of increasing tenure. Even with a four or five year commitment, some officers will be allowed to leave the service early, and as such a cohort will never be as large as upon entering. Early and late promotions can result in members of a year group holding differing ranks, which can be accounted for by adjusting year groups, or maintaining a separate data element for rank. One consideration in designing an optimal force profile is to provide consistent or constant opportunities to each year group. This complements the Army notion of fairness and supports the need to have a smoothly aging Army.

The number of officers that choose to stay each year is expressed as a survival distribution. The survival distribution is the probability of surviving to year $n+1$, conditioned upon surviving to year n. These statistics are calculated annually and aggregation of 30 year groups provides a the survival distribution for a given year. A stationary survival distribution would suggest that the opportunities for each year group of officers would remain relatively constant, as economic factors and force structure changes do provide variability in opportunities inside and outside the Army and result in evolving survival distributions. Variability in the accession cohorts subject to changing requirements and opportunities will provide an source of variability to survival distributions.

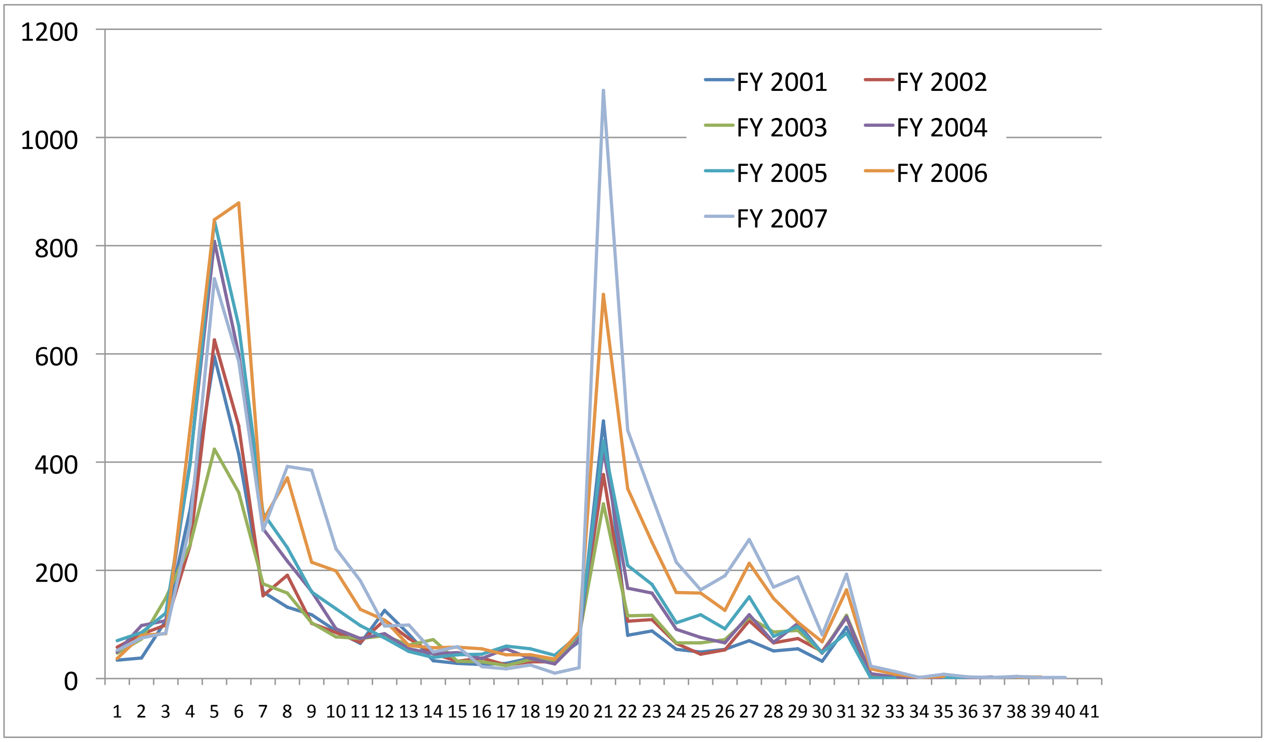

Figure 3.1: Officer Losses by Years of Service, FY 2001-07

Figure 3.1 displays seven historical loss distributions by years of service from 2001-2007. This data suggests a stationary distribution for losses by year group. The sharp peaks in Figure 3.1 correspond to the ending of initial obligation for officers between four and six years of service and retirement at 20 years of service, the first point of retirement pension vesting. Either continuation distributions or loss rates can be used to model force losses, and both methods have the advantage of tying the current population size to the losses during each period.

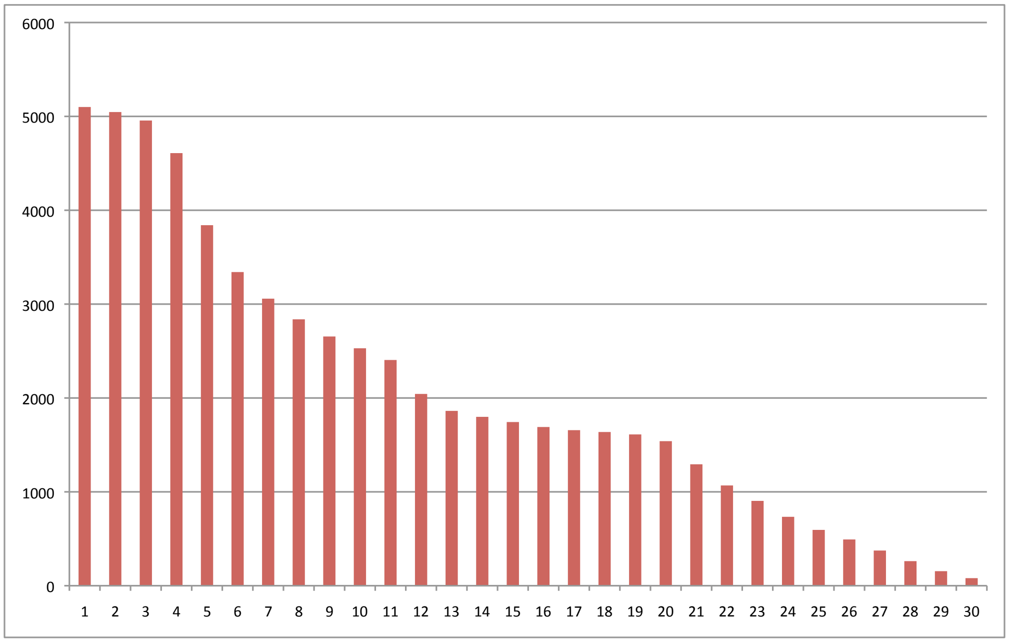

Figure 3.2: An Army Officer Force Profile, Number of Officers Remaining by Years of Service

With support for the assumption of a stationary survival distribution, a projected accession cohort can be aged by the survival distributions to approximate the size of the total force at steady state. In such a two-parameter model, the survival probabilities and accession cohort can be adjusted to move the total force distribution towards a desired distribution. Figure 3.2 represents a cohort force profile, based upon 1988-89 continuation distributions and 5100 annual officer accessions. Although 1988-89 might seem like an arbitrary timeframe to be chosen in 2008-2009, it is the last time period when the Army was in something close to steady state. The period of the 90's was one of officer drawdowns after the fall of the Berlin Wall and the end of the Cold War. The period since 2001 has been dominated by the Global War on Terror (GWOT) and force structure growth. By combining an example of steady state survival distributions with the current cohort planning figure, we gain some insight into a possible steady-state officer profile. The current planning figure for the Army's accession mission is 5100 new officers each year (Haight and Lewis 2008).

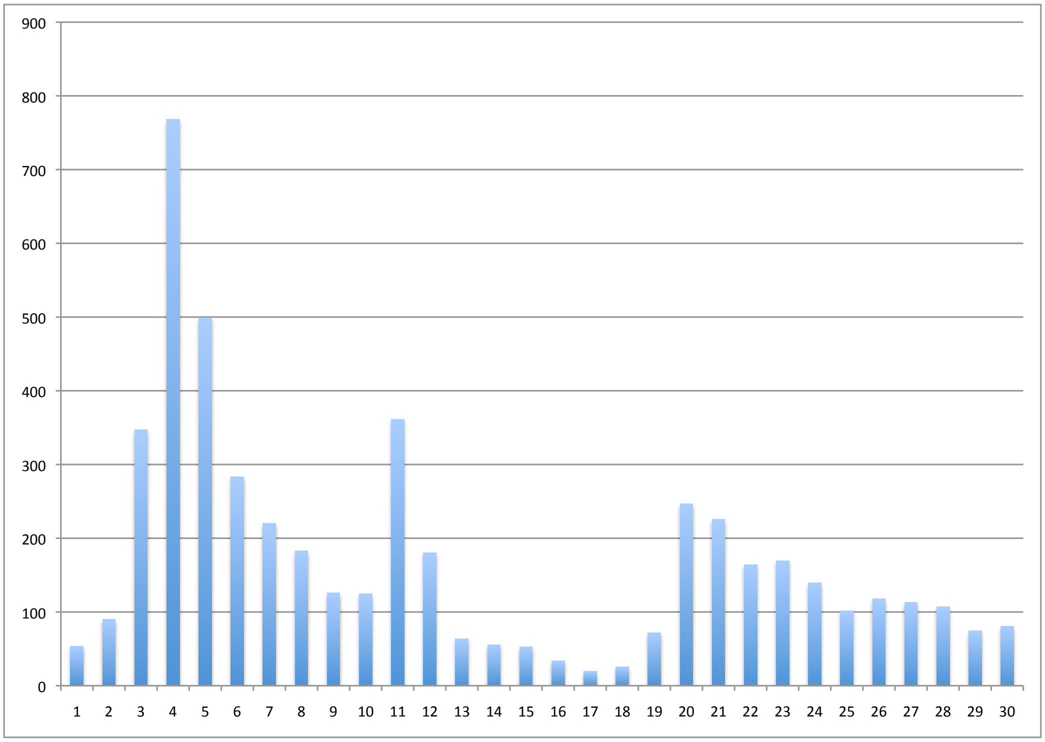

Figure 3.3: Officer Loss Profile, Number of Officer Losses Per Year

Figure 3.3 displays the loss distribution for an individual cohort, implied by the survival distributions from Figure 3.2. An Army at steady-state with cohorts exhibiting this behavior would require 5100 new officers each year and requires 5100 losses, officers leaving the system, to achieve a stable distribution.

This type of cohort model requires rank to be added after the time in service distribution has been determined, and this years-of-service distribution would equate to 15102 LTs, 20343 CPTs, 15730 MAJs, 8054 LTCs and 2693 COLs, and a end-strength of 61,924 officers under a 100% cohort promotion in primary zone of consideration assumption. Such a force could be maintained by a promotion policy to have a field grade population in accordance with the DOPMA constraints for a force between 60,000 and 65,000 officers. Under the assumption that officers would comprise 15% of the total force, this would equate to 400 K to 435 K Active Army, significantly less than the Congressionally mandated end-strength of 525,400 for FY 2009, for which this officer distribution comprise 12% the total force.

A prescription for promotion is not a integral component of the model, and merely looking at the expected inventory based upon size of incoming cohorts assumes the Army will be able to attain a steady state. It also does not incorporate any additional information that might provide insight into determining which officers might be choosing to stay. Additional modeling fidelity based upon rank and specialty will help to provide insight into the heterogeneity of the officer population.

3.3.3 Rank Modeling

A slightly different approach to discovering the optimal force profile is to first partition the officer corps by rank. Whereas in cohort modeling aging is a natural process, when partitioning by rank, promotion policies must be articulated to define transition probabilities. The traditional approach has been to model using a network model such as Figure 3.4. A network flow model would be used to allow for prescription of promotions and losses.

In an optimal force profile, losses and promotions should both be stationary distributions. The network flow model lends to modeling losses as opposed to con-

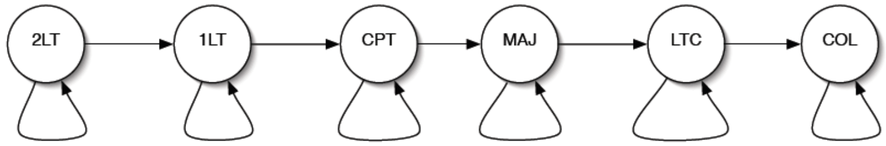

Figure 3.4: Network for Rank State Space

tinuation rates. At each time period, a subset of each rank will be chosen to be promoted, and other subset designated as losses. When modeling by rank alone, an upper bound or a combination of upper and lower bounds would be required to describe promotion opportunities from state to state. Adding accessions at 2LT, and the network flow model would provide an approximation of the personnel distribution.

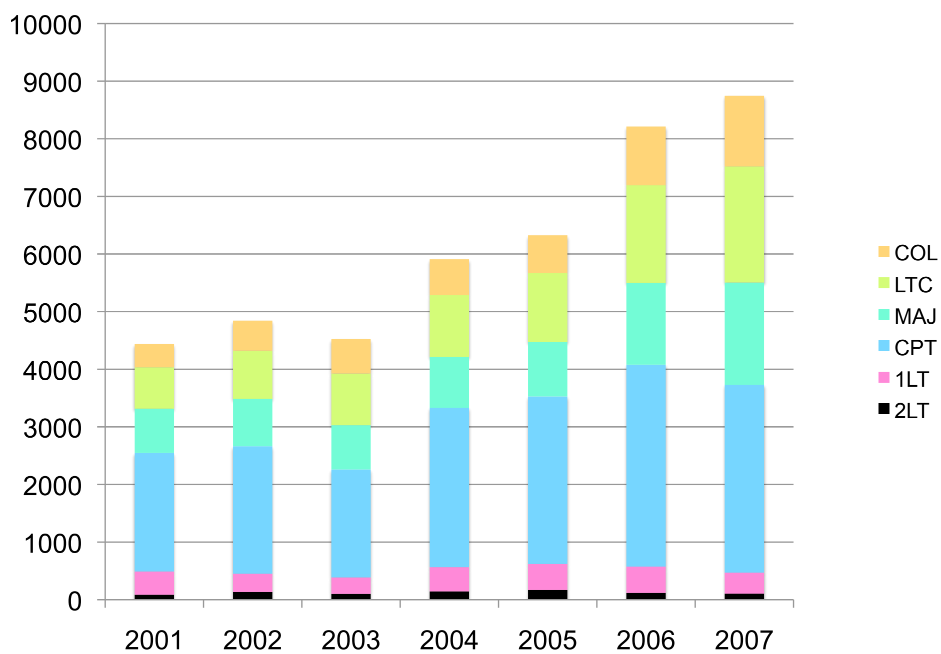

Figure 3.5: Annual Officer Losses by Rank

Whereas Figure 3.1 supports the assumption of a stationary distribution over the cohort, over the same time period, Figure 3.5 does not appear to describe a stationary distribution by rank. To gain modeling fidelity, we would like to combine the two paradigms, that of modeling cohorts and rank. Using both rank and cohort would result in 240 possible year of service and rank combinations. However, all possible combinations are not feasible, and many do not add fidelity to our model.

In order to provide a model that will describe the current system and suggest an optimal force profiles, we will aggregate in a different manner than previously considered in the literature. Combining rank, specialty, and cohort, we add additional rank states to model the time spent in grade between promotions opportunities. We are also able to aggregate all officers that have been passed over for promotion, a measure of the robustness of a given force profile and the flexibility afforded to manpower planners.

The network structure is shown in Figure 3.6. Starting with accessions in state 1, and using natural aging and annual promotion opportunities, the accession cohorts are distributed over the 29 states during their military career. States 3, 6, 7, 13, 20 and 26 represent states where personnel have been passed over for promotion and strength can accumulate. When younger cohorts lack sufficient strength, these states provided additional flexibility to the system. The on-time or due-course promotions, e.g. 2-4, 19-21, have a lower bound on the promotion percentage from these states, while the early promotions, e.g. 18-21, have a upper bound on the promotion percentage. This limits the number of promotable officers in any given model year and mirrors the way Army promotion boards consider officer cohorts.

An advantage of using a 29-state model for rank is this logical interpretation of

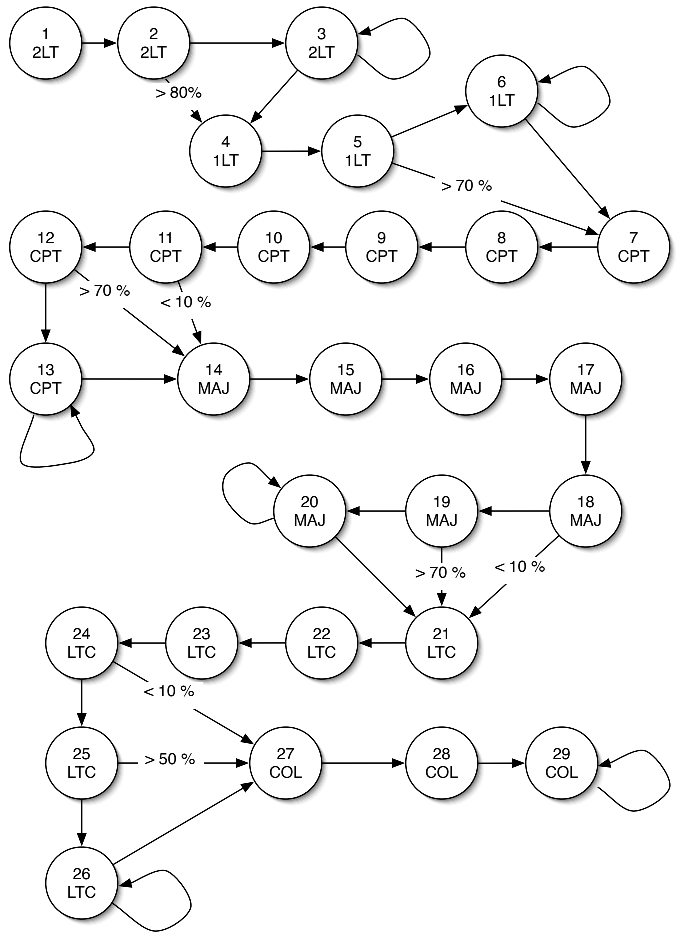

Figure 3.6: Network for a 29 Rank State Space

Army promotion policy. The structure of Figure 3.6 captures all promotion states, i.e. below the zone, in the zone and above the zone. This allows the targeting of possible promotions at the states where they would be applied by Army policy. This methodology allows the incorporation of specialty to the model using 1,247 states per period, assuming 43 modeled specialties, as opposed to 7,740 per period required to add time in service and specialty to Figure 3.4.

3.4 A Linear Programming Officer Manpower Model

We describe a linear programming manpower model that exploits the network structure in Figure 3.6 to capture rank and cohort fidelity while incorporating specialty. This model allows the explorations of different controls to develop alternative manpower plans to meet Army requirements. This is the first model in the literature to incorporate officer individual specialty, while aggregating the total officer demand from the requirements of the Personnel Management Authorization Document (PMAD) over the planning horizon.

Each of the data elements and decision variables are in units of individuals. Thus all elements of the objective function are in the same units, alleviating the need for scaling in the objective function. The network structure and the smoothing constraints work together to minimize end-effects, as will be seen in the numerical results. Also, the model provides parameters to weigh the different portions of the objective function as well as to prioritize ranks or specialties.

The model starts with a personnel stock vector describing the number of officers in the system. This personnel vector, X, is indexed by

g	grade $\in \{1 \dots 29\}$
s	specialty
t	time $\in \{0 \dots T\}$

with $t = 0$ providing the initial starting conditions for the strength of the force.

Thus the stock vector is $X_{s,g,t}$ records the number of officers in specialty s at grade g in time (or period) t. The initial personnel vector, X_0 has dimension 29×43. The final period of the projection horizon is represented by T, which will range from 2 to 9 years. Specialties are defined as listed in Figure 3.8.

The data or endogenous variables are

$$
\begin{array}{ll}
D_{s,r,t} & \text{PMAD} \\
L_{s,g,t} & \text{losses.}
\end{array}
$$

The PMAD contains the Army force structure requirements has time dimension equal to the projection horizon and the program will be run over increasing finite time horizon, from two to nine years. The PMAD is indexed by specialty and rank, $r \in \{\text{LT, CPT, MAJ, LTC, COL}\}$, which is mapped to grade, g, according to Figure 3.7. Losses, reflecting the decision of an officer to retire or resign their commission, are empirically estimated for the first two years of the projection horizon from 84 months of data, capturing losses from 2001-2007 by grade and specialty. The loss data was provided by PRS, and is graphically depicted in Figure 3.1. The model requires the PMAD to be input over the entire planning horizon, but has the flexibility to shape the force with reductions in all years for which estimated losses are not provided.

The decision variables are accessions, promotions, and reductions. These con-

Grade	Rank	
1..6	LT	All Lieutenants
7..13	CPT	Captains
14..20	MAJ	Majors
21..26	LTC	Lieutenant Colonels
27..29	COL	Colonels

Figure 3.7: Model Grade to PMAD Rank Mapping

trol variables are represented by

$A_{s,g,t}$	accessions
$R_{s,g,t}$	reductions
$P_{s,g,t}$	promotions.

Accessions are used both to indicate initial accessions at 2LT and any other gains or additions to $X_{s,g,t}$. Reductions are used to decrement any $X_{s,g,t}$, in addition to the losses entered as data, and the combination of an accession and a reduction, $A_{i,g,t}$ and $R_{j,g,t}$, represents a transfer from specialty $s = j$ to $s = i$ given grade g and time in service t. The 43 Army specialties incorporated in the model are listed in Figure 3.8.

