



**NAVAL
POSTGRADUATE
SCHOOL**

MONTEREY, CALIFORNIA

THESIS

**DETERMINING OPTIMAL EVACUATION DECISION
POLICIES FOR DISASTERS**

by

Jason C. Crews

March 2012

Thesis Advisor:
Second Reader:

Emily M. Craparo
David L. Alderson

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) 19-3-2012			2. REPORT TYPE Master's Thesis		3. DATES COVERED (From — To) 2011-02-15—2012-03-30	
4. TITLE AND SUBTITLE Determining Optimal Evacuation Decision Policies For Disasters					5a. CONTRACT NUMBER	
					5b. GRANT NUMBER	
					5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Jason C. Crews					5d. PROJECT NUMBER	
					5e. TASK NUMBER	
					5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943					8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Department of the Navy					10. SPONSOR/MONITOR'S ACRONYM(S)	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited						
13. 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 not applicable.						
14. ABSTRACT Decision making in the face of uncertainty is a difficult task, and this is exacerbated when the decision is irreversible, it involves a near-term deadline, and/or the cost of a bad decision is high. Deciding whether to stay or evacuate from an impending natural disaster is difficult for all of these reasons. This thesis explores the evacuation decision as a Markov decision problem. We develop a generic disaster model to explore the tensions and tradeoffs in the decision to evacuate and use a dynamic programming algorithm to determine optimal decision policies for the decision maker. We explore how these policies are affected by evacuation costs as well as disaster uncertainty.						
15. SUBJECT TERMS						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 81	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (include area code)	

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

DETERMINING OPTIMAL EVACUATION DECISION POLICIES FOR DISASTERS

Jason C. Crews
Lieutenant, United States Navy
B.S., University of South Carolina, 2005

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

**NAVAL POSTGRADUATE SCHOOL
March 2012**

Author: Jason C. Crews

Approved by: Emily M. Craparo
Thesis Advisor

David L. Alderson
Second Reader

Robert F. Dell
Chair, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Decision making in the face of uncertainty is a difficult task, and this is exacerbated when the decision is irreversible, it involves a near-term deadline, and/or the cost of a bad decision is high. Deciding whether to stay or evacuate from an impending natural disaster is difficult for all of these reasons. This thesis explores the evacuation decision as a Markov decision problem. We develop a generic disaster model to explore the tensions and tradeoffs in the decision to evacuate and use a dynamic programming algorithm to determine optimal decision policies for the decision maker. We explore how these policies are affected by evacuation costs as well as disaster uncertainty.

THIS PAGE INTENTIONALLY LEFT BLANK

Table of Contents

1	INTRODUCTION	1
2	BACKGROUND AND CONTRIBUTIONS	3
2.1	Overview	3
2.2	Mitigating Earthquake Risks	3
2.3	Probabilistic Evacuation Decision Model for River Floods	4
2.4	Dynamically Modeling Hurricane Evacuation Decisions	5
2.5	Evacuation Decisions and Hurricane Track Uncertainty	6
2.6	Thesis Contributions	6
3	MODELING THE DISASTER PROCESS	9
3.1	Special Case: Deterministic Horizontal Motion ($\Delta H = 1$)	12
3.1.1	Impact of Vertical Volatility on Final Particle State	13
3.2	General Disaster Model	16
3.2.1	Horizontal Volatility	16
3.2.2	Expected TTG	18
3.3	Calculating the Hit Probability (P_{hit})	20
3.4	P_{hit} versus Vertical Volatility	24
3.5	Chapter Summary	25
4	DECISION MAKING IN DISASTER SITUATIONS	29
4.1	Cost-To-Go Function	30
4.2	Creating Decision Cost Functions	30
4.3	How Evacuation Costs Affect Decision Policies	33
4.3.1	Constant Evacuation Cost	33
4.3.2	Increasing Evacuation Cost	33
4.4	Impact of Particle Volatility on Decision Policies	37
4.5	Chapter Summary	38

