NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

PLANNING FOR AN ADAPTIVE EVADER WITH APPLICATION TO DRUG INTERDICTION OPERATIONS

by

Philip Gift

September 2010

Thesis Advisor:	Johannes O. Royset
Second Reader:	Moshe Kress

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			*Form Approved OMB No. 0704-0188*

1. AGENCY USE ONLY *(Leave blank)*	**2. REPORT DATE** September 2010	**3. REPORT TYPE AND DATES COVERED** Master's Thesis

4. TITLE AND SUBTITLE Planning for an Adaptive Evader with Application to Drug Interdiction Operations	**5. FUNDING NUMBERS**
6. AUTHOR(S) Philip Gift	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000	**8. PERFORMING ORGANIZATION REPORT NUMBER**
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A	**10. SPONSORING/MONITORING AGENCY REPORT NUMBER**

11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. IRB Protocol Number _________N/A________

12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited	**12b. DISTRIBUTION CODE**

13. ABSTRACT (maximum 200 words)

In an effort to impede the flow of drugs from South America, a Coalition Force headed by Joint Interagency Task Force (JIATF) – South allocates its assets to detect and interdict drug smuggling vessels such as the self-propelled semi-submersible (SPSS) used by a Drug Trafficking Organization (DTO). In this thesis, we develop an interdiction model to place the Coalition Force assets optimally. We also develop a model – known as the Adaptive Evader Model – for a DTO that is able to learn the placement of the Coalition Force assets. This model is akin to the multi-armed bandit problem. We create two algorithms for the Adapting Evader Model. One algorithm uses an optimal learning policy and the other uses a heuristic learning policy. We also create an algorithm for the interdiction model using the Cross-Entropy method. Finally, we construct a case study that we use to draw some insights about how a DTO, that is capable of learning, reacts to different optimal plans. This information can be used by the Coalition Force to more effectively allocate their limited number of assets during drug interdiction operations.

14. SUBJECT TERMS Stochastic model, Dynamic programming, Stochastic optimization, Multi-armed bandit problem, Cross Entropy Method, Interdiction model, Search and detection, Self-propelled semi-submersible, SPSS			**15. NUMBER OF PAGES** 99
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	**18. SECURITY CLASSIFICATION OF THIS PAGE** Unclassified	**19. SECURITY CLASSIFICATION OF ABSTRACT** Unclassified	**20. LIMITATION OF ABSTRACT** UU

NSN 7540-01-280-5500 Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

PLANNING FOR AN ADAPTIVE EVADER WITH APPLICATION TO DRUG INTERDICTION OPERATIONS

Philip D. Gift
Lieutenant, United States Navy
B.S., United States Naval Academy, 2002

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

NAVAL POSTGRADUATE SCHOOL
September 2010

Author: Philip Gift

Approved by: Johannes O. Royset
Thesis Advisor

Moshe Kress
Second Reader

Robert F. Dell
Chairman, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

In an effort to impede the flow of drugs from South America, a Coalition Force headed by Joint Interagency Task Force (JIATF)—South allocates its assets to detect and interdict drug smuggling vessels such as the self-propelled semi-submersible (SPSS) used by a Drug Trafficking Organization (DTO). In this thesis, we develop an interdiction model to place the Coalition Force assets optimally. We also develop a model—known as the Adaptive Evader Model—for a DTO that is able to learn the placement of the Coalition Force assets. This model is akin to the multi-armed bandit problem. We create two algorithms for the Adapting Evader Model. One algorithm uses an optimal learning policy and the other uses a heuristic learning policy. We also create an algorithm for the interdiction model using the Cross-Entropy method. Finally, we construct a case study that we use to draw some insights about how a DTO, that is capable of learning, reacts to different optimal plans. This information can be used by the Coalition Force to more effectively allocate their limited number of assets during drug interdiction operations.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

a	Discount factor
AEW	Airborne early warning
AOR	Area of Responsibility
b	Variable in the linear law; $b = \frac{v_j}{u}$
C	Stopping criterion for the CE method; the number of consecutive times that the same asset allocation for each route has the highest probability
CE	Cross-Entropy
c	Iteration count for the CE method
d_i	Width of route i in nm
DTO	Drug trafficking organization
$f(X;\alpha^0,\beta^0)$	Optimization function in the Asset Allocation Model
G	Number of plans chosen with a random mechanism for the CE method
G^l	Number of plans chosen with a random mechanism in iteration l for the CE method
g	Current simulation of a random mechanism for the CE method
GAO	United States Government Accountability Office
h	Variable in the linear law; $h = \frac{\left(d_i - w_{i,j}\right)}{w_{i,j}}$
i	A route
IACM	Interagency Assessment of Cocaine Movement
j	A specific asset type for the interdictor

JIATF	Joint Interagency Task Force
k	A specific asset allocation
l	An iteration of a random mechanism for the CE method
m	Total number of different asset types for the interdictor
MPA	Maritime patrol aircraft
n	Total number of routes
NDIC	National Drug Intelligence Center
ONDCP	Office of National Drug Control Policy
p_i	Actual probability that an evader transits route i successfully
$p_i(X)$	Actual probability that an evader transits route i successfully given an interdiction plan X
$p_i(X^{l,g})$	Actual probability that an evader transits route i successfully in iteration l for interdiction plan $X^{l,g}$ in iteration g in the CE Asset Allocation Algorithm
pmf	probability mass function
r_j	Number of available assets for the interdictor of type j
s	Track spacing in the barrier patrol
SOUTHCOM	Southern Command
SPSS	Self-Propelled Semisubmersible
t	Time (one time period / one try of a route)
T	Total number of time periods / times that the routes can be chosen
u	Speed of an evader's vessel
v_j	Speed of an interdictor's asset of type j
$w_{i,j}$	Sweep width for an interdictor's asset of type j on route i
X	An interdiction plan

$X_{i,j}$	Interdiction plan with the number of assets of type j on route i
$X^{l,g}$	g^{th} interdiction plan for iteration l for the Cross-Entropy method
$X_{i,j}^{l,g}$	Number of assets of type j on route i in interdiction plan $X^{l,g}$.
$\widehat{X}$	Interdiction plan chosen by CE Asset Allocation Algorithm
$\widehat{X}_{i,j}$	$\arg\max_k \psi_{i,j,k}^{l}$ for the CE Asset Allocation Algorithm
z	Variable used in the inverse cube law; $z = \sqrt{\frac{\pi}{2}}\frac{w_{i,j}}{s}$
α	A parameter in the beta distribution that represents the number of successes (represents a vector for all of the routes)
α_i	A parameter in the beta distribution that represents the number of successes on route i
α_i^0	The initial parameter in the beta distribution that represents the believed initial number of successes on route i
β	A parameter in the beta distribution that represents the number of failures (represents a vector of all of the routes)
β_i	A parameter in the beta distribution that represents the number of failures on route i
β_i^0	The initial parameter in the beta distribution that represents the believed initial number of failures on route i
ε_0	Parameter for the ε-decreasing strategy
ε_t	Updated value for the ε-decreasing strategy
η_j	Number of interdiction assets of type j for a barrier patrol
$v(\alpha,\beta)$	Gittins Index with α and β
γ	Value calculated from user input, ρ, times D to compare the function values for the CE method

λ	Value between one and zero for the smoothing function for the CE method
$\phi_{i,j}$	Probability of detecting a vessel with asset type j on route i
ψ^l	Random mechanism for iteration l
$\psi^l_{i,j,k}$	Probability for asset type j on route i to have k assets for iteration l
$\tilde{\psi}^l$	Interim random mechanism for iteration l
ρ	Percent of plans to use to update the parameters for the CE method
$\theta_i(t)$	Believed probability of success on route i at time t
$\theta(t)$	Vector of believed probability of successes on the routes at time t

EXECUTIVE SUMMARY

This thesis focuses on a portion of the drug problem; the shipment of cocaine from South and Central America into the United States. The major suppliers of these illicit drugs are the Colombian drug trafficking organizations (DTOs). One of the vessels that they are now using to traffic the cocaine into the United States is the self-propelled semi-submersible (SPSS). The SPSSs are a major issue since they are difficult to detect and can carry a large amount of cargo. To stop the flow of drugs, the United States has partnered with other nations to create a Coalition Force, which is headed by Joint Interagency Task Force (JIATF)—South. Since the Coalition Force has a limited number of assets, they need to utilize them effectively and efficiently.

In this thesis, we develop an interdiction model to allocate the Coalition Force (interdictor) assets optimally. We also develop a model for a DTO (evader) that is able to learn the placement of the interdictor assets. This model is referred to as the Adapting Evader Model and is akin to the multi-armed bandit problem. In both models, there are a given number of "smuggling" routes available. The interdictor allocates his assets on these routes. The evader uses these same routes to transport his SPSS.

In the Adapting Evader Model, the evader, at each decision point, chooses one route with the goal of maximizing his discounted number of successful "smuggling runs." Discounting the number of successes means that it is more beneficial to have a success sooner than later. By varying the amount we discount, we change the patience of the evader. As the discount gets closer to 1, the evader becomes more patient and is more willing to explore the routes, instead of just using the route which he believes has the highest probability of success. After trying a route, the evader learns if the SPSS successfully traversed the route without being caught. The evader uses the knowledge about which routes produce successes versus failures to choose the next route. He continues trying routes for an infinite time horizon. We create two algorithms for the Adapting Evader Model. One algorithm is exact and uses an optimal learning policy based on Gittins indices and the other uses a heuristic learning policy.

The interdiction model is an integer nonlinear program, which minimizes the evader's discounted number of successes by optimally allocating the interdictors assets. We create an algorithm for solving the interdiction model using the Cross-Entropy method.

Finally, we construct a case study that we use to draw some insights about how a DTO, that is capable of learning, reacts to different optimal plans. In the case study, there are five routes and the interdictor has eight assets, which are split into three different types each with a different probability of detecting an SPSS. The goal of the case study is to gain a better understanding of how the Coalition Force can more effectively allocate their limited number of assets to impede the effectiveness of the SPSS. One insight is that the interdiction plan is not dependent on the algorithm the DTO uses to choose his routes. The exact and heuristic algorithms give about the same interdiction plan. Other attributes such as the DTO's patience and prior belief of success on the routes play a more substantial role in designing the interdiction plan. Another insight relates to the time it takes the DTO to find the route with the highest probability of success, which we refer to as the best route. In the Coalition Force's worst-case scenario, a patient DTO that chooses the routes optimally and does not have a prior belief of success, it takes about forty tries to find the best route. To keep the DTO from utilizing the knowledge about the best route, the Coalition Force needs to change their interdiction plan before they believe the DTO has found the best route. By changing this plan before the DTO can take advantage of the knowledge of the best route, the Coalition Force can increase the number of SPSSs that they interdict. The models and algorithms developed in this thesis provide a means for the Coalition Force to allocate their limited number of assets effectively and efficiently.

ACKNOWLEDGMENTS

I would like to thank my advisor, Professor Johannes O. Royset, for his expertise and guidance. He made writing this thesis a positive learning experience. I would also like to thank my second reader, Professor Moshe Kress, for his feedback and positive encouragement.

Most importantly, I would like to thank my family for their support and patience during this thesis process. My wife, Laura, has been a wonderful mom to our son and daughter. I would also like to thank James and Tess who always made sure that I took time off from work.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

A. SCENARIO OVERVIEW

Illegal drugs are a large problem in the United States. The National Drug Intelligence Center (NDIC), in their National Drug Threat Assessment 2009, states that more than 35 million people in 2007 used illicit drugs or abused prescription drugs (NDIC, 2009). The Assessment further states, "in September 2008 there were nearly 100,000 inmates in federal prisons convicted and sentenced for drug offenses, representing more than 52 percent of all federal prisoners." The major suppliers of the illicit drugs are the Mexican and Colombian drug trafficking organizations (DTOs). They "generate, remove, and launder between $18 billion and $39 billion in wholesale drug proceeds annually" (NDIC, 2009). The federal government has allocated more than $14 billion in 2009 for drug interdiction, counterdrug law enforcement, international counterdrug assistance, and drug treatment and prevention (NDIC, 2009).

1. Trafficking of Cocaine

A more specific drug problem is that associated with cocaine. "Analysis of law enforcement reporting as well as national drug threat, availability, demand, and treatment data indicates that cocaine trafficking is the greatest drug threat to the United States" (NDIC, 2009). Almost all of the cocaine transported to the United States started in the Andean countries of Colombia, Peru, and Bolivia (GAO, 2008). According to the Interagency Assessment of Cocaine Movement (IACM) in 2007 between 545 and 707 metric tons of cocaine departed South America for the United States (NDIC, 2009). Cocaine can be transported from South America by air, land, or sea. Recently, the majority is transported by sea (GAO, 2008). There are also currently no reports of the cocaine being transported by land (ONDCP, 2008). Figure 1 shows the flow of cocaine from South America in 2007 as estimated using confirmed, substantiated, and higher-confidence suspect events in the Consolidated Counterdrug Database (NDIC, 2009).

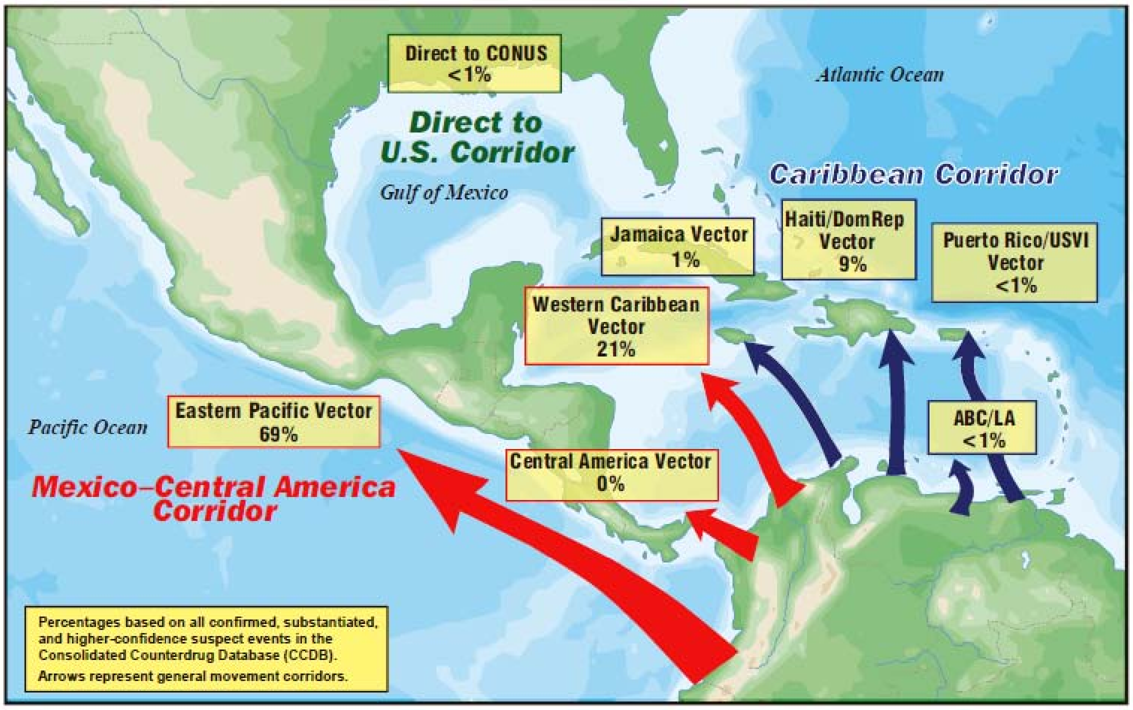

Source: Interagency Assessment of Cocaine Movement.

Figure 1. Documented Cocaine Flow Departing South America, in January-December of 2007 (From NDIC, 2009)

One of the more problematic vessels that the DTOs are using to transport their drugs is the self-propelled semi-submersible (SPSS). The SPSS is a boat with a low draft. It combines the capacity of a fishing vessel with the small surface size of a go-fast, which makes it difficult to find. The majority of the SPSS vessels are between 25-65 feet in length, with speeds up to 13 knots, can hold 4–5 crewmembers, and carry up to 10 metric tons (with an average of 3-6 metric tons) of illicit cargo for distances up to 5000 NM (with refueling) and 2500 NM without refueling (USCG, 2009).

The Coast Guard estimates that currently 32 percent of all maritime cocaine that flows through the Eastern Pacific is by SPSS. This is expected to increase since only 23 SPSS events occurred in the 6.5 years before September 2007, but for 9 months after September 2007, there have been 62 SPSS events (USCG, 2009). Figure 2 is a picture of a boarded SPSS. Figure 3 is a picture of an SPSS under construction.

Figure 2. An SPSS being boarded (From USCG, 2009)

Figure 3. An SPSS under construction (From USCG, 2009)

2. Stopping the Trafficking of Cocaine

To defend the United States from drug trafficking, the Department of Defense is the leading federal branch of government. Specifically, the Joint Interagency Task Force (JIATF) – South is the United States Southern Command (SOUTHCOM) agency that heads all of the interagency and partner nations counter drug operations in the Caribbean

Sea, Gulf of Mexico, and Eastern Pacific. For counter drug operations in this Area of Responsibility (AOR), there are a variety of American and foreign assets that are used for detection, monitoring and interdiction.

The United States Navy and Coast Guard, as well as partner nations such as Britain, France, Netherlands, Canada, and Colombia, use ships to help patrol the AOR (JIATF, 2009). The ships include United States Navy frigates (see Figure 4), United States Coast Guard Cutters, and partner nation frigates. For the interdiction of the suspected vessels, there is a Coast Guard Law Enforcement Detachment embarked on the US ships, and sometimes on the partner nation ships (JIATF, 2009). Aboard some of the United States Navy ships there are helicopter squadron detachments to help with detection and monitoring (see Figure 5) (JIATF, 2009). To compliment the Coalition Force's ship operations, the United States and partner nations also use land-based aircraft such as maritime patrol aircraft (MPA) and airborne early warning (AEW) aircraft (see Figure 6).

Figure 4. United States Navy Frigate (USS John L. Hall, FFG – 32) (From Jane's, Oliver Hazard Perry class, 2010)

Figure 5. United States Navy Helicopter (MH-60R) (From Jane's, Sikorsky MH-60R Seahawk, 2010)

Figure 6. United States Navy P-3C Orion Airborne Early Warning (AEW) Aircraft (From Jane's, Lockheed P-3C Orion, 2010)

Currently, the United States is making some significant strides in the battle against cocaine trafficking. In the past year, according to the NDIC in 2009, there has been a decrease in the availability of domestic cocaine. The reason is unclear, but they state that one of the factors is most likely "several exceptionally large cocaine seizures made while the drug was in transit toward the United States" (NDIC, 2009). Since, as Coast Guard Rear Adm. Joseph L. Nimmich, Commander, JIATF-South stated, "Every time we turn around, the smugglers are extraordinarily creative, extraordinarily adaptive," the United States needs to continue to improve on the way it interdicts the drug smuggler's vessels (SOUTHCOM, 2009). One of the ways to improve is with smarter methods of employing the Coalition Force's search and interdiction platforms.

B. SCOPE, GOAL, AND BENEFITS OF STUDY

This thesis focuses on a portion of the drug problem; the shipment of cocaine from South America into the United States, using the SPSS. The SPSS is a good smuggling vessel; it is difficult to detect with its low profile and it is able to carry a large cargo. The SPSS also has the capability of smuggling other contraband such as weapons of mass destruction. For this reason, the goal of this thesis is to find a better way to allocate the Coalition Force assets to increase the number of interdictions of SPSSs.

This thesis develops a model, which produces the optimal allocation of Coalition Force assets against a DTO that is capable of learning the location of those assets. This model will have operational value when used as a decision aid. It also provides a means to delve into the decision process of a DTO choosing where to use his SPSSs.

C. THESIS ORGANIZATION

This thesis is organized with Chapter I presenting an overview of the problem and the main players. Chapter II describes the scenario and lists other references that have delved into this topic. The models and algorithms used in this thesis are presented in Chapter III. A numerical case study is presented and analyzed in Chapter IV. Chapter V summarizes this thesis with conclusions. In Appendix A, there are tables of pre-computed Gittins Indices and values to estimate Gittins Index. Appendix B describes how to calculate the number of possible interdiction plans. Appendix C describes how the probability of success on each route is computed for the numerical case study in Chapter IV. Appendix D is a graphical representation of the results from the Case Study.

II. PROBLEM DESCRIPTION

A. SITUATION

An evader (i.e., DTO) and an interdictor (i.e., Coalition Force) are operating against each other in an environment consisting of n routes. A route consists of a starting location, a path, and a destination. For instance, Figure 1 shows an environment with seven routes ($n = 7$). Both the evader and interdictor know the routes and they cannot change them.