The mathematical program is designed to minimize the absolute deviation from the PMAD and to place smoothing controls on the spectrum of manpower decisions. The objective function minimizes deviation from structure within specialty and minimizes the yearly changes in accessions, promotions and reductions. The mathematical program's objective function has three components, the first two

24	Telecom Systems Engineer
30	Information Operations
34	Strategic Intelligence
37	Psychological Operations
40	Space Operations
43	Human Resources
45	Comptroller
46	Public Affairs
47	Academy Professor
48	Foreign Area Officer
49	ORSA
50	Force Management
52	Nuclear and Counterproliferation
53	Information Systems Management
57	Simulation Operations
59	Strategic Plans and Policy
AC	Acquisition
AD	Air Defense Artillery
AG	Adjutant General
AR	Armor
AV	Aviation
CA	Civil Affairs
CM	Chemical
EN	Engineer
FA	Field Artillery
FI	Finance
IN	Infantry
MI	Military Intelligence
MP	Military Police
MS	Medical Service
OD	Ordinance
QM	Quartermaster
SC	Signal Corps
SF	Special Forces
TC	Transportation

Figure 3.8: Army Specialties

components are deviation from PMAD targets

$$\begin{aligned}
&\sum_{t,s}\left|\sum_{g\in\{1...6\}} X_{s,g,t} - D_{s,LT,t}\right| \\
&+\sum_{t,s}\left|\sum_{g\in\{7...13\}} X_{s,g,t} - D_{s,CPT,t}\right| \\
&+\sum_{t,s}\left|\sum_{g\in\{14...20\}} X_{s,g,t} - D_{s,MAJ,t}\right| \\
&+\sum_{t,s}\left|\sum_{g\in\{21...26\}} X_{s,g,t} - D_{s,LTC,t}\right| \\
&+\sum_{t,s}\left|\sum_{g\in\{27...29\}} X_{s,g,t} - D_{s,COL,t}\right| \qquad (3.1)
\end{aligned}$$

and smoothing constraints on manpower controls between adjacent year pairs,

$$\gamma\sum_{s,g,t}|A_{s,g,t} - A_{s,g,t+1}| + \lambda\sum_{s,g,t}|R_{s,g,t} - R_{s,g,t+1}| + \eta\sum_{s,g,t}|P_{s,g,t} - P_{s,g,t+1}|\,. \qquad (3.2)$$

With (3.1) and (3.2), a linear combination of operating strength deviation and the variability of accessions, promotions and reductions, the weighting parameters γ, λ and η can be adjusted to balance the contribution of each of the annual difference in each of the manpower controls. In addition to the smoothing constraints, we will minimize reductions, as reductions involve either reclassification of an officer, or firing of an officer. Thus the third component of the objective function minimizing reductions in each period is given by

$$\sum_{s,g,t} R_{s,g,t}. \qquad (3.3)$$

The network structure from Figure 3.6 results in flow constraints $\forall s \in$ Specialties across the projection horizon, $t \in \{1 \ldots T-1\}$,

$$
\begin{aligned}
X_{s,1,t+1} &= A_{s,1,t} \\
X_{s,g,t+1} &= X_{s,g-1,t} - L_{s,g-1,t} - R_{s,g-1,t} \qquad g \in \{2,5\} \\
X_{s,g,t+1} &= \sum_{j=g-1}^{g} \left(X_{s,j,t} + P_{s,j,t} - L_{s,j,t} - R_{s,j,t}\right) \qquad g \in \{3,6\} \\
X_{s,g,t+1} &= \sum_{j=g-2}^{g-1} P_{s,j,t} \qquad g \in \{4,7\} \\
X_{s,g,t+1} &= \sum_{j=g-3}^{g-1} P_{s,j,t} \qquad g \in \{14,21,27\} \\
X_{s,g,t+1} &= X_{s,g-1,t} - L_{s,g-1,t} - R_{s,g-1,t} + A_{s,g-1,t} \qquad g \in \{8,9,10,11,15,16,17,18,22,23,24\} \\
X_{s,g,t+1} &= X_{s,g-1,t} - P_{s,g-1,t} - L_{s,g-1,t} - R_{s,g-1,t} + A_{s,g-1,t} \qquad g \in \{12,19,25\} \\
X_{s,g,t+1} &= \sum_{j=g-1}^{g} \left(X_{s,j,t} - P_{s,j,t} - L_{s,j,t} - R_{s,j,t} + A_{s,j,t}\right) \qquad g \in \{13,20,26\} \\
X_{s,28,t+1} &= X_{s,27,t} - L_{s,27,t} - R_{s,27,t} + A_{s,27,t} \\
X_{s,29,t+1} &= \sum_{j=g-1}^{g} \left(X_{s,j,t} - L_{s,j,t} - R_{s,j,t} + A_{s,j,t}\right).
\end{aligned}
$$

The DOPMA constraints on field grade officer strengths are modeled as $\forall t \in \{1 \ldots T\}$

$$
\begin{gathered}
\sum_{s,g=27\ldots29} X_{s,g,t} \leq 0.05 \sum_{s,g} X_{s,g,t} \\
\sum_{s,g=21\ldots26} X_{s,g,t} \leq 0.11 \sum_{s,g} X_{s,g,t} \\
\sum_{s,g=14\ldots20} X_{s,g,t} \leq 0.24 \sum_{s,g} X_{s,g,t}
\end{gathered}
$$

which limit COL, LTC and MAJ to 5%, 11% and 24% of the total officer corps for each period in the projection. Transfers are modeled using an accession and a reduction pair at a given rank. All field grade accessions are transfers from another specialty, and thus provide an additional lower bound on reductions $\forall g \in \{14 \ldots 29\}$ and $t \in \{1 \ldots T\}$

$$\sum_s R_{s,g,t} \geq \sum_s A_{s,g,t}.$$

Promotion constraints are modeled to allocate upper bounds to early promotion states, and lower bounds to due-course promotion states. $\forall s \in$ Specialties across the projection horizon, $t \in \{1 \ldots T\}$,

$$\begin{aligned}
P_{s,g,t} &= 0 \qquad g \in \{1, 4, 14, 15, 16, 17, 21, 22, 23, 27, 28, 29\} \\
P_{s,2,t} &\geq 0.8X_{s,2,t} \\
P_{s,g,t} &\geq 0.7X_{s,g,t} \qquad g \in \{5, 12, 19\} \\
P_{s,25,t} &\geq 0.5X_{s,25,t} \\
P_{s,g,t} &\leq 0.1X_{s,g,t}, \qquad g \in \{11, 18, 24\} \qquad (3.4)
\end{aligned}$$

$\forall g \in$ Grades, and $\forall s \in$ Specialties across the projection horizon, $t \in \{1 \ldots T\}$,

$$P_{s,g,t} \leq X_{s,g,t}$$

requiring that the population must be larger than the number chosen for promotion.

The next several constraints are properly categorized as goals. Initial officer accessions at 2LT are floored by the constraint for $\forall t \in \{1 \ldots T\}$

$$\sum_{s} A_{s,1,t} \geq 4000$$

and distributed using a pair of bounds $\forall t \in \{1 \ldots T\}$

$$0.75X_{s,1,0} \leq A_{s,1,t} \leq 1.25X_{s,1,0}$$

which keep the annual accession mission for each of the branch within an acceptable band. Minimum fill constraints are also defined specialties and $t \in \{1 \ldots T\}$

$$\begin{aligned}
\sum_{g=1...6} X_{s,g,t} &\geq \alpha_{s,LT} D_{s,LT,t} \\
\sum_{g=7...13} X_{s,g,t} &\geq \alpha_{s,CPT} D_{s,CPT,t} \\
\sum_{g=14...20} X_{s,g,t} &\geq \alpha_{s,MAJ} D_{s,MAJ,t} \\
\sum_{g=21...26} X_{s,g,t} &\geq \alpha_{s,LTC} D_{s,LTC,t} \\
\sum_{g=27...29} X_{s,g,t} &\geq \alpha_{s,COL} D_{s,COL,t}
\end{aligned}$$

where α is a percentage fill requirement by rank and specialty. Losses are input data for the first two years of the model horizon. The following years the reductions variable allows the model to choose the optimal losses to achieve a steady state distribution. The floors for losses were taken from the empirical distribution and are applied $\forall t \in \{1 \ldots T\}$

$$
\begin{aligned}
\sum_{s,g=1...3} R_{s,g,t} + \sum_{s,g=1...3} L_{s,g,t} &\geq 87 \\
\sum_{s,g=4...6} R_{s,g,t} + \sum_{s,g=4...6} L_{s,g,t} &\geq 260 \\
\sum_{s,g=7...13} R_{s,g,t} + \sum_{s,g=7...13} L_{s,g,t} - \sum_{s,g=7...13} A_{s,g,t} &\geq 1462 \\
\sum_{s,g=14...20} R_{s,g,t} + \sum_{s,g=14...20} L_{s,g,t} - \sum_{s,g=14...20} A_{s,g,t} &\geq 466 \\
\sum_{s,g=21...26} R_{s,g,t} + \sum_{s,g=21...26} L_{s,g,t} - \sum_{s,g=21...26} A_{s,g,t} &\geq 652 \\
\sum_{s,g=27...29} R_{s,g,t} + \sum_{s,g=27...29} L_{s,g,t} - \sum_{s,g,27...29,t} A_{s,g,t} &\geq 368.
\end{aligned}
$$

The number of accessions adjust the total losses to account for transfers within specialties. The last goal ensures the overall percentage of the Army End Strength allowed in the officer corps, $\forall t \in \{1 \ldots T\}$

$$\sum_{s,g} X_{s,g,t} \leq 0.15(525400)$$

is kept below 15%.

3.5 Results and Programming Implications

The linear program is designed to produce an optimal force profile, subject to all the constraints, in accordance with the stated objective function. This solution is, however, a set of 29 x 43 dimensional objects, for each time period. The decision variables of accessions, reductions, and promotions each prescribe a operational plan for manpower planners. The operational implications of such a profile, and the acceptability of the solution, are dependent on the input data, as well as choices of weighting functions and prioritization.

Years	Objective Function Value	Time (minutes)	Constraints (#)	Variables (#)	Non-Zero Coef (#)
2	67950.2	00:13.4	7664	9842	42528
3	89062.2	00:21.2	10990	14258	66331
4	110174.2	00:44.9	14316	18674	90134
5	131286	01:24.1	17642	23090	113937
6	152398.2	02:17.1	20968	27506	137740
7	173510.2	05:57.2	24294	31922	161543
8	194622.2	29:27.6	27620	36338	185346
9	215878.95	35:18.9	30946	40754	209149

Table 3.1: OPL 6.1 / CPLEX 11.2 Run Data

The model was run over an increasing time horizon from two to nine years of projection. Optimal objective function values are presented in Table 3.1 along with model sizes and the running time on a Dell Optiplex GX 240 PC with a 1.70GHZ Pentium 4 and 512 MB RAM. The optimization model code is presented in Appendix B and was solved using CPLEX 11.2. The two-year projection solved in 13.39 seconds with 9842 variables and 7664 constraints, while the nine-year projection required 35 minutes and 18.9 seconds to solve the model with $40,754$ variables and $30,946$ constraints. Figure 3.9 indicates that the solution time grows exponentially, yet the growth in the objective function value, number of variables and constraints grow linearly with increasing time horizon.

The structure of this experiment results in eight solutions for the first projection year, decreasing to only one solution for the ninth projection year. The stability of the solution displayed with the increasing time horizon indicates limited

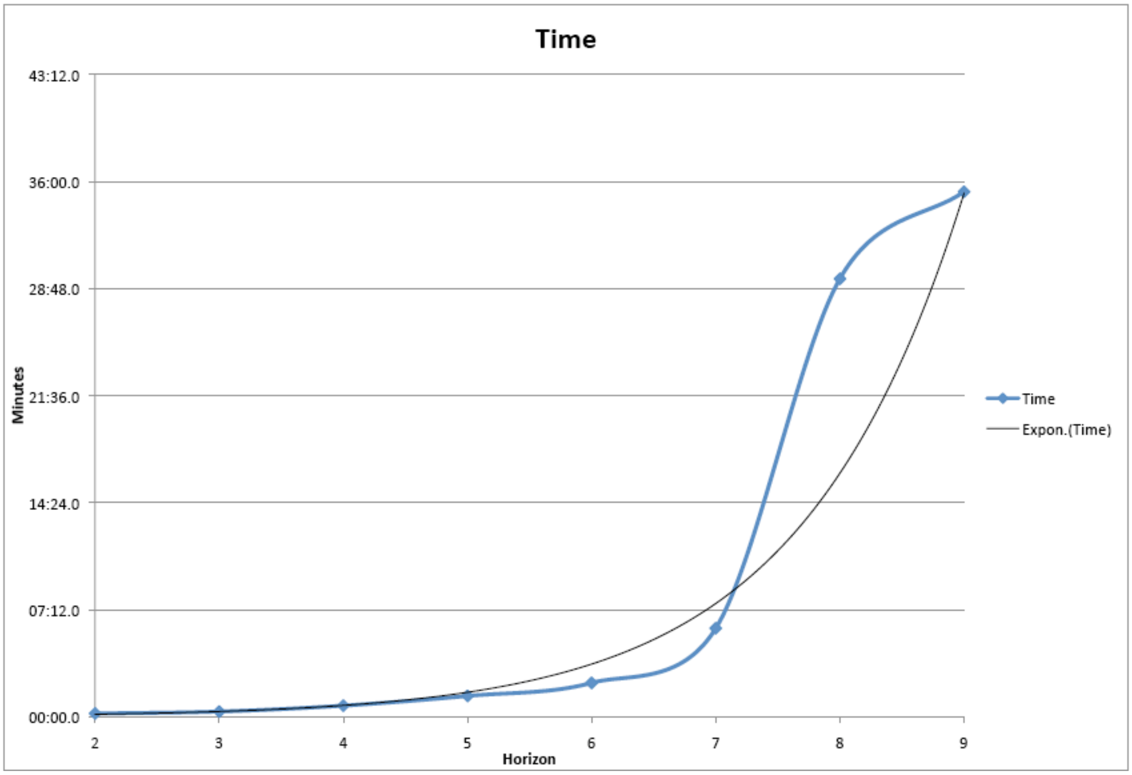

Figure 3.9: Solution Time vs. Projection Horizon

end effects as the modeling horizon is increased. The multi-objective function and smoothing constraints contribute to the stability of the model and will be further explored in the following sections.

Looking at the value of the optimal objective function in Figure 3.10, the intercept displays the cost of the misalignment of the force in the initial program and the linear increase highlights the tradeoff between the smoothing constraints and the deviation from structure over the increasing projection horizon.

Could the optimal objective function value be zero in steady state? If the objective function was comprised of only (3.1) and (3.2), a perfectly aligned force could have objective function value of zero. However, with the current formulation,

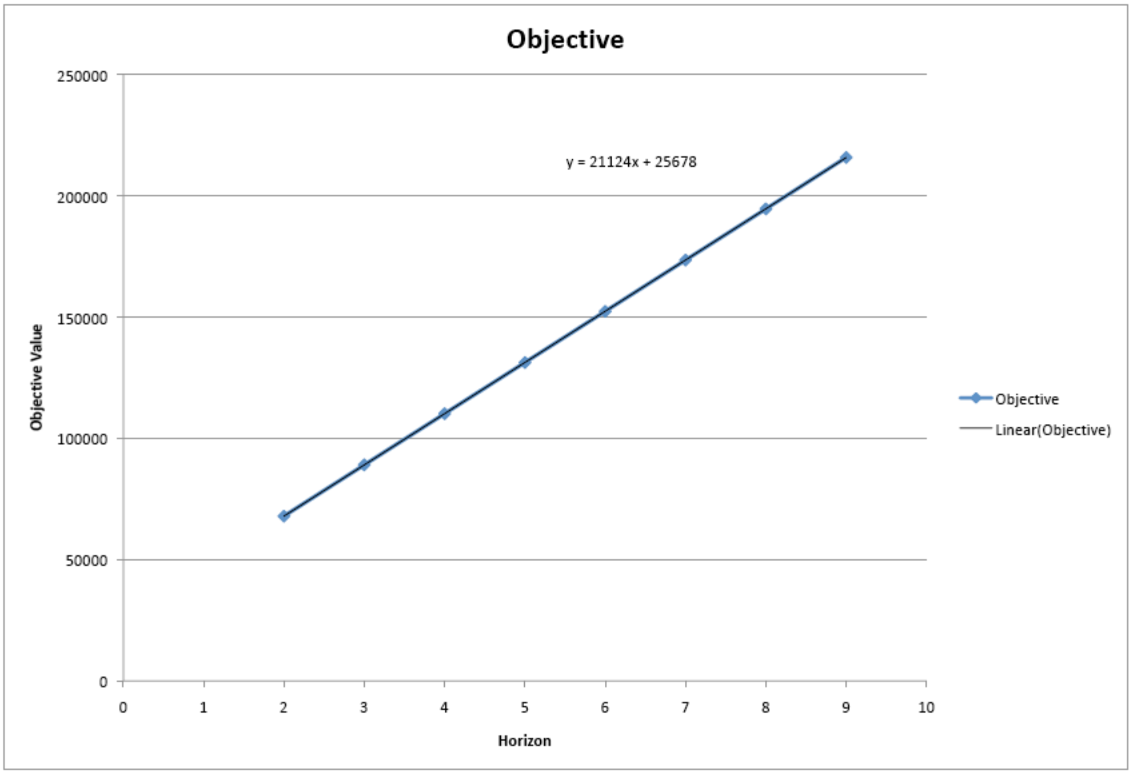

Figure 3.10: Objective Value vs. Projection Horizon

the inclusion of (3.3) ensures that the objective function will be positive. First, although reductions are minimized, there is a requirement for reductions as a part of the Career Field Designation (CFD), a process of redistribution of company grade officers into field grade specialties. This structural shift between CPT and MAJ will necessitate some transfers, which in this formulation are composed of an accession and a reduction. These logically should be loosely bounded from above by the number of accession at 2LT in steady-state, as the steady state solution cannot redistribute more MAJ each year than are accessed as 2LT.

Reductions are also floored at 3,295 and in steady state need to be at least as large as accessions. Thus, although losses are floored at 3,295, they should be expected to be greater than 4,000, based upon yearly accessions. If steady-state could be achieved with regard to perfect structure match, (3.1), and controls (3.2), the formulation would require 3,295 losses plus the transfers required for the CFD process, creating a feasible floor on the objective.

3.5.1 Changing Requirements

There are three general scenarios in which Army manpower planners find themselves: grow, maintain, or shrink the Army. With political pressure to reduce the U.S. budget deficit, additional pressure is applied by DOD to all the services, and the Army in particular, to reduce their budgets. The scenario described here is one of initial strength that is greater than requirements, in the aggregate. The first year accessions, i.e. initial accessions and transfers, in the two year model shown in Table 3.3, indicate an initial misalignment. With only a two-year horizon, the model has to use more force management controls to counter the large deviation between the larger officer inventory and smaller requirements. As the horizon lengthens, the solution becomes more stable, as evident from Table 3.2 and Table 3.3: the effect of the smoothing constraints are apparent in both accessions and promotions.

The values of the smoothing parameters determine the relative weighting between the portions of the objection function values. As all the elements of the objective function are in units of officers, we initially choose equal weighting. With equal weighting, for the 2-8 year projections, each year's promotions were identical as the model smoothed differences between annual promotions. In the 9 year

	2	3	4	5	6	7	8	9
1	18497	15972	15647	15539	15538.6	15539	15538.6	15482.4
2	18497	15972	15647	15539	15538.6	15539	15538.6	15393.8
3		15972	15647	15539	15538.6	15539	15538.6	15393.8
4			15647	15539	15538.6	15539	15538.6	15393.8
5				15539	15538.6	15539	15538.6	15393.8
6					15538.6	15539	15538.6	15393.8
7						15539	15538.6	15393.8
8							15538.6	15393.8
9								15393.8

Table 3.2: Forecasted Officer Personnel Promotions

	2	3	4	5	6	7	8	9
1	19311	7140.34	6713.4	6621.8	7273.83	7283.6	7696.75	7699.62
2	19311	7140.34	6713.4	6621.8	7273.83	7283.6	7696.75	7699.62
3		7140.34	6713.4	6621.8	7273.83	7283.6	7696.75	7699.62
4			6713.4	6621.8	7273.83	7283.6	7696.75	7699.62
5				6621.8	7273.83	7283.6	7696.75	7699.62
6					7273.83	7283.6	7696.75	7699.62
7						7283.6	7696.75	7699.62
8							7696.75	7699.62
9								7699.62

Table 3.3: Forecasted Officer Personnel Accessions

projection, only the first year differed. Current constraints only penalize the one year difference in promotions, and as such, the first year and the last year are only included in one difference, encouraging the model to make changes in either the first or last years of the projection. The agreement between year n, i.e. the last year of the projection, and $n-1$ at all time horizons is evidence of minimal end-effects in promotions as the model horizon is increased. Rounding to units of officers, yearly promotions do decrease with the increasing time horizon, indicating a move toward steady-state distribution that requires fewer promotions than required in the near term. This is consistent with the greater initial population than requirements.