5	CONCLUSIONS	45
5.1	Summary	45
5.2	Future Work	46
5.3	Final Thoughts	47
	LIST OF REFERENCES	47
	Appendices	
A	RELATIONSHIP BETWEEN VERTICAL VOLATILITY AND INITIAL HORIZONTAL POSITION	51
B	SENSITIVITY OF OPTIMAL EVACUATION POLICIES TO STATE SPACE DISCRETIZATION	57
	INITIAL DISTRIBUTION LIST	61

List of Figures

Figure 3.1	Disaster Model Schematic	9
Figure 3.2	Two Sample Particle Trajectories	11
Figure 3.3	Disaster Model (Special Case)	12
Figure 3.4	Three Sample Particle Trajectories ($VV = 0, VV = 1, VV = 5$)	14
Figure 3.5	Gridded State Space with numbered Bins	16
Figure 3.6	Grid of Histograms	17
Figure 3.7	Extension of the Random Walk Particle Model	18
Figure 3.8	Sample Particle Trajectory ($VV = 3, HV = 3$)	19
Figure 3.9	Explanation of Volatility	20
Figure 3.10	Two Sample Particle Trajectories	21
Figure 3.11	Two Similar Histograms	22
Figure 3.12	Creating a Probability Map	23
Figure 3.13	Large Probability Matrix (Map)	24
Figure 3.14	Particle Trajectory with P_{hit} data	26
Figure 3.15	P_{hit} versus Vertical Volatility	27
Figure 4.1	Cost-To-Evacuate Functions	31
Figure 4.2	Decision Matrix for Constant Evacuation Cost	32
Figure 4.3	Decision Matrices for Linearly Increasing Evacuation Cost	35
Figure 4.4	Heatmap (Decision Matrix) for Linear Increasing Evacuation Functions	36

Figure 4.5	Heatmap (Decision Matrix) for Convex Increasing Evacuation Functions	39
Figure 4.6	Heatmap (Decision Matrix) for Concave Increasing Evacuation Functions	40
Figure 4.7	Heatmap (Decision Matrix) for Combinations of Linear and Convex Evacuation Functions.	41
Figure 4.8	Decision Matrices for Changing HV	42
Figure 4.9	Volatility and Dynamic Programming	43
Figure A.1	Final V_{posit} Distributions: $H_{posit}(0) \leq 30$ and $H_{posit}(0) \leq 90$	54
Figure A.2	Final V_{posit} Distributions: $H_{posit}(0) \leq 150$ and $H_{posit}(0) \leq 270$	55
Figure B.1	Decision Matrices for a Linearly Increasing Evacuation Cost	58
Figure B.2	Decision Matrices for a Convex Increasing Evacuation Cost	59

List of Tables

Table 1.1	Saffir-Simpson Hurricane Scale	2
Table 4.1	Cost Matrix	31
Table A.1	Number of Steps (k') with $VV = n$	53

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

$2D$	Two Dimensional
C_e	Evacuation Cost
$\#Bins$	Number of Bins
ΔH	Horizontal Step
ΔV	Vertical Step
DP	Dynamic Programming
E	Evacuation Indicator
H_{posit}	Horizontal Position
HV	Horizontal Volatility
J^*	Cost-To-Go Function
NOAA	National Oceanic and Atmospheric Administration
NHC	National Hurricane Center
P_{hit}	Hit Probability
TTG	Time-To-Go
V_{posit}	Vertical Position
VV	Vertical Volatility

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

Natural disasters effect every level of our society, from individual citizens to high-ranking government officials. These destructive events affect our livelihood, our economy, and the security of our nation. Unfortunately, natural disasters carry some level of uncertainty in that it is hard to predict when or where a disaster will occur and how destructive a disaster will be when and if it does occur.

Decision making in the face of uncertainty is a difficult task. This task is particularly difficult when it involves one or more of the following characteristics: 1) The decision maker is faced with a near-term deadline, 2) The stakes are high for any incorrect decision, and/or 3) The decision is irreversible. Unfortunately, the decision of whether or not to evacuate from a natural disaster has all of these characteristics.