The evader travels these routes with his vessels. The vessels go one at a time across a route in hopes of a successful trip. A successful trip means that the vessel makes it to its destination. Associated with route i, i=1, 2, …, n, is the probability p_i that a trip on this route completed successfully. The evader does not know p_i and thus must try to learn these probabilities. The only way for the evader to learn is to try the route and experience the outcome. An outcome will be either a success or a failure. Each try of a route takes a time period (e.g., a week). At the beginning of time period t a vessel becomes available. The evader will then choose a route, say route i, for the vessel to traverse. By the end of the time period, the evader knows if the vessel succeeds or fails. With this knowledge about route i, the evader updates his belief, $\theta_i(t)$, on route i for time period t. $\theta_i(t)$ is a random variable that represents the evader's belief about the probability of success of traversing route i at time period t. The evader does not update his belief (pdf of $\theta_i(t)$) on any of the other routes for time period t. With the updated belief, the evader now uses his beliefs on the routes ($\theta_i(t+1)$, i=1, 2, …, n) to choose a route to use for the next vessel during the following time period ($t+1$). This process continues for an infinite time horizon. The evader has the option to explore different routes and risk a vessel now to gain knowledge for later use. The evader aims to find a best route so he can maximize his expected number of discounted successes. A best route is the route with a probability of success that is equal to or greater than the respective probabilities of the other routes. It is possible to have multiple best routes. By finding the best route as soon as possible, the evader maximizes his expected number of discounted successes. The willingness of the evader to search for the best route is dependent on the evader's patience. This patience can be represented using a discount factor (Fudenberg & Levine, 1998). The discount factor decreases the value of a success to the evader as t increases. At each time period, the value of a success to the evader is exponentially discounted. We denote the discount factor by a, which is between zero and one. As a gets closer to one, the evader gets more patient. A more patient evader will explore the routes longer than an impatient evader.

The interdictor determines the probabilities of success on the routes that the evader is trying to learn. He does this by allocating his limited number of assets on the n routes. The assets assigned to route i yields a specific p_i. The interdictor has m different types of assets ($j = 1, 2, \ldots, m$), which have different probabilities of detecting a vessel. The corresponding probability of detecting a vessel for asset j on route i is $\phi_{i,j}$. An asset's probability of detecting a vessel is independent of the other assets' probability of detecting a vessel. Let $X_{i,j}$ be the number of assets of type j on route i and let X be the vector with components $X_{i,j}$ for all $i = 1, 2, \ldots, n$ and $j = 1, 2, \ldots, m$. We refer to X as an interdiction plan. We denote the number of assets of type j available by r_j. We let $p_i(X)$ be the probability of success along a route given an interdiction plan X, which we compute using $\phi_{i,j}$; see Appendix A. The interdictor aims to minimize the evader's expected number of discounted successes over the infinite time horizon.

B. LITERATURE REVIEW

Using mathematical models to allocate resources that interdict the flow of contraband is not new. Wood (1993) uses network interdiction models[1] to help decide

---

[1] An interdiction model is also known as an attacker-defender model. This model is a Stackelberg game (von Stackelberg, 1952). In a Stackelberg game two sides play sequentially. The interdictor goes first followed by the evader. For a more detailed explanation and extension of these kind of models, see Brown et al. (2006).

where to allocate a limited number of resources to interdict drugs in South America. He develops flexible integer programming models to minimize the maximum flow that can be pushed through a capacitated network.

More recently, Pfeiff (2009) uses an interdiction model to help the Coalition Force allocate their assets to catch the most SPSSs, in both the Eastern Pacific and the Caribbean. The solution to his model is a mixed strategy for the interdictor and a least risk path for the evader.

Washburn and Wood (1995) develop game-theoretic models for the drug flow along a network of roads and rivers in parts of South America. They use a network interdiction model to solve this two-person zero-sum game, with the solution being the maximum expected equilibrium.

Dimitrov et al. (2009) use dynamic programming in their model to show the optimal places to build stationary radiation detectors on a transportation network. The nuclear-material smuggler in the model is adaptive, but also has full knowledge of the detector locations and evasion probabilities.

Bailey et al. (1994) model the interaction of the United States Coast Guard cutters and smuggling vessels using Monte Carlo simulation. The evader chooses its route using a sequence of finite horizon dynamic programs. The dynamic programs take into account an evader that is “forced to combine his short-run profit goals with his need to gain future information about the configuration of the cutters” (Bailey et al., 1994). This configuration of the interdictors becomes a discrete-time Markov chain from the evader’s point of view.

Caulkins et al. (1993) model how interdicting cocaine shipments affect smuggling costs using dynamic programming and Monte Carlo simulation. The evader has three modes of transportation: air, sea, and land. Each mode can use a set number of routes. The routes are treated as generic since the routes in the model are not associated with any particular geographic boundaries (Caulkins et al., 1993). The evader uses a heuristic algorithm to choose a route in the resulting model, which is a multi-armed bandit problem.[2] The heuristic algorithm randomly chooses a route based on time-weighted estimates of the probability of interdiction. The interdictor in this model is only represented by a probability of interdiction on each route and not by any physical asset.

The first four papers differ from this thesis by the modeling of the evader. Wood (1993) and Pfeiff (2009) both have an evader that does not know nor is able to learn the interdictor's plan. Washburn & Wood (1995) and Dimitrov et al. (2009) both have an evader that already knows the interdictor's plan in some sense. In Washborn & Wood (1995), the interdictor's plan is a mixed strategy. In Dimitrov et al. (2009), the interdict's plan is the location of the detectors. The last two papers differ from this thesis by the modeling of the interdictor. Bailey et al. (1994) and Caulkins et al. (1993) both have an evader that is adaptive and capable of learning the interdictor's plan. However, the interdictor's plan, in both these papers, is determined by the user and is not optimized. This thesis develops models that find the optimal interdiction plan against an evader that is adaptive and capable of learning this plan.

---

[2] The first person to pose the bandit problem was W. Thompson in 1933 (Berry & Fristedt, 1985). Thompson (1933, 1935) uses the Beta distribution with Bayesian updating to maximize the expected number of successes in the first T pulls for two arms. There was little interest in the bandit problem until 1952 when H. Robbins wrote a paper on it. Since Robbins, many people have explored the multi-armed bandit problem and its extensions, which are laid out in Berry and Fristedt's book (1985).

III. MODEL AND ALGORITHM DEVELOPMENT

A. OPTIMIZATION MODELS FOR THE EVADER AND INTERDICTOR

1. The Evader Optimization Model

We develop a model for an evader that is able to learn the placement of the interdictor's assets which we refer to as the Adapting Evader Model. In our model, the evader's belief about the probability of success of the routes is represented by a beta distribution, which is also in Thompson (1933, 1935); see comment on page 10. The beta distribution is able to account for the two parts associated with a belief. The first part is the believed probability of success and the second part is the uncertainty associated with that probability. A beta distribution can represent these parts using its parameters α and β. The believed probability is the mean of the beta distribution and is calculated by

$$\frac{\alpha}{\alpha+\beta} \tag{3.1}$$

The confidence of the belief is represented by the variance of the beta distribution and is calculated by

$$\frac{\alpha\beta}{\left(\alpha+\beta\right)^2\left(\alpha+\beta+1\right)} \tag{3.2}$$

$\theta_i(t)$ has a beta distribution with parameters α_i and β_i for time period t. Using Bayesian updating, if route i is used and the trip is a success, then α_i is increased by one, but if the trip is a failure then β_i is increased by one. The Bayesian updating is possible since the beta distribution is a conjugate pair with the Bernoulli distribution. To start the Bayesian update, the evader needs to have an initial belief on route i. We denote the initial belief on route i by α_i^0 and β_i^0, which is an input to the model. The initial belief gives modeling flexibility. It allows the modeling of different prior disposition about the routes. The Adapting Evader Model follows below and takes the form of a multi-armed bandit problem.

Adapting Evader Model

<u>*Index*</u>

i	*routes* $(i = 1, 2, \ldots, n)$

<u>*Data*</u>

a	*discount factor*
α_i^0	*initial value of alpha on route i*; $\alpha^0 = \left(\alpha_1^0, \alpha_2^0, \ldots, \alpha_n^0\right)$
β_i^0	*initial value of beta on route i*; $\beta^0 = \left(\beta_1^0, \beta_2^0, \ldots, \beta_n^0\right)$
p_i	*actual probability of success on route i*

<u>*State*</u>

α_i	*current value of alpha on route i*
β_i	*current value of beta on route i*

<u>*Function*</u>

$$
\begin{aligned}
\mathrm{R}(\alpha,\beta) = \max\Big\{ & p_1\left[a + a\mathrm{R}\left(\left(\alpha_1 + 1, \alpha_2, \ldots, \alpha_n\right), \beta\right)\right] \\
& + \left(1 - p_1\right)\left[0 + a\mathrm{R}\left(\alpha, \left(\beta_1 + 1, \beta_2, \ldots, \beta_n\right)\right)\right], \\
& p_2\left[a + a\mathrm{R}\left(\left(\alpha_1, \alpha_2 + 1, \ldots, \alpha_n\right), \beta\right)\right] \\
& + \left(1 - p_2\right)\left[0 + a\mathrm{R}\left(\alpha, \left(\beta_1, \beta_2 + 1, \ldots, \beta_n\right)\right)\right], \\
& \ldots, \\
& p_n\left[a + a\mathrm{R}\left(\left(\alpha_1, \alpha_2, \ldots, \alpha_n + 1\right), \beta\right)\right] \\
& + \left(1 - p_n\right)\left[0 + a\mathrm{R}\left(\alpha, \left(\beta_1, \beta_2, \ldots, \beta_n + 1\right)\right)\right]\Big\}
\end{aligned}
\tag{3.3}
$$

<u>*Formulation*</u>

Determine $\mathrm{R}\left(\alpha^0, \beta^0\right)$

The Adapting Evader Model is a dynamic program with Bellman's equation given by Equation 3.3. The current reward in Equation 3.3 is the expected discounted success for that time period, which is $p_i(a)$ for a success and $(1 - p_i)(a)$ for a failure. The future rewards are the discounted successes for a later time period, which is $p_i\left[a\mathrm{R}\left(\left(\alpha_1, \alpha_{i-1}, \alpha_i + 1, \alpha_{i+1}, \ldots, \alpha_n\right), \beta\right)\right]$ for a success and

$(1-p_i)\left[a\mathrm{R}\left(\alpha,\left(\beta_1,\beta_{i-1},\beta_i+1,\beta_{i+1},\ldots,\beta_n\right)\right)\right]$ for a failure. To solve the Adaptive Evader Model using the backward-recursion dynamic programming algorithm is difficult and we instead utilize a procedure for determining an optimal policy based on Gittins indices[3] and a heuristic policy as we describe in Section B.

2. The Interdictor Optimization Model

We develop an interdiction model that finds the optimal allocation of the interdictor's assets, which we refer to as the Asset Allocation Model. The solution to the objective function in our model is the minimum of the maximum expected number of discounted successes as defined by the Adaptive Evader Model. The Asset Allocation Model consists of minimizing this maximum by selecting an interdiction plan as described next.

[3] Gittins & Jones (1974) solve the multi-armed bandit problem optimally using Gittins Indices. Gittins Indices are optimal under the following assumptions: exponential discounting, an infinite time horizon, and beliefs about routes that are not chosen cannot change. Since then, others such as Whittle (1980), Varaiya et al. (1985), Tsitsiklis (1986), Gittins (1989), Weber (1992), Tsitsiklis (1994), Bertsimas & Niño-Mora (1996), and Dacre et al. (1999) also have proven the optimality of Gittins index.

Asset Allocation Model

Index

i — *routes* $(i = 1, 2, \ldots, n)$

j — *asset types* $(j = 1, 2, ..., m)$

Data

a — *discount factor*

α_i^0 — *initial value of alpha on route i*; $\alpha^0 = \left(\alpha_1^0, \alpha_2^0, \ldots, \alpha_n^0\right)$

β_i^0 — *initial value of beta on route i*; $\beta^0 = \left(\beta_1^0, \beta_2^0, \ldots, \beta_n^0\right)$

r_j — *total number of assets of type j*

Decision Variables

$X_{i,j}$ — *number of assets of type j on route i*

Functions

$p_i\left(X\right)$ — *actual probability of success on route i given interdiction plan X*; $X = \left(X_{1,1}, X_{2,1}, \ldots, X_{2,1}, X_{2,2}, \ldots, X_{n,m}\right)$

$f(X; \alpha^0, \beta^0) = \mathsf{R}\left(\alpha^0, \beta^0\right)$ — *expected number of discounted successes given interdiction plan X and* α^0 *and* β^0 *as defined by* $\mathsf{R}\left(\alpha^0, \beta^0\right)$ *in Adapting Evader Model with* $p_i\left(X\right)$ *replacing* p_i *for all i*

Formulation

$$\min_X \quad f(X; \alpha^0, \beta^0)$$

s.t.

$$\sum_i X_{i,j} = r_j \qquad \forall j$$

$$X_{i,j} \in \left\{0, 1, ..., r_j\right\} \qquad \forall i, j$$

The Asset Allocation Model minimizes the optimization function of the Adaptive Evader Model with the optimal interdiction plan.

B. OPTIMIZATION ALGORITHMS

1. The Evader Optimization Algorithm

This thesis develops two algorithms for the Adaptive Evader Model. The first uses an optimal policy, which we denote Gittins Choice Algorithm. The second uses a heuristic policy, which we denote Decreasing Choice Algorithm. In both of these algorithms, we let T be the maximum number of time periods with $t = 1, 2, \ldots, T$. We set T to be a large number to represent an infinite time horizon. In the Decreasing Choice Algorithm, we let ε_0 be a user defined parameter and ε_t be a variable that changes with time.

a. *Optimal Evader Optimization Algorithm*

We use Gittins Index to solve the evader optimization model optimally. By using Gittins Index the evader's n-dimensional problem, which is choosing a route at each time period to achieve the highest expected number of discounted successes in an infinite horizon, becomes n one-dimensional problems (Varaiya et al., 1985). Gittins (1989) provides an equation to estimate the Index and tables of pre-computed Indices. The equation and tables depend directly on a, since a dictates how much the evader wants to explore versus exploit. The equation and tables for $a = 0.50$ and $a = 0.99$ are given in Appendix B.

To use Gittins Indices, a Gittins Index is calculated for each of the n routes. The evader chooses to use the route with the highest index. After learning if the vessel completes its trip successfully, the evader then recalculates the Index for the chosen route. The process continues for an infinite time horizon.

As an illustration of Gittins Index, consider an evader that has two routes to choose from ($n = 2$) and is patient ($a = 0.99$). The evader's prior belief on Route 1 is manifested in $\alpha = 1$ and $\beta = 1$. The evader's prior belief on Route 2 is manifested in $\alpha = 30$ and $\beta = 10$. If the evader chooses a route based on expected probability of success alone (Equation 3.1), then he would choose Route 2 since 0.75 is greater than 0.5. Using Gittins Index, in this example, the evader chooses Route 1 since its Gittins index is larger. Since there is more uncertainty around Route 1 compared to Route 2, Route 1 has a higher potential to have a high probability of success. In this instance, using Gittins Index, the evader will explore instead of exploit, since he chooses the route with more uncertainty over the route with a higher expected probability of success. Using the same example with an impatient evader ($a = 0.50$) changes the outcome. The evader chooses to exploit instead of exploring and selects Route 2.

Below is the Gittins Choice Algorithm, which solves the Adapting Evader Model. The policy determined by the Gittins Choice Algorithm is optimal, but the optimal value (expected number of discounted successes) is estimated using Monte Carlo simulation.

Gittins Choice Algorithm

1. Set initial conditions

   T : Maximum time; $t = 1$; a : discount factor; n : routes ($i = 1, 2, \ldots, n$)

   α_i^0 , β_i^0 : Evader's initial belief about the probability of success on the route i .

   p_i : Actual probability of a success on route i .

2. Calculate Gittins Index for all of the routes using the initial beliefs.
3. Choose the route with the highest index (route i).
4. Randomly decide if the run was a success using the actual probability on route i (p_i).

   If run was a success, increase the optimal value by the value a^t .

5. Update the belief on route i .

   If run was a success increase α_i by 1. Else, increase β_i by 1 for route i .

6. Recalculate Gittins Index for route i .
7. Increase t by 1.
8. If $t > T$, then stop. Else, go to step 3.

b. Heuristic Evader Optimization Algorithm

The evader could choose from many heuristic policies. For instance, the evader could just explore at every time period by randomly choosing a route. At the other extreme, the evader could exploit at every time period; he would choose the route with the highest expected probability of success. We consider a heuristic policy that is something in between, the ε-decreasing strategy; see Vermorel & Mohri (2005).

In each time period, using the ε-decreasing strategy, the evader randomly decides whether to explore (with probability ε_t) or exploit (with probability $1-\varepsilon_t$) where:

$$\varepsilon_t = \min\left\{1, \frac{\varepsilon_0}{t}\right\} \tag{3.4}$$

and ε_0 is a non-negative value set by the user. ε_0 is the evader's patience. A higher value for ε_0 means the evader will explore for a longer time before he starts to exploit. If $\varepsilon_0 = 0$ then the evader will always exploit and choose the route with the highest believed probability of success. The believed probability of success on each route is the mean of the beta distribution and is calculated by Equation 3.1. ε_t is monotonically decreasing as time (t) increases. This means that as time increases the evader is more likely to exploit than to explore. There will also be a point (depending on the value of ε_0) when the evader will exploit with a very high probability.

Below is the Decreasing Choice Algorithm, which uses a heuristic policy for the Adapting Evader Model. This algorithm uses Monte Carlo simulation to evaluate the value of the policy.

Decreasing Choice Algorithm

1. Set initial conditions
   - T: Maximum time; $t = 1$; a: discount factor
   - ε_0: Non-negative value to represent the evader's patience.
   - n: Routes ($i = 1, 2, \ldots, n$)
   - α_i^0, β_i^0: Evader's initial belief about the probability of success on the route i.
   - p_i: Actual probability of a successful run on route i.
2. Calculate the believed probability of success on the routes using Equation 3.1.
3. Calculate ε_t using Equation 3.4.
4. Randomly choose a number (δ), where $\delta \in [0,1]$.

   If $\delta > \varepsilon_t$ then choose the route with highest believed probability of success. Else, randomly choose a route.

5. Randomly decide if the run was a success using the actual probability on route i (p_i).
   If run was a success, then increase the optimal value by the value a^t.
6. Update the belief on route i.
   If run was a success, then increase α_i by 1. Else, increase β_i by 1 for route i.
7. Recalculate the believed probability of success on route i using Equation 3.1.
8. Increase t by 1.
9. If $t > T$, then stop. Else go to step 3.

2. The Interdictor Optimization Algorithm

We solve the Asset Allocation Model by the Cross-Entropy (CE) method (Rubinstein & Kroese 2004, Allon et al., 2004). The CE method is an iterative method that generates a sequence of sets of interdiction plans according to a specified random mechanism. We denote the random mechanism by ψ^l for iteration l. This random mechanism generates a user specified number of interdiction plans G^l in iteration l. As in Allon et al. (2004), we do not change the number of interdiction plans between iterations and drop the superscript henceforth. We let g =1, 2, …, G be the interdiction plan counter. We then let $X^{l,g}$ be the g^{th} interdiction plan generated by ψ^l in iteration l.

The random mechanism induces a probability mass function (pmf) for each asset type and each route, except Route n. Route n is excluded since it receives the assets that have not already been allocated to the previous routes. We denote the probability for route i to have k assets of type j in iteration l by $\psi^l_{i,j,k}$. The random mechanism consists of the collection of these pmf, i.e. $\psi^l = \left\{\psi^l_{i,j,k} : 1 \le i \le n, 1 \le j \le m, 0 \le k \le r_j\right\}$. The initial pmf for each route is uniform, which means that the probability for route i to have any number of assets (within the limit of the total number of assets available) is equally likely. The pmf for Route 1 and asset type j in iteration l is $\psi^l_{1,j,0}$, $\psi^l_{1,j,1}$, …, ψ^l_{1,j,r_j}, where r_j is the number of assets of type j. Table 1 is an example of an initial pmf for four routes and five assets for asset type j. The hyphen in Table 1 on Route 4 delineates that there is no pmf on Route 4, since the number of assets that have not been allocated on Routes 1, 2, or 3 will be allocated on Route 4.

Route	Number of Assets on the Route of Type *j*: 0	1	2	3	4	5
1	1/6	1/6	1/6	1/6	1/6	1/6
2	1/6	1/6	1/6	1/6	1/6	1/6
3	1/6	1/6	1/6	1/6	1/6	1/6
4	-	-	-	-	-	-

Table 1. An example of an initial pmf for four routes ($n = 4$) and five assets ($r_j = 5$)

The possible number of assets that can be allocated on a route is dependent on the placement of assets on previous routes. To generate the plan for asset of type j in iteration l, we randomly choose the number of assets for Route 1 according to its pmf. Given the number of assets of type j on Route 1 is y, we construct a renormalized distribution for Route 2 as shown in Equation 3.5

$$\frac{\psi^{l}_{2,j,k}}{\psi^{l}_{2,j,0}+\psi^{l}_{2,j,1}+\ldots+\psi^{l}_{2,j,r_j-y}} \qquad \text{for } \forall j; k = 0,1,\ldots,r_j - y \tag{3.5}$$

and generate the number of assets on Route 2. Given the number of assets of type j on Route 1 and Route 2 is y, we construct a renormalized distribution for Route 3 as shown in Equation 3.6

$$\frac{\psi^{l}_{3,j,k}}{\psi^{l}_{3,j,0}+\psi^{l}_{3,j,1}+\ldots+\psi^{l}_{3,j,r_j-y}} \qquad \text{for } \forall j; k = 0,1,\ldots,r_j - y \tag{3.6}$$

and generate the number of assets on Route 3. We continue this procedure for all of the other routes and the other asset types to compute $X^{l,g}_{i,j}$, which is the number of assets of type j on route i in interdiction plan $X^{l,g}$. Table 2 is an example of a renormalized distribution for Route 3, from the previous example of Table 1. The black squares in Table 2 denote asset allocations that are no longer possible since some assets are allocated on previous routes. Notice that the probabilities on Route 3 still add up to one.