Often in manpower optimization models the models solution will vary considerable between the n and $n-1$ year of projection. This has made truncation of the last several periods of a projection a popular method for addressing end effects. In our model, we do not observe that nature of end effects. Rather, as the projection horizon lengthens, we observe a change varying over the entire solution.

Values of several of the decision variables are presented in Tables 3.2-3.3. The years of the projection are annotated on the left and the number of projection years in the model run are recorded across the top of the table. The required officer accessions are shown in Table 3.3. When combined with the total losses to the system, $\sum L + \sum R - \sum A$, as shown in Table 3.4, they equate to the debits for each year of the projection. Traditionally, losses are viewed as the number of officers that exit the system. However, in this model we have both losses, statistically estimated, and reductions, a decision variable. The comparision of Table 3.3 and Table 3.4 show a steady state 2LT accession mission of $4,000$ officers and equal losses to the system. Table 3.2 displays promotions moving towards steady state. The stability of the number of promotions in each projection year is also an indicator of the effectiveness of the underlying network structure.

	2	3	4	5	6	7	8	9
1	21631	21630.7	21631	21631	21630.7	21631	21630.7	21658.8
2	21644	8796	8796	8796	8796	8796	8796	8796
3		4000	4000	4000	4000	4000	4000	4000
4			4000	4000	4000	4000	4000	4000
5				4000	4000	4000	4000	4000
6					4000	4000	4000	4000
7						4000	4000	4000
8							4000	4000
9								4000

Table 3.4: Forecasted Officer Personnel Debits

The accessions in Table 3.3 represent two components, transfers between specialties at higher grades, as well as initial 2LT accessions. Most of the accessions are actually transfers, transfers from the CFD process as well transfers as to correct structural imbalances. They are required by the lack of company grade positions in many of the specialties as well as changing structural distribution of requirements between branches. In this respect, even in a perfectly aligned force, there will be a requirement for transfers between the company and field grade levels, as well as the required accessions at 2LT.

3.5.2 Accessions

An annual debate centers around the number of accessions. If more new officers are accessed than available jobs , only a portion of the formative years will be spent in jobs directly relating to the officer's specialty. When only a portion of the force are able to have career enhancing jobs, or alternatively, each officer spends less of their time in career enhancing jobs, job satisfaction declines. In the Army, declining job satisfaction as a 1LT translates into higher CPT losses, as the initial four year service obligation requires the majority of officers to serve until after their promotion to CPT.

The distribution of accessions by branch is captured in Figure 3.11. Interestingly, although only one year in the 9 year projection had accession different from 4000, year two at 5742, the distribution over the branches has considerable variability. The losses in different year groups and sizes of previous accession cohort changes the distribution over the branches in the second and following years. The largest variability was in the largest combat arms branches, armor, infantry and field artillery. These are the core jobs of the Army, and extra accessions in these specialties are tied to higher field grade populations that can absorb variability.

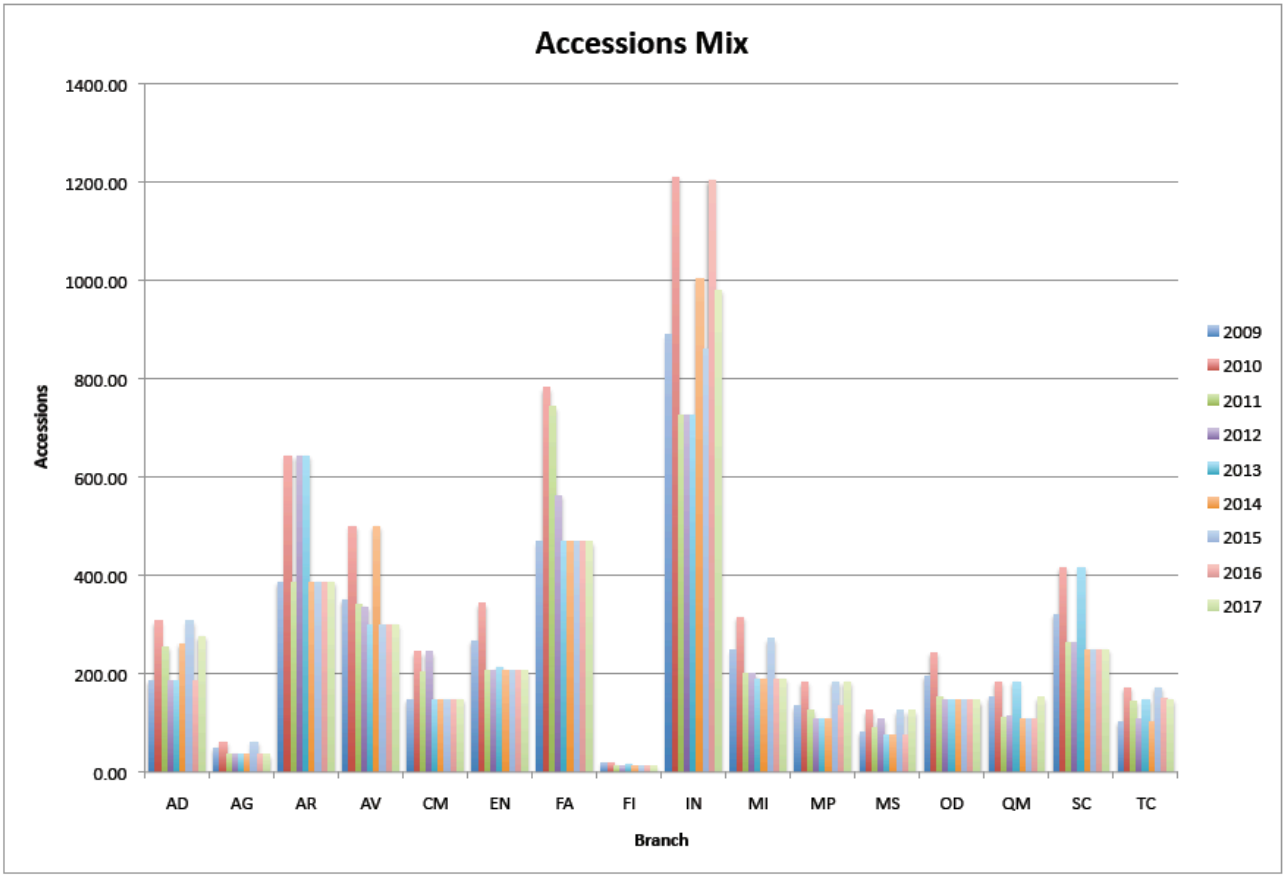

Figure 3.11: Accessions by Branch for Each Projection Year

Current accession planning is focused at 5,100 accession annually (Haight and Lewis 2008). However, our optimal force profile required only 4000 accessions in all but the second year of the projection. One of the motivations for over accessing is the shortages at higher ranks combined with the bar to lateral entry. However, rather than over-accessing, another option is to find a way to commission or access at the point where the officers are needed. If the shortage is at CPT or MAJ, staff positions could be identified that require more specialized non-combat skills and the Army could directly hire against those positions. This would require a change to the bar to lateral entry as well as cultural modification, but could ease shortages at field grade ranks as well as properly sizing the 2LT and 1LT populations for the structure.

Within the optimization, adjustment can be made to increase or decrease accessions. One method is to increase the floor on accessions. The other is to add a goal to the objective and penalize the difference from the goal. Another modification that could be made to the optimization is to add rank substitutability constraints, and in this way directly assess 2LT against field grade shortages. This however, would require increasing the specialties to which initial accessions are authorized. In this manner, if the Army were short Public Affairs Officers, a 2LT could be used to replace a CPT or MAJ. Although counter to current Army culture, an argument can be made for 2LT with advanced education or civilian experience to serve in staff positions alongside MAJs. The teamwork of experienced MAJs in the Army methods and traditions could be complemented by the specialized experience of highly educated 2LTs in some specialties, e.g. network engineering, public affairs officer, comptroller.

3.5.3 Losses

Optimization models can make use of the expected number of losses each year, as well as cohort sizes, to choose optimal values of all the decision variables simultaneously. Using loss numbers estimates in lieu of rates in the near term has the advantage of incorporating additional information on retirements or resignations that might be known to the manpower planner over a subset of the planning horizon. Applying such techniques over the entire planning horizon would be susceptible to small cell size problems and, using this approach, it would be increasing difficult to maintain model feasibility over a longer planning horizon.

One of the obvious insights of this manpower model is more insightful than at first glance. At steady state, losses must equal accessions. The interesting question is where the losses should occur. By allowing the manpower model to allocate losses past the first two years of the model, the optimal loss behavior is revealed. Another modeling technique, that of rate-based losses, results in proportional losses to cohort size. Although initially attractive, the assumption of proportional losses may not hold when specialties are under-strength, i.e. increased availability of career enhancing jobs, or over-strength. With our focus on finding an optimal force profile, rather than prescribing all losses, we instead allow the model to prescribe where losses should optimally be taken.

The choice to allow the model to determine losses was an attempt to reach a steady-state solution in the process of finding an optimal force profile. Over the horizon of the President's Budget, the first two years of the projection, the manpower

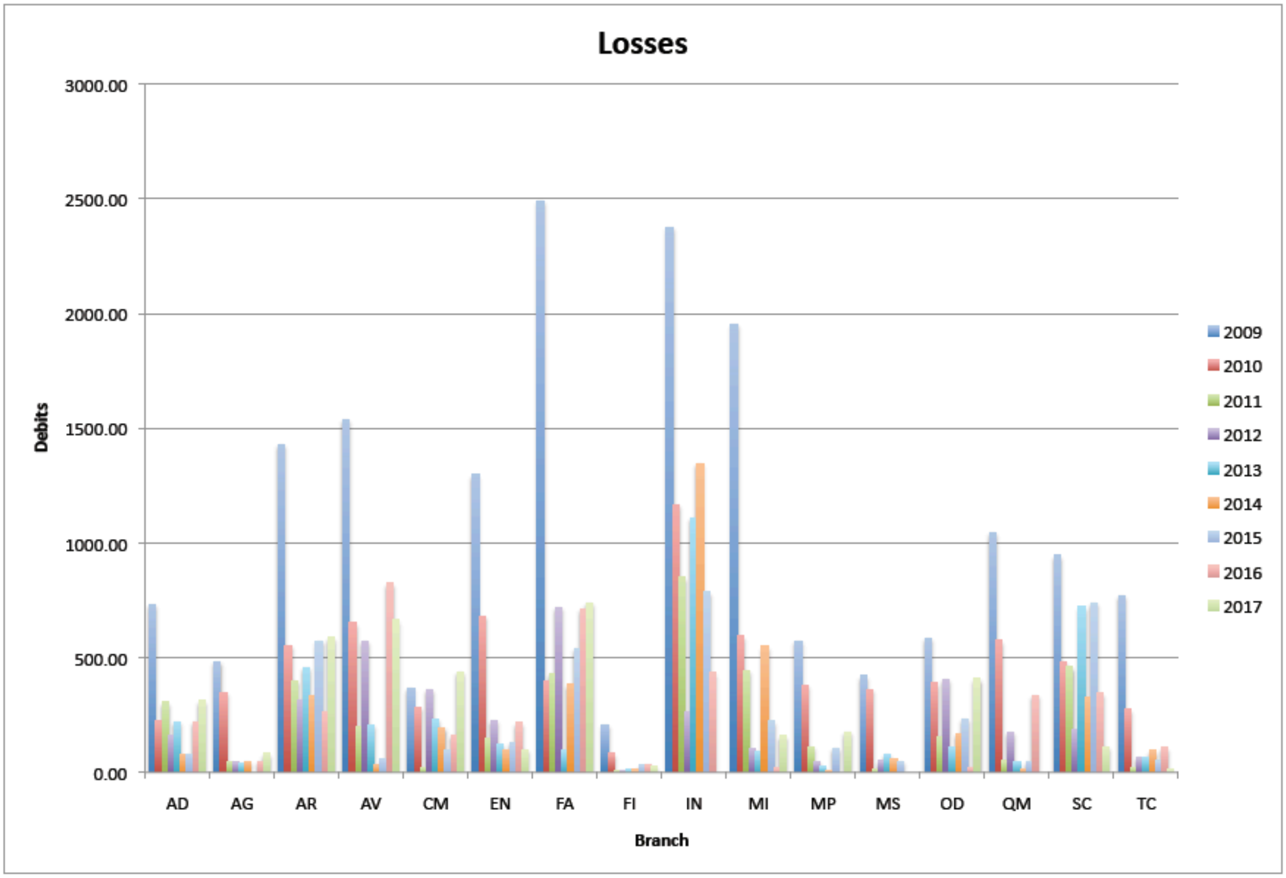

Figure 3.12: Branch Losses/Debits by Projection Year

planners have very few tools to change an officer's behavior or shift the behavior of an entire year-group. Following the same line of reasoning, over the entire POM, the Army has more control to influence loss behavior. The Army has just implemented a bonus system to affect just such a change in the entire cohort behavior for CPTs in the 1999-2004 cohorts (Haight and Lewis 2008). Additional policies can be invoked to exert more control over loss rates over a longer planning horizon which could be implemented in the final years of the POM.

Our focus on finding an optimal force profile also motivated the choice of modeling CFD as a combination of an accession and a reduction, allowing the model to prescribe where the officers should be taken from to fill the increasing number of specialties at the field grade level. Figure 3.4 illustrates both higher losses in the first year of the model due to the misalignment with structure, but also indicates the branches, e.g. infantry, field artillery, that must provide a large portion of the officers required for CFD process.

3.5.4 Promotions

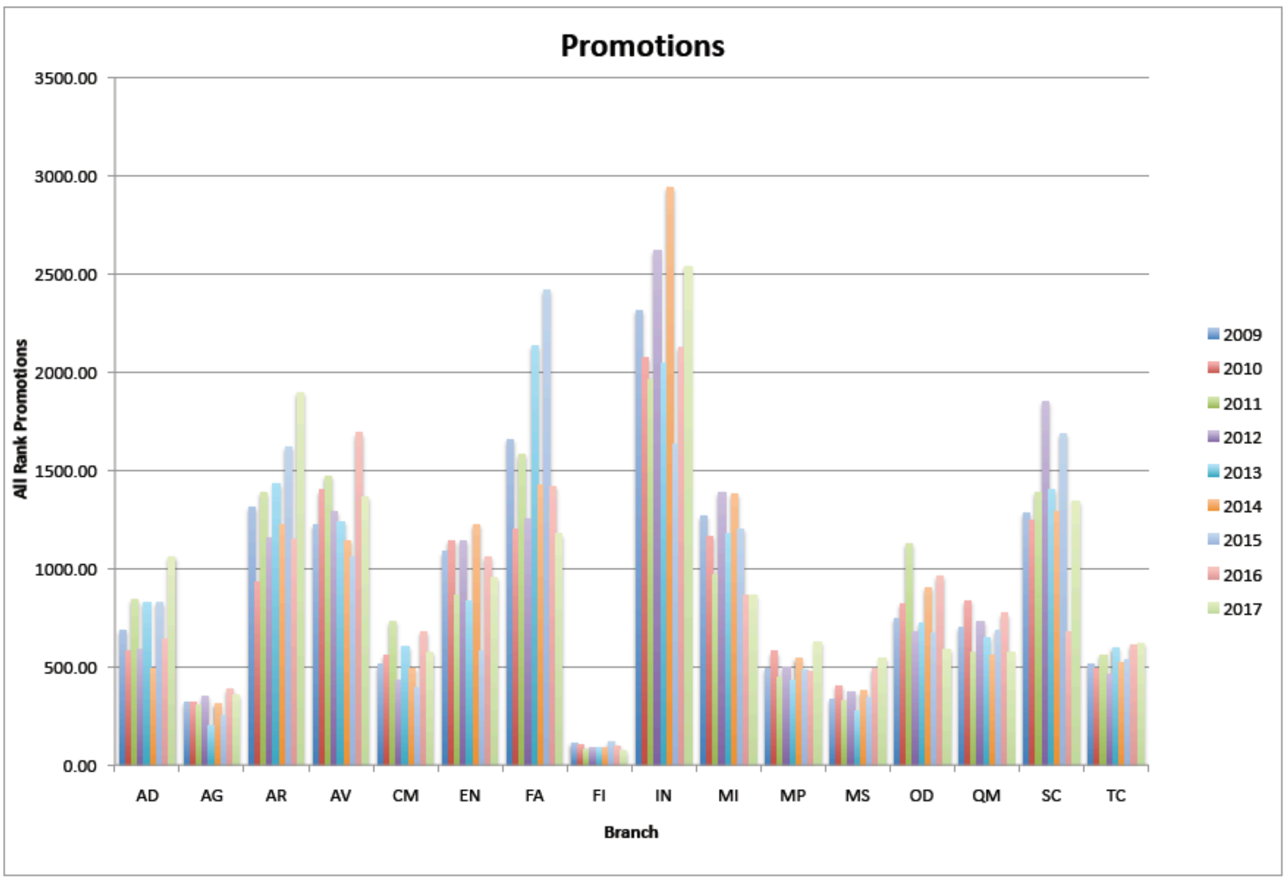

Figure 3.13: Branch Promotions by Projection Year

Model promotions provide a glimpse of the viability of the different career paths. Where Table 3.2 displays that the model finds steady-state promotions in the aggregate, Figure 3.13 displays the distribution of promotions across specialties. Increasing promotions in a specialty would indicate a specialty in which there is growing opportunities, where decreasing promotions would indicate the over saturation of the specialty. The distribution indicates that the optimization found a solution that has a healthy distribution of promotions across the initial entry branches.

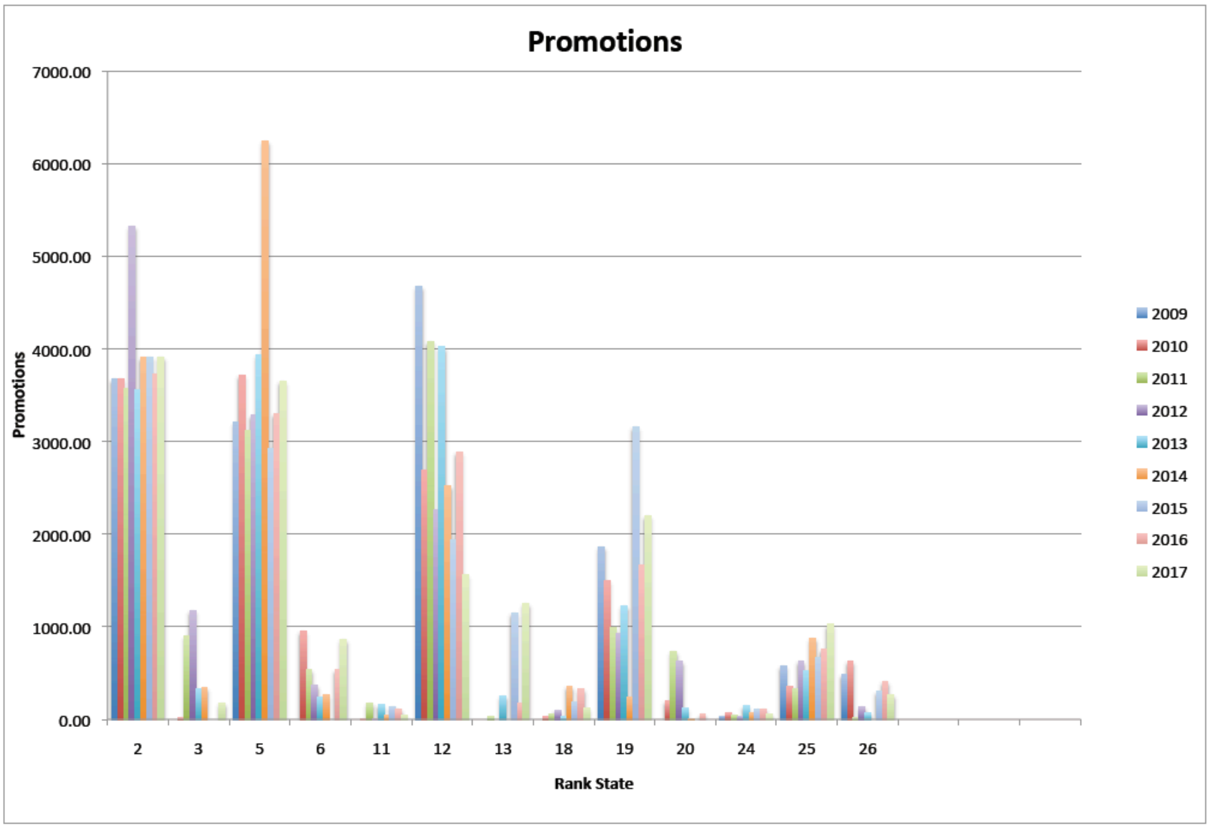

Figure 3.14: Rank State Promotions by Projection Year

Another approach to examining the promotion prescription is to look at the distribution across the states that are eligible for promotions. Figure 3.14 displays the opportunity for early, due-course and late promotions. The largest promotions are in the due-course promotions, i.e. states 2, 5,12, 19, 25, as expected. These states all have a proportionally defined lower bounds, (3.4), to ensure that due-course promotions meet the Congressional guidance within DOPMA. The promotions in late promotion states, especially 13, indicate where additional promotions could be shifted to early promotions. The early promotions in described in states 11, 18, and 24, have an upper bound in the formulation, while the model has the most flexibility in the late promotion states.