This thesis explores the evacuation decision problem for a generic disaster process that idealizes natural disasters such as hurricanes, wildfires, and flooding. We create a stochastic model that simulates disaster movement and describe how this model is used to explore the tensions and tradeoffs in the decision to to evacuate. In this model, a “particle” moves through a discrete and bounded state space toward a “target.” The particle represents the disaster, while the target represents an decision maker who is at risk from the impending disaster.

As the disaster progresses, the decision maker must decide whether to “stay” or “evacuate.” In our model, the decision to evacuate is irreversible and the decision to stay is reversible, meaning that the decision maker may decide to evacuate at a later time. Each decision carries some potential consequences that are expressed as a cost. The decision maker can evacuate and incur an evacuation cost at the time of this decision. The evacuation cost is described via an evacuation cost function (C_e). The decision maker may evacuate at any time before the disaster strikes; however, if the decision maker has not evacuated when the disaster strikes, the decision maker incurs a cost based on the outcome of the disaster. No cost is incurred if the decision maker stays and the disaster does not strike the target; however, the decision maker incurs a significant cost if he or she stays and the disaster strikes the target.

We implement a dynamic programming (DP) algorithm to find optimal decision policies for the decision maker. We assume that the decision maker seeks to minimize the expected cost that will be incurred once the disaster outcome is known. A top-down version of DP is used to calculate an optimal decision (stay or evacuate) for all achievable states as well as the expected

cost associated with each state.

An evacuation policy describes regions of the state space in which the optimal decision is to evacuate. We show that when C_e is constant, the decision maker can always wait until the last time step to evacuate. Then, by exploring the effects of an increasing C_e , we see that decision policies may have multiple distinct evacuation regions and that the shape and size of these regions is affected by the size and shape of the evacuation cost function (C_e).

Lastly, we describe the impact of particle (disaster) uncertainty on the size and shape of the evacuation region. We show that the evacuation region expands as disaster uncertainty increases, thus indicating that it is prudent to respond even to distant threats when the uncertainty about the threat's future states is high.

This thesis lays the framework for future work by developing a disaster model and providing initial results on optimal decision policies and how these policies are impacted by uncertainties about the disaster process. We recommend some ideas for future work in this area and describe how these extensions will provide a more detailed understanding of the decisions surrounding disaster evacuations.

Acknowledgements

First and foremost, I acknowledge my Lord and Savior Jesus Christ who gives me the strength to do all things. I would like to thank all of the faculty and staff of the Naval Postgraduate School Operations Research Department for providing me with the skills to tackle the world's problems. I especially thank my thesis advisor, Dr. Emily Craparo, for her patience, support, and wisdom. She made my "thesis journey" a memorable experience. I would also like to thank Dr. David Alderson for his guidance. Last but not least, I would like to thank my lovely wife LeAnne and my two children Ethan and Alyssa. Thanks for standing by me throughout my career in the United States Navy.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 1:

INTRODUCTION

Decision making in the face of uncertainty is a difficult task. This task is particularly tough when it involves one or more of the following characteristics: 1) The decision maker is faced with a near-term deadline, 2) The stakes are high for any incorrect decision, and/or 3) the decision is *irreversible*, in the sense that it significantly reduces the variety of choices available in the future (Henry, 1974).

Consider a law enforcement officer that must decide whether or not to fire his or her weapon at a suspect. Once the officer fires the weapon, this decision cannot be reversed and he or she must face the consequences of that decision. If the officer shoots and kills a suspect that turns out not to have been a real threat, that officer faces civil and potential criminal penalties. If the officer does not shoot and the suspect is a real threat, the lives of the officer and others may be endangered. The uncertainty about the suspect, combined with both the irreversible nature of the action and the high stakes, make this decision difficult.