	Number of Assets on the Route of Type ***j***					
Route	0	1	2	3	4	5
1		1				
2			2			
3	1/3	1/3	1/3			
4	-	-	-	-	-	-

Table 2. An example of a renormalized distribution for route three with a total of four routes ($n = 4$) and five assets ($r_j = 5$)

After all of the assets are allocated on the routes, we compute the probability of success for each route. We denote the probability of success on each route in iteration l for the g^{th} interdiction plan $X^{l,g}$ by $p_i(X^{l,g})$, i =1, 2, …, n; see Appendix A, Section B for calculating these probabilities. We use the solution from either the Gittins Choice Algorithm or the Decreasing Choice Algorithm with these probabilities as the corresponding function values. The function values are then used to update the random mechanism, as we describe next. To update the random mechanism with the best interdiction plans, we only use the top 100ρ percent of function values, where $\rho \in [0,1]$ and is set by the user. Since we want to minimize the function value, the best interdiction plans are those associated with the lowest function values. We denote γ as the function value for the interdiction plan that has the $\lceil \rho G \rceil$ lowest function value, where $\lceil \; \rceil$ means to round up to the next integer. To update $\psi^{l}_{i,j,k}$, we use the percentage of times that asset k on route i is used in the plans that have a function value less than γ. Equation 3.7 shows the mathematical formula to update the probabilities:

$$\tilde{\psi}^{l}_{i,j,k} = \frac{\sum_{g=1}^{G} \left(I_{\left\{ f\left(X^{l,g};\alpha^0,\beta^0\right) \geq \gamma \right\}} I_{\left\{ X^{l,g}_{i,j} = k \right\}} \right)}{\sum_{g=1}^{G} I_{\left\{ f\left(X^{l,g};\alpha^0,\beta^0\right) \geq \gamma \right\}}} \qquad \forall i, j, k \qquad (3.7)$$

where $I_{\{Boolean\}}$ is an expression that is a one when the Boolean expression is true and zero when it is false. We compute the random mechanism, ψ^{l+1}, for iteration $l+1$ by

$$\psi^{l+1} = \lambda\tilde{\psi}^{l} + (1-\lambda)\psi^{l} \tag{3.8}$$

where $\tilde{\psi}^{l}$ is the collection $\left\{\tilde{\psi}^{l}_{i,j,k} : 1 \le i \le n, 1 \le j \le m, 0 \le k \le r_j\right\}$ given in Equation 3.7 and λ is a smoothing parameter. λ is a value between zero and one, given by the user, with λ between 0.7 and 0.9 giving the best results (Allon et el., 2004).

The stopping criterion of the CE method is based on the convergence of the sequence of ψ^{l}, which in the long term will converge to a degenerate matrix under certain assumptions (Allon et el., 2004). Our stopping criterion is when the same asset allocation for each route has the highest probability for C consecutive iterations. The CE Asset Allocation Algorithm stated in detail, below.

<u>CE Asset Allocation Algorithm</u>

1. Set initial conditions
   - n: Number of routes (i = 1, 2, …, n)
   - m: Number of asset types (j = 1, 2, …, m)
   - r_j: Total number of assets of type j (k =0, 1, …, r_j)
   - $\phi_{i,j}$: Probability for asset of type j detecting a vessel on route i
   - G: Number of plans chosen for each random mechanism
   - C: Number of times that the highest probability for each route for each asset in the probability distribution is the same before terminating (stopping criteria)
   - ρ: Percentage of plans that have the lowest reward
   - λ: A smoothing parameter for Equation 3.8
2. Set c =1, g =1, and l =1
3. Construct the random mechanism ($\psi^{l}_{i,j,k}$) using a uniform distribution for each asset type (j) on each route (i) distributed across the asset allocations (k)
4. Choose $X^{l,g}$ according to ψ^{l}
5. Calculate $p_i(X^{l,g})$ for all i using Equation A.1 with $p_i(X^{l,g})$ replacing $p_i(X)$
6. Calculate the value from the Gittins Choice Algorithm or Decreasing Choice Algorithm with $p_i\left(X^{l,g}\right)$ replacing p_i for all i
7. Increase g by 1
8. If $g < G$ then go back to step 4. Else, solve for γ and increase l by 1
9. Update ψ^{l} according to Equation 3.8

10. If $\arg\max_k \ \psi^{l-1}_{i,j,k} = \arg\max_k \ \psi^{l}_{i,j,k}$ for each asset type (j) on each route (i) then increase c by 1. Else, reset $c = 1$.
11. If $c < C$ then reset $g = 1$ and go to step 4. Else, stop and output the interdiction plan $\widehat{X} = \left(\widehat{X}_{1,1}, \widehat{X}_{1,2}, \ldots, \widehat{X}_{n,m}\right)$, where $\widehat{X}_{i,j} = \arg\max_k \ \psi^{l}_{i,j,k}$.

IV. NUMERICAL RESULTS

A. CASE STUDY

In this chapter, we present a case study that addresses a specific scenario, the inputs into the algorithms, the results of multiple runs, and the insights resulting from this analysis. The evader in our scenario is a DTO who is using SPSSs to transport his drugs out of Colombia by the sea. The interdictor in our scenario is the Coalition Force who is using its assets to interdict the SPSSs.

1. Routes

The environment consists of five drug shipment routes from Colombia. Figure 7 shows the location of the routes graphically. Route 1 goes to the Galapagos Islands and then up to Central America. Route 2 goes through the Eastern Pacific along the coast to Central America. Route 3 goes through the Caribbean Sea along the coast to Central America. Route 4 goes through the Caribbean Sea into the Gulf of Mexico to Mexico. Route 5 goes through the Caribbean Sea to Cuba.

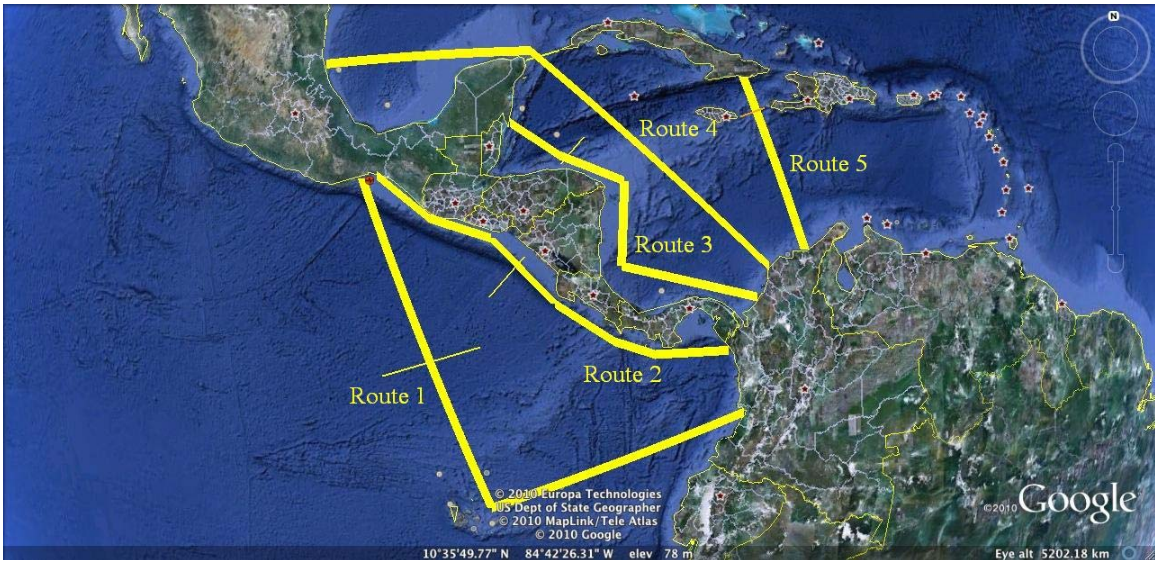

Figure 7. A map showing the five routes (From Google Earth).

2. Evader

We solve the Adapting Evader Model with the Gittins Choice Algorithm and also the Decreasing Choice Algorithm, which yield an optimal and heuristic policy, respectively. The DTO is either patient ($a = 0.99$ and $\varepsilon_0 = 10$) or impatient ($a = 0.50$ and $\varepsilon_0 = 0$). When the DTO is patient, he explores more often. For instance, using the Decreasing Choice Algorithm with $\varepsilon_0 = 10$ the DTO randomly explores for at least ten time periods. If the DTO is impatient then a success now significantly outweighs the value of a success later and thus there is no time to explore; the DTO always exploits. The DTO may have a predetermined belief which is manifested in α^0 and β^0 of the beta distribution. A predetermined belief is when the DTO initially is more likely to choose one route over another. Table 3 shows the α^0 and β^0 for the DTO with a predetermined belief, while Table 4 shows the α^0 and β^0 for a DTO without a predetermined belief.

Route	**Initial Alpha**	**Initial Beta**	**Beta Distribution**	
			Mean	**Variance**
1	15	5	0.75	0.0089
2	45	5	0.90	0.0018
3	10	10	0.50	0.0119
4	5	5	0.50	0.0227
5	10	30	0.25	0.0046

Table 3. Initial belief on the routes for a DTO with a predetermined belief

Route	**Initial Alpha**	**Initial Beta**	**Beta Distribution**	
			Mean	**Variance**
1	1	1	0.50	0.0833
2	1	1	0.50	0.0833
3	1	1	0.50	0.0833
4	1	1	0.50	0.0833
5	1	1	0.50	0.0833

Table 4. Initial belief on the routes for a DTO without a predetermined belief

There are eight different cases obtained from the three characteristics – decision process, patience, and predetermined belief—of the DTO, which are shown in Table 5.

Case	Decision Process	Patience	Predetermined Belief
1	Optimal	Patient	Yes
2	Optimal	Patient	No
3	Optimal	Impatient	Yes
4	Optimal	Impatient	No
5	Heuristic	Patient	Yes
6	Heuristic	Patient	No
7	Heuristic	Impatient	Yes
8	Heuristic	Impatient	No

Table 5. Cases for the DTO using the Gittins Choice Algorithm for the optimal decision process and the Decreasing Choice Algorithm for the heuristic decision process

3. Interdictor

The Coalition Force has three asset types (j =1, 2, 3) for their use: four Navy P-3 aircraft ($j=1$), two Navy E-2 aircraft ($j=2$), and two Coast Guard Cutters ($j=3$). There are 15,750 possible interdiction plans for the eight assets allocated on the five routes. Appendix C explains how to compute the total number of plans.

Each individual asset type has a probability of detecting an SPSS, which depends upon the route it is searching. The assets perform a barrier search on their assigned route. To calculate the probability of detecting an SPSS we use the inverse cube law for an aircraft and the linear law for a ship. These probabilities are listed in Table 6. For a more detailed explanation of these probabilities, see Appendix A.

	Asset Type (j)		
Route (i)	P-3 ($j=1$)	E-2 ($j=2$)	CG Cutter ($j=3$)
1	0.1184	0.2903	0.0420
2	0.2341	0.5434	0.0836
3	0.1279	0.3146	0.0456
4	0.1279	0.3146	0.0456
5	0.1595	0.6854	0.0569

Table 6. The probabilities of each Coalition Force asset detecting an SPSS on each route for the case study ($\phi_{i,j}$)

B. NUMERICAL RESULTS

The results depicted below are obtained from running the CE Asset Allocation Algorithm with either the Gittins Choice Algorithm or the Decreasing Choice Algorithm described in Chapter III, using VBA for Microsoft Excel. All of the simulations use the same seed of 300 for the random number generator. The values for the parameters for the Gittins Choice Algorithm and the Decreasing Choice Algorithm are: $T = 298$; a=0.99 for a patient DTO and a=0.50 for an impatient DTO; ε_0=10 for a patient DTO and ε_0=0 for an impatient DTO; α_i^0 and β_i^0 are in Table 3 and Table 4. We use a value of 298 for T, because at time period 298 a discounted success is only worth 0.05. The values in the parameters for the CE Asset Allocation Algorithm are: n=5; m=3; r_1=4, r_2=2, and r_3=2; $\phi_{i,j}$ are in Table 6; $G = 1000$; $C = 20$; $\rho = 0.01$; and $\lambda = 0.9$.

1. Optimal Interdiction Plans

Tables 7 – 14 give the interdiction plans from the CE Asset Allocation Algorithm for the cases from Table 5. For instance, Plan 1 is the interdiction plan that the CE Asset Allocation Algorithm develops for Case 1. We use the Gittins Choice Algorithm in the CE Asset Allocation Algorithm to choose Plans 1 – 4. We use the Decreasing Choice Algorithm in the CE Asset Allocation Algorithm to choose Plans 5 – 8. We note that a patient DTO (as in Cases 1, 2, 5, and 6) would have had about 94 discounted successes if he is successful for every time period. We also note that an impatient DTO (as in Cases 3, 4, 7, and 8) would have about 1 discounted success if he is successful for every time period. We denote this value as the maximum possible number of discounted successes.

Table 7 gives the interdiction plan corresponding to Case 1. The majority of the assets are located on Routes 1 and 2. The only other route being searched is Route 4 with an E-2.

	Route				
	1	**2**	**3**	**4**	**5**
Navy P-3	4	0	0	0	0
Navy E-2	0	1	0	1	0
Coast Guard Cutter	1	1	0	0	0
Probability of Success: $p_i(X)$	0.5788	0.4184	1.0000	0.6854	1.0000

Table 7. Interdiction plan for Case 1 with an expected number of discounted successes of 50.1

Table 8 gives the interdiction plan corresponding to Case 2. The assets are equally dispersed amongst all of the routes.

	Route				
	1	**2**	**3**	**4**	**5**
Navy P-3	0	1	0	1	2
Navy E-2	1	0	1	0	0
Coast Guard Cutter	0	1	0	1	0
Probability of Success: $p_i(X)$	0.7097	0.7018	0.6854	0.8324	0.7064

Table 8. Interdiction plan for Case 2 with an expected number of discounted successes of 72.6

Table 9 gives the interdiction plan corresponding to Case 3. All of the assets are located on Routes 1 and 2.

	Route				
	1	**2**	**3**	**4**	**5**
Navy P-3	2	2	0	0	0
Navy E-2	0	2	0	0	0
Coast Guard Cutter	2	0	0	0	0
Probability of Success: $p_i(X)$	0.7134	0.1223	1.0000	1.0000	1.0000

Table 9. Interdiction plan for Case 3 with an expected number of discounted successes of 0.123

Table 10 gives the interdiction plan corresponding to Case 4. The assets are dispersed amongst all of the routes except Route 4.

	Route				
	1	**2**	**3**	**4**	**5**
Navy P-3	2	2	0	0	0
Navy E-2	0	0	1	0	1
Coast Guard Cutter	0	0	0	0	2
Probability of Success: $p_i(X)$	0.77729	0.58656	0.68539	1	0.54486

Table 10. Interdiction plan for Case 4 with an expected number of discounted successes of 0.745

Table 11 gives the interdiction plan corresponding to Case 5. The assets are dispersed amongst all of the routes except Route 5.

	Route				
	1	**2**	**3**	**4**	**5**
Navy P-3	4	0	0	0	0
Navy E-2	0	1	0	1	0
Coast Guard Cutter	1	0	1	0	0
Probability of Success: $p_i(X)$	0.5788	0.4566	0.9544	0.6854	1.0000

Table 11. Interdiction plan for Case 5 with an expected number of discounted successes of 59.9

Table 12 gives the interdiction plan corresponding to Case 6. The assets are equally dispersed amongst all of the routes.

	Route				
	1	**2**	**3**	**4**	**5**
Navy P-3	1	1	0	0	2
Navy E-2	0	0	1	1	0
Coast Guard Cutter	1	1	0	0	0
Probability of Success: $p_i(X)$	0.8446	0.7018	0.6854	0.6854	0.7064

Table 12. Interdiction plan for Case 6 with an expected number of discounted successes of 72.5

Table 13 gives the interdiction plan corresponding to Case 7. All of the assets are located on Routes 1 and 2.

	Route				
	1	**2**	**3**	**4**	**5**
Navy P-3	3	1	0	0	0
Navy E-2	0	2	0	0	0
Coast Guard Cutter	1	1	0	0	0
Probability of Success: $p_i(X)$	0.6565	0.1463	1.0000	1.0000	1.0000

Table 13. Interdiction plan for Case 7 with an expected number of discounted successes of 0.151

Table 14 gives the interdiction plan corresponding to Case 8. The assets are equally dispersed amongst all of the routes.

	Route				
	1	**2**	**3**	**4**	**5**
Navy P-3	1	1	0	1	1
Navy E-2	0	1	1	0	0
Coast Guard Cutter	1	0	0	0	1
Probability of Success: $p_i(X)$	0.8446	0.3497	0.6854	0.8721	0.7927

Table 14. Interdiction plan for Case 8 with an expected number of discounted successes of 0.745

Tables 9 and 13 are very similar with all of the assets assigned to the first two routes. Table 7 also has all of the assets assigned to the first two routes except for one asset on route four. All three of these cases have a DTO that has a predetermined belief. Table 11 also has a DTO with a predetermined belief, but the assets are dispersed over four out of the five routes. Most of the assets are on Route 1. Tables 8, 12, and 14 are similar with the assets evenly dispersed over the routes. The DTO in each of these cases does not have a predetermined belief. In Table 10, the DTO does not have predetermined belief, and the assets are allocated on all of the routes except for one.

2. Comparing the Plans

To better understand which characteristics of the DTO drives the CE Asset Allocation Algorithm to choose an interdiction plan, we create sixty-four sub cases. Each sub case is a combination of a case and a plan. We use the Gittins Choice Algorithm and the Decreasing Choice Algorithm to estimate their expected number of successes with 1000 replications. See Appendix D for the full case study results.

Since the sub cases do not all have the same maximum possible value for the number of discounted successes, we standardize the values. We use percentage of discounted successes as defined by the fraction

$$\frac{\textit{number of discounted successes}}{\textit{maximum possible number of discounted successes}} \tag{5.1}$$

To explore the results, we group the sub cases in pairs. In each pair, the sub cases have one out of the three characteristics of the DTO different, and the other two the same. There are eight pairs for each set of analysis. To be 95% confident in our findings, involving eight combinations, we use a confidence level of $100\%(1-0.025/8)$ for the individual sub cases.

a. Optimal Versus Heuristic Decision Process by the DTO

The sub cases in Figure 8 are grouped by the DTO's decision process. In one sub case the DTO chooses routes optimally and in the other sub case the DTO chooses routes heuristically. The other two characteristics, which are the DTO's patience and initial belief, are the same within the group. For instance, Case 1 is a patient DTO that chooses the routes optimally and has a predetermined belief. Plan 1 is the interdiction plan against Case 1. Case 5 is a patient DTO that chooses the routes heuristically and has a predetermined belief. Plan 5 is the interdiction plan against Case 5.

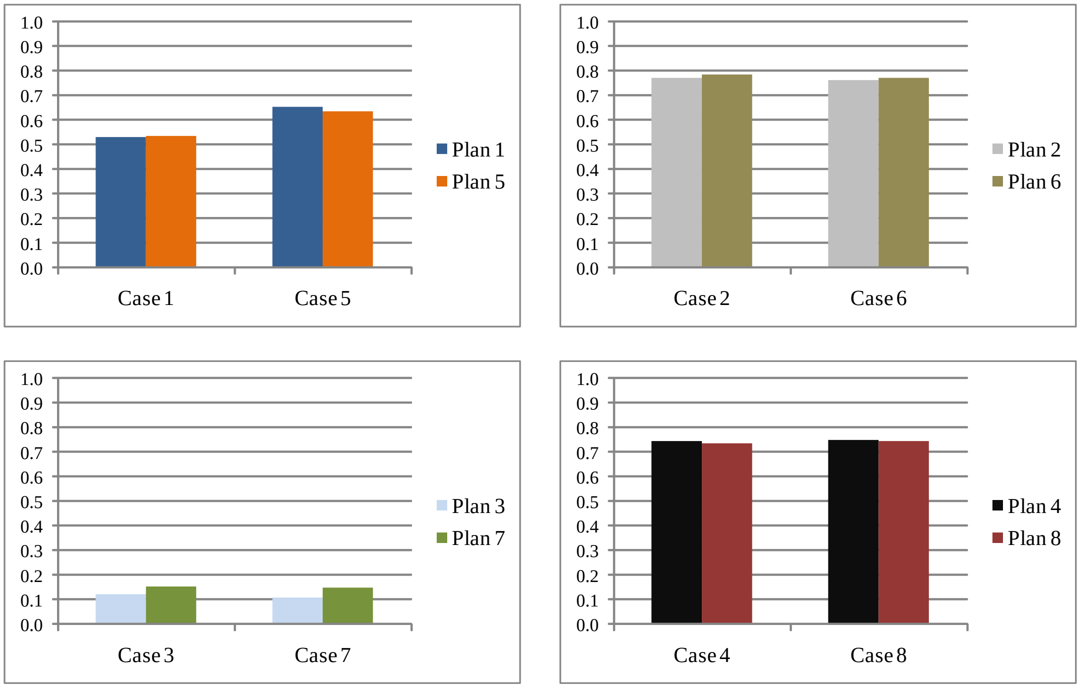

Figure 8. Pairings of optimal versus heuristic sub cases

All of the pairs have a close percentage of discount successes, with four out of the eight pairs statistically indistinguishable. Those four pairs are Case 1 with Plans 5 and 6, Case 3 with Plans 3 and 7, Case 4 with Plans 4 and 8, and Case 8 with Plans 4 and 8. When comparing the interdiction plans, the plans chosen by the CE Asset Allocation Algorithm where the case only differs by the decision process are very similar; see Tables 7 – 14. This is indicative that the decision process of the DTO does not play a role when the CE Asset Allocation Algorithm chooses the interdiction plan.

b. Patient Versus Impatient DTO

The sub cases in Figure 9 are grouped by the DTO's patience. In one sub case the DTO is patient and in the other sub case the DTO is impatient. The other two characteristics, which are the DTO's decision process and initial belief, are the same within the group.