3.5.5 Force Profiles

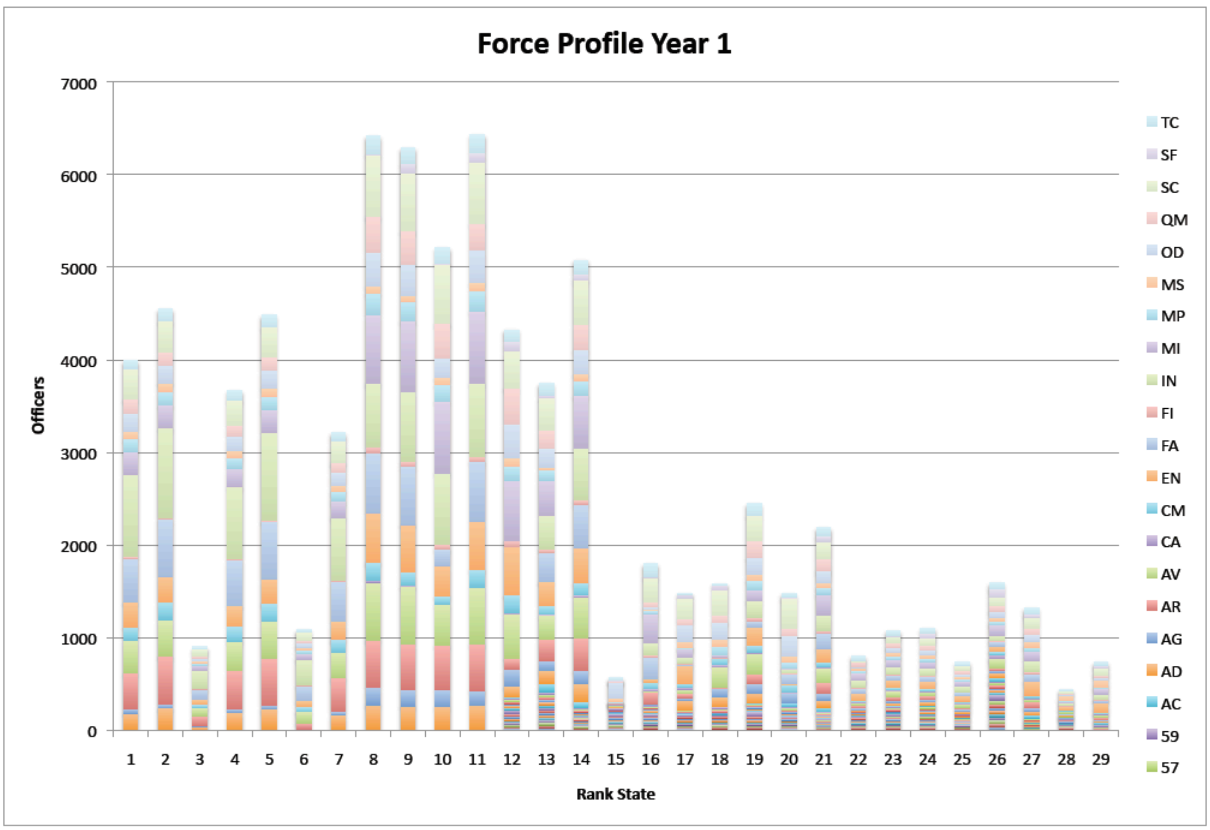

Figure 3.15: Force Profile for Projection Year 1

The force profile for the first, fifth and ninth year of the 9 year projection are displayed in Figures 3.15-3.17. The company grade component of the force has three semi-absorbing states, i.e. 3, 6, 13. Interestingly, both states 3 and 6, the last promotions to 1LT and CPT, tend towards zero. However, state 13, the CPT passed over for MAJ, grows throughout the horizon. This suggests that promotion rates to MAJ have considerable flexibility, due to the pool of available manpower in state 13, but also indicates that there is a significant demand at CPT, and as a result the model has not chosen eliminate these passed over officers though a reduction. This optimal company grade structure suggests opportunities for direct assignment at the company grade level to make up for the shortage in manpower that the model identifies across the model horizon which results in the growth in state 13. This interesting disconnect between requirements at the company grade and field grade structure is apparent in the optimal solution, but much more difficult to identify solely though structure documents.

At the field grade levels, both state 20 and state 26 decrease from higher populations in the initial population. The populations in the COL states are very stable, maintaining a stable distribution over states 27, 28 and 29. It is an interesting result from this model that the proposed structure is stable across all ranks, with the exception of CPT. The middle grades are always the most challenging for the Army, and the model clearly indicates a need for additional manpower. The possible alternative of direct assignments at CPT could smooth the flow between 1LT and MAJ and appears a viable alternative.

After examining the characteristics of these optimal force profiles, constraints can be updated, additional constraints can be added, and objective function weightings revisited. One suggestion that is not supported by this research, is the need

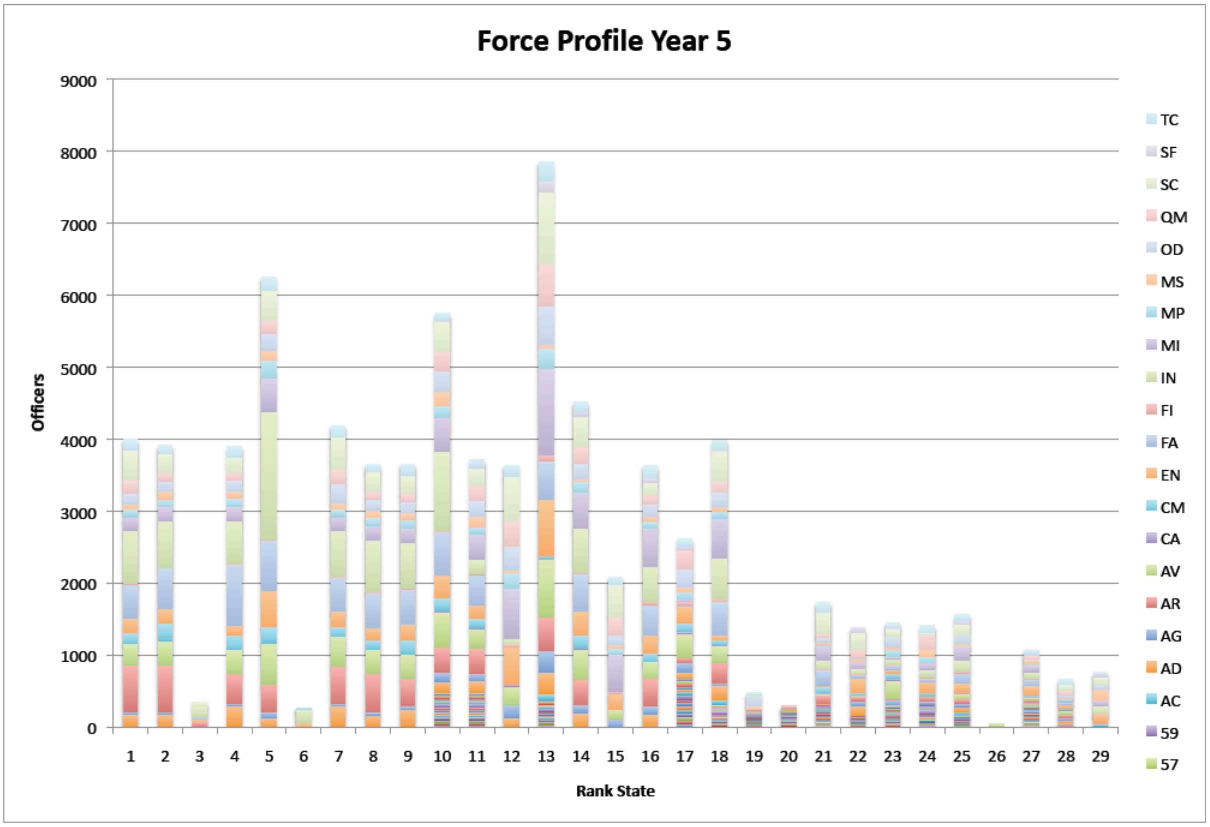

Figure 3.16: Force Profile for Projection Year 5

for longer field grade careers for Army officers. This would be characterized in the model as increasing populations at states 20, 26 and 29, which is not evident in the model. Alternatively, the research supports offering incentives to ensure that CPTs continue to serve only to the promotion window for MAJ. The Army has considered bonuses to get officers to stay until they reach MAJ. This research suggests that an optimal Army force profile contains a large number of officers passed over for MAJ. This could be achieved by modifying the retirement vesting policies to incentivize a higher portion of CPTs to stay in the Army until the 10-year point and increasing the competition for promotion to MAJ.

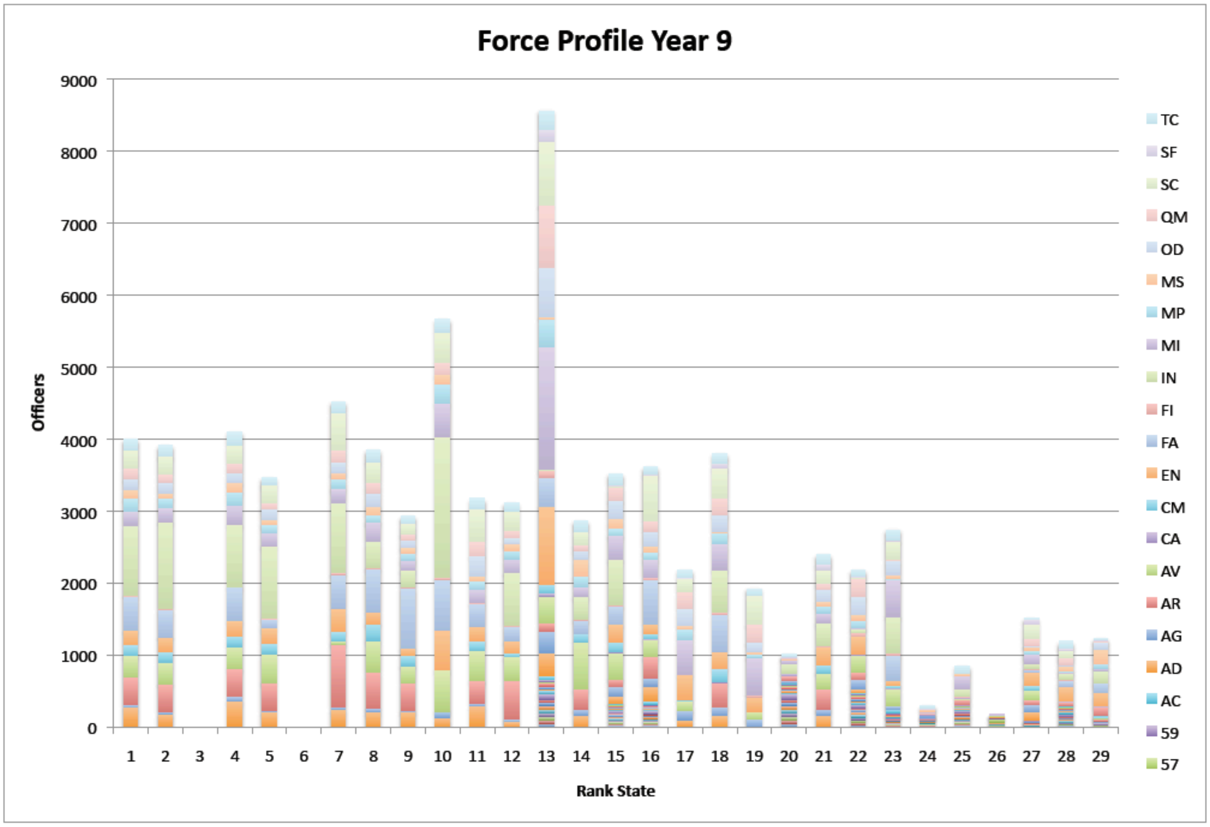

Figure 3.17: Force Profile for Projection Year 9

3.6 Future Work

The purpose of this research was to create a model to explore the optimal force profile for the U.S. Army. As such, the solution should approach a steady-state distribution and allow enough fidelity that modification to the optimal policy can be achieved by modifying weighting on the objective function and constraints, or by shifting goals from constraints to the objective function.

The process of shifting the employment of this model from a tool to suggest an optimal force profile to a forecasting tool is one of adjusting weighting and objectives until manpower analysts are satisfied with the forecasts of the model. The model is built with the flexibility to change deviations from structure by rank and specialty, or by the 29-state grade and specialty, as well as changing the weighting of each of the force management levers. In addition, constraints, e.g. DOPMA or total force officer percentage, can be easily shifted into the objective function to transition from constraints to goals. This model can be extended by changing the projection horizon from annual to either quarterly or monthly projections by applying a distribution to the annual accession mission, and estimating losses over the shortened time step.

The model also provides a formulation stable enough to be modified for use in answering policy questions involving additional stochastic elements. Alternative force structures can be compared as well as different potential loss profiles to identify strengths and weaknesses of potential courses of action and solutions that are optimal over a distribution of possible future structures.

One possible extension of this manpower model is a principle agent model, making use of a model of individual behavior and this manpower model to explore differences in ability and natural sorting. The common network structure upon which the models in this chapter and the previous chapter are built would facilitate integration and provide possible insight into potential distributions of ability at the field grade ranks based upon different promotion, signaling and compensation policies. Several potential policies changes could be explored to ensure that the aggregate optimal personal decisions agree with a system optimal policy.

This is the first linear programming model for Army manpower to incorporate specialty over the entire planning horizon. We were able to achieve this though a new network structure that focused on time in grade rather than modeling time in service as has been previously implemented in models focusing on enlisted manpower. The limited end effects and stability of this model make it an excellent candidate for extension and inclusion into the suite of models employed by the Army G1.

Chapter 4

Gradient Estimation and Mountain Range Options

4.1 Introduction

Simulation and gradient estimation research have found a fertile ground in financial engineering applications. Often the pay-out function for complex securities and derivatives do not permit a closed form analytical solution, and therefore must be priced by Monte Carlo simulation or other numerical methods.

This chapter focuses on developing gradient estimates for a relatively recent form of options known as mountain range products. These products were first introduced by Société Générale in 1998 and have a mixture of the characteristics of basket options, barrier options and other path dependent options (Quessette 2002). We will first look at pricing several mountain range products and then turn to gradient estimation in order to complete calculation of the Greeks, whose derivation can be critical to hedging these products.

As we calculate the Greeks, we will examine indirect gradient estimation, pathwise or Infinitesimal Perturbation Analysis (IPA), and Likelihood Ratio/ Score Function (LR/SF), to include evaluating the strengths and weaknesses of each method. We will look to prior work on calculating the Greeks for Asian options, barrier options, and basket options for insight into efficient gradient estimation for mountain range products.

We find that IPA and LR/SF methods are efficient methods of gradient estimation for Mountain Range products at a considerably reduced computation cost compared with the commonly used indirect methods, e.g. finite differences or bumping, which provide biased estimators and are less effective at estimating the derivatives for these products. We also find that the most commonly calculated Greek, delta, has value zero by construction in models of Mountain Range options where the stock process is driven by multiplicative distributions.

We will proceed by first introducing gradient estimation in a stochastic simulation setting, and provide a worked out example for European call options. Section 4.3 will explain the mountain range products and we then develop gradient estimators for Everest and Atlas options in Section 4.4. We close with a numerical example in Section 4.5 and results in Section 4.6.

4.2 Gradient Estimation

The price of each the mountain range options is dependent upon the various parameters of the model. Calculating the derivatives of each function with respect to the parameters of interest will result in better understanding of how sensitive each option is to system changes, and will be accomplished by gradient estimation.

Fu and Hu (1995) first applied gradient estimation techniques to option pricing focusing on IPA estimation for both European and American options. Broadie and Glasserman (1996) applied both IPA and LR/SF methods to European and Asian options. Within financial engineering, Glasserman (2004) has become the standard reference. For broader application, Fu (2006) provides an detailed survey of gradient estimation within stochastic simulation and explains how the direct techniques are much more efficient than "brute-force" finite difference methods. Fu (2008) provides a more gentle introduction with a focus on application to option pricing.

We begin with $J(\theta)$, a performance measure dependent upon θ, and desire to calculate

$$\frac{dJ(\theta)}{d\theta}.$$

Often θ appears in the performance measure, but is may also be introduced though the random variables of the system. Assuming we lack detailed information concerning the distribution of J, we turn to simulation as an effective technique to estimate the distribution of $J(\theta)$. We carry out a large number of independent trials and calculate the sample performance L, applying the law of large numbers to justify using the sample mean to estimate

$$J(\theta) = E[L(\theta)] = E[L(X_1, X_2, ..., X_n)] \tag{4.1}$$

where $X = X_1, X_2, ..., X_n$ are dependent on θ. Expanding the expected value of our sample performance,

$$E[L(X)] = \int y dF_L(y) = \int L(x) dF_X(x), \tag{4.2}$$

where F_L is the distribution of L and F_X is the distribution of the input random variables X, can assist in illuminating the importance of the different forms of θ dependence.

Looking to make sense of the right hand side of (4.2),

$$E[L(X)] = \int_0^1 L(X(\theta; u))du \tag{4.3}$$

$$E[L(X)] = \int_{-\infty}^{\infty} L(x)f(x;\theta)dx \tag{4.4}$$

where f is the probability density function (p.d.f.) of X, provide two different integral representations, highlighting a structural difference in the forms of θ dependence. The θ dependance can be path-wise from the input random variables, as shown in (4.3), or in the distribution or measure of the input random variable F_X, as in (4.4).

Gradient estimation seeks to estimate

$$\frac{dE[L(\theta)]}{d\theta}$$

and we will describe indirect methods before addressing direct methods: IPA and LR/SF.

4.2.1 Indirect Methods

Indirect methods can can be used for sensitivity analysis or simulation optimization. An intuitive approach corresponds to the notion of the derivative and changing the value of the parameter "a little" and observing the change in the performance measure. A one-sided forward difference is the most straightforward implementation, but can have a large variance and results in a biased estimator. With e_i, the unit vector in the ith direction, and c_i, the scalar perturbation in the ith direction,

$$\frac{\hat{J}(\theta + c_i e_i) - \hat{J}(\theta)}{c_i} \tag{4.5}$$

expresses a forward difference gradient estimator. Choosing a very small c_i can result in noisy estimators, but larger values results in increasing bias of the estimator, which illustrates the trade off between variance and bias motivating the a priori choice of c_i. This method is referred to as bumping a parameter in finance, e.g. bumping the spot. The forward difference estimator requires a second simulation of the same dimension as the initial simulation.

A slightly more sophisticated approach is the two-sided symmetric difference

$$\frac{\hat{J}(\theta + c_i e_i) - \hat{J}(\theta - c_i e_i)}{2c_i}, \tag{4.6}$$

which has increased accuracy over (4.5), making use of a central difference approximation. A two-sided estimator requires two additional simulations whereas (4.5) requires only one additional simulation: the increased accuracy comes at the price of increased computational burden.

An additional gradient estimation technique designed for stochastic approximation is the simultaneous perturbation estimator. Where the first two methods are well suited for sensitivity analysis, the simultaneous perturbation estimator is especially well suited to simulation optimization, such as making small changes to several positions simultaneously, e.g. in an investment portfolio optimization using simultaneous perturbation stochastic approximation. Letting Δ be a d-dimensional vector of perturbations, the simultaneous perturbations estimator is

$$\frac{\hat{J}(\theta + c\Delta) - \hat{J}(\theta - c\Delta)}{2c_i\Delta_i}. \tag{4.7}$$

A slight modification, the random directions gradient estimator, shifts Δ into the numerator and results in a estimator,

$$\frac{(\hat{J}(\theta + c\Delta) - \hat{J}(\theta - c\Delta))\Delta_i}{2c_i}, \tag{4.8}$$

which allows the use of the normal distribution for the perturbation sequence (Fu 2006).

Where (4.5) and (4.6) are well suited to sensitivity analysis, analyzing the change to the performance measure when only one item has changed, the simultaneous changes in (4.7) and (4.8) help speed convergence in simulation optimization.

4.2.2 Direct Methods

Where all indirect methods require additional simulations to estimate the gradient, direct methods exploit the structure of the performance measure and stochastics of the system to calculate the derivative estimate directly during the initial simulation. Returning to (4.1) we examine the dependance on the parameter θ and categorize the performance measure as either having either sample (4.9) or measure dependency (4.10),

$$\begin{aligned} \frac{dE[L(X)]}{d\theta} &= \int_0^1 \frac{dL(X(\theta;u))}{d\theta} du && (4.9) \\ &= \int_{-\infty}^{\infty} L(x) \frac{df(x;\theta)}{d\theta} dx. && (4.10) \end{aligned}$$

When the θ dependence occurs in the input random variable as seen in (4.3), pathwise estimates can be employed. When the θ dependence occurs in the distribution (measure) of the input random variable F_X as in (4.4), the likelihood ratio method is appropriate.

4.2.2.1 Pathwise

Pathwise estimates require integrability conditions which are easily satisfied when the performance measure is continuous with respect to the given parameter. Two different notions of derivatives will be necessary for implementing direct gradient estimates. The notion of the derivative required for pathwise gradient estimation,

$$\frac{dE[L(X)]}{d\theta} = \int_0^1 \frac{dL(X(\theta;u))}{d\theta} du,$$

is one where the derivative of a random variable defined by

$$\frac{dX(\theta,\omega)}{d\theta} = \lim_{\Delta\theta \to 0} \frac{X(\theta+\Delta\theta,\omega) - X(\theta,\omega)}{\Delta\theta} \text{ w.p.1.}$$

The family of random variables parameterized by θ are defined on a common probability space such that $X(\theta) \sim F(\cdot;\theta)$ s.t. $\forall\theta \in \Theta, X(\theta)$ is differentiable w.p.1.

Assuming the interchange of differentiation and expectation is permissible, by application of the chain rule we obtain

$$\frac{dE[L(X)]}{d\theta} = \int_0^1 \frac{dL(X(\theta;u))}{d\theta} du = \int_0^1 \frac{dL}{dX}\frac{dX(\theta)}{d\theta} du$$

from which the estimator is

$$\frac{dL}{dX}\frac{dX(\theta)}{d\theta}. \tag{4.11}$$

The condition for applicability of the pathwise estimator is the uniform integrability of $\frac{dL}{dX}\frac{dX(\theta)}{d\theta}$, which can be asserted by the Lebesgue Dominated Convergence Theorem.

4.2.2.2 Likelihood Ratio

The LR/SF method makes use of distributional differentiation.

$$\frac{dE[L(X)]}{d\theta} = \int_{-\infty}^{\infty} L(x)\frac{df(x;\theta)}{d\theta} dx \tag{4.12}$$

$$= \int_{-\infty}^{\infty} L(x) \frac{d \ln f(x;\theta)}{d\theta} f(x) dx$$

and the estimator is

$$L(x) \frac{d \ln f(x;\theta)}{d\theta},$$

where $\frac{d \ln f(x;\theta)}{d\theta}$ is the score function of statistics.

4.2.3 European Call Example

We calculate the sensitivities for European Call options to changes in spot (delta) and volatility (vega) as an illustrative example. The payoff of the European Call struck at strike K which expires at T, with a risk free interest rate of r is

$$V_T(K) = e^{-rT}(S_T - K)^+ \tag{4.13}$$

where $(\cdot)^+$ represents $\max(\cdot, 0)$. This is the standard hockey stick payoff which is always continuous with respect to the value of S_T, the value of the stock at expiry.

A standard Brownian motion is a stochastic process, W_t, where $W_0 = 0$ and $t \to W_t$ is continuous on $[0, T]$ w.p.1 and has independent increments $W_t - W_s$ which are distributed $N(0, t - s)$ or equivalently $W_t - W_s$ has mean 0 and variance $t - s$. A stochastic process S_t is a geometric Brownian motion if $\log(S_t)$ is a Brownian motion with initial value $\log S_0$. In the Black-Scholes-Merton model, the price of the underlying stock, S_t, follows geometric Brownian motion and therefore has a lognormal distribution for any fixed t.

We specify a geometric Brownian motion as

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t$$

with W_t, a standard Brownian motion. With a dividend yield of δ, spot (initial or current stock price) S_0, volatility σ, and drift μ replaced by $r-\delta$,

$$S_t = S_0 e^{(r-\delta-\sigma^2/2)t+\sigma W_t}$$

represents the risk-neutral dynamics of the stock price. Letting Z represent a standard normal random variable, this can be simulated though

$$S_T = S_0 e^{(r-\delta-\sigma^2/2)T+\sigma\sqrt{T}Z}. \tag{4.14}$$

The price of the Europen Call, $C_T(K)$, is the expected value of the discounted payoff under the risk-neutral measure, Q,

$$C_T(K) = E^Q[V_T(K)] = e^{-rT}E^Q[(S_T-K)^+]. \tag{4.15}$$

The price of a European Call has a closed-form solution, which is presented along with Monte Carlo price estimates from $10,000$ sample paths in Table 4.1, for a stock with $S_0 = 90, 100$ and 110, $K = 100$, $r-\delta = 0.07$, volatility, $\sigma = 0.25$, and $T = 0.2$. Calculation of the standard errors, as reported in Table 4.1 requires additional code and is contained in Appendix B.

4.2.3.1 Pathwise

We will first illustrate calculating vega, $\frac{dV_T}{d\sigma}$, by the pathwise method of gradient estimation. From (4.11), the IPA estimator for vega is

$$\frac{dV_T}{d\sigma} = \frac{dV_T}{dS_T}\frac{dS_T}{d\sigma}.$$

From the stock dynamics in (4.14), we calculate

$$\begin{aligned} \frac{dS_T}{d\sigma} &= S_T(-\sigma T + \sqrt{T}Z) \\ &= \frac{S_T}{\sigma}\left(\ln\left(\frac{S_T}{S_0}\right) - \left(r - \delta + \frac{1}{2}\sigma^2\right)T\right). \end{aligned} \tag{4.16}$$

Next we evaluate $\frac{dV_T}{dS_T}$, by taking the derivative of (4.13) with respect to S_T. As (4.13) is equal to zero when $S_T \leq K$, V_T is only sensitive to changing values of S_T when $S_T \geq K$, which can be captured by

$$\frac{dV_T}{dS_T} = e^{-rT}1_{S_T \geq K}, \tag{4.17}$$

where $1_{\{\}}$ is the set indicator function. Substituting (4.16) and (4.17) result in the IPA estimator

$$e^{-rT}\left[\frac{S_T}{\sigma}\left(\ln\left(\frac{S_T}{S_0}\right) - \left(r - \delta + \frac{1}{2}\sigma^2\right)T\right)\right]1_{S_T \geq K}$$

for vega.

Returning again to (4.11), the IPA estimator for delta is

$$\frac{dV_T}{dS_0} = \frac{dV_T}{dS_T}\frac{dS_T}{dS_0}.$$

Using (4.17) and calculating $\frac{dS_T}{dS_0}$ results in

$$e^{-rT}\frac{S_T}{S_0}1_{S_T \geq K}.$$

4.2.3.2 Likelihood Ratio

The LR/SF method requires the differentiation of the distribution of the stock path. The first step to implementing this method is finding the distribution of X, the stock path. To find the distribution, we express the stock path from (4.14) as a function of a normal random variable,

$$S_T = g(Z),$$

calculate the inverse

$$g^{-1}(x) = \frac{1}{\sigma\sqrt{T}}\left(\ln\frac{x}{S_0} - \left(r - \delta - \frac{\sigma^2}{2}\right)T\right)$$

and calculate the Jacobian

$$\left|\frac{dg^{-1}(x)}{dx}\right| = \frac{1}{x\sigma\sqrt{T}}.$$

As Z is a standard normal random variable, $X \sim f(x)$ with

$$\begin{aligned} f(x) &= \left|\frac{dg^{-1}(x)}{dx}\right| N(g^{-1}(x)) \\ &= \frac{1}{x\sigma\sqrt{2\pi T}} exp\left(-\frac{1}{2}\left[\frac{1}{\sigma\sqrt{T}}\left(\ln\frac{x}{S_0} - \left(r - \delta - \frac{\sigma^2}{2}\right)T\right)\right]^2\right). \end{aligned} \tag{4.18}$$

Combining (4.15) and the distribution of the stock process, we can write

$$E[V_T] = \int_0^\infty e^{-rT}(x-K)^+ f(x)dx.$$

To find the LR/SF estimator for vega, assuming the order of integration and expectation can be interchanged, we obtain

$$\frac{dE[V_T]}{d\sigma} = \int_0^\infty e^{-rT}(x-K)^+ \frac{df(x)}{d\sigma}dx.$$

Simplifying using the identity $\frac{df}{d\sigma}/f = \frac{d\ln f}{d\sigma}$

$$\frac{dE[V_T]}{d\sigma} = \int_0^\infty e^{-rT}(x-K)^+ \frac{d\ln f(x)}{d\sigma} f(x)dx,$$

and we find the LR/SF estimator for vega to be

$$e^{-rT}(x-K)^{+}\frac{d\ln f(x)}{d\sigma}. \tag{4.19}$$

Using the distribution (4.18), the score function of (4.19) can be implemented in simulation by

$$\frac{T\left(4r^2T-8rT\delta+4T\delta^2-4\sigma^2-T\sigma^4\right)+8T(-r+\delta)\ln\left[\frac{x}{S_0}\right]+4\ln\left[\frac{x}{S_0}\right]^2}{4T\sigma^3}$$

or by substituting $\hat{r}=r-\delta$

$$\frac{4\hat{r}^2T-\sigma^2\left(4+T\sigma^2\right)}{4\sigma^3}-\frac{2\hat{r}\ln\left[\frac{x}{S_0}\right]}{\sigma^3}+\frac{\ln\left[\frac{x}{S_0}\right]^2}{T\sigma^3}.$$

To find the LR/SF estimator for delta changes the score function from (4.19) to

$$e^{-rT}(x-K)^{+}\frac{d\ln f(x)}{dS_0}$$

for which the score function can be implemented as

$$\frac{\ln\left[\frac{x}{S_0}\right]-T\left(r-\delta-\frac{\sigma^2}{2}\right)}{S_0T\sigma^2}.$$

4.2.4 Numerical Experiment

From the results in Table 4.1, we observe that the IPA estimate is closest to the true value of delta and vega at each spot. We also observe that the bump

Spot	90	(Std Err)	100	(Std Err)	110	(Std Err)
Call Price	1.2198		5.1259		12.3271	
Monte Carlo	1.241056	0.035826	5.144484	0.074739	12.379637	0.109306
Delta	0.2219		0.5684		0.8443	
Bump	0.22375	0.004601	0.573811	0.005436	0.847988	0.004056
Pathwise	0.22328	0.007846	0.573682	0.004407	0.846242	0.001895
LR/SF	0.228012	0.008227	0.574714	0.012631	0.838723	0.013245
Vega	11.946		17.446		11.4347	
Bump	13.087574	0.294961	20.36339	0.320319	16.156308	0.406102
Pathwise	12.082765	0.275682	17.49492	0.302503	11.492773	0.398022
LR/SF	12.52041	0.707849	18.32197	1.144887	12.463275	1.625017

Table 4.1: European Call Options Gradient Estimation

estimate is closer to the true value for delta than LR/SF, but LR/SF is closer to the true value for vega. The IPA has smaller standard errors than the LR/SF for each estimate. The effectiveness of the bump method for calculating delta might suggest the extra work to calculate direct gradient estimators is unnecessary, but the method is clearly less effective at calculating vega. We will show that it is also much less efficient at estimating the gradients for mountain range options.

4.3 Mountain Range Options

Very little has been written about mountain range options. Quessette (2002) describes several new exotic derivatives to include the mountain range options. He discusses the additional risks involved in products that are traded over-the-counter as well as the importance of understanding the model sensitivities and importance of model calibration. Vega stability, sensitivity of price to the volatility of the stock, as well as forward skew and multi-underlying skew issues are discussed, but computational details are not provided. This discussion of different modeling techniques resulting in different prices highlights the model dependence in exotic derivatives. Overhaus (2002) details the Himalaya options, explaining many of the different potential payoff functions and highlighting the main difficulty in pricing such products: correlation of the stocks within the basket. A basket is the term used to annotate a set of stocks upon which an option is based. By grouping stocks into a basket, it is possible to aggregate performance of multiple assets while reducing transaction costs. Overhaus (2002) suggests finite difference methods for calculation of delta, sensitivity to spot, but acknowledges that discontinuities occurring as stocks exit the basket will be problematic. He suggests measure changes can be used to develop more robust gradient estimates, but does not provide details. Meaney (2007) expands the treatment of Himalayan options by looking at volatility smile. He proposes a mixture of densities technique to gain greater fidelity in the volatility modeling of Himalayan options, providing a two-stock example.

Several of the mountain range products have path dependencies, e.g. Altiplano and Himalayan, while all are variations on basket options written on the level of the returns. As each of the following products are traded over-the-counter, each payoff must be modeled to match the individual contract that has been written, but the following payoff formulations capture the general flavor of each option variety.

4.3.1 Everest

Mount Everest is the highest point on earth and in the Himalayan Mountain range. Curiously, the Everest option is the pay-out on the worst performer in a basket, normally of 10-25 stocks, with 10-15 year maturity (Quessette 2002).

Given n stocks $S_1, S_2, \ldots, S_n$ in a basket, the payoff for an Everest options is

$$\min_{i=1\ldots n}\left(\frac{S_i^T}{S_i^0}\right). \tag{4.20}$$

4.3.2 Atlas

Atlas was a figure in Greek mythology that supported the world on his back. The Atlas option is a call on the mean of a basket of stocks, a basket that has some of the best and worst performers removed (Quessette 2002). The number of stocks at the top and bottom of the basket to be removed can be varied.

To properly describe the Atlas option, we need to introduce the notion of order statistics. Given n stocks $S_1, S_2, \ldots, S_n$ in a basket, we define$R_{(1)}^t = \min\{\frac{S_1^t}{S_1^0}, \frac{S_2^t}{S_2^0}, \ldots, \frac{S_n^t}{S_n^0}\}$, and $R_{(n)}^t = \max\{\frac{S_1^t}{S_1^0}, \frac{S_2^t}{S_2^0}, \ldots, \frac{S_n^t}{S_n^0}\}$ and $R_{(i)}^t$ is the ith smallest return, so that $R_{(1)}^t \leq R_{(2)}^t \leq \ldots \leq R_{(n)}^t$. The Atlas option removes a fixed number of stocks from the basket with n_1 and n_2 detailing the number of stocks to be removed from the minimum and maximum of the ordering. Given two numbers n_1, n_2 where $n_1 + n_2 < n$, and n stocks $S_1, S_2, \ldots, S_n$ in a basket with strike K, the payoff for the Atlas option is

$$\left(\sum_{j=1+n_1}^{n-n_2} \frac{R_{(j)}^T}{n-(n_1+n_2)} - K\right)^+. \tag{4.21}$$

For example, a strike of $K = 1$ would be in the money when the average of the remaining stocks in the basket had a return of at least 1. Likewise, a $K = 1.1$ would be in the money when the average of the stocks remaining in the basket increased in value by 10% over the period.

4.3.3 Altiplano/Annapurna

Altiplano is the highest plateau outside of Asia, located in the central Andes in South America and Annapurna is a mountain range in the Himalayas. The Altiplano pays a coupon if none of the stocks in the basket hits or goes above a certain level before expiration. In the event that one or more stocks hits the critical level or barrier, the Altiplano is a call option on the basket of stocks. When the limit is a floor rather than a ceiling for the stocks in the basket, the option is an Annapurna.

Given n stocks $S_1, S_2, \ldots, S_n$ in a basket, a coupon amount, C, a limit L, and strike K, with barrier period beginning at t_1 and ending at t_2, the payoff for an Altiplano option is

$$\begin{array}{cl} C & if \;\; \max\left(\frac{S_i^t}{S_i^0}\right) \leq L \quad \forall i, \forall t \in \{t_1, t_2\} \\ \left(\sum_{j=1}^{n} \frac{S_j^T}{S_j^0} - K\right)^+ & \text{otherwise.} \end{array}$$

The payoff for the Altiplano option will be discontinuous, having a jump when any of the stocks in the basket hits the barrier. Each of the parameters of these options, e.g. the coupon, level, strike, are all in units of returns. Imagine if the option coupon would pay 1.05, with $L = 1.05$ and $K = 0$. The Altiplano option would have a continuous payoff which would be constant until a stock had a positive return, at which time the Altiplano would jump to pay-out based upon the returns in the basket. Unless all the stocks in the basket simultaneously broke the upper limit, the value of the basket would decrease in price. As K increases, the pay-out in this option would have a larger drop as a stock breaks the boundary. Likewise, with a $C = 0.5$, and $K = 0.5$ with boundary at $L = 1$, the option would decrease in value as the stocks hit the boundary. With the three variables, a wide variety of payoffs can be constructed, which will be discontinuous w.p.1.

4.3.4 Himalaya

The Himalayan mountain range contains all the 8,000+ meter peaks on Earth. The Himalayan option is a call option on the average of the maximum performer in the basket each period. Each period, the top performer is removed from the basket, and its return over the last period is paid as a coupon. The option can be created so that the number of periods are equal or less than the number of stocks in the basket (Overhaus 2002).

To describe the removal of stocks from the basket, we will introduce additional set notation. We will define the set of returns from the stocks in the basket, and then incrementaly decrease the number of returns remaining in the basket to be chosen at each subsequent time period. Define

$$\begin{aligned} \mathcal{R}_i &= \left\{ \frac{S_1^t}{S_1^0}, \frac{S_2^t}{S_2^0}, \ldots, \frac{S_n^t}{S_n^0} \right\} \\ i_1^* &= \arg\max \mathcal{R}_1 \\ i_2^* &= \arg\max \mathcal{R}_2 \setminus \left\{ \frac{S_{i_1^*}^{t_1}}{S_{i_1^*}^0} \right\} \end{aligned}$$

$$i_3^* \quad = \quad \arg\max \mathcal{R}_3 \setminus \left\{ \frac{S_{i_1^*}^{t_1}}{S_{i_1^*}^0}, \frac{S_{i_2^*}^{t_2}}{S_{i_2^*}^0} \right\},$$

continuing until i_n^* is chosen from a set $\mathcal{R}_n$ which only contains a single element.

Given n stocks $S_1, S_2, \ldots, S_n$ in a basket, and a number of time points, $\{t_0, t_1, t_2, \ldots T\}$, we construct

$$\{R_{i_1^*}^{t_1}, R_{i_2^*}^{t_2}, \ldots R_{i_n^*}^{T}\}.$$

The payoff for the Himalayan option would be

$$\left(\sum_{j=1}^{n} \left(R_{i_j^*}^{t_j} - 1 \right) \right)^+$$

if globally floored or

$$\sum_{j=1}^{n} \left(R_{i_j^*}^{t_j} - 1 \right)^+ \tag{4.22}$$

if locally floored. In locally floored options, each coupon is positive by construction, and results in higher payoff than the globally floored option. In both of these representations, by construction, $R_{i_j^*}^{t_j}$ is the highest performing stock during the time interval $(t_0, t_j]$ remaining in the basket. Alternatively, the option could pay a coupon based upon the highest performing stock over the preceding time interval $(t_{j-1}, t_j]$, as opposed to paying all returns from time t_0. When the Himalayan option is constructed to pay on only a subset of the assets in the basket, the locally floored and globally floored payoffs are represented as

$$\left(\sum_{j=1}^{M}\left(R_{i_j^*}^{t_j}-1\right)\right)^+$$

or

$$\sum_{j=1}^{M}\left(R_{i_j^*}^{t_j}-1\right)^+$$

respectively, with $M < n$.

The Himalayan option can be constructed using the many different features highlighted here to provide a complex set of payoffs. The defining feature of a stock exiting the basket each period will make the payoff discontinuous with respect to the stock prices and stock returns in the basket.

4.4 Derivation of Gradient Estimators

To describe the basket of stocks of interest and to price each of the mountain range options will require multi-dimensional processes. To further explore these Mountain Range options, we will derive IPA and LR/SF estimators in the Black-Scholes-Merton model of geometric Brownian motion. Of the Mountain Range options, the Everest and Atlas both have continuous payoff functions, making them suitable for IPA, as well as LR/SF gradient estimation. We will derive IPA and LR/SF gradient estimators for these two of the Mountain Range options.

We define the correlated multi-dimentional geometric Brownian motion diffusion model by a system of stochastic differential equations,

$$\frac{dS_i(t)}{S_i(t)} = \alpha_i dt + \sigma_i dW_i(t), \quad i = 1, 2, ..., n \tag{4.23}$$

where $W_i(t)$ is a standard one-dimentional Brownian motion. Defining the correlation between $W_i(t)$ and $W_j(t)$ as $\rho_{i,j}$, we create a covariance matrix, $\Sigma_{i,j} = \rho_{i,j}\sigma_i\sigma_j$. Using the Cholesky factorization for $\Sigma = AA^T$, and we can rewrite (4.23) as

$$\frac{dS_i(t)}{S_i(t)} = \alpha_i dt + \sum_{j=1}^{n} A_{ij} dW_j(t).$$

Thus we can replace the drift α_i with $r - \frac{1}{2}\sigma_i^2$ and simulate the risk-neutral dynamics of the stocks in each of the Mountain Range options by

$$S_i(t_k + 1) = S_i(t_k) e^{\left(r - \frac{1}{2}\sigma_i^2\right)(t_{k+1} - t_k) + \sqrt{t_{k+1} - t_k} \sum_{j=1}^{n} A_{ij} Z_{k+1,j}}. \tag{4.24}$$

4.4.1 Everest Options Sensitivities

From the Everest option payoff function (4.20), we can notice the lack of a strike and the payoff based upon a minimum return of the stocks in the basket. When pricing under the risk-neutral measure, a change in measure ensures that the martingale condition on the stock price process is satisfied. If the final price of a stock S_i^T is such that it has the minimum return in the basket, any change in S_i^T will change the valuation of the basket, until the point where it no longer has the minimum return, at which point a different stock, S_j^T would be the new minimum. The payoff is thus a continuous and monotonically non-decreasing linear function of S_i^T while it is minimum, and therefore will be a piecewise linear function of the stocks in the basket, with a potential kink at each point when the stock with minimum return changes.