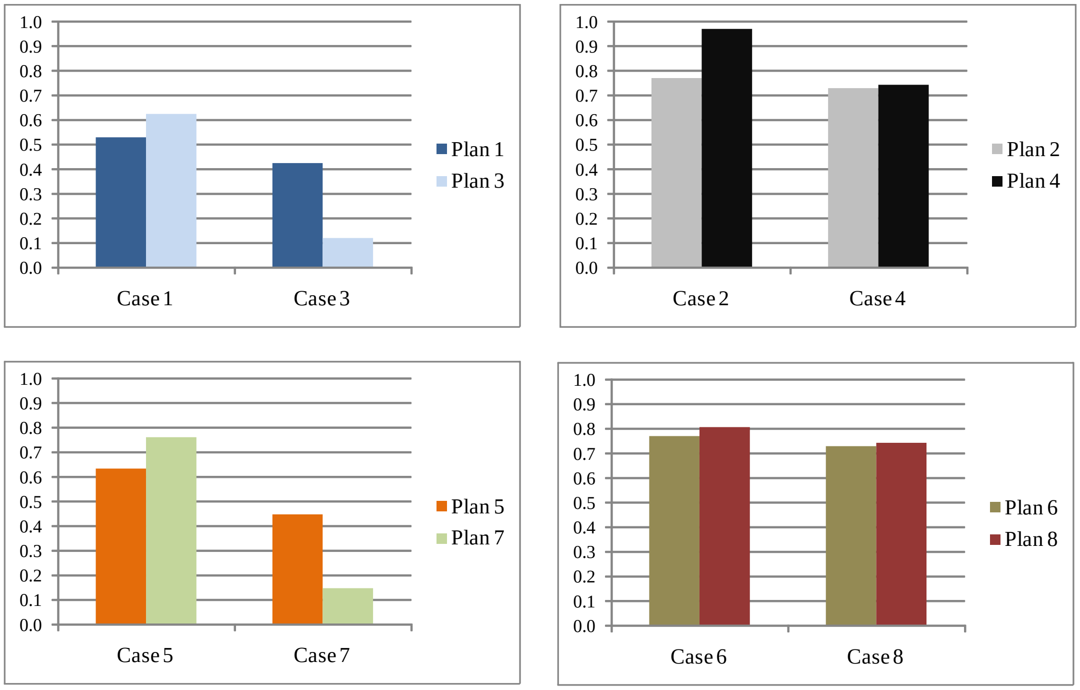

Figure 9. Pairings of patient versus impatient sub cases

Two out of the eight pairs are statistically indistinguishable. Those two are Case 4 with Plans 2 and 4, and Case 8 with Plans 6 and 8. When comparing the interdiction plans, the plans chosen by the CE Asset Allocation Algorithm where the case only differs by the patience are close with a couple assets different; see Tables 7 – 14. The patience of the DTO thus plays a role when the CE Asset Allocation Algorithm chooses the interdiction plan.

c. Predetermined Belief Versus No Predetermined Belief

The sub cases in Figure 10 are grouped by the DTO's initial belief. In one sub case the DTO has a predetermined belief and in the other sub case the DTO does not have a predetermined belief. If a DTO has a predetermined belief then he starts the scenario more willing to use some routes over other routes. A DTO with no predetermined belief starts the scenario willing to try every route equally. The other two characteristics, which are the DTO's decision process and patience, are the same within the group.

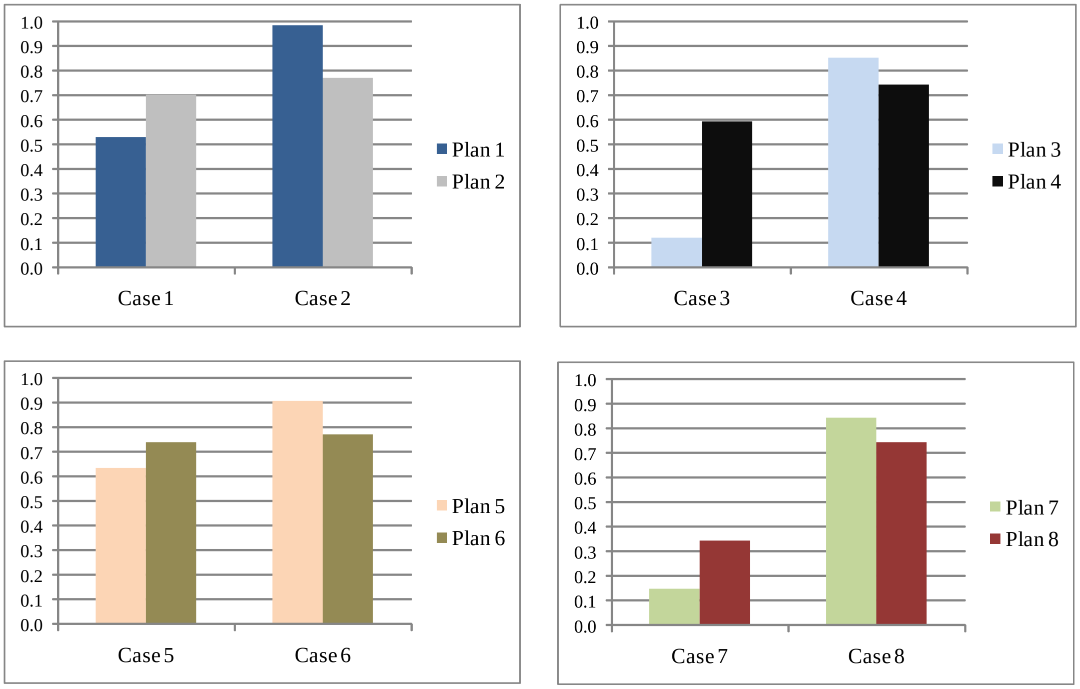

Figure 10. Pairings of predetermined belief versus no predetermined belief

None of the eight pairs are statistically indistinguishable. When comparing the interdiction plans, the plans chosen by the CE Asset Allocation Algorithm where the case only differs by the initial belief are very different; see Tables 7 – 14. The plans against a DTO with a predetermined belief allocate most if not all of the assets on Route 1 and Route 2. Since a DTO with a predetermined belief will first try these routes, the Coalition Force gets the most benefit with assets on these routes. The plans against a DTO that does not have a predetermined belief disperses the assets among all of the routes, since the DTO has an equal chance at initially trying any of the routes. The initial belief for the DTO plays a role when the CE Asset Allocation Algorithm chooses the interdiction plan.

3. The DTO’s Perspective

In the following analysis, we compare the eight cases, which are defined in Table 5, from the DTO’s perspective. With eight cases, there are $8 \cdot 7/2=28$ combinations. To be 95% confident in our findings involving twenty-eight combinations, we use a confidence level of 100%$(1-0.025/28)$ for the individual sub cases. To get the average number of discounted successes for Case 1, we take an average over the percentages of discounted successes for the sub cases that involve Case 1.

a. Discounted Successes

Figure 11 displays the average percentage of discounted successes by case. The only cases that are statistically indistinguishable are Case 8 with Case 4 and Case 3 with Case 7.

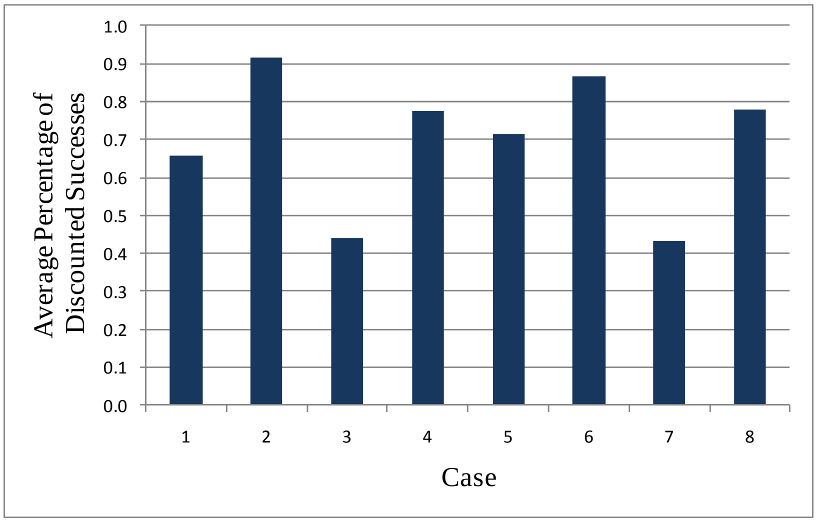

Figure 11. Average Percentage of Discounted Successes by Case

Case 2 is the best case for the DTO. Case 2 is a patient DTO that chooses the routes optimally and does not have a predetermined belief. Exploring the data further, the top four cases have one common feature; the DTO does not have a predetermined belief. This allows the DTO to be flexible and find the best route. If this DTO is playing against a plan created to combat a predetermined belief then he can easily find the uncovered routes. On the other hand, if this DTO is playing against a plan that was created without a predetermined belief then the limited interdiction assets are spread across the routes, which limit the probability of detecting a vessel on each route.

b. Finding the Best Route

Figure 12 is the average time period it took the DTO to find the best route. For each case, we take an average of the time periods that the DTO finds the best route of the associated eight sub cases. The only cases that are statistically indistinguishable are Case 3 with Case 7 and Case 7 with Case 1.

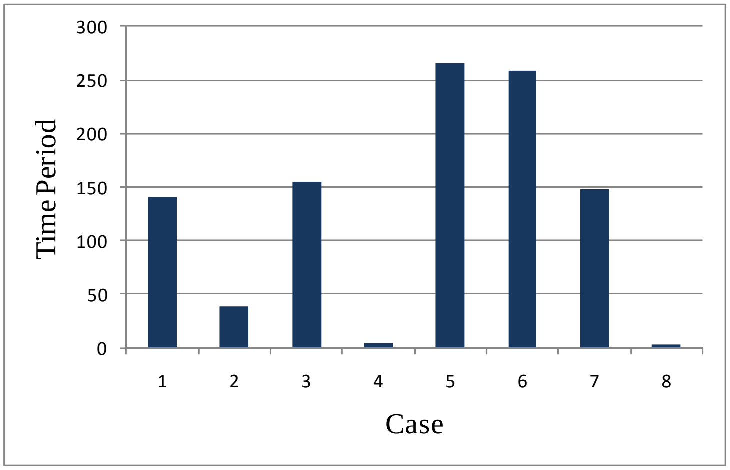

Figure 12. The average time the DTO found the best route given he found the best route

Even if the DTO finds the best route quickly when he finds the best route, does not mean he will have a high average percentage of discounted success. Case 2 has the highest average percentage of discounted successes and a very low time to find the best route. Case 6, though, has the second highest average percentage of discounted successes and the second highest time to find the best route. This disparity has two possible explanations. The first possible explanation is that even if the best route was not found does not mean the DTO did not find a very good route. The DTO would thus not do as well as he could have if he found the best route, but could still do very well. The second possible explanation is the collection of the data. The DTO did not always find the best route. There is no data for those cases, so the information is skewed. To help mitigate this problem Figure 13 shows the average percentage of the number of times the DTO finds the best route.

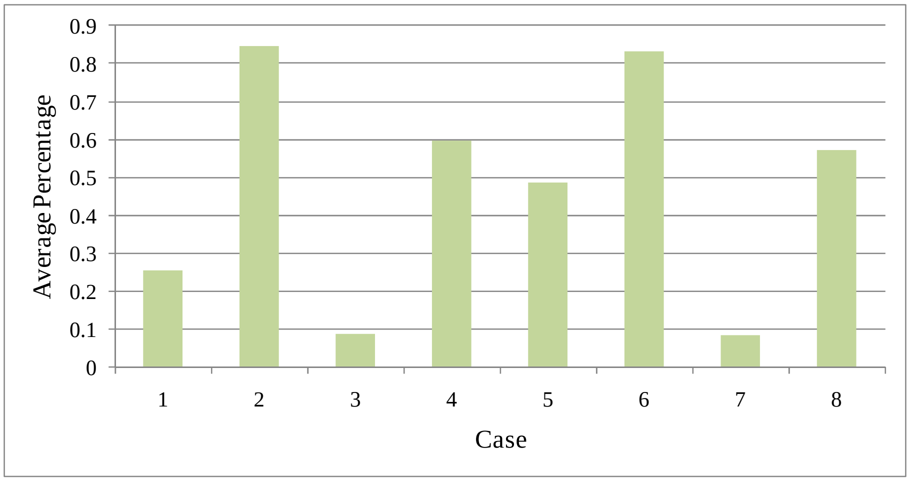

Figure 13. Average Percentage the DTO finds the best route

As Figure 14 shows, though, the more times that the DTO finds the best route, even if he finds it late, the higher the percentage of discount successes. Figure 14 also confirms that Case 2 is the best for the DTO.

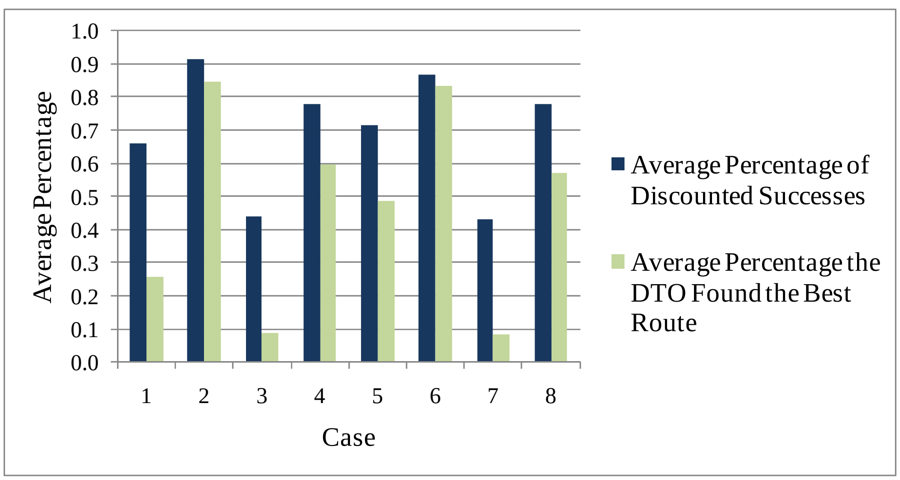

Figure 14. Trend of discount successes with finding the best route

C. SUMMARY

While planning, the Coalition Force needs to take into account a DTO that is capable of learning the location of assets. The pace of the DTO leaning the location of the Coalition Force is dependent on the DTO's characteristics. By determining which characteristics of the DTO to center on, the Coalition Force can focus its intelligence gathering efforts and ultimately make a more informed decision on the interdiction plan.

The primary characteristic to focus on is the DTO's initial belief. If the DTO has a predetermined belief and the Coalition Force can ascertain that belief, then they may be able to ignore some possible routes and focus their efforts in a particular region. This is currently being seen, with the majority the SPSSs found in the Eastern Pacific.

The other characteristic to be aware of is the patience of the DTO. Does the DTO feel that he quickly needs to make successful runs? If he does, then he will be more willing to use the routes that he currently knows about. For instance, a patient DTO will be willing to try the SPSS in the Caribbean Sea, while an impatient DTO will stick with using the Eastern Pacific.

The one aspect about the DTO that does not need to be considered is how he chooses his routes, optimally or heuristically. The choice will affect the number of discounted successes, but not the development of the interdiction plan.

The length of time that the interdiction plan should be used varies greatly and is not an easy question to answer. One way to try to answer it is with a worst-case scenario. The DTO's best case (and thus worst case for the Coalition Force) is Case 2. Case 2 finds the best route on average in about 40 time periods. From the time that the DTO finds best route until the Coalition Force changes its interdiction plan, the DTO will maximize his expected number of successes. To keep the DTO from maximizing his expected number of successes the Coalition Force needs to change their interdiction plan. By changing the allocation of the Coalition Force's assets before the DTO can take advantage of learning the best route, the Coalition Force can increase the number of SPSSs that they interdict.

THIS PAGE INTENTIONALLY LEFT BLANK

V. CONCLUSION AND RECOMMENDATIONS

A. CONCLUSION

This thesis focuses on a portion of the drug trafficking problem. It looks at the shipment of cocaine from South America into the United States using SPSS. The SPSSs are difficult to detect, and they can carry a huge amount of cargo. Currently the cargo is only drugs, but there is potential for the cargo to be other contraband items. The Coalition Force needs to develop better ways to stop the use of the SPSS.

In this thesis, we construct models and then develop algorithms to gain insight on a DTO that is capable of learning the placement of interdiction assets. We first use our algorithms to develop interdiction plans given different assets based on a DTO that is capable of learning and adapting. We then use these algorithms to analyze the how well a DTO performs against the plans. By analyzing how a DTO, that is capable of learning, reacts to different optimal plans, we can gain a better understanding of how the Coalition Force can more effectively allocate their limited number of assets to impede the effectiveness of the SPSS. One such insight is the interdiction plan is not dependent on the technique the DTO uses to choose his routes. Other attributes such as his patience and prior belief of success on the routes play a more substantial role in designing the interdiction plan. Another insight is the time it takes the DTO to find the route with the highest probability of success. In the Coalition Force's worst-case scenario, a patient DTO that chooses the routes optimally and does not have a prior belief of success takes about forty tries. This is the time the interdictor should change his interdiction plan.

B. SUGGESTED WORK AHEAD

Since the DTO has more than one vessel that he uses, an extension of the thesis should allow the evader to choose multiple vessels. Each vessel costs a different amount to build, has a different probability of detection, and a different payout if the vessel is successful.

In a possible extension, the DTO's decisions could be split into two phases. The first phase is the planning phase. In the planning phase, the DTO decides which vessels to build. The second phase is the deployment phase. In each time period of this phase, the DTO sends a vessel over a route to get the maximum expected number of discounted successes.

Allowing the DTO to have multiple vessels will let us have an insight into which vessel is the most beneficial to the DTO. The Coalition Force can then focus their efforts on stopping that vessel to best impede the drug trade.

APPENDIX A. PROBABILITY OF DETECTION

This appendix presents the mathematical tools used to calculate the probability of detection for each asset and the probability of success for the evader on each route. It also describes how the probabilities were calculated for the case study.

A. BARRIER PATROL

Since evader's vessels are traversing the routes, the interdictor's assets are going to perform a barrier patrol across their assigned routes. The assumptions for a barrier patrol are the evader's probability of detection does not change and the time the evader transits the barrier is unknown, but equally likely throughout the time interval under consideration (Wagner et al., 1999). To calculate the probability of detection some parameters need to be known. One of which is width of the route that is being patrolled which is d_i. The speed of the interdictor's asset of type j (v_j) and the speed of the evader's vessel (u) also need to be known. We denote the sweep width associated with the interdictor's asset on route i as $w_{i,j}$. Depending on the method chosen to calculate the probability of detection, more parameters may need to be known.

1. Inverse Cube Law

If the speed of the interdictor's asset of type j is much greater than the evader's speed then the inverse cube law can be used to calculate the asset's probability of detecting the evader. This is common when the interdictor asset is an aircraft and the evader's vessel is a ship. For the inverse cube law, two other parameters are used. The first is the number of assets that are in the barrier patrol (η). The last parameter needed is the track spacing (s). The probability of detecting an evader on route i with an interdiction asset of type j is

$$\phi_{i,j} = 2\int_0^z \varphi(y)dy \qquad \text{(B.1)}$$

where φ is the standardized, normal probability density function with mean zero and variance one, and (see (Wagner et al., 1999))

$$z = \sqrt{\frac{\pi}{2}} \frac{w_{i,j}}{s} \tag{B.2}$$

and

$$s = \sqrt{\frac{v_j + u}{v_j - u}} \left(\frac{d_i u}{\eta_j v_j} \right) \tag{B.3}$$

2. Linear Law

If the speed of the interdictor's asset of type j is about the same as the evader's speed then the linear law can be used to calculate the asset's probability of detecting the evader. This is common when the interdictor asset is a ship and the evader's vessel is a ship. The probability of detecting an evader on route i with an interdiction asset of type j is (see (Wagner et al., 1999))

$$\phi_{i,j} = \begin{cases} 1 - \left(h - \frac{\sqrt{b^2 + 1} - 1}{2} \right)^2 \frac{1}{h(h+1)}, & \text{if } b \le 2\sqrt{h(h+1)} \\ 1, & \text{otherwise} \end{cases} \tag{B.4}$$

where

$$b = \frac{v_j}{u} \tag{B.5}$$

and

$$h = \frac{(d_i - w_{i,j})}{w_{i,j}} \tag{B.6}$$

B. PROBABILITY OF DETECTING A VESSEL ON A ROUTE

Once the interdictor has decided on an interdiction plan X, then the actual probability of detecting a vessel on each route needs to be computed. We assume the probability of detecting a vessel is independent among the interdictor's assets. We also assume that when a vessel is detected, it is captured. Using these two assumptions, the probability that an evader is transits route i successfully for an interdiction plan X is

$$p_i(X) = \prod_{j=1}^{m} (1 - \phi_{i,j})^{X_{i,j}} \qquad \forall i \tag{A.1}$$

C. CASE STUDY'S PROBABILITIES

1. Routes

There are five routes for the case study presented in Chapter IV. Each of these routes has a width associated with it. These routes are in either the Eastern Pacific or the Caribbean Sea. Table 15 displays this information for the routes. The widths are in nautical miles.