The pathwise gradient estimator is

$$\begin{aligned}
\frac{dJ_T}{d\theta} &= \sum_{i=1}^{n} \frac{dJ_T}{dS_i^T}\frac{dS_i^T}{d\theta} \\
\frac{dJ_T}{d\theta} &= \sum_{i=1}^{n} \left(\frac{1}{S_i^0}\frac{dS_i^T}{d\theta} - \frac{S_i^T}{(S_i^0)^2}\frac{dS_i^0}{d\theta} \right) 1_{\frac{S_i^T}{S_i^0} \leq \frac{S_j^T}{S_j^0}, \forall j \neq i},
\end{aligned}$$

and a sufficient condition for the interchange of differentiation and expectation is that the payoff is a continuous function with respect to the parameter, θ. As

$$\frac{dJ_T}{dS_0^T} = \sum_{i=1}^{n} \left(\frac{1}{S_i^0}\frac{S_i^T}{S_i^0} - \frac{S_i^T}{(S_i^0)^2} \right) = 0,$$

the Everest option has a delta equal to zero with respect to each stock in the basket. By directly substituting for S_i^T in (4.20), we obtain

$$\min_{i=1\ldots n} e^{\left(r - \frac{1}{2}\sigma_i^2\right)T + \sqrt{T}AZ},$$

which clearly shows that the price of the Everest option has no dependence on the initial spot of the stocks in the basket under geometric Brownian motion, and also holds under any multiplicative distribution. For all other Greeks, $\frac{dS_i^0}{d\theta} = 0$, and correspondingly

$$\frac{dJ_T}{dS_i^T} = \frac{1}{S_i^0}\frac{dS_i^T}{d\theta}1_{\frac{S_i^T}{S_i^0}\le\frac{S_j^T}{S_j^0},\forall j\neq i}.$$

The IPA estimator for vega with respect to each of the stocks in the basket is

$$\begin{aligned}\frac{dJ_T}{d\sigma_i} &= \frac{1}{S_i^0}\frac{dS_i^T}{d\sigma_1}1_{\frac{S_i^T}{S_i^0}\le\frac{S_j^T}{S_j^0},\forall j\neq i}\\ &= -\sigma_i T\frac{S_i^T}{S_i^0}1_{\frac{S_i^T}{S_i^0}\le\frac{S_j^T}{S_j^0},\forall j\neq i}.\end{aligned}$$

The IPA for rho and theta are, respectively,

$$\begin{aligned}\frac{dJ_T}{dr} &= \sum_{i=1}^{n}\frac{1}{S_i^0}\frac{dS_i^T}{dr}1_{\frac{S_i^T}{S_i^0}\le\frac{S_j^T}{S_j^0},\forall j\neq i}\\ &= \sum_{i=1}^{n}T\frac{S_i^T}{S_i^0}1_{\frac{S_i^T}{S_i^0}\le\frac{S_j^T}{S_j^0},\forall j\neq i}.\end{aligned}$$

$$\begin{aligned}\frac{dJ_T}{dT} &= \sum_{i=1}^{n}\frac{1}{S_i^0}\frac{dS_i^T}{dT}1_{\frac{S_i^T}{S_i^0}\le\frac{S_j^T}{S_j^0},\forall j\neq i}\\ &= \sum_{i=1}^{n}\left(r-\frac{1}{2}\sigma_i^2+\frac{AZ}{\sqrt{T}}\right)\frac{S_i^T}{S_i^0}1_{\frac{S_i^T}{S_i^0}\le\frac{S_j^T}{S_j^0},\forall j\neq i}.\end{aligned}$$

For the LR/SF gradient estimator, we need the payoff and the score function. The payoff function can be viewed as coming for free, as any simulation of the price of the Everest option will require implementation of the payoff function. In addition, once we determine the score function, it is also a reusable component, allowing straightforward computation of LR/SF gradient estimates for options with the same stochastics but different payoff functions. For the Everest option, we start with

$$\min_{i=1\ldots n}\left(\frac{S_i^T}{S_i^0}\right)\frac{d\ln f(S_1^T,S_2^T,\ldots S_n^T;\theta)}{d\theta}.$$

as the form of the LR/SF gradient estimator. According our model of stock dynamics, (4.23), returns are log-normally distributed,

$$\begin{aligned} S_i^T &\sim exp(X) \\ X &\sim N(\mu(\theta), T\Sigma), \end{aligned} \tag{4.25}$$

and the score function will be the derivative of (4.25) with respect to the parameter of interest. The score function is

$$\frac{d\ln f(S_1^T,S_2^T,\ldots S_n^T;\theta)}{d\theta}=\frac{(S_t-\mu(\theta))^T\Sigma^{-1}}{\sqrt{T}}\frac{d\mu(\theta)}{d\theta}.$$

For vega, rho and Theta, with $\mu_i=\log S_i^0+(r-\frac{1}{2}\|\Sigma_i\|^2)T$, we obtain

$$\begin{aligned} \frac{d\mu_i}{d\sigma} &= -\|\Sigma_i\|T, \\ \frac{d\mu_i}{dr} &= T \\ \frac{d\mu_i}{dT} &= r-\frac{\|\Sigma_i\|^2}{2} \end{aligned}$$

The score function will be used in the LR/SF gradient estimator for each of the

Mountain Range options, and simulating in accordance with (4.24)

$$\frac{(S_t-\mu(\theta))^T\Sigma^{-1}}{\sqrt{T}}\frac{d\mu}{d\theta}=\frac{Z^TA^{-1}}{\sqrt{T}}\frac{d\mu}{d\theta}. \tag{4.26}$$

4.4.2 Atlas Options Sensitivities

The Atlas option has the interesting feature of trimming both the highest and lowest performers in the basket and then incorporating a strike on the average of the remaining returns. The pathwise gradient estimator is

$$\begin{aligned}\frac{dJ_T}{d\theta} &= \sum_{i=1}^{n}\frac{dJ_T}{dS_i^T}\frac{dS_i^T}{d\theta}1_{1+n_1\leq i\leq n-n_2}\\ \frac{dJ_T}{dS_i^T} &= \frac{1}{(n-(n_1+n_2))S_{(i)}^0}1_{\sum_{j=1+n_1}^{n-n_2}\frac{R_{(j)}^T}{(n-(n_1+n_2))}>K}\\ &\quad\times 1_{1+n_1\leq i\leq n-n_2}\end{aligned}$$

The differentiability of the Atlas option will mirror the Everest Option, except that the Atlas option will not be differentiable at the point where $\frac{R_{(j)}^T}{(n-(n_1+n_2))}>K$. However, this is a point, so (4.21) is differentiable w.p.1. Implementing the gradient estimator for Atlas vega takes the form

$$\frac{dJ_T}{d\sigma_{(i)}}=\frac{1}{(n-(n_1+n_2))S_{(i)}^0}1_{\sum_{j=1+n_1}^{n-n_2}\frac{R_{(j)}^T}{(n-(n_1+n_2))}>K}\frac{dS_{(i)}^T}{d\sigma_{(i)}}1_{1+n_1\leq i\leq n-n_2}.$$

The IPA estimators for rho and theta are

$$\frac{dJ_T}{dr}=\sum_{i=1}^{n}\frac{1}{(n-(n_1+n_2))S_{(i)}^0}1_{\sum_{j=1+n_1}^{n-n_2}\frac{R_{(j)}^T}{(n-(n_1+n_2))}>K}\frac{dS_{(i)}^T}{dr}1_{1+n_1\leq i\leq n-n_2}.$$

$$\frac{dJ_T}{dT} = \sum_{i=1}^{n} \frac{1}{(n-(n_1+n_2))S^0_{(i)}} 1_{\sum_{j=1+n_1}^{n-n_2} \frac{R^T_{(j)}}{(n-(n_1+n_2))} > K} \frac{dS^T_{(i)}}{dT} 1_{1+n_1 \leq i \leq n-n_2}.$$

with $\frac{dS_i^T}{d\sigma}$, $\frac{dS_i^T}{dr}$, $\frac{dS_i^T}{dT}$ as previously derived.

The LR/SF gradient estimator makes use of the same score function developed for the Everest options, (4.26). The estimator can be implemented as

$$\left(\frac{\sum_{j=1+n_1}^{n-n_2} R^T_{(j)}}{n-(n_1+n_2)} - K \right)^+ \frac{Z^T A^{-1}}{\sqrt{T}} \frac{d\mu}{d\theta}.$$

4.4.3 Altiplano, Annapurna and Himalayan Option Sensitivities

IPA will not be applicable for gradient estimation for Altiplano, Annapurna and Himalayan options. We discussed removing the discontinuity in the derivative at K for the Atlas option, and the same holds here. However, the barrier option characteristics of the Altiplano/Annapurna option and the changing basket size of the Himalaya option both make the IPA estimator biased, as the price paths are discontinuous. Even though these options do not have continuous price paths, the LR/SF gradient estimator will still be applicable.

For the Altiplano/Annapurna option, the three parameters of strike, coupon payment, and barrier level all add to the complexity of this option. With the combination of $C = 1$, $K = 4$, and $L = 1$, on a basket of 4 stocks, the option would pay a coupon value of 1, unless the return of any of the stocks went above 1 during, at which point the option would pay the return over 4 for the entire basket. Interestingly, the coupon payment can set so that the payment has either a positive or a negative jump upon hitting the boundary.

The Himalayan option will have also have a discontinuous payoff as each period a stock is removed from the basket. However, a strength of the LR/SF gradient estimator is that once the proper pricing algorithm has been implemented, the score from the derivative of the measure can be combined with the pay-out from any option to complete the LR/SF estimator.

4.5 Numerical Example

We conducted a series of experiments by Monte Carlo simulation to price Everest, Atlas, Annapurna, and Himalaya options. Each of the experiments was conducted on a four stock basket. The correlation of the stocks in the basket was estimated from physical correlation, using 256 days of data from 1 June 2006 though 8 June 2007 provided by the Wharton Research Data Services (WRDS). The selected stocks are from the financial sector, Citigroup Inc. (C), Freddie Mac (FRE), J.P. Morgan Chase and Company (JPM) and Lehman Brothers (LEH), and each of the stocks had a zero dividend yield over the period. The values used in the simulation are displayed in Tables 4.2 and 4.3.

	Spot	Volatility
C	54	0.096313
FRE	66	0.092968
JPM	53	0.113978
LEH	61	0.316459

Table 4.2: Simulation Stock Parameter Values

	C	FRE	JPM	LEH
C	0.096313	0.043048	0.070818	0.097052
FRE	0.043048	0.092968	0.056524	0.07439
JPM	0.070818	0.056524	0.113978	0.122116
LEH	0.097052	0.07439	0.122116	0.316459

Table 4.3: Simulation Covariance Matrix Values

4.5.1 Mountain Range Greek Estimates

First we implemented the Everest Option over a one-year time horizon, with the risk-free rate set at 4%. We simulated 10,000 paths, calculated direct gradient estimates, central difference (CD), and forward difference (FD) estimates for vega, rho and theta. The results of the IPA and LR/SF estimators are provided in Table 4.4. The IPA estimators had small standard errors, and agreed with the results from CD and FD. The LR/SF agreed with the IPA and indirect methods for all but theta. The standard errors for each of the estimators were an order of magnitude less than the estimated value.

GBM Everest Option Gradient Estimation

Greek	IPA	StErr	Central diff	St Err	Forward Diff	St Err	LR/SF	Std Err
rho	0.741860	0.0026	0.741860	0.0026	0.741897	0.0026	0.693955	0.037
theta	-0.118181	0.0010	-0.118181	0.0010	-0.118177	0.0010	-0.057471	0.0023
vega								
	IPA	StErr	Central diff	St Err	Forward Diff	St Err	LR/SF	St Err
C	-0.112123	0.0016	-0.112123	0.0016	-0.112131	0.0016	-0.096837	0.012
FRE	-0.060290	0.0012	-0.060280	0.0012	-0.060299	0.0012	-0.052115	0.0092
JPM	-0.037741	0.00097	-0.037751	0.00097	-0.037766	0.00098	-0.031285	0.011
LEH	-0.083076	0.0014	-0.083076	0.0014	-0.083084	0.0014	-0.095874	0.0096

Table 4.4: GBM Everest Option Gradient Estimation

Similar results to those from the first experiment were observed for the Atlas options. Once again, we simulated 10,000 paths, calculated direct gradient estimates, central difference estimates, and forward difference estimates. We set

$K = 1.0$, $n_1 = 1$ and $n_2 = 1$, so the Atlas option was an option on the two remaining stocks in the basket, simulated over a one-year horizon.

In vega estimates, LR/SF and IPA produced similar results to CD and FD. LR/SF, FD, and CD were all within standard errors for the estimate of rho, and IPA, FD and CD were all within standard errors for estimates of theta. The LR/SF tended to have larger standard errors than the other methods, but the ease of implementation and wider applicability are the benefits of this direct method.

GBM Atlas Option Gradient Estimation

Greek	IPA	StErr	Central diff	St Err	Forward Diff	St Err	LR/SF	Std Err
rho	0.427622	0.0066	0.524983	0.0065	0.525715	0.0065	0.519970	0.016
theta	-0.050868	0.0012	-0.050950	0.0016	-0.050933	0.0012	-0.038693	0.0011
vega								
	IPA	StErr	Central diff	St Err	Forward Diff	St Err	LR/SF	St Err
C	-0.060735	0.0013	-0.060872	0.0013	-0.060939	0.0013	-0.062522	0.0038
FRE	-0.038250	0.00079	-0.038311	0.00079	-0.038357	0.00079	-0.035196	0.0033
JPM	-0.047935	0.00082	-0.048039	0.00083	-0.048114	0.00083	-0.048716	0.0039
LEH	-0.048748	0.0011	-0.048829	0.0011	-0.048863	0.0011	-0.048678	0.0034

Table 4.5: GBM Atlas Option Gradient Estimation

When estimating the Greeks for the Atlas and Everest options, the CD and FD methods produced similar results. The additional simulations required for estimation of these indirect gradients with a priori choice of step size did result in similar estimates. To calculate the CD with 10,000 stock paths cost an additional 20,000 simulations per calculated Greek, for a total of 120,000 additional stock paths. In the next section we will address the difference in step size and the errors for the CD method.

The next two options, Annapurna and Himalaya, are both path dependent and have discontinuous payoffs, and consequently the IPA method is not applicable.

Once again, we simulated 10,000 paths, calculated LR/SF gradient estimates, central difference estimates, and forward difference estimates, however, we simulated the monthly stock price for a four month period.

GBM Annapurna Option Gradient Estimation

Greek	Central diff	St Err	Forward Diff	St Err	LR/SF	Std Err
rho	0.670548	0.0077	0.670554	0.0077	0.288962	0.024
theta	-0.358369	0.018	-0.35836	0.017	-0.059088	0.0059
vega						
	Central diff	St Err	Forward Diff	St Err	LR/SF	St Err
C	-0.092154	0.0011	-0.092155	0.0011	-0.008374	0.0075
FRE	-0.049504	0.00058	-0.049505	0.00058	-0.029182	0.0055
JPM	-0.046857	0.00054	-0.046857	0.00054	-0.032924	0.0064
LEH	-0.070909	0.00083	-0.070909	0.00083	-0.026962	0.0059

Table 4.6: GBM Annapurna Option Gradient Estimation

For the Annapurna option, we set the boundary level, $L = 0.75$ and the coupon payment, $C = 0.75$. The strike, $K = 4.0$, was set such that the option would be in the money if the sum of returns were positive on average. In all of the estimates, the LR/SF estimate was considerably lower than the CD and FD estimates. The apparent discrepancy between the CD/FD and LR/SF for the Annapurna could result from either over estimation by CD/FD or under estimation by the LR/SF due to the pathwise dependence in the payoff structure, but we have yet to explain this discrepancy.

The Himalaya gradient estimates are displayed in Table 4.7. For the payoff function we used a local floor, as described in (4.22), and included each of the four stocks over the four month period of the option. The estimates agree for three of the four stocks vega estimates and theta but the LR/SF estimates for rho were smaller than the CD and FD estimates.

GBM Himalaya Option Gradient Estimation

Greek	Central diff	St Err	Forward Diff	St Err	LR/SF	Std Err
rho	0.404421	0.0034	0.404421	0.0034	0.356934	0.011
theta	-0.092976	0.0035	-0.108016	0.0039	-0.102731	0.0029
vega						
	Central diff	St Err	Forward Diff	St Err	LR/SF	St Err
C	-0.050899	0.00057	-0.050899	0.00057	-0.041537	0.0038
FRE	-0.030806	0.00038	-0.030806	0.00038	-0.035096	0.0030
JPM	-0.031068	0.00038	-0.031068	0.00038	-0.034695	0.0036
LEH	-0.041222	0.00049	-0.041222	0.00049	-0.041042	0.0030

Table 4.7: GBM Himalaya Option Gradient Estimation

In all of the Greeks for these path-dependent options, the CD and FD estimates agreed. However, the implementation for the CD and FD required significant additional coding. The implementation of the LR/SF requires two lines of code for each Greek in addition to required code to calculate the price of the option. For the Annapurna option, when calculating each of the indirect options, an additional overhead is incurred to calculate if any of the perturbed values broke the lower limit. For the Himalaya option it is required to keep track of which stocks have been removed from the basket on each perturbed price path required to calculate each Greek.

The additional modeling complexity is an additional cost above and beyond the the additional 20,000 simulations per calculated Greek. In pricing European options, the calculation of CD and FD estimates are trivial, but the complexity increases dramatically when pricing Mountain Range options. In addition, the a priori choice of step size may will require a tradeoff between bias and variance, and the step choices that are consistent with the model variables may not result in the lowest standard errors. In the next section we will address the difference in step size and the errors for the CD method.

4.5.2 Central Difference Errors

The standard errors for the Everest and Atlas options displayed in Figure 4.1 and Figure 4.2. These results agree with the tradeoff between bias and variance: smaller step size has larger errors but smaller bias (Fu 2006). When deciding upon a step size for the CD estimate, a starting guess would be based upon the order of magnitude of the parameter. For an interest rate of 0.04 or a volatility of 0.15, a guess of 0.001 would seem to be consistent with the model parameters. These step sizes provide reasonable errors for the Atlas and Everest options.

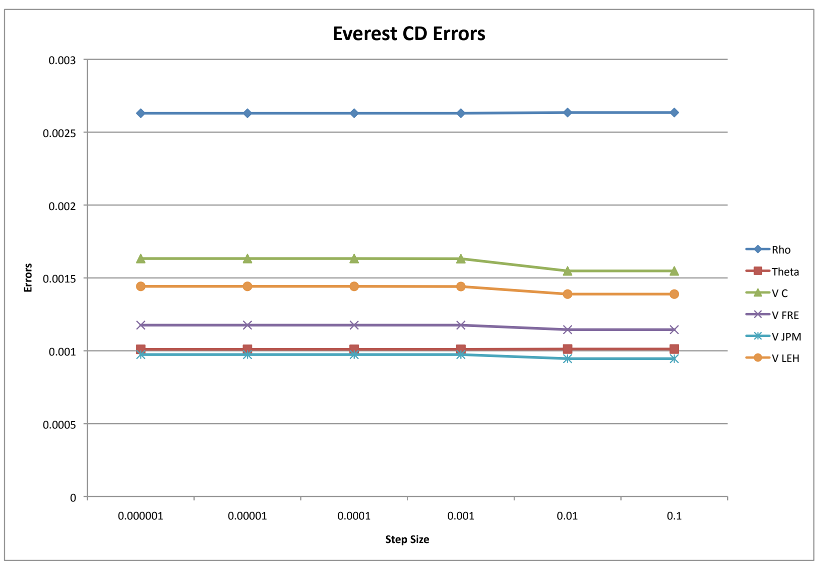

Figure 4.1: Standard Errors for Differing Step Sizes for Everest Options

Looking at the standard errors for the path dependent options provides insight into the important of the choice of step size. Using the same logic presented before, the a priori choice of 0.001 for a step size would result in larger standard errors

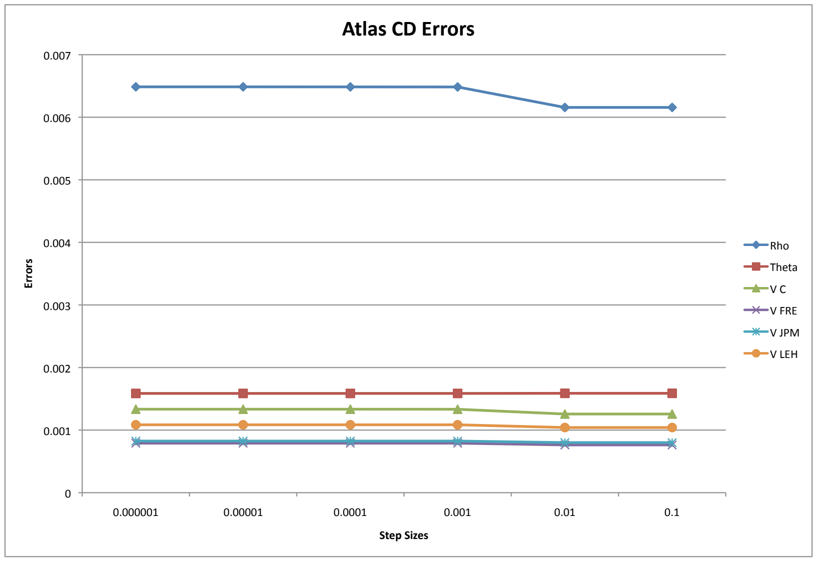

Figure 4.2: Standard Errors for Differing Step Sizes for Atlas Options

than necessary. As seen in Figure 4.3, the change in error from 0.01 to 0.0001 is increasing as expected. However, the smallest step sizes have unexpectedly low standard errors. In addition, as shown in Figure 4.4, the direction of the change in the value of the estimated Greek with a smaller step size is not consistent. The value of vega at different step sizes is very stable, where the value of theta varies considerably with smaller step sizes.

The standard errors for the Himalaya options are shown in Figure 4.5. The errors show similar behavior with the errors for the Annapurna options. However, the estimated values of the Greeks are more stable for the Himalaya option than for the Annapurna option. In both options, the value of theta decreased when a smaller step size was selected.

When pricing vanilla options, the computational complexity of the FD and CD methods seem trivial. However, with more complex options, the implementation of

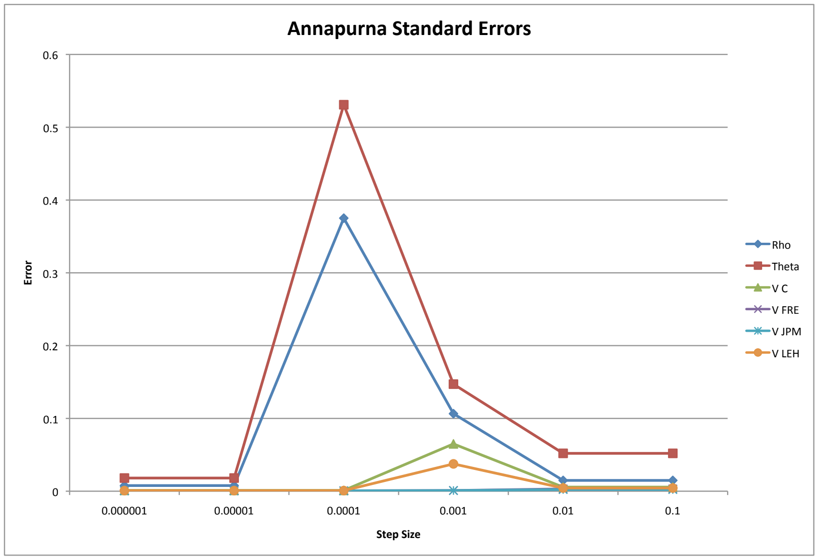

Figure 4.3: Standard Errors for Differing Step Sizes for Annapurna Options

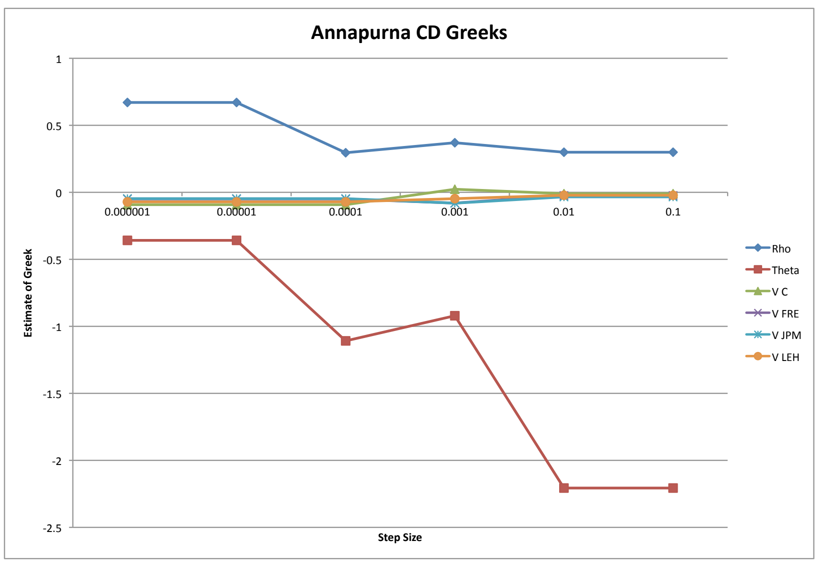

Figure 4.4: Differing Step Sizes for Annapurna Options

these indirect estimators results in additional modeling complexity, in addition to the additional simulation costs. After adding the additional modeling complexity and the uncertainty concerning the proper step size to the additional simulations required for CD, the benefits of direct gradient estimation become increasingly apparent.

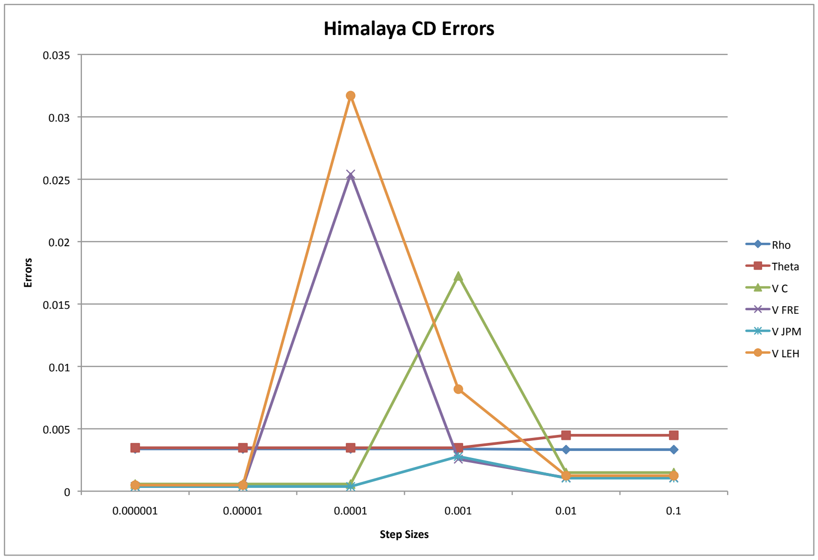

Figure 4.5: Standard Errors for Differing Step Sizes for Himalaya Options

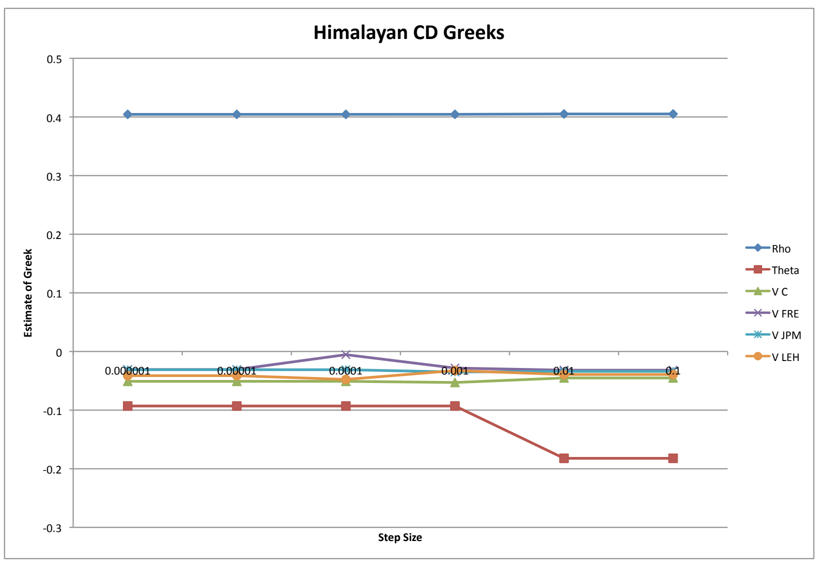

Figure 4.6: Differing Step Sizes for Himalaya Options

4.6 Results and Future Work

We found that the Mountain Range options provide an interesting example of efficient use of gradient estimation in a pricing context that lacks a closed-form solution and requires Monte Carlo or another numerical estimation technique to price the asset. We derived IPA estimators that accurately and efficiently estimated the gradients for Atlas and Everest options. Our results have shown that when IPA is applicable, it is the gradient estimate technique of choice for Mountain Range options.

The ease of implementation and mathematical simplicity makes indirect gradient estimation an attractive technique, but these methods can have very high variance and are known to result in biased estimators. Moreover, an additional 120,000 sample paths beyond the 10,000 needed to price the option itself are required to implement central and forward difference for the six Greeks estimated for these Mountain Range options. In addition to the increase in model complexity for the path-dependent options, there is an increased simulation cost for any implementation of indirect gradient estimators. As the number of Greeks needed increases, direct gradient estimators become the clear choice for Mountain Range options.

The commonality of each of these options is the payoff based upon return rather than stock level. Many options, e.g. spread, Asian, European, American, pay off on the basis of the level of the stock. However, for any options that rely on returns, estimates of delta based upon geometric Brownian motion and other multiplicative process will be a constant, 0.

We explored an application where Monte Carlo simulation is required in order to price the options, and direct gradient estimation provide an efficient technique for gradient estimation. This work has illuminated several additional research extensions. First, we could compute results for additional models of stock dynamics. Looking at models with variable volatility would be well suited to exploration by simulation. We would look at gradient estimates for Lévy and Variance Gamma processes. Also, exploring the differences between simulation of non-multiplicative models and our current results would add to the understanding of Mountain Range options.

Additional work in creating weak derivative estimates could compare results for lognormal distributions with the results from the normal distribution (Fu 2006). In addition, developing weak derivative estimators provide another comparison measure for LR/SF when IPA is not applicable. The research could also be extended by examining hedging strategies for these options, to determine if Mountain Range options can be perfectly hedged by static or dynamic positions. Any of the Greeks required for a hedging strategy can be implemented as a straightforward extension of this work.

Appendix A
OPL Code for Army Officer Specialty Model

```
/*********************************************
 * OPL 6.1+.1 Model
 * Author: ahall
 * Creation Date: Nov 13, 2008 at 3:22:52 PM
 *********************************************/
//Data

{string} Specialties = ...;
{string} Ranks = ...;

int NbPeriods = ...;
int NbStates = ...;
int NbLTS = ...;
int Nb2LTS = ...;
int NbProjection = NbPeriods -1;

range Periods = 1..NbPeriods;
range Projection = 1..(NbPeriods - 1);
range Difference = 1 ..(NbPeriods - 2);
range Rank_States = 1.. NbStates;
range FG = 8..NbStates;
range LTS = 1..NbLTS;
range SLT = 1..3;
range FLT = 4..6;
range CPT = 7..13;
range MAJ = 14..20;
range LTC = 21..26;
range COL = 27..29;


int X[Rank_States][Specialties] = ...;
int D[Projection][Ranks][Specialties] = ...;
int Losses[Projection][Specialties][Rank_States] = ...;

//Variables
//dvar int+ Accessions[Projection][Specialties][Rank_States];
//dvar int+ Promotions[Projection][Rank_States][Specialties];
//dvar int+ Reductions[Projection][Rank_States][Specialties];
```

```
//dvar int+ Personnel[Periods][Rank_States][Specialties];

dvar float+ Accessions[Projection][Specialties][Rank_States];
dvar float+ Promotions[Projection][Rank_States][Specialties];
dvar float+ Reductions[Projection][Rank_States][Specialties];
dvar float+ Personnel[Periods][Rank_States][Specialties];
//*******************************************************************
//Objective

minimize

//  Penalize all ranks OSD

 sum(p in Projection, s in Specialties)abs(sum (r in 1..6)Personnel[p+1][r][s]
 -   D[p]["1LT"][s])+
 sum(p in Projection, s in Specialties)abs(sum (r in 7..13)Personnel[p+1][r][s]
  -  D[p]["CPT"][s])+
 sum(p in Projection, s in Specialties)abs(sum (r in 14..20)Personnel[p+1][r][s]
 - D[p]["MAJ"][s])+
 sum(p in Projection, s in Specialties)abs(sum (r in 21..26)Personnel[p+1][r][s]
  - D[p]["LTC"][s])+
 sum(p in Projection, s in Specialties)abs(sum (r in 27..29)Personnel[p+1][r][s]
  - D[p]["COL"][s])

// Smoothing Accessions
 +
 sum(d in Difference)
 abs(sum(s in Specialties, r in Rank_States)Accessions[d][s][r]-
 sum(s in Specialties, r in Rank_States)Accessions[d+1][s][r])

// Smoothing Promotions
 +
 sum(d in Difference)
 abs(sum(s in Specialties, r in Rank_States)Promotions[d][r][s]-
 sum(s in Specialties, r in Rank_States)Promotions[d+1][r][s])
+
 // Minimize Reductions
 sum(p in Projection, s in Specialties, r in Rank_States)Reductions[p][r][s]
 +
 // Smoothing Reductions
 sum(d in Difference)
 abs(sum(s in Specialties, r in Rank_States)Reductions[d][r][s]-
 sum(s in Specialties, r in Rank_States)Reductions[d+1][r][s])
 ;
//**********************************************************************
```

```
//Goals and Constraints

subject to{

 // Constraints: Promotions
    forall (s in Specialties, p in Projection)
     Promotions[p][1][s] == 0;
    forall (s in Specialties, p in Projection)
     Promotions[p][4][s] == 0;
    forall (r in 7..10, s in Specialties, p in Projection)
     Promotions[p][r][s] == 0;
    forall (r in 14..17, s in Specialties, p in Projection)
     Promotions[p][r][s] == 0;
    forall (r in 21..23, s in Specialties, p in Projection)
     Promotions[p][r][s] == 0;
    forall (r in 27..29, s in Specialties, p in Projection)
     Promotions[p][r][s] == 0;

// In the Zone Promotions

  forall(s in Specialties, p in Projection)
      Promotions[p][2][s] >= (0.8)*Personnel[p][2][s];
  forall(s in Specialties, p in Projection)
      Promotions[p][5][s] >= (0.7)*Personnel[p][5][s];
    forall(s in Specialties, p in Projection)
      Promotions[p][12][s] >= (0.7)*Personnel[p][12][s];
  forall(s in Specialties, p in Projection)
      Promotions[p][19][s] >= (0.7)*Personnel[p][19][s];
  forall(s in Specialties, p in Projection)
      Promotions[p][25][s] >= (0.5)*Personnel[p][25][s];

// Below the Zone Promotions
   forall (s in Specialties, y in Projection)
     Promotions[y][11][s] <= (0.10)*Personnel[y][11][s];
  forall (s in Specialties, y in Projection)
     Promotions[y][18][s] <= (0.10)*Personnel[y][18][s];
   forall (s in Specialties, y in Projection)
     Promotions[y][24][s] <= (0.10)*Personnel[y][24][s];
   forall (r in Rank_States, s in Specialties, y in Projection)
     Promotions[y][r][s] <= Personnel[y][r][s];

// Constraints:  Personnel
// Aggregate flow constraints

// FY 08 Constraints - Initial Conditions
```

```
  forall (s in Specialties, r in Rank_States)
      Personnel[1][r][s] == X[r][s];

// Projection Flow Constraints

 forall (s in Specialties, p in Projection)
      Personnel[p+1][1][s] == Accessions[p][s][1];

 forall (s in Specialties, p in Projection)
      Personnel[p+1][2][s] == Personnel[p][1][s] - Losses[p][s][1]
       - Reductions[p][1][s];

 forall (s in Specialties, p in Projection)
      Personnel[p+1][3][s] == Personnel[p][2][s] + Personnel[p][3][s]
      - Promotions[p][2][s] - Promotions[p][3][s] - Losses[p][s][2]
       - Losses[p][s][3]- Reductions[p][2][s]- Reductions[p][3][s];

 forall (s in Specialties, p in Projection)
      Personnel[p+1][4][s]== Promotions[p][2][s] + Promotions[p][3][s];

 forall (s in Specialties, p in Projection)
      Personnel[p+1][5][s] == Personnel[p][4][s] - Losses[p][s][4]
      - Reductions[p][4][s];

 forall (s in Specialties, p in Projection)
      Personnel[p+1][6][s] == Personnel[p][5][s] + Personnel[p][6][s]
      - Promotions[p][5][s] - Promotions[p][6][s] - Losses[p][s][5]
       - Losses[p][s][6]- Reductions[p][5][s] - Reductions[p][6][s];

 forall (s in Specialties, p in Projection)
        Personnel[p+1][7][s] == Promotions[p][5][s] +
         Promotions[p][6][s] + Accessions[p][s][7];

 forall (r in 8..11, s in Specialties, p in Projection)
      Personnel[p+1][r][s] == Personnel[p]  [r-1][s]-
      Losses[p][s][r-1]- Reductions[p][r-1][s]
       + Accessions[p][s][r-1];

 forall (s in Specialties, p in Projection)
      Personnel[p+1][12][s] == Personnel[p][11][s] -
       Promotions[p][11][s] - Losses[p][s][11]
      - Reductions[p][11][s] + Accessions[p][s][11];

 forall (s in Specialties, p in Projection)
      Personnel[p+1][13][s] == Personnel[p][12][s] +
```

```
        Personnel[p][13][s]- Promotions[p][12][s]
        - Promotions[p][13][s] - Losses[p][s][12] -
         Losses[p][s][13]- Reductions[p][12][s]
        - Reductions[p][13][s] + Accessions[p][s][12] + Accessions[p][s][13];

forall (s in Specialties, p in Projection)
         Personnel[p+1][14][s] == Promotions[p][11][s] + Promotions[p][12][s] +
          Promotions[p][13][s];

forall ( r in 15..18, s in Specialties, p in Projection)
       Personnel[p+1][r][s] == Personnel[p][r-1][s] -
       Losses[p][s][r-1]- Reductions[p][r-1][s]
       + Accessions[p][s][r-1];

forall (s in Specialties, p in Projection)
       Personnel[p+1][19][s] == Personnel[p][18][s] - Promotions[p][18][s]
       - Losses[p][s][18]- Reductions[p][18][s] + Accessions[p][s][18];

forall (s in Specialties, p in Projection)
       Personnel[p+1][20][s] == Personnel[p][19][s] + Personnel[p][20][s]
       - Promotions[p][19][s] - Promotions[p][20][s] - Losses[p][s][19] -
       Losses[p][s][20]- Reductions[p][19][s]- Reductions[p][20][s]
       + Accessions[p][s][19] + Accessions[p][s][20];

forall (s in Specialties, p in Projection)
         Personnel[p+1][21][s] == Promotions[p][18][s]
          + Promotions[p][19][s] + Promotions[p][20][s];

forall (r in 22..24, s in Specialties, p in Projection)
       Personnel[p+1][r][s] == Personnel[p][r-1][s]
        - Losses[p][s][r-1] - Reductions[p][r-1][s]
       + Accessions[p][s][r-1];

forall (s in Specialties, p in Projection)
       Personnel[p+1][25][s] == Personnel[p][24][s]
       - Promotions[p][24][s] - Losses[p][s][24]
       - Reductions[p][24][s] + Accessions[p][s][24];

forall (s in Specialties, p in Projection)
       Personnel[p+1][26][s] == Personnel[p][25][s]
       + Personnel[p][26][s]- Promotions[p][25][s]
        - Promotions[p][26][s] - Losses[p][s][25] -
        Losses[p][s][26]- Reductions[p][25][s]-
        Reductions[p][26][s] + Accessions[p][s][25] + Accessions[p][s][26];
```

```
 forall (s in Specialties, p in Projection)
        Personnel[p+1][27][s] == Promotions[p][24][s] + Promotions[p][25][s] +
         Promotions[p][26][s];

 forall (s in Specialties, p in Projection)
      Personnel[p+1][28][s] == Personnel[p][27][s] - Losses[p][s][27] -
      Reductions[p][27][s] + Accessions[p][s][27];

 forall (s in Specialties, p in Projection)
      Personnel[p+1][29][s] == Personnel[p][29][s] +
       Personnel[p][28][s]- Losses[p][s][28]
      - Losses[p][s][29] - Reductions[p][28][s]- Reductions[p][29][s] +
      Accessions[p][s][28] + Accessions[p][s][29];


// DODMA Field Grade Requirements
 forall (p in Projection)
  sum (s in Specialties, r in 27..29)Personnel[p+1][r][s] <= 0.05 *
sum(s in Specialties, r in Rank_States)Personnel[p+1][r][s];
 forall (p in Projection)
  sum (s in Specialties, r in 21..26)Personnel[p+1][r][s] <= 0.11
 * sum(s in Specialties, r in Rank_States)Personnel[p+1][r][s];
 forall (p in Projection)
  sum (s in Specialties, r in 14..20)Personnel[p+1][r][s] <= 0.24
 * sum(s in Specialties, r in Rank_States)Personnel[p+1][r][s];

// ***************************************************************************
// Goals: Accession bounds

 forall ( y in Projection)
   sum (s in Specialties) Accessions[y][s][1] >= 4000;

 forall (y in Projection, s in Specialties)
  Accessions[y][s][1] <= 1.25*X[1][s];

 forall (y in Projection, s in Specialties)
  Accessions[y][s][1] >= 0.75*X[1][s];

//  Goals: Minimum Fill Constraints

 forall (p in Projection, s in Specialties)
   sum (r in 1..6)Personnel[p+1][r][s] >= 0.5 * D[p]["1LT"][s];
 forall (p in Projection, s in Specialties)
  sum (r in 7..13)Personnel[p+1][r][s] >= 0.5 * D[p]["CPT"][s];
 forall (p in Projection, s in Specialties)
```

```
  sum (r in 14..20)Personnel[p+1][r][s] >= 0.5 * D[p]["MAJ"][s];
 forall (p in Projection, s in Specialties)
  sum (r in 21..26)Personnel[p+1][r][s] >= 0.5 * D[p]["LTC"][s];
 forall (p in Projection, s in Specialties)
  sum (r in 27..29)Personnel[p+1][r][s] >= 0.5* D[p]["COL"][s];


// Reduction in Force Toggle

 forall (r in FG, p in Projection)
sum( s in Specialties)Reductions[p][r][s] >=
sum(s in Specialties)Accessions[p][s][r];

// Losses

forall (p in Projection)
sum(s in Specialties, r in SLT)Reductions[p][r][s]
+ sum(s in Specialties, r in SLT)Losses[p][s][r] >= 87;

forall (p in Projection)
sum(s in Specialties, r in FLT)Reductions[p][r][s]
+ sum(s in Specialties, r in SLT)Losses[p][s][r] >= 260;


forall (p in Projection)
sum(s in Specialties, r in CPT)Reductions[p][r][s]
+ sum(s in Specialties, r in CPT)Losses[p][s][r]
- sum(s in Specialties, r in CPT)Accessions[p][s][r] >= 1462;

forall (p in Projection)
sum(s in Specialties, r in MAJ)Reductions[p][r][s] +
sum(s in Specialties, r in MAJ)Losses[p][s][r]
- sum(s in Specialties, r in MAJ)Accessions[p][s][r] >= 466;

forall (p in Projection)
sum(s in Specialties, r in LTC)Reductions[p][r][s]
+ sum(s in Specialties, r in LTC)Losses[p][s][r]
- sum(s in Specialties, r in LTC)Accessions[p][s][r] >= 652;

forall (p in Projection)
sum(s in Specialties, r in COL)Reductions[p][r][s]
+ sum(s in Specialties, r in COL)Losses[p][s][r]
- sum(s in Specialties, r in COL)Accessions[p][s][r] >= 368;

// Percentage of Total Force
```

```
forall(p in Projection)
sum(s in Specialties, r in Rank_States)Personnel[p+1][r][s]
<= 0.15 * 525400;

};

execute OUTPUT {

  var ofile = new IloOplOutputFile("Plan10lpc.csv");

  for (var w in Projection){
  for (var x in (SLT)){
  for (var y in Specialties){
  ofile.writeln(NbProjection,",",w, ",",x, ",2LT,", y,
  ",",Personnel[w+1][x][y], ",", Accessions[w][y][x],
  ",",Promotions[w][x][y],",", Reductions[w][x][y],
  ",",Losses[w][y][x]);
  }}}
  for (var w1 in Projection){
  for (var x1 in (FLT)){
  for (var y1 in Specialties){
  ofile.writeln(NbProjection,",",w1, ",",x1, ",1LT,", y1,
  ",",Personnel[w1+1][x1][y1],",", Accessions[w1][y1][x1],
   ",",Promotions[w1][x1][y1],",", Reductions[w1][x1][y1],
   ",",Losses[w1][y1][x1]);
  }}}
  for (var w2 in Projection){
  for (var x2 in (CPT)){
  for (var y2 in Specialties){
  ofile.writeln(NbProjection,",",w2, ",",x2, ",CPT,", y2,
  ",",Personnel[w2+1][x2][y2],",", Accessions[w2][y2][x2],
   ",",Promotions[w2][x2][y2],",", Reductions[w2][x2][y2],
   ",",Losses[w2][y2][x2]);
  }}}
  for (var w3 in Projection){
  for (var x3 in (MAJ)){
  for (var y3 in Specialties){
  ofile.writeln(NbProjection,",",w3, ",",x3, ",MAJ,", y3,
  ",",Personnel[w3+1][x3][y3],",", Accessions[w3][y3][x3],
   ",",Promotions[w3][x3][y3],",", Reductions[w3][x3][y3],
   ",",Losses[w3][y3][x3]);
  }}}
  for (var w4 in Projection){
  for (var x4 in (LTC)){
  for (var y4 in Specialties){
```

```
  ofile.writeln(NbProjection,",",w4, ",",x4, ",LTC,", y4,
  ",",Personnel[w4+1][x4][y4],",", Accessions[w4][y4][x4],
   ",",Promotions[w4][x4][y4],",", Reductions[w4][x4][y4],
   ",",Losses[w4][y4][x4]);
  }}}
  for (var w5 in Projection){
  for (var x5 in (COL)){
  for (var y5 in Specialties){
  ofile.writeln(NbProjection,",",w5, ",",x5, ",COL,", y5,
  ",",Personnel[w5+1][x5][y5],",", Accessions[w5][y5][x5],
   ",",Promotions[w5][x5][y5],",", Reductions[w5][x5][y5],
   ",",Losses[w5][y5][x5]);
  }}}

  ofile.close();
}
```