Route (i)	Width (d_i)	Location
1	300	Eastern Pacific
2	150	Eastern Pacific
3	100	Caribbean Sea
4	100	Caribbean Sea
5	80	Caribbean Sea

Table 15. Width and area for each route in the case study

As seen in Figure 15, the two locations have different sea surface winds. For this thesis, the Eastern Pacific has winds that are approximately 10 knots (5.14 meters/second) and the Caribbean Sea has winds that are approximately 15 knots (7.72 meters/second).

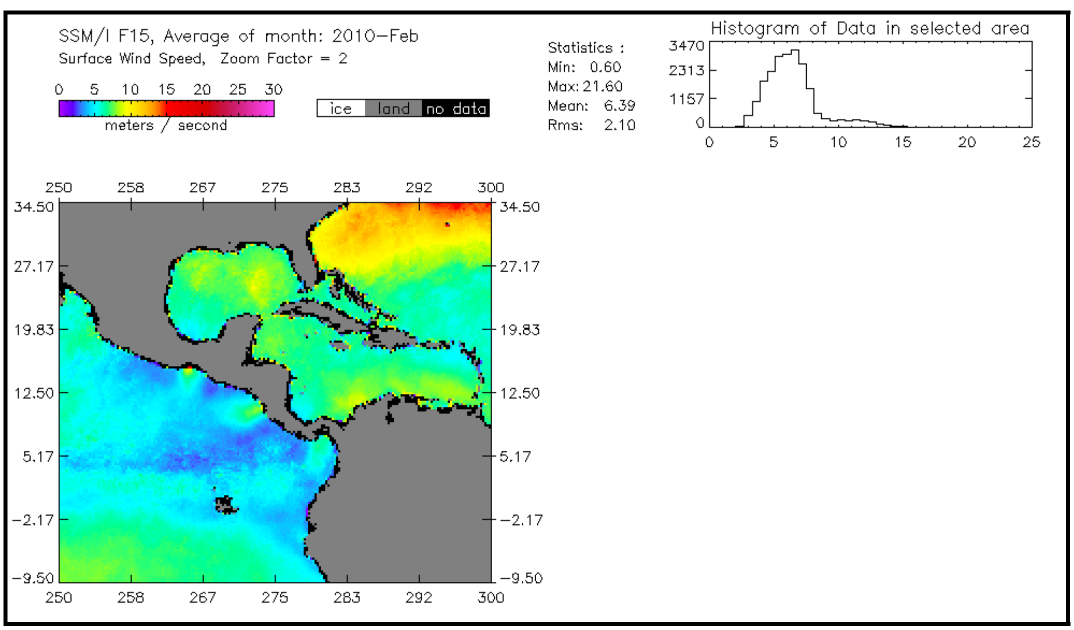

Figure 15. Average winds for February 2010 (From Remote Sensing System, 2010)

2. Evader

In the case study, the DTO has an unlimited supply of SPSSs that he wants to use to ship his drugs. The SPSS is going to traverse the route at 6 knots ($u = 6$), which is the same assumption that Pfeiff (2009) makes. This speed is a little less than half of the maximum speed of an SPSS.

3. Interdictor

In the case study, the coalition force has three asset types ($m = 3$). There are four Navy P-3 aircraft ($j = 1$), two Navy E-2 aircraft ($j = 2$), and two Coast Guard Cutters ($j = 3$). The Navy P-3s will search for the SPSSs at a speed of 180 knots ($v_1 = 180$), which is the same assumption that Pfeiff (2009) makes. They will not be on station searching for the whole time due to fuel requirements, crew rest, and limited flight hours. For this reason the number of aircraft per P-3 will be a fraction ($\eta_1 = 1/7$), since the aircraft will be on station for only a fraction of the time. The P-3s that are deployed in JIATF-South's AOR are equipped with the APS-137 radar. Since the coalition force are looking for SPSSs, which have a small cross section above the water, the sweep widths for the P-3s are 8.6 NM for the Eastern Pacific routes and 3.1 NM for the Caribbean Sea routes. This means $w_{1,1} = w_{2,1} = 8.6$ and $w_{3,1} = w_{4,1} = w_{5,1} = 3.1$. These numbers are pulled from Table 16 for the "4 to 10 person life raft" and winds "to 10" and "to 15."

Table H-27 Sweep Widths for Forward-Looking Airborne Radar (AN/APS-137)

16 Nautical Mile Radar Range Scale (Sweep Width in Nautical Miles)										
	On scene Surface Winds (kts)									
Object Type	*< 5*	*to 10*	*to 15*	*to 20*	*to 25*	*to 35*	*to 45*	*to 55*	*to 65*	*> 65*
4 to 10 person life raft	12.1	8.6	3.1	0	0	0	0	0	0	0
17 to 25 foot recreational boat	13.6	11.9	8.2	2.8	0	0	0	0	0	0
26 to 35 foot recreational boat	16.6	16.3	15.4	14.2	12.6	9.5	3.9	0	0	0
36 to 50 foot recreational boat	21.0	20.7	19.9	18.9	17.5	14.7	9.8	3.5	0	0

32 Nautical Mile Radar Range Scale (Sweep Width in Nautical Miles)										
	On scene Surface Winds (kts)									
Object Type	*< 5*	*to 10*	*to 15*	*to 20*	*to 25*	*to 35*	*to 45*	*to 55*	*to 65*	*> 65*
17 to 25 foot recreational boat	17.4	15.7	12.0	6.6	0	0	0	0	0	0
26 to 35 foot recreational boat	22.1	21.7	20.9	19.7	18.1	14.9	9.3	2.1	0	0
36 to 50 foot recreational boat	29.0	28.7	27.9	26.9	25.5	22.7	17.8	11.5	3.8	0

Table 16. APS-137 sweep widths (From USCG Addendum, 2004)

The Navy E-2s will search for the SPSSs at a speed of 180 knots ($v_2 = 180$),which is the same assumption that Pfeiff (2009) makes. The E-2s will also not be on station searching for the whole time due to the same limitations as the P-3s. For this reason the number of aircraft per E-2 will also be a fraction ($\eta_2 = 1/7$). The E-2s are equipped with the APS-145 radar. Using a factor of 2.5 to the P-3's sweep width, we obtain the E-2's sweep width since the APS-145 antenna area is roughly 20 times the APS-137 antenna area as explained by Pfeiff (2009). The sweep widths for the E-2s are 21.5 NM for the Eastern Pacific routes and 7.8 NM for the Caribbean Sea routes. This means $w_{1,2} = w_{2,2} = 21.5$ and $w_{3,2} = w_{4,2} = w_{5,2} = 7.8$.

We assume that the Coast Guard Cutter will search for the SPSSs at a speed of 15 knots ($v_3 = 15$). Coast Guard Cutters are equipped with the SPS-73 radar. Looking at Table 17, the sweep width for a raft with a reflector (such as the exhaust tubing on the SPSS) and in moderate rain (worst case scenario) is 4.7 NM and 1.7 NM for winds "to 10" and "to 15" respectively. This means $w_{1,3} = w_{2,3} = 4.7$ and $w_{3,3} = w_{4,3} = w_{5,3} = 1.7$.

Table H-25 Sweep Widths and Recommended Settings for AN/SPS-73 Radar (4-10 person life rafts with and without radar reflectors)

		SWEEP WIDTHS FOR AN/SPS-73 RADAR (Nautical Miles)			
WEATHER	OBJECT TYPE	On scene Surface Winds (kts)			
		<5	to 10	to 15	>15
No Rain or Drizzle	Raft w/ reflector	10.6	8.6	5.8	unknown
	Raft w/o reflector	5.1	2.5	0.9	nil
Moderate Rain	Raft w/ reflector	8.3	4.7	1.7	unknown
	Raft w/o reflector	4.3	1.5	0.3	nil

RECOMMENDED SETTINGS		
	Range Scale:	6 NM range scale
	Pulse Width:	M1 pulse width (AUTO)
	STC:	Zero
	FTC: for rain	Less than 80% for no rain, at least 80%
	Persistence:	No higher than 15
	Interference Rejection:	ON at 100%

Table 17. APS-137 sweep widths (From USCG Addendum, 2004)

Substituting the above values into Equation B.1 for the P-3 and E-2 and into Equation B.4 for the Coast Guard Cutter returns the probability of detection for the assets on each route, as seen in Table 6.

APPENDIX B. GITTINS INDEX TABLES

This appendix presents the equation that estimates Gittins Index (Equation 3.4) and pre-calculated Indices for $a = 0.50$ (Tables 18 – 19) and $a = 0.99$ (Tables 21 – 23).

Equation 3.4 is valid when $a = 0.50$ and $\alpha > 20$ and $\beta > 20$. It is also valid when $a = 0.99$ and $\alpha + \beta > 40$. In Equation 3.4 $v(\alpha, \beta)$ is Gittins Index.

$$\left[v(\alpha,\beta) - \alpha*(\alpha+\beta)^{-1}\right]^{-1} = (\alpha+\beta)*\alpha^{-1/2}*\beta^{-1/2}*\left(A + B*(\alpha+\beta) + C*(\alpha+\beta)^{-1}\right)*(1-a)^{1/2} \quad \text{(A.1)}$$

The values for A, B, C are also in Gittins' book (Gittins, 1989). They are Table 20 for $a = 0.50$ and Table 24 for $a = 0.99$.

α \ β	1	2	3	4	5	6	7	8	9	10
1	0.559	0.3758	0.2802	0.2223	0.1837	0.1563	0.1358	0.12	0.1074	0.0972
2	0.706	0.5359	0.4298	0.3577	0.3058	0.2668	0.2364	0.2122	0.1923	0.1758
3	0.7772	0.6289	0.5258	0.4512	0.3947	0.3504	0.3149	0.2859	0.2616	0.2411
4	0.8199	0.6899	0.5937	0.5201	0.4626	0.4163	0.3783	0.3465	0.3196	0.2965
5	0.8485	0.7333	0.6441	0.5736	0.5165	0.4697	0.4306	0.3974	0.3688	0.344
6	0.8691	0.7658	0.6832	0.6161	0.5606	0.514	0.4745	0.4407	0.4113	0.3855
7	0.8847	0.7911	0.7144	0.6507	0.5971	0.5515	0.5121	0.4781	0.4482	0.4218
8	0.8969	0.8114	0.7399	0.6795	0.6279	0.5835	0.5447	0.5107	0.4807	0.4541
9	0.9067	0.828	0.7611	0.7038	0.6543	0.6111	0.5732	0.5396	0.5096	0.4828
10	0.9148	0.8419	0.7791	0.7247	0.6771	0.6353	0.5982	0.5652	0.5355	0.5087
11	0.9216	0.8537	0.7946	0.7428	0.6971	0.6566	0.6205	0.588	0.5587	0.5321
12	0.9274	0.8638	0.808	0.7586	0.7147	0.6755	0.6403	0.6085	0.5797	0.5534
13	0.9324	0.8726	0.8197	0.7726	0.7304	0.6925	0.6582	0.6271	0.5987	0.5728
14	0.9376	0.8804	0.8301	0.785	0.7444	0.7077	0.6743	0.6439	0.6161	0.5906
15	0.9405	0.8872	0.8393	0.7961	0.757	0.7215	0.689	0.6593	0.6321	0.6069
16	0.9439	0.8933	0.8476	0.8061	0.7684	0.734	0.7024	0.6734	0.6467	0.622
17	0.9469	0.8988	0.855	0.8152	0.7788	0.7454	0.7147	0.6864	0.6602	0.6359
18	0.9496	0.9037	0.8618	0.8235	0.7883	0.7559	0.726	0.6984	0.6727	0.6488
19	0.952	0.9081	0.8679	0.831	0.797	0.7656	0.7365	0.7095	0.6843	0.6609
20	0.9542	0.9122	0.8736	0.8379	0.805	0.7745	0.7461	0.7198	0.6952	0.6721

Table 18. Pre-Calculated Gittins Index when $a = 0.50$ Part 1 (From Gittins, 1989)

α \ β	11	12	13	14	15	16	17	18	19	20
1	0.0888	0.0816	0.0756	0.0703	0.0657	0.0617	0.0582	0.055	0.0522	0.0496
2	0.1619	0.15	0.1397	0.1307	0.1228	0.1158	0.1095	0.1039	0.0988	0.0942
3	0.2236	0.2084	0.1951	0.1833	0.1729	0.1636	0.1553	0.1477	0.1409	0.1346
4	0.2764	0.2589	0.2434	0.2297	0.2174	0.2064	0.1964	0.1873	0.179	0.1714
5	0.3223	0.3032	0.2862	0.2709	0.2572	0.2448	0.2335	0.2232	0.2138	0.2051
6	0.3627	0.3423	0.3242	0.3078	0.2931	0.2796	0.2673	0.2561	0.2457	0.2362
7	0.3983	0.3773	0.3583	0.3411	0.3255	0.3113	0.2982	0.2862	0.2751	0.2648
8	0.4302	0.4086	0.3891	0.3713	0.3551	0.3402	0.3265	0.3139	0.3022	0.2913
9	0.4587	0.4369	0.417	0.3988	0.3821	0.3668	0.3526	0.3395	0.3273	0.3159
10	0.4845	0.4625	0.4424	0.424	0.407	0.3913	0.3767	0.3632	0.3506	0.3389
11	0.5079	0.4859	0.4657	0.447	0.4298	0.4139	0.3991	0.3853	0.3724	0.3603
12	0.5294	0.5073	0.487	0.4683	0.451	0.4349	0.4199	0.4058	0.3927	0.3804
13	0.549	0.527	0.5068	0.488	0.4706	0.4544	0.4392	0.4251	0.4118	0.3992
14	0.567	0.5453	0.5251	0.5063	0.4888	0.4726	0.4573	0.4431	0.4296	0.417
15	0.5837	0.5621	0.542	0.5233	0.5059	0.4896	0.4743	0.4599	0.4464	0.4337
16	0.599	0.5777	0.5578	0.5393	0.5219	0.5055	0.4902	0.4758	0.4622	0.4494
17	0.6133	0.5922	0.5726	0.5541	0.5368	0.5205	0.5052	0.4908	0.4772	0.4643
18	0.6266	0.6058	0.5863	0.568	0.5509	0.5347	0.5194	0.5049	0.4913	0.4784
19	0.639	0.6185	0.5992	0.5811	0.5641	0.548	0.5327	0.5183	0.5047	0.4917
20	0.6506	0.6304	0.6113	0.5934	0.5765	0.5605	0.5454	0.531	0.5174	0.5045

Table 19. Pre-Calculated Gittins Index when $a = 0.50$ Part 2 (From Gittins, 1989)

$\frac{\alpha}{(\alpha + \beta)}$	A	B	C
0.025	9.2681	9.2848	0.5004
0.05	6.8028	6.8134	0.1667
0.1	5.1912	5.1849	-0.033
0.2	4.2665	4.2244	-0.029
0.3	4.2057	3.9207	-0.463
0.4	4.3423	3.8095	-0.633
0.5	4.271	3.8637	-0.007
0.6	5.0689	3.8097	-0.839
0.7	5.9458	3.921	-1.102
0.8	7.6929	4.2241	-2.383
0.9	11.906	5.1853	-5.885
0.95	17.93	6.8137	-15.3
0.975	24.176	9.3011	21.783

Table 20. Pre-Calculated A, B, and C when $a = 0.50$ (From Gittins, 1989)

α \ β	1	2	3	4	5	6	7	8	9	10	11	12	13
1	0.8699	0.7005	0.5671	0.4705	0.3969	0.3415	0.2979	0.2632	0.235	0.2117	0.1922	0.1756	0.1614
2	0.9102	0.7844	0.6726	0.5806	0.5093	0.4509	0.4029	0.3633	0.3303	0.302	0.2778	0.2571	0.2388
3	0.9285	0.8268	0.7308	0.649	0.5798	0.5225	0.4747	0.4337	0.3986	0.3679	0.3418	0.3187	0.2982
4	0.9395	0.8533	0.7696	0.6952	0.6311	0.5756	0.5277	0.4876	0.452	0.4208	0.3932	0.3685	0.3468
5	0.947	0.8719	0.7973	0.7295	0.6697	0.6172	0.571	0.53	0.4952	0.464	0.4359	0.4108	0.3882
6	0.9525	0.8857	0.8184	0.7561	0.6998	0.6504	0.6061	0.5665	0.5308	0.5002	0.4722	0.4469	0.4239
7	0.9568	0.8964	0.835	0.7773	0.7249	0.6776	0.6352	0.597	0.5625	0.531	0.5034	0.4782	0.4551
8	0.9603	0.9051	0.8485	0.7949	0.7456	0.7004	0.6599	0.623	0.5895	0.5589	0.5307	0.5057	0.4827
9	0.9631	0.9122	0.8598	0.8097	0.7631	0.7203	0.6811	0.6456	0.613	0.5831	0.5556	0.5302	0.5072
10	0.9655	0.9183	0.8693	0.8222	0.7781	0.7373	0.6997	0.6653	0.6337	0.6045	0.5776	0.5527	0.5295
11	0.9675	0.9234	0.8775	0.8331	0.7912	0.7522	0.7161	0.6826	0.6521	0.6236	0.5973	0.5728	0.5501
12	0.9693	0.9278	0.8846	0.8426	0.8027	0.7653	0.7307	0.6984	0.6685	0.6408	0.615	0.5911	0.5687
13	0.9709	0.9138	0.8909	0.8509	0.8129	0.7771	0.7437	0.7125	0.6833	0.6564	0.6312	0.6077	0.5856
14	0.9722	0.9352	0.8964	0.8584	0.822	0.7877	0.7554	0.7253	0.697	0.6706	0.646	0.6229	0.6012
15	0.9735	0.9383	0.9014	0.8651	0.8302	0.7972	0.766	0.7369	0.7094	0.6835	0.6595	0.6369	0.6155
16	0.9746	0.9411	0.9059	0.8711	0.8377	0.8059	0.7758	0.7474	0.7208	0.6956	0.6719	0.6498	0.6289
17	0.9756	0.9436	0.9099	0.8766	0.8444	0.8137	0.7846	0.7571	0.7312	0.7067	0.6834	0.6617	0.6412
18	0.9765	0.9458	0.9136	0.8816	0.8506	0.821	0.7928	0.766	0.7408	0.7169	0.6942	0.6278	0.6527
19	0.9774	0.9479	0.9169	0.8861	0.8563	0.8276	0.8003	0.7743	0.7497	0.7264	0.7043	0.6832	0.6634
20	0.9781	0.9498	0.92	0.8903	0.8615	0.8338	0.8072	0.7821	0.7579	0.7352	0.7136	0.693	0.6734
21	0.9788	0.9516	0.9228	0.8942	0.8663	0.8394	0.8137	0.7891	0.7657	0.7435	0.7223	0.7021	0.6829
22	0.9795	0.9532	0.9255	0.8978	0.8708	0.8447	0.8197	0.7958	0.773	0.7512	0.7305	0.7107	0.6918
23	0.9801	0.9547	0.9279	0.9011	0.875	0.8497	0.8253	0.802	0.7797	0.7584	0.7382	0.7188	0.7003
24	0.9807	0.9561	0.9302	0.9043	0.8789	0.8543	0.8306	0.8078	0.7861	0.7652	0.7454	0.7264	0.7082
25	0.9812	0.9574	0.9323	0.9072	0.8825	0.8586	0.8355	0.8133	0.7921	0.7717	0.7522	0.7336	0.7157
26	0.9817	0.9587	0.9343	0.9099	0.8859	0.8626	0.8401	0.8185	0.7977	0.7778	0.7586	0.7403	0.7228
27	0.9822	0.9598	0.9362	0.9124	0.8891	0.8664	0.8445	0.8234	0.803	0.7835	0.7647	0.7468	0.7295
28	0.9827	0.9609	0.9379	0.9148	0.8921	0.87	0.8486	0.828	0.808	0.7889	0.7705	0.7529	
29	0.9831	0.9619	0.9396	0.9171	0.895	0.8734	0.8525	0.8323	0.8128	0.7941	0.7761		
30	0.9835	0.9629	0.9411	0.9192	0.8977	0.8766	0.8562	0.8364	0.8174	0.799			
31	0.9839	0.9638	0.9426	0.9213	0.9002	0.8797	0.8597	0.8403	0.8217				
32	0.9842	0.9646	0.944	0.9232	0.9027	0.8826	0.863	0.8441					
33	0.9846	0.9655	0.9453	0.925	0.9049	0.8853	0.8662						
34	0.9849	0.9662	0.9466	0.9267	0.9071	0.8879							
35	0.9852	0.967	0.9478	0.9284	0.9092								
36	0.9855	0.9677	0.9489	0.93									
37	0.9858	0.9684	0.95										
38	0.986	0.969											
39	0.9866												

Table 21. Pre-Calculated Gittins Index when $a = 0.99$ Part 1 (From Gittins, 1989)

α \ β	14	15	16	17	18	19	20	21	22	23	24	25	26
1	0.1491	0.1384	0.129	0.1206	0.1132	0.1066	0.1006	0.0952	0.0903	0.0858	0.0817	0.078	0.0745
2	0.2228	0.2086	0.196	0.1847	0.1746	0.1654	0.157	0.1494	0.1425	0.1361	0.1302	0.1248	0.1198
3	0.2799	0.2637	0.2491	0.2359	0.2239	0.213	0.2031	0.194	0.1856	0.1778	0.1707	0.164	0.1578
4	0.3274	0.3097	0.2938	0.2792	0.2659	0.2539	0.2428	0.2325	0.2231	0.2142	0.2061	0.1985	0.1914
5	0.3677	0.3491	0.3324	0.317	0.3028	0.2898	0.2778	0.2666	0.2564	0.2468	0.2379	0.2295	0.2217
6	0.403	0.3839	0.3663	0.3501	0.3355	0.3218	0.3092	0.29794	0.2864	0.2762	0.2666	0.2577	0.2493
7	0.434	0.4145	0.3967	0.3801	0.3648	0.3505	0.3374	0.3252	0.3138	0.3031	0.293	0.2836	0.2747
8	0.4615	0.442	0.4238	0.407	0.3914	0.3769	0.3633	0.3505	0.3387	0.3277	0.3173	0.3074	0.2981
9	0.4862	0.4666	0.4484	0.4314	0.4156	0.4009	0.387	0.3741	0.3619	0.3503	0.3396	0.3295	0.3199
10	0.5083	0.4889	0.4707	0.4537	0.4378	0.4228	0.4088	0.3957	0.3833	0.3716	0.3605	0.35	0.3402
11	0.5288	0.5091	0.4911	0.4741	0.4581	0.4431	0.429	0.4157	0.4031	0.3913	0.3801	0.3694	0.3593
12	0.5477	0.528	0.5096	0.4929	0.4769	0.4619	0.4477	0.4343	0.4216	0.4096	0.3983	0.3875	0.3772
13	0.565	0.5456	0.5273	0.51	0.4943	0.4793	0.4651	0.4517	0.4389	0.4268	0.4153	0.4044	0.3941
14	0.5808	0.5617	0.5436	0.5265	0.5103	0.4955	0.4814	0.4679	0.4551	0.443	0.4314	0.4204	0.4099
15	0.5955	0.5766	0.5587	0.5418	0.5258	0.5105	0.4965	0.4831	0.4703	0.4582	0.4466	0.4355	
16	0.6091	0.5904	0.5728	0.556	0.5402	0.5251	0.5106	0.4974	0.4846	0.4725	0.4609		
17	0.6217	0.6033	0.5859	0.5693	0.5536	0.5387	0.5244	0.5107	0.4981	0.486			
18	0.6336	0.6154	0.5982	0.5818	0.5662	0.5514	0.5372	0.5237	0.5107				
19	0.6446	0.6267	0.6097	0.5935	0.5781	0.5634	0.5494	0.5359					
20	0.655	0.6373	0.6205	0.6045	0.5893	0.5747	0.5608						
21	0.6647	0.6473	0.6308	0.615	0.5999	0.5854							
22	0.6738	0.6568	0.6404	0.6248	0.6099								
23	0.6825	0.6657	0.6496	0.6342									
24	0.6908	0.674	0.6582										
25	0.6986	0.6821											
26	0.7059												

Table 22. Pre-Calculated Gittins Index when $a = 0.99$ Part 2 (From Gittins, 1989)

α \ β	27	28	29	30	31	32	33	34	35	36	37	38	39
1	0.0713	0.0684	0.0656	0.0631	0.0607	0.0585	0.0564	0.0545	0.0526	0.0509	0.0493	0.0478	0.0463
2	0.1151	0.1107	0.1067	0.1029	0.0994	0.096	0.0929	0.09	0.0872	0.0846	0.0821	0.0797	
3	0.1521	0.1467	0.1416	0.1369	0.1325	0.1283	0.1244	0.1206	0.1171	0.1138	0.1106		
4	0.1848	0.1786	0.1727	0.1673	0.1621	0.1572	0.1526	0.1483	0.1441	0.1402			
5	0.2143	0.2075	0.201	0.1949	0.1891	0.1837	0.1785	0.1736	0.169				
6	0.2414	0.234	0.2269	0.2203	0.214	0.208	0.2024	0.1971					
7	0.2663	0.2583	0.2509	0.2439	0.2372	0.2308	0.2248						
8	0.2894	0.2811	0.2732	0.2658	0.2587	0.252							
9	0.3109	0.3023	0.2941	0.2864	0.279								
10	0.3309	0.3221	0.3136	0.3056									
11	0.3497	0.3405	0.332										
12	0.3675	0.3582											
13	0.3842												

Table 23. Pre-Calculated Gittins Index when $a = 0.99$ Part 3 (From Gittins, 1989)

$\frac{\alpha}{(\alpha + \beta)}$	A	B	C
0.025	6.1145	1.7222	36.772
0.05	10.308	1.6483	-119.1
0.1	13.424	1.608	-122.1
0.2	17.192	1.582	-139.3
0.3	20.299	1.5709	-168.5
0.4	21.514	1.5716	-152.5
0.5	22.958	1.5795	-155.1
0.6	24.962	1.5745	-150.2
0.7	27.292	1.5796	-148.9
0.8	31.032	1.5892	-151.9
0.9	38.459	1.62	-126.6
0.95	49.51	1.6793	-79.08
0.975	65.911	1.7798	0.1

Table 24. Pre-Calculated A, B, and C when $a = 0.99$ (From Gittins, 1989)

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX C. TOTAL NUMBER OF POSSIBLE PLANS

The interdictor's goal is to minimize the evader's expected number of successful trips. He does this by implementing an interdiction plan. The total number of possible plans for the interdictor to choose from is dependent on the number of assets and the number of routes. The number of possible plans could be small if there are not many assets or routes. For instance with one asset and n routes then there are only n possible plans. With any number of assets, but only one route then there is one possible plan. The number of possible plans, though, explodes as the number of routes and assets increases.

A. USING A TREE DIAGRAM

To compute the total number of plans, first consider each asset type separately. For instance, we will look at computing the number of plans for only one asset type, j. The two parameters that will affect the quantity of the plans for this asset type are the number of assets (r_j) and the number of routes (n). When figuring out how many assets to place on the first route there are $r_j + 1$ possible choices. This extra choice is because the interdictor could place zero assets on the first route. The options for the assets on the second route are $r_j + 1$ minus the assets placed on the first route and so forth. When looking at the number of routes, there are $n-1$ degrees of freedom, since the n^{th} route will get whatever assets are left. This type of setup lends itself to a tree diagram, which is $n-1$ deep. The number of plans will be the number of leaves. Figures 16 – 18 are tree diagrams for two assets ($r_j = 2$) with two, three, and four routes, respectively.

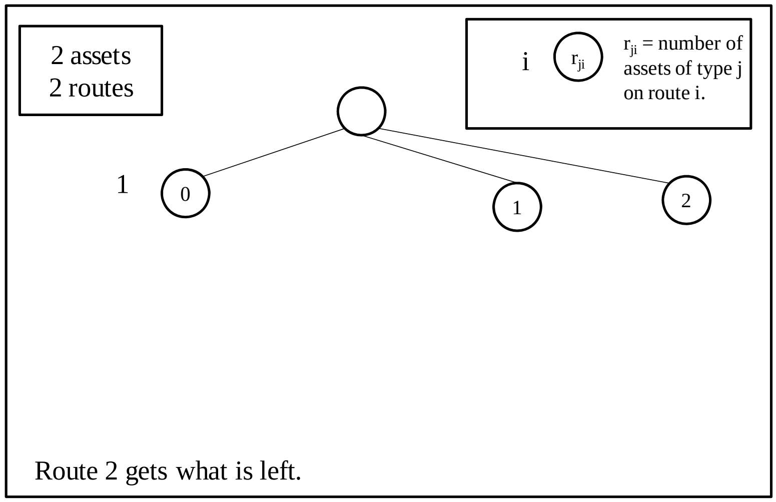

Figure 16. There are three possible plans with two assets and two routes

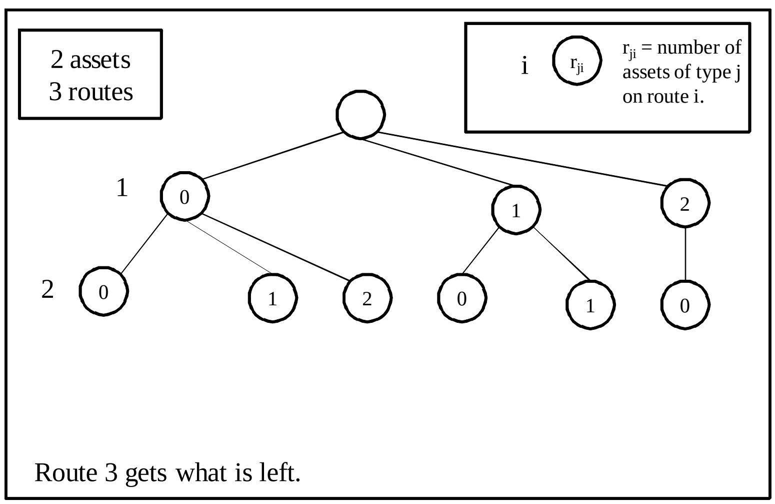

Figure 17. There are six possible plans with two assets and three routes

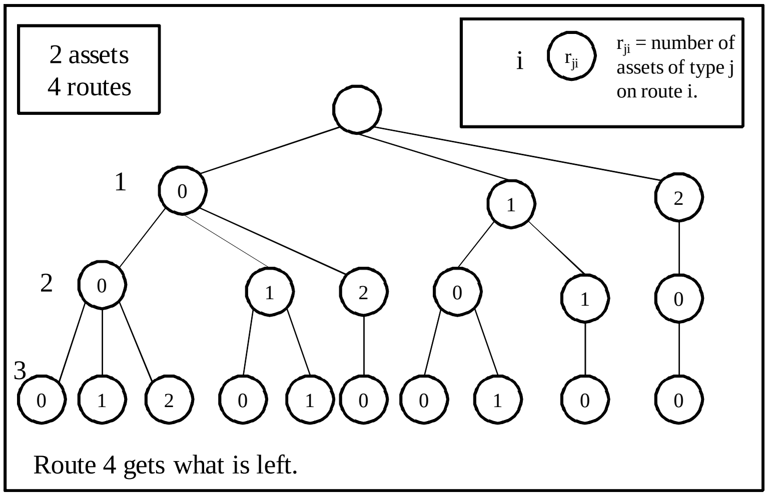

Figure 18. There are ten possible plans with two assets and three routes

B. USING AN ITERATIVE TABLE

Another way to calculate the number of plans is by setting up a table. Table 25 is an example of this for two assets. The rows of the table are the number of routes, starting at two. The column headers are the number of possible choices ($r_j + 1$ to one). The first row is populated with zeros except for the first column ($r_j + 1$) which has a one. The rest of the first column ($r_j + 1$) is also populated with ones. To fill in the rest of the table is a top down, left to right process. Since the first row and first column are already filled in the next block to populate is for three routes and column r_j. To fill this block in, add the number above the block with the number on the left of the block. Repeat this process for the rest of the table until it is complete. To calculate the number of possible plans for a route, multiply the number in the row for that route with its column header and add those products together. For instance with four routes and two assets, (as seen in Table 25), multiply the first column header by the number in the first column ($3 \cdot 1$), then add that to the answer when the second column header is multiplied to the second column ($2 \cdot 2$), and finally add that to the answer when the third column header is multiplied to the third column ($1 \cdot 3$). The final answer is 10. Therefore, with four routes and one asset type with two assets there are only ten possible plans for the interdictor.

	Number of Possible Choices			
Total Routes	3	2	1	**Total Plans**
2	1	0	0	3
3	1	1	1	6
4	1	2	3	10
5	1	3	6	15
6	1	4	10	21
7	1	5	15	28

Table 25. The number of plans for two assets

C. MULTIPLE ASSET TYPES

If there is more than one asset type then each asset type is considered separately then those numbers are multiplied together to get the number of plans. For example, an interdictor is trying to cover four routes ($n = 4$) with four asset types ($m = 4$) each with a possible different number of assets available (r_j). With $r_1 = 2$, there are ten possible plans when just considering $j = 1$. With $r_2 = 4$, there are thirty-five possible plans when just considering $j = 2$. With r_3 and r_4 both equal to one, there are four possible plans for each asset type. To get the total number of possible defense plans of 5,600, multiply 10 by 35 by 4 by 4.

APPENDIX D. CASE STUDY DATA

Appendix D contains the results of the sixty-four sub cases. Each figure is a box plot[4] of one case and the corresponding eight plans. The box plots show the maximum value, the 95th percentile upper bound, the mean, the 95th percentile lower bound, and the minimum value of the discounted number of successes for the sub cases.

A. DISCOUNTED NUMBER OF SUCCESSES

Figure 19 gives the discounted number of successes for Case 1. The 95% confidence interval is small for each plan against Case 1. The two best plans for the Coalition Force are Plan 1 and Plan 5 with a discounted number of successes of about 50. We note that the CE Asset Allocation Algorithm suggests Plan 1 for Case 1. The two worst plans for the Coalition Force are Plan 6 and Plan 8 with a discounted number of successes of about 72.

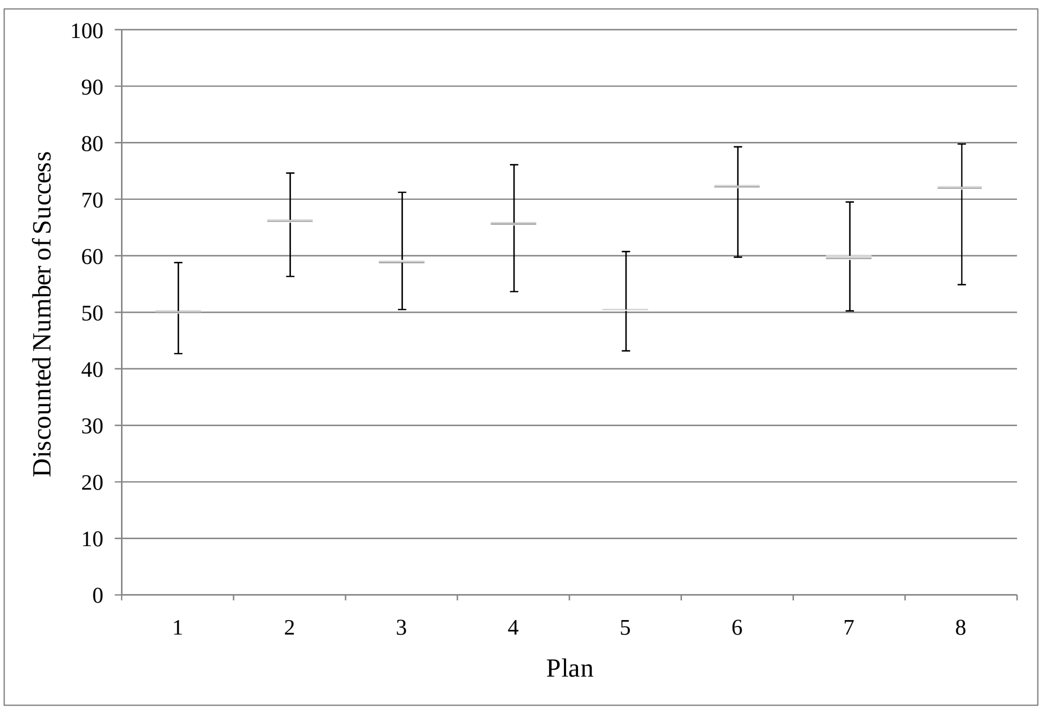

Figure 19. Discounted Number of Successes for Case 1

[4] Box plots were created in Excel with the help of bloggpro (2007).

Figure 20 gives the discounted number of successes for Case 2. The 95% confidence interval is small for each plan against Case 2. The two best plans are Plan 2 and Plan 6 with a discounted number of successes of about 73. We note that the CE Asset Allocation Algorithm suggests Plan 2 for Case 2.

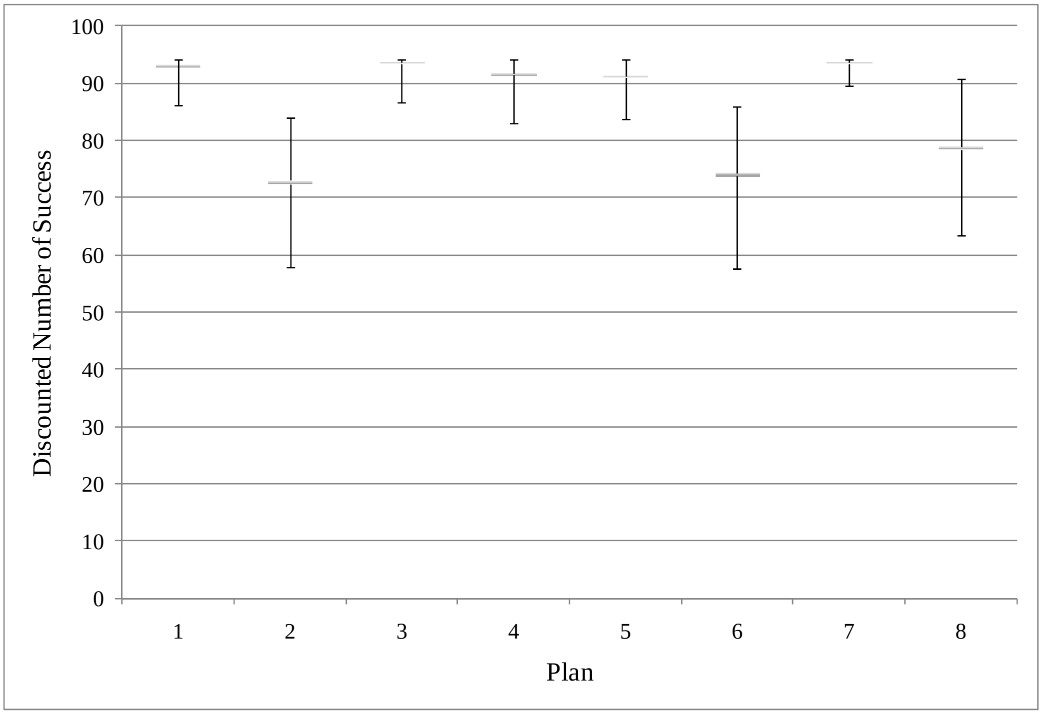

Figure 20. Discounted Number of Successes for Case 2

Figure 21 gives the discounted number of successes for Case 3. The two best plans are Plan 3 and Plan 7 with a discounted number of successes of about 0.14. We note that the CE Asset Allocation Algorithm suggests Plan 3 for Case 3. The two worst plans are Plan 2 and Plan 6 with a discounted number of successes of about 0.70.

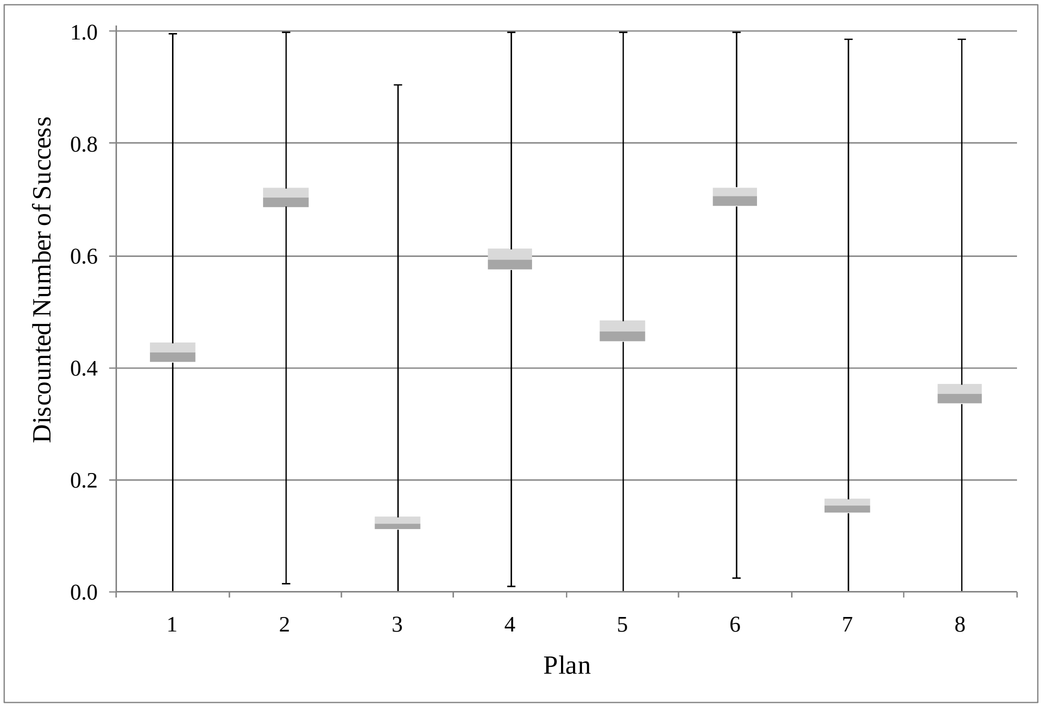

Figure 21. Discounted Number of Successes for Case 3

Figure 22 gives the discounted number of successes for Case 4. There is a huge variation in the minimum and maximum values, but on average all of the plans perform relatively the same with a discounted number of successes of about 0.80.