Appendix B
C++ Code for European Call Options

```
double stock_process(double Z, double spot, double vol, double T){
double next;
next = spot*exp((rate - 0.5 *sig1*sig1)*T)*exp(sqrt(T)*vol*Z);
return next;
}
```

Figure B.1: C++ Implementation of geometric Brownian Motion

```
double Euro_price(int sim_limit)
{
for(sim_num = 0; sim_num < sim_limit; ++sim_num)
{
uniform1 = mcg();  uniform2 = mcg();
Z1 = boxmuller(uniform1, uniform2);
stock1 = stock_process(Z1, spot1, sig1, h);
price += ( (stock1 > K) ? exp(-rate*h)*(stock1 - K) : 0)/double(sim_limit);
}
return price;
}
```

Figure B.2: Monte Carlo Call Option Price

```
delta +=((stock1 > K) ? exp(-rate*h)*(stock1/spot1) : 0)/sim_limit;
vega  +=((stock1 > K) ?  exp(-rate*h)*(stock1/sig1)*(log(stock1/spot1)
- (rate + 0.5*sig1*sig1)*h) : 0)/sim_limit;
```

Figure B.3: C++ Implementation of IPA estimators

```
delta_LR +=((stock1 > K) ? exp(-rate*h)*(stock1 - K)
*(log(stock1/spot1) -delta_d)/(spot1*h*sig1*sig1) : 0)/sim_limit;
vega_LR +=((stock1 > K) ? exp(-rate*h)*(stock1 - K)
*((log(stock1/spot1)*log(stock1/spot1)/(h*sig1*sig1*sig1) -
(2*rate*log(stock1/spot1))/(sig1*sig1*sig1)
+ (rate*rate*h -sig1*sig1-0.25*h*sig1*sig1*sig1*sig1)
/(sig1*sig1*sig1))): 0)/sim_limit;
```

Figure B.4: C++ Implementation of LR/SF estimators

```
double boxmuller(double rand_u1, double rand_u2)
{
double R; double V;
R = - 2 * log(rand_u1);
V = 2 * 3.141592653589793 * rand_u2;
gauss1 = sqrt(R)*cos(V);  gauss2 = sqrt(R)*sin(V);
return gauss1;
}
```

Figure B.5: Box-Muller Algorithm for generating Normal Random Variates

Bibliography

Abrams, John W. 1957. Military applications of operational research. *Operations Research* **5**(3) 434–440.

Adda, J., R. Cooper. 2003. *Dynamic Economics*. MIT Press, Cambridge, Mass.

Arguden, R. Yilmaz. 1988. There is no free lunch: Unintended effects of the new military retirement system. *Journal of Policy Analysis and Management* **7**(3) 529–541.

Asch, Beth J., James R. Hosek, Jeremy Arkes, C. Christine Fair, Jennifer Sharp, Mark Totten. 2002. *Military recruiting and retention after the fiscal year 2000 military pay legislation*. Rand, Santa Monica, CA.

Asch, Beth J., James R. Hosek, John T. Warner. 2007. *New Economics of Manpower in the Post-Cold War Era*, *Handbook of Defense Economics*, vol. 2, chap. 32. Elsevier, 1075–1138. URL `http://ideas.repec.org/h/eee/hdechp/2-32.html`.

Asch, Beth J., John T. Warner. 2001. A theory of compensation and personnel policy in hierarchical organizations with application to the united states military. *Journal of Labor Economics* **19**(3) 523–562.

Bartholomew, David J. 1982. *Stochastic Models for Social Processes*. 3rd ed. Wiley Series in Probability and Mathematical Statistics, Wiley, Chichester ; New York.

Bartholomew, David J., Andrew F. Forbes, S. I. McClean. 1991. *Statistical Techniques for Manpower Planning*. 2nd ed. Wiley Series in Probability and Mathematical Statistics. Applied Probability and Statistics, Wiley, Chichester; New York.

Bausch, Dan O., Gerald G. Brown, Danny R. Hundley, Stephen H. Rapp, Richard E. Rosenthal. 1991. Mobilizing Marine Corps officers. *Interfaces* **21**(4) 26–38.

Berkovec, James, Steven Stern. 1991. Job exit behavior of older men. *Econometrica* **59**(1) 189–210.

Bertsekas, Dimitri P. 1995. *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, Mass.

Black, Fischer, Myron Scholes. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* **81**(3) 637-654.

Black, Matthew, Robert Moffitt, John T. Warner. 1990. The dynamics of job separation: The case of federal employees. *Journal of Applied Econometrics* **5**(3) 245–262.

Boudreau, John W. 2004. 50th Anniversary Article: Organizational behavior, strategy, performance, and design in management science. *Management Science* **50**(11) 1463–1476.

Bres, E. R., D. Burns, A. Charnes, W. W. Cooper. 1980. A goal programming model for planning officer accessions. *Management Science* **26**(8) 773–783.

Broadie, M., Paul Glasserman. 1996. Estimating security price derivatives using simulation. *Management Science* **42**(2) 269–285.

Burkhauser, R. V., J. S. Butler, G. Gumus. 2004. Dynamic programming model estimates of social security disability insurance application timing. *Journal of Applied Econometrics* **19**(6) 671–685.

Cashbaugh, D., A. Hall, M. Kwinn, T. Sriver, N. Womer. 2007. Manpower and personnel. L. Rainey, A. Loerch, eds., *Methods for Conducting Military Operational Analysis*. Military Operations Research Society.

Charnes, A., William W. Cooper, Richard J. Niehaus. 1978. *Management science approaches to manpower planning and organization design*. Studies in the management sciences v. 8, North-Holland Publishing Co., Amsterdam ; New York.

Chopra, Sunil, William Lovejoy, Candace Yano. 2004. Five decades of operations management and the prospects ahead. *Management Science* **50**(1) 8–14.

Collins, Roger W., Saul I. Gass, Edward E. Rosendarl. 1983. The ASCAR model for evaluating military manpower policy. *Interfaces* **13**(3) 44–53.

DACMC, Defense Advisory Committee on Military Compensation. 2006. The military compensation system: Completing the transition to an all-volunteer force.

Dailey, John T. 1958. Prediction of first-cruise reenlistment rate. *Operations Research* **6**(5) 686–692.

Daula, Thomas, Robert Moffitt. 1995. Estimating dynamic models of quit behavior: The case of military reenlistment. *Journal of Labor Economics* **13**(3) 499–523.

Davies, G. S. 1976. Consistent recruitment in a graded manpower system. *Management Science* **22**(11) 1215–1220.

Davies, G. S. 1982. Control of grade sizes in a partially stochastic Markov manpower model. *Journal of Applied Probability* **19**(2) 439–443.

Edwards, John S. 1983. A survey of manpower planning models and their application. *The Journal of the Operational Research Society* **34**(11) 1031–1040.

Feichtinger, Gustav. 1976. On the generalization of stable age distributions to gani-type person-flow models. *Advances in Applied Probability* **8**(3) 433–445.

Feichtinger, Gustav, Alexander Mehlmann. 1976. The recruitment trajectory corresponding to particular stock sequences in Markovian person-flow models. *Mathematics of Operations Research* **1**(2) 175–184.

Feuer, M. J., A. P. Schinnar. 1984. Sensitivity analysis of promotion opportunities in graded organizations. *Journal of the Operational Research Society* **35**(10) 915–922.

Fisher, Franklin M., Anton S. Morton. 1967a. The costs and effectiveness of reenlistment incentives in the Navy. *Operations Research* **15**(3) 373–387.

Fisher, Franklin M., Anton S. Morton. 1967b. Reenlistments in the U.S. Navy: A cost effectiveness study. *The American Economic Review* **57**(2) 32–38.

Flynn, James. 1975a. Retaining productive units: A dynamic programming model with a steady state solution. *Management Science* **21**(7) 753–764.

Flynn, James. 1975b. Steady state policies for a class of deterministic dynamic programming models. *SIAM Journal on Applied Mathematics* **28**(1) 87–99.

Flynn, James. 1979. Steady state policies for deterministic dynamic programs. *SIAM Journal on Applied Mathematics* **37**(1) 128–147.

Fry, M. J., M. J. Magazine, U. S. Rao. 2006. Firefighter staffing including temporary absences and wastage. *Operations Research* **54**(2) 353–365.

Fu, M. C. 2006. Gradient estimation. S. G. Henderson, B. L. Nelson, eds., *Simulation*. Handbooks in Operations Research and Management Science, Elsevier, Amsterdam, The Netherlands, 575–616. Chapter 19.

Fu, Michael C. 2008. What you should know about simulation and derivatives. *Naval Research Logistics* **55**(8) 723-736.

Fu, Michael C., Jian-Qiang Hu. 1995. Sensitivity analysis for Monte Carlo simulation of option pricing. *Probability in the Engineering and Informational Sciences* **9**(3) 417–446.

Gass, Saul I. 1991. Military manpower planning models. *Computers & Operations Research* **18**(1) 65–73.

Gass, Saul I., Roger W. Collins, Craig W. Meinhardt, Douglas M. Lemon, Marcia D. Gillette. 1988. The army manpower long-range planning system. *Operations Research* **36**(1) 5–17.

Glasserman, Paul. 2004. *Monte Carlo methods in financial engineering*. Applications of mathematics 53, Springer, New York.

Gotz, G. A., John J. McCall. 1984. *A Dynamic Retention Model for Air Force Officers: Theory and Estimates*. RAND Corp.

Gotz, Glenn A., John J. McCall. 1983. Sequential analysis of the stay/leave decision: U.S. Air Force officers. *Management Science* **29**(3) 335–351.

Green, Linda V., Peter J. Kolesar, Joao Soares. 2001. Improving the sipp approach for staffing service systems that have cyclic demands. *Operations Research* **49**(4) 549–564.

Grinold, R. C., R.E. Stanford. 1974. Optimal control of a graded manpower system. *Management Science* **20**(8) 1201-1216.

Grinold, Richard C. 1976. Manpower planning with uncertain requirements. *Operations Research* **24**(3) 387–399.

Grinold, Richard C., Kneale T. Marshall. 1977. *Manpower planning models*. Publications in Operations Research Series, North-Holland Pub. Co., New York.

Haight, Matthew, William Lewis. 2008. Officer forecasting model. BAAFEX 2008: British American Australian Forecasting Exchange.

Hall, Andrew O. 2004. Validation of the enlisted grade model gradebreaks. *WSC '04: Proceedings of the 36th conference on Winter simulation*. Winter Simulation Conference, 921–925.

Harrison, J. Michael, David M. Kreps. 1979. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory* **20**(3) 381–408.

Hartley, Keith, Todd Sandler. 2007. *Handbook of defense economics, Volume 2*. Elsevier, Amsterdam ; New York.

Henning, Charles A. 2006. Army officer shortages : background and issues for congress. Tech. rep., Congressional Research Service, the Library of Congress, Fort Belvoir, VA.

Henry, Todd M., A. Ravi Ravindran. 2005. A goal programming application for army officer accession planning. *INFOR* **43**(2) 111–119.

Holder, A. 2005. Navy personnel planning and the optimal partition. *Operations Research* **53**(1) 77–89.

Holz, Betty W., James M. Wroth. 1980. Improving strength forecasts: Support for Army manpower management. *Interfaces* **10**(6) 37–52.

Jaquette, D. L., G. R. Nelson. 1976. Optimal wage rates and force composition in military manpower planning. *Management Science* **23**(3) 257–266.

Klingman, Darwin, Nancy V. Phillips. 1984. Topological and computational aspects of preemptive multicriteria military personnel assignment problems. *Management Science* **30**(11) 1362–1375.

Krass, Iosif A., Mustafa Pinar, Theodore J. Thompson, Stavros A. Zenios. 1994. A network model to maximize Navy personnel readiness and its solution. *Management Science* **40**(5) 647–661.

Lewis, M. W., K. R. Lewis, B. J. White. 2006. Guided design search in the interval-bounded sailor assignment problem. *Computers & Operations Research* **33**(6) 1664–1680.

Martel, A., W. Price. 1981. Stochastic programming applied to human resource planning. *The Journal of the Operational Research Society* **32**(3) 187–196.

Martel, Alain, Ahmad Al-Nuaimi. 1973. Tactical manpower planning via programming under uncertainty. *Operational Research Quarterly (1970-1977)* **24**(4) 571–585.

McCall, John, Neil Wallace. 1969. A supply function of first-term re-enlistees to the Air Force. *The Journal of Human Resources* **4**(3) 293–310.

Meaney, Joseph. 2007. Dealing with the volatility smile of himalayan options. Michele Vanmaele, Griselda Deelstra, Ann De Schepper, Jan Dhaene, Huguette Reynaerts, Wim Schoutens, Paul Van Goethem, eds., *5th Actuarial and Financial Mathematics Day*. Universa Press, Flemish Royal Academy of Arts and Sciences for Belgium.

Mehlmann, Alexander. 1977. Markovian manpower models in continuous time. *Journal of Applied Probability* **14**(2) 249–259.

Merton, Robert C. 1973. Theory of rational option pricing. *The Bell Journal of Economics and Management Science* **4**(1) 141–183.

Miranda, Mario J., Paul L. Fackler. 2002. *Applied Computational Economics and Finance*. MIT Press, Cambridge, Mass.

Nilakantan, K., B. G. Raghavendra. 2005. Control aspects in proportionality Markov manpower systems. *Applied Mathematical Modelling* **29**(1) 85–116.

Nilakantan, K., B. G. Raghavendra. 2008. Length of service and age characteristics in proportionality Markov manpower systems. *IMA J Management Math* **19**(3) 245–268.

Office of the Secretary of Defense. 2008. *Defense Budget Materials*. http://www.defenselink.mil/comptroller/defbudget/fy2008/index.html.

Overhaus, Marcus. 2002. Himalaya options. *Risk* **15**(3) Deutsche Bank Masterclass, Risk.net (March 2002).

Price, W. L. 1978. Solving goal-programming manpower models using advanced network codes. *The Journal of the Operational Research Society* **29**(12) 1231–1239.

Price, W. L., W. G. Piskor. 1972. The application of goal programming to manpower planning. *INFOR* **10**(3) 221–231.

Puterman, Martin L. 2005. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley-Interscience, Hoboken, NJ.

Quessette, Richard. 2002. New products, new risks. *Risk* **15**(3). Deutsche Bank Masterclass, Risk.net (March 2002).

Reeves, Gary R., Randall C. Reid. 1999. A military reserve manpower planning model. *Computers & Operations Research* **26**(12) 1231–1242.

Rostker, Bernard. 2006. *I Want You! : the Evolution of the All-Volunteer Force*. RAND, Santa Monica, CA.

Rostker, Bernard, Office of the Assistant Secretary of Defense/Force Management and Personnel, Defense Research Institute. 1993. *The Defense Officer Personnel Management Act of 1980 : a retrospective assessment*. Rand, Santa Monica, CA.

Roth, Alvin E., Elliott Peranson. 1999. The redesign of the matching market for American physicians: Some engineering aspects of economic design. *The American Economic Review* **89**(4) 748–780.

Rust, J., C. Phelan. 1997. How social security and medicare affect retirement behavior in a world of incomplete markets. *Econometrica* **65**(4) 781-831. 0012-9682.

Rust, John P. 1987. A dynamic programming model of retirement behavior. *NBER Working Paper Series*.

Senate, United States. 2009. Senator Everett Mckinley Dirksen Dies. URL http://www.senate.gov/artandhistory/history/minute/Senator Everett Mckinley Dirksen Dies.htm.

Shrimpton, Dan, Alexandra M. Newman. 2005. The US Army uses a network optimization model to designate career fields for officers. *Interfaces* **35**(3) 230–237.

Stanford, Robert E. 1980. The effects of promotion by seniority in growth-constrained organizations. *Management Science* **26**(7) 680–693.

Stanford, Robert E. 1985. A stochastic model of personnel mobility: Application of reliability theory in manpower system analysis. *Operations Research* **33**(1) 109–125.

Stigler, George J. 1962. Information in the labor market. *The Journal of Political Economy* **70**(5) 94–105.

Tien, James M., Angelica Kamiyama. 1982. On manpower scheduling algorithms. *SIAM Review* **24**(3) 275–287.

Vajda, S. 1978. *Mathematics of Manpower Planning*. Wiley, Chichester, Eng. ; New York.

Warner, D. Michael, Juan Prawda. 1972. A mathematical programming model for scheduling nursing personnel in a hospital. *Management Science* **19**(4) 411–422.

Warner, John T., Matthew S. Goldberg. 1984. The influence of non-pecuniary factors on labor supply: The case of Navy enlisted personnel. *The Review of Economics and Statistics* **66**(1) 26–35.

Warner, John T., Saul Pleeter. 2001. The personal discount rate: Evidence from military downsizing programs. *The American Economic Review* **91**(1) 33–53.

Wilson, Norman A. B. 1969. *Manpower research: the proceedings of a conference held under the Aegis of the N.A.T.O. Scientific Affairs Committee in London from 14th-18th August, 1967*. English Universities P., London,.

Yadavalli, V. S. S., R. Natarajan. 2001. A semi-Markov model of a manpower system. *Stochastic Analysis & Applications* **19**(6) 1077–1086.

Zanakis, Stelios H., Martin W. Maret. 1980. A Markov chain application to manpower supply planning. *The Journal of the Operational Research Society* **31**(12) 1095–1102.

Zanakis, Stelios H., Martin W. Maret. 1981. A Markovian goal programming approach to aggregate manpower planning. *The Journal of the Operational Research Society* **32**(1) 55–63.