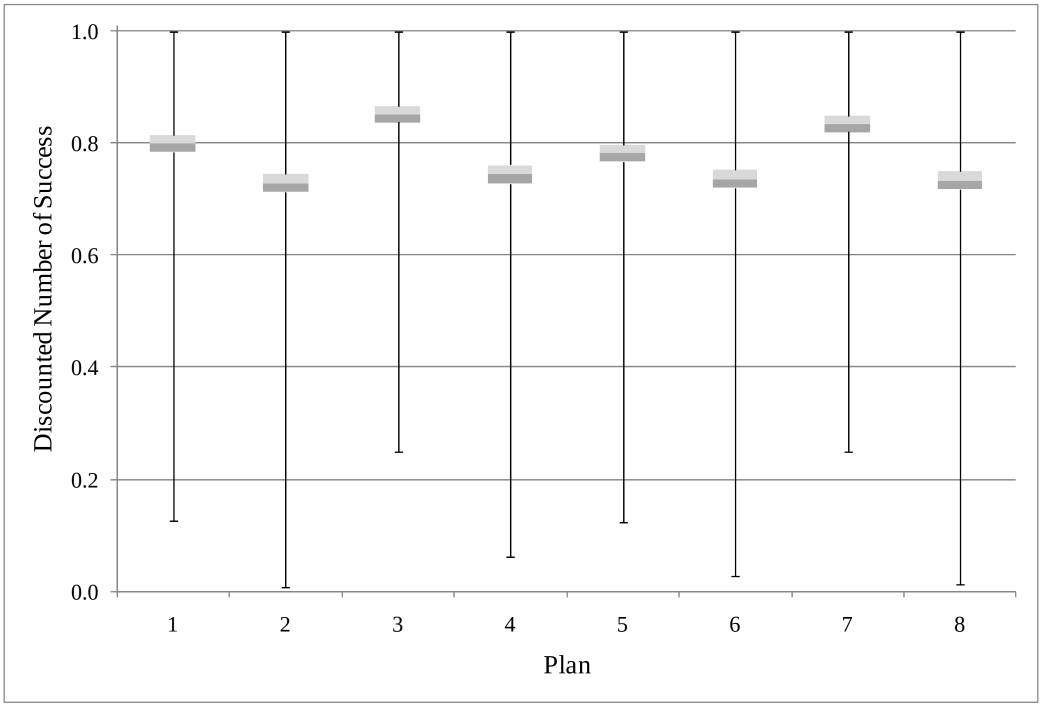

Figure 22. Discounted Number of Successes for Case 4

Figure 23 gives the discounted number of successes for Case 5. The 95% confidence interval is small for each plan against Case 5. The two best plans are Plan 1 and Plan 5 with a discounted number of successes of about 60. We note that the CE Asset Allocation Algorithm suggests Plan 5 for Case 5. The worst plan is Plan 7 with a discounted number of successes of about 60.

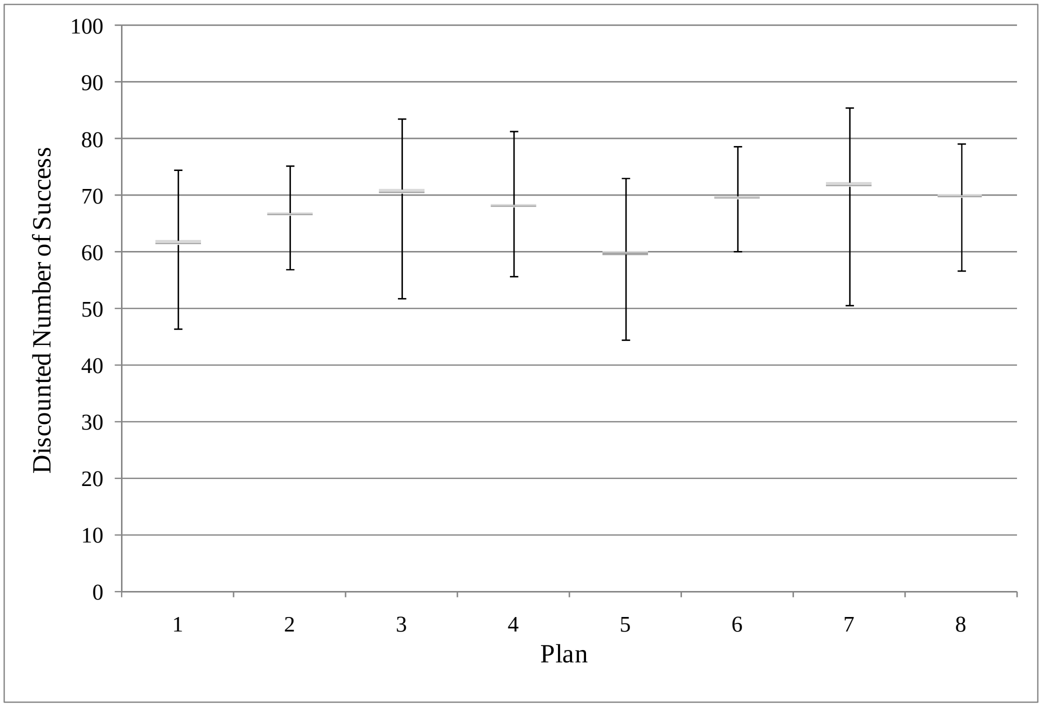

Figure 23. Discounted Number of Successes for Case 5

Figure 24 gives the discounted number of successes for Case 6. The 95% confidence interval is small for each plan against Case 6. The two best plans are Plan 2 and Plan 6 with a discounted number of successes of about 72. We note that the CE Asset Allocation Algorithm suggests Plan 6 for Case 6.

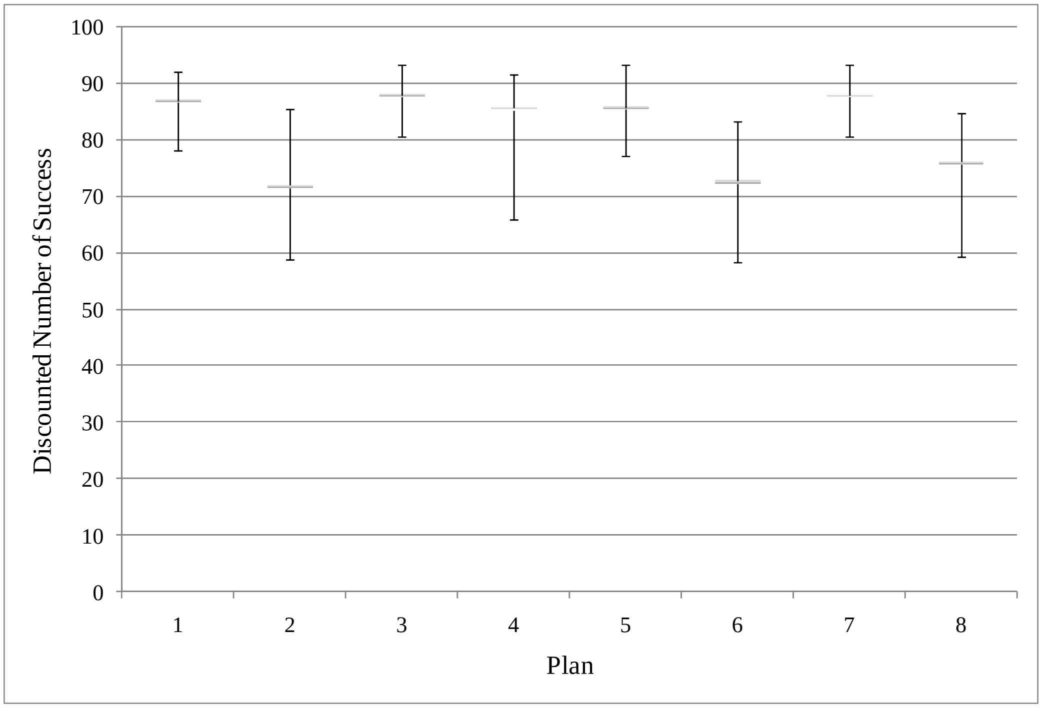

Figure 24. Discounted Number of Successes for Case 6

Figure 25 gives the discounted number of successes for Case 7. The two best plans are Plan 3 and Plan 7 with a discounted number of successes of about 0.13. We note that the CE Asset Allocation Algorithm suggests Plan 7 for Case 7. The two worst plans are Plan 2 and Plan 6 with a discounted number of successes of about 0.69.

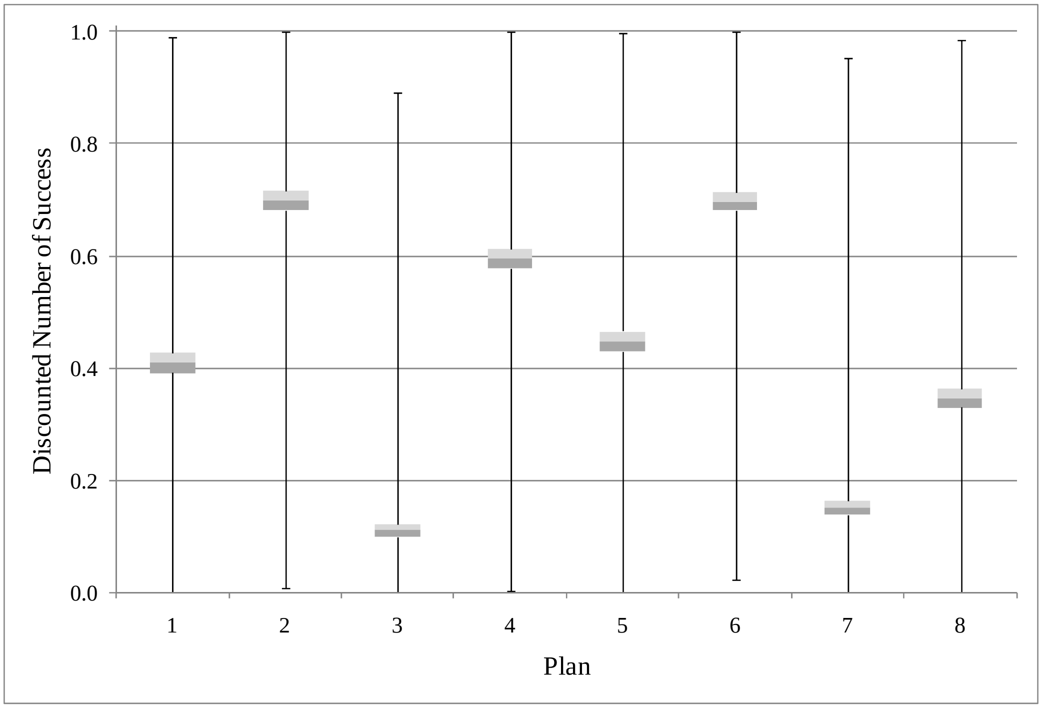

Figure 25. Discounted Number of Successes for Case 7

Figure 26 gives the discounted number of successes for Case 8. All of the plans perform relatively the same, except Plan 3 and Plan 7, which have a discounted number of successes of about 0.84.

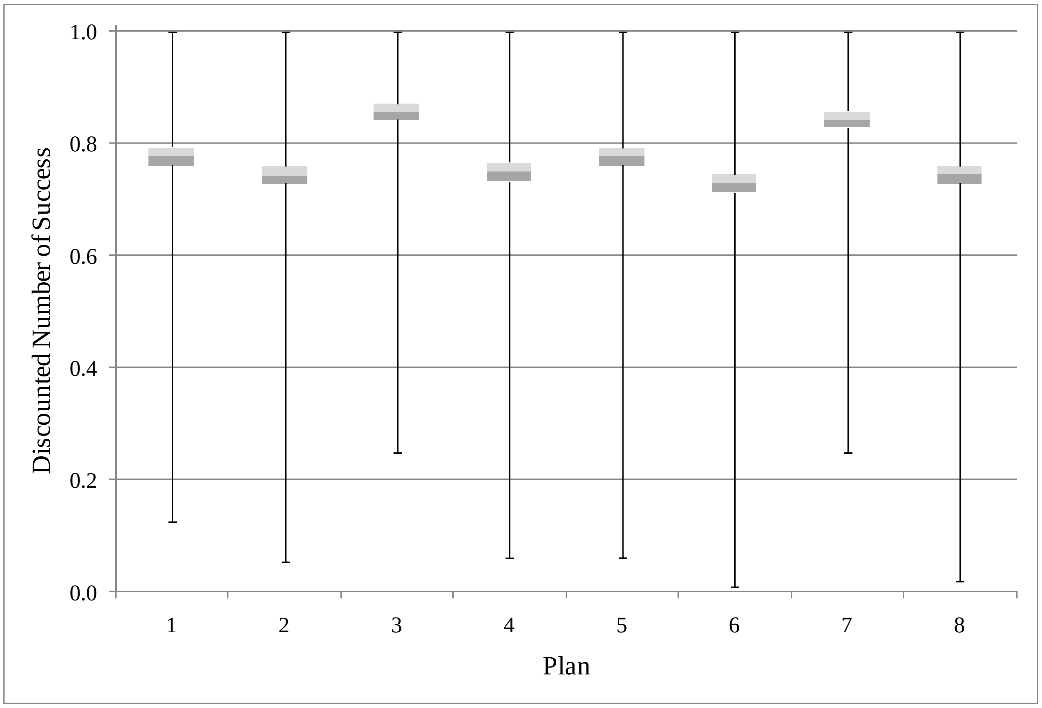

Figure 26. Discounted Number of Successes for Case 8

B. THE TIME THE DTO FOUND THE BEST ROUTE

The results for the time that the DTO finds the best route for the sixty-four sub cases are in Figures 27–34. The minimum time period is 1, if the DTO finds the best route with his first decision and continues to use the same route. The maximum time period is 298, since both the Gittins Choice Algorithm, and Decreasing Choice Algorithm run for 298 time periods. If the DTO does not find the best route then there is no result to collect. Therefore, each sub case will have a different number of observations.

Figure 27 gives the time the DTO finds the best route for Case 1. The time period that the DTO finds the best route is divided into three groups. The DTO either finds the best route in about time period 200 (Plan 1 and Plan 4), time period 100 (Plan 3, Plan 6 and Plan 7, or did not find the best route (Plan 2, Plan 5, and Plan 8).

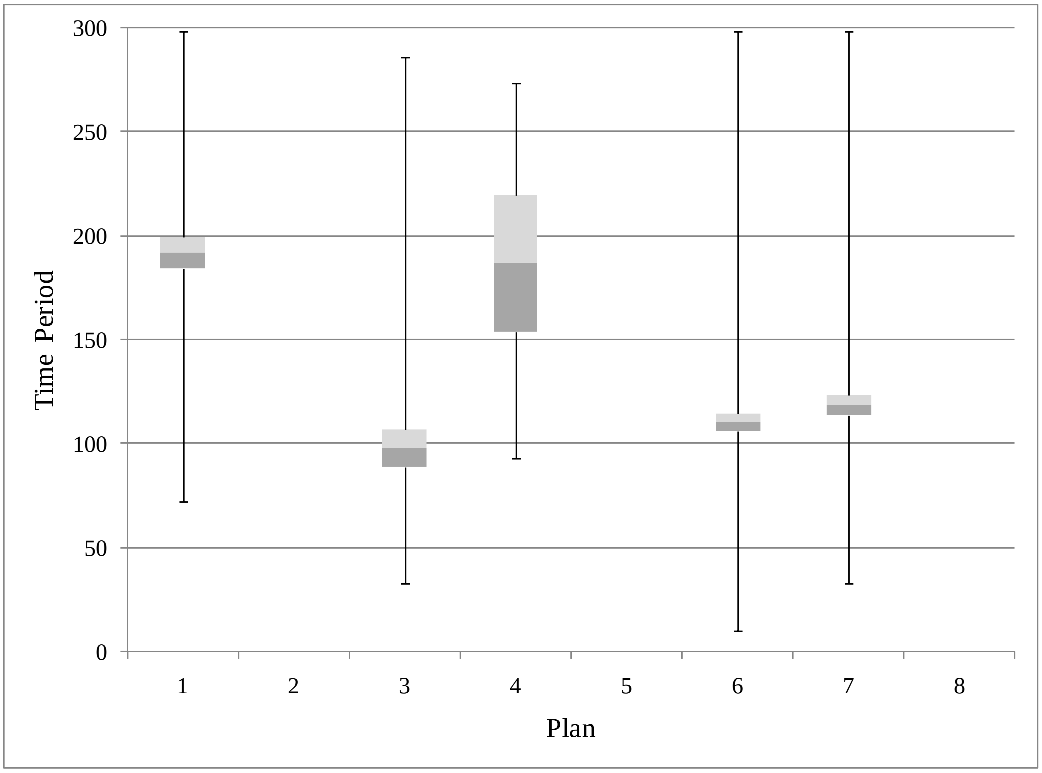

Figure 27. The time the DTO found the best route given he found the best route for Case 1

Figure 28 gives the time the DTO finds the best route for Case 2. The time period that the DTO finds the best route is divided into two groups. The DTO either finds the best route in about time period 7 (Plan 1, Plan 3, Plan 4, Plan 5, and Plan 7) or he finds the best route in time period 80 (Plan 2, Plan 6, and Plan 8).

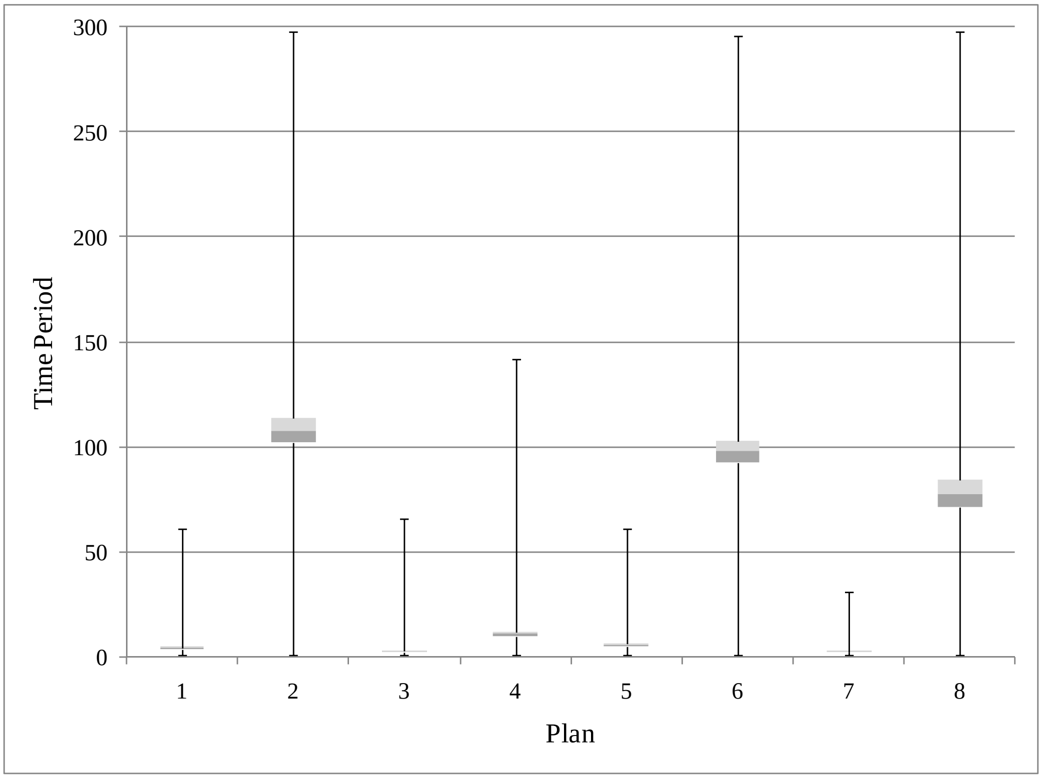

Figure 28. The time the DTO found the best route given he found the best route for Case 2

Figure 29 gives the time the DTO finds the best route for Case 3. The DTO rarely finds the best route in Case 3. Against Plan 1, Plan 6 and Plan 7 does the DTO find the best route.

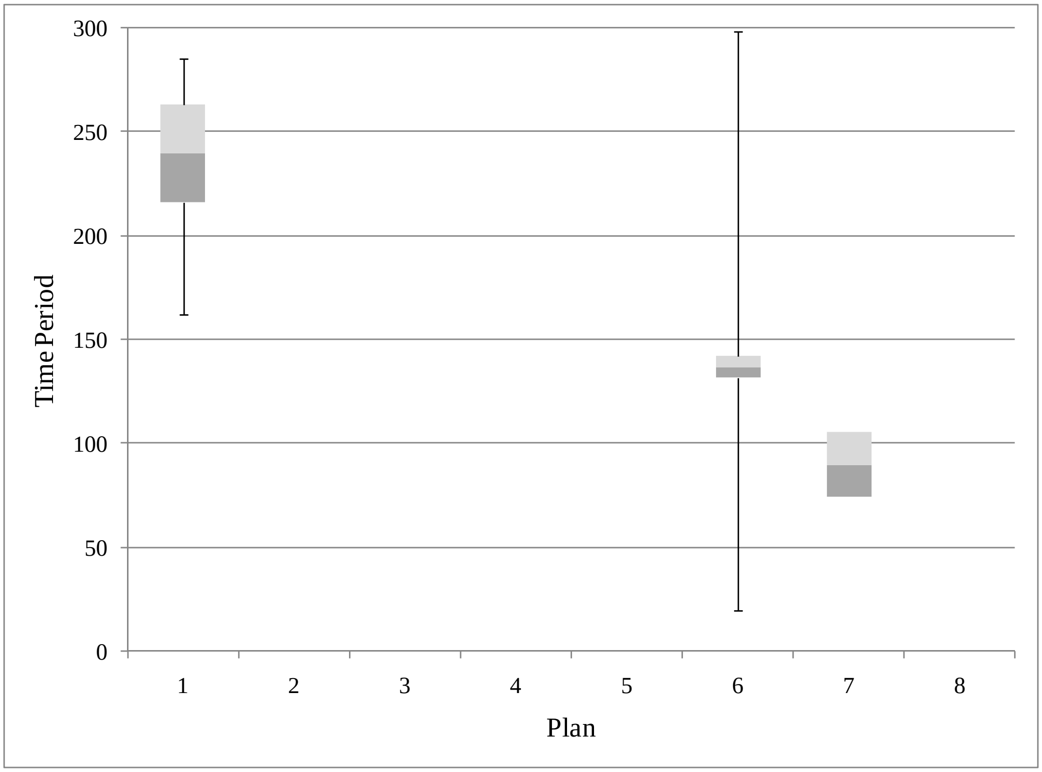

Figure 29. The time the DTO found the best route given he found the best route for Case 3

Figure 30 gives the time the DTO finds the best route for Case 4. The DTO always finds the best route in Case 4 at approximately time period 10.

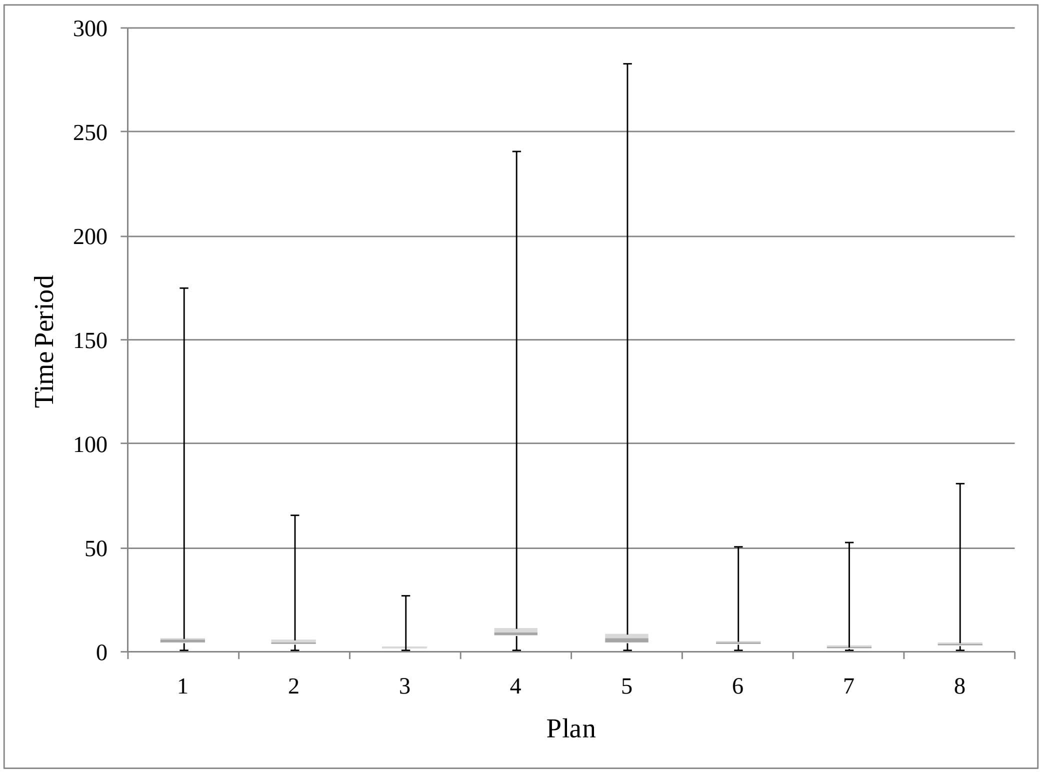

Figure 30. The time the DTO found the best route given he found the best route for Case 4

Figure 31 gives the time the DTO finds the best route for Case 5. The DTO finds the best route in all of the plans, but it takes him on average to time period 265.

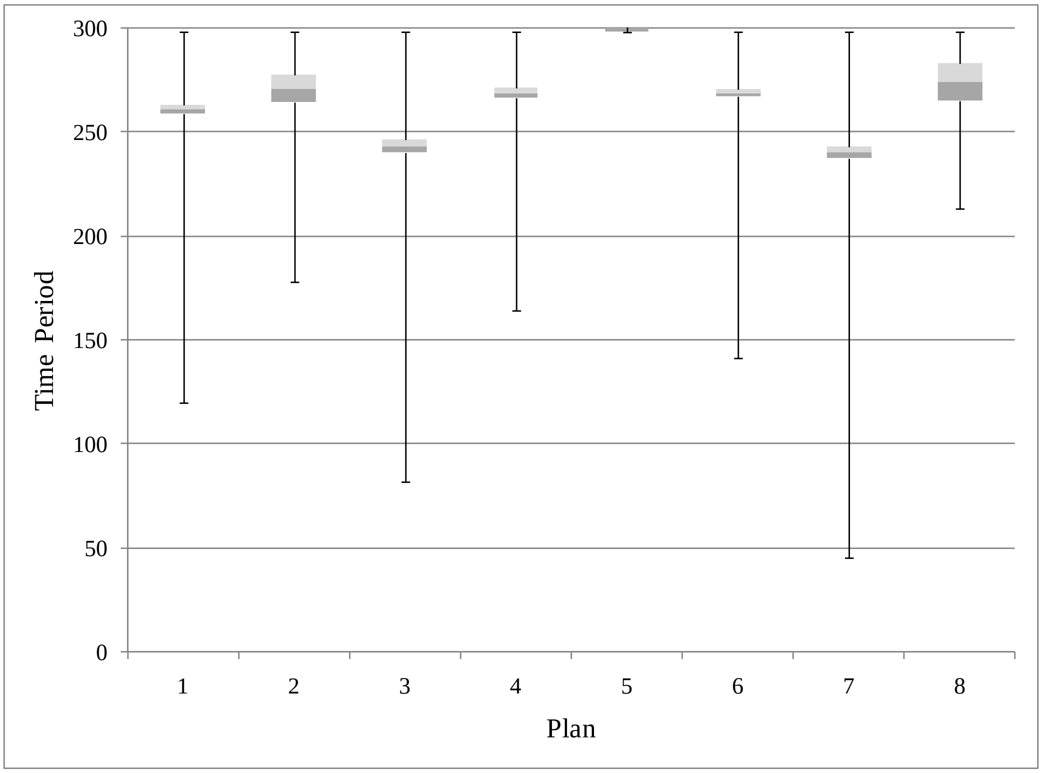

Figure 31. The time the DTO found the best route given he found the best route for Case 5

Figure 32 gives the time the DTO finds the best route for Case 6. The DTO finds the best route in all of the plans, but it takes him on average to time period 258.

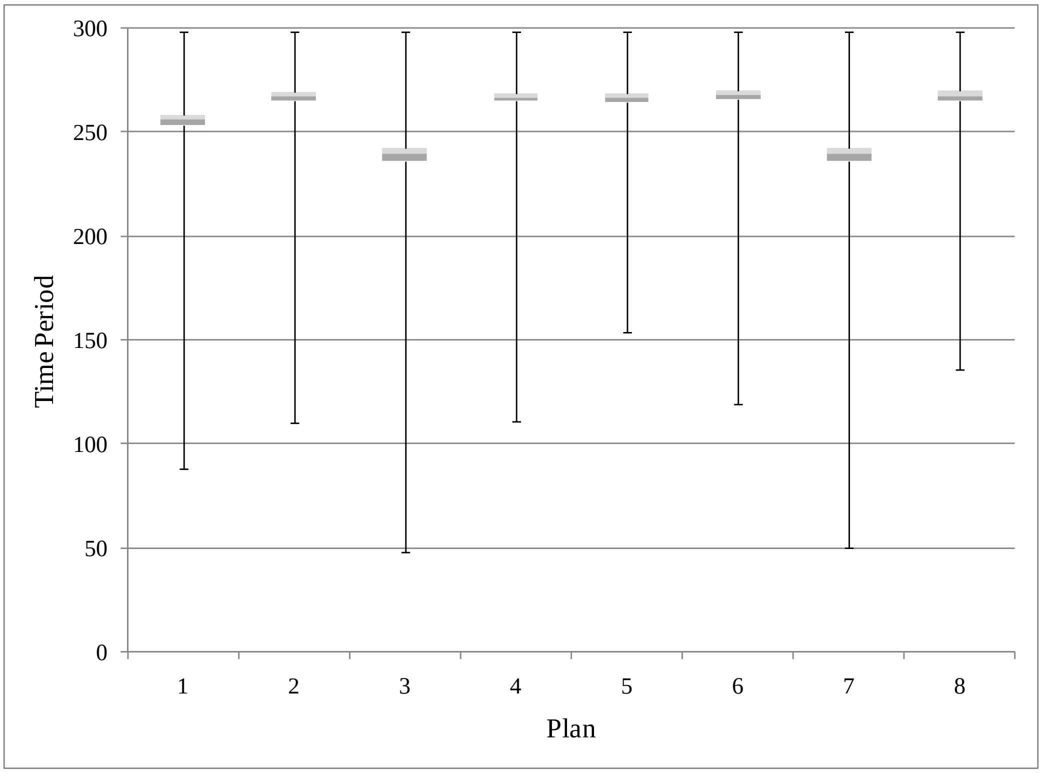

Figure 32. The time the DTO found the best route given he found the best route for Case 6

Figure 33 gives the time the DTO finds the best route for Case 7. The DTO rarely finds the best route in Case 7. He only finds the best route against Plan 1, Plan 6 and Plan 7.

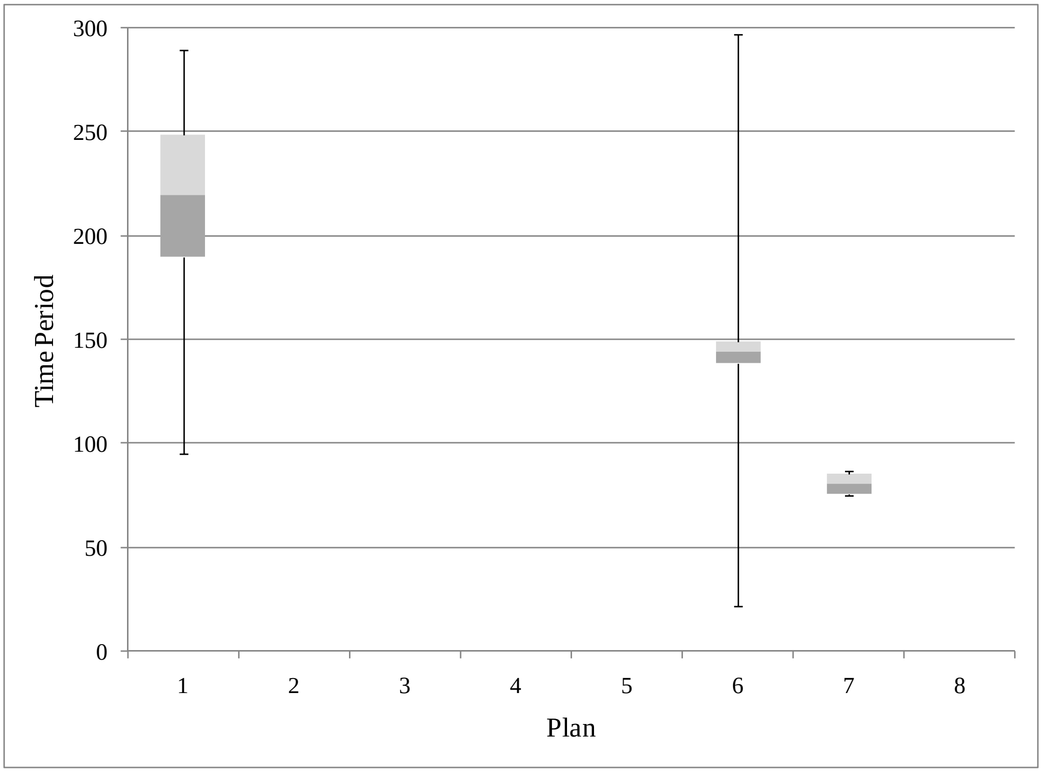

Figure 33. The time the DTO found the best route given he found the best route for Case 7

Figure 34 gives the time the DTO finds the best route for Case 8. The DTO finds the best route in all of the plans, but it takes him on average to time period 10.

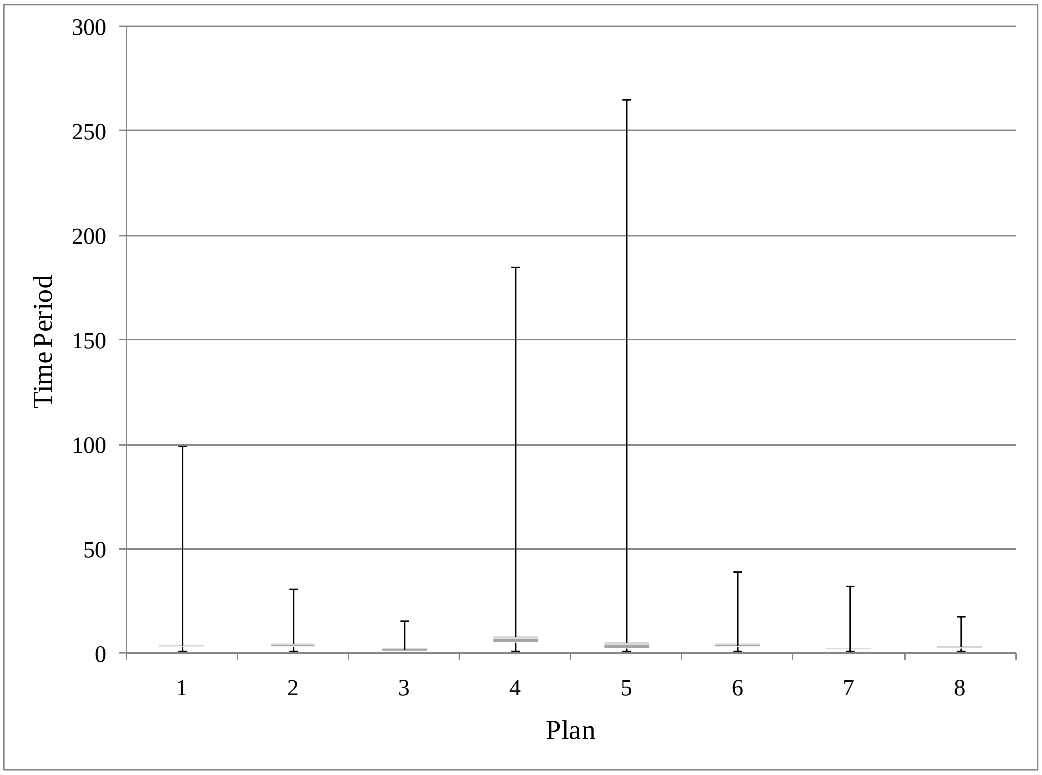

Figure 34. The time the DTO found the best route given he found the best route for Case 8

LIST OF REFERENCES

Allon, G., D. P., Kroese, T. Raviv, & R. Y. Rubinstein. (2004). Application of the cross-entropy method to the buffer allocation problem in a simulation-based environment. Annals of Operations Research, Vol 134 No 1, 137–151, 2005.

Bailey, M. P., R. F. Dell, & K. D. Glazebrook. 1994. Simulatio-Based Dynamic Optimization: Planning United States Coast Guard Law Enforcement Patrols. In Proceedings of the 1994 Winter Simulation Conference. 392–398.

Berry, D. A. & B. Fristedt, 1985. Bandit Problems: Sequential Allocation of Experiments. Chapman and Hall, London.

Bertsimas, D. & J. Niño-Mora, 1996. Conservation Laws, Extended Polymatroid and Multiarmed Bandit Problems; A Polyhedral Approach to Indexabel Systems. Mathematics of Operations Research, Vol 21 No 2, 257–306.

Bloggpro. 2007. Retrieved September 1, 2010 from http://www.bloggpro.com/box-plot-for-excel-2007/.

Brown, G., M. Carlyle, J. Salmeron & K. Wood. 2006. Defending Critical Infrastructure. Interfaces, Vol 36 No 6, 530–544.

Caulkins, J.P., G. Crawford & P. Reuter. 1993. Simulation of Adaptive Response: A Model of Drug Interdiction. Mathematical and Computer Modeling. Vol 17 No 2, 37–52.

Dacre, M., K. Glazebrook, & J. Niño-Mora, 1999. The Achievable Region Approach to the Optimal Control of Stochastic Systems. Journal of the Royal Statistcal Society. Series B (Statistical Methodology), Vol 61 No 4, 747–791.

Dimitrov, N.B., D. P. Michalopoulos, D. P. Morton, M. V. Nehme, F. Pan, E. Popova, E. A. Schneider, & G. G. Thoreson. 2009. Network Deployment of Radiation Detectors with Physics-Based Detection Probability Calculations. Annals of Operations Research.

Fudenberg, Drew & David K. Levine, 1998. The Theory of Learning in Games. The MIT Press, Cambridge, MA.

GAO, 2008. United States Government Accountability Office. Drug Control, Cooperation with Many Major Drug Transit Countries Has Improved, but Better Performance Reporting and Sustainability Plans are Needed. Retrieved June 17, 2010 from https://www.hsdl.org/?view&doc=101763&coll=limited.

Gittins, J. 1989. Multi-armed Bandit Allocation Indices. John Wiley & Sons, New York, NY.

Gittins, J. & D. Jones, 1974. A dynamic allocation index for the sequential design of experiments. Progress in Statistics, ed. J. Gani, North-Holland, Amsterdam, 241–266.

Google Earth, 2010. Retrieved August 25, 2010 from http://earth.google.com/.

Jane's, Oliver Hazard Perry class. Retrieved September17, 2010 from http://www2.janes.com/janesdata/yb/jfs/jfs_3535.htm#img.

Jane's, Lockheed P-3C Orion. Retrieved September17, 2010 from http://www2.janes.com/janesdata/yb/jfs/jfs_3559.htm#img.

Jane's, Sikorsky MH-60R Seahawk. Retrieved September17, 2010 from http://www2.janes.com/janesdata/yb/jfs/jfs_5703.htm#img.

JIATF. 2009. Joint Interagency Task Force-South. Retrieved December 4, 2009 from http://www.jiatfs.southcom.mil/.

NDIC. 2009. The National Drug Intelligence Center. National Drug Threat Assessment 2009. Retrieved December 4, 2009 from http://www.justice.gov/ndic/pubs31/31379/index.htm.

ONDCP, 2008. Office of National Drug Control Policy. Cocaine Smuggling in 2007. Retrieved June 17, 2010 from http://www.ncjrs.gov/ondcppubs/publications/pdf/cocaine_smuggling07.pdf.

Pfeiff, D. 2009. Optimizing Employment of Search Platforms to Counter Self-Propeleed Semi-Submersibles. Master's Thesis, Naval Postgraduate School. Monterey, CA.

Remote Sensing System, 2010. Retrieved August 30, 2010 from http://www.remss.com/ssmi/ssmi_browse.html.

Robbins, H. 1952. Some aspects of the sequential design of experiments. Bulletin of the American Mathematical Society. Vol 58, 527–535.

Rubinstein, R. Y. & D. P. Kroesse, 2004. The Cross-Entropy Method: A Unified Apporach to Monte Carlo Simulation, Randomized Optimization and Machine Learning. Springer Verlag.

SOUTHCOM. 2009. United States Southern Command. Retrieved December 4, 2009 from http://www.southcom.mil/.

Thompson, W. R. 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. Biometrika. Vol 25, 275–294.

Thompson, W. R. 1935. On the theory of apportionment. American Journal of Mathematics. Vol 57, 450–456.

Tsitsiklis, J. 1986. A Lemma on the Multiarmed Bandit Problem. IEEE Transactions, Automatic Control. Vol 31 No 6, 576–577.

Tsitsiklis, J. 1994. A Short Proof of the Gittins Index Theorem. The Annals of Applied Probability. Vol 4 No 1, 194–199.

USCG. 2009. United States Coast Guard. SPSS Background. Retrieved December 4, 2009 from http://www.uscg.mil/Comdt/all_hands/docs/SPSS%20Background.pdf.

USCG Addendum, 2009. U.S. Coast Guard Addendum to the United States National Search and Rescue Suppliment (NSS).

Varaiya, P., J. Walrand, & C. Buyukkoc. 1985. Extensions of the Multiarmed bandit problem: The discounted case. IEEE Transactions, Automatic Control, Vol 30 No 5, 426–439.

Vermorel, J. & M. Mohri. 2005. Multi-Armed Bandit Algorithms and Empirical Evaluation. In Proceedings of the 16th European Conference on Machine Learning: ECML. 437–448.

von Stackelberg, H. 1952. The Theory of the Market Economy (translated from German). William Hodge & Co., London, UK.

Washburn, A. & K. Wood. 1995. Two-Person Zero-Sum Games for Network Interdiction. Operations Research, Vol 43 No 2, 243–251.

Weber, R. 1992. On the Gittins Index for Multiarmed Bandits. The Annals of Applied Probability, Vol 2 No 4, 1024–1033.

Whittle, P. 1980. Multiarmed Bandits and the Gittins Index. Journal of the Royal Statistical Society, Series B (Methodoligical), Vol 42 No 2, 143–149.

Wagner, D., W. Mylander, & T. Sanders. 1999, Naval Operations Analysis, 3rd ed., Naval Institute Press, Annapolis, Maryland.

Wood, K. 1993. Deterministic Network Interdiction. Mathematical and Computer Modeling, Vol 17 No 2, 1–18.

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California

3. Professor Johannes O. Royset
   Department of Operations Research
   Naval Postgraduate School
   Monterey, California

4. Professor Moshe Kress
   Department of Operations Research
   Naval Postgraduate School
   Monterey, California