In emergency situations, people also must make time-critical decisions based on uncertain information. Their situational awareness can sometimes be improved through additional observations of the threat; however, these observations delay action. A resident of a coastal city faces a difficult decision when encountering an approaching hurricane. If the hurricane makes landfall, the consequences of delayed action can be catastrophic because it could result in significant property damage, injury, or death. However, hasty action can also be costly. Both the individual and government could unnecessarily incur significant costs if everyone evacuates and the hurricane does not strike that region.

Unfortunately, this situation is very common in the United States. From 1851-2004, 273 hurricanes have made landfall on mainland United States (Blake, Rappaport, & Landsea, 2007). Of these 273 hurricanes, 92 have been considered major hurricanes (Blake *et al.*, 2007). A major hurricane has wind speeds greater than 110 mph and corresponds to a category 3, 4, or 5 on the Saffir-Simpson Hurricane Scale, shown in Table 1.1 (Blake *et al.*, 2007). So on average, there have been approximately 17.7 hurricanes per decade and 6 major hurricanes per decade since 1851.

Table 1.1: Saffir-Simpson Hurricane Scale (Blake *et al.*, 2007).

Category	Winds(MPH)	Damage
1	74-95	Minimal
2	96-110	Moderate
3	111-130	Extensive
4	131-155	Extreme
5	>155	Catastrophic

Hurricanes can inflict a significant amount of damage on the affected regions. In 2005, hurricanes striking the mainland United States caused 1,525 deaths and damages costing \$115 billion (Blake *et al.*, 2007). Hurricane Katrina, which struck southeast Louisiana on August 29, 2005, was the most costly hurricane to strike the mainland United States; damage caused by Katrina is estimated at \$81 billion (Blake *et al.*, 2007).

The costs associated with hurricane preparation and evacuation are significant to both evacuees and government organizations. For an evacuee, there are costs associated with home preparation, missed work and other opportunities, as well as the travel, lodging, food costs associated with the evacuation itself. At the government level, the cost of evacuating a region is often estimated at about \$1 million per mile of coastline Regnier (2008). Moreover, Regnier (2008) states, “a rough estimate of the average annual cost of false alarm evacuations is over \$1 billion.”

Decisions related to natural disasters are challenging. The uncertainties associated with these phenomena make it difficult to predict which regions will be affected. The short lead times and the potential for great devastation necessitate timely decisions, while the costs associated with a false alarm require well-informed decisions. Logistical factors make it impossible to delay the decision until just before the crisis happens. Clearly, good decision policies are required to make timely decisions that save both lives and money in the context of a natural disaster.

This thesis explores the evacuation decision problem for a generic disaster process. Chapter 2 provides a literature review to support the efforts of this study. Chapter 3 describes a model for the generic disaster process and describes how this model is used to explore the tensions and tradeoffs in the decision to evacuate. Then Chapter 4 describes a method to determine optimal decision policies for the decision maker. Finally, Chapter 5 summarizes the efforts of this thesis and recommends areas for future work.

CHAPTER 2:

BACKGROUND AND CONTRIBUTIONS

2.1 Overview

There is a wide variety of research that studies evacuations from natural disasters. This research includes efforts that seek to improve disaster forecast models, evacuation routes, and evacuation decision policies. A small subset of this research area incorporates either decision theory, dynamic modeling, or dynamic programming techniques. This chapter includes a review of four academic papers that either seek to determine an optimal evacuation policy or model the costs associated with a particular disaster. Each of these papers is relevant to and supports the efforts of this thesis. Each discussion includes a comparison between the methods used in each article and the methodology of this thesis.

2.2 Mitigating Earthquake Risks

Large earthquakes have the potential to cause significant structural damage and a large number of fatalities. In 1995, Awaji island and the Kobe area of Japan experienced a 6.8-magnitude earthquake that left 6,000 people dead (Tamura, Yamamoto, Tomiyama, & Hatano, 2000). Events such as this have motivated researchers to determine ways to mitigate the risks associated with these events. Some research seeks to predict earthquakes, but this area of study faces many difficulties (Tamura *et al.*, 2000). Preparation is another approach to mitigating the affects of an earthquake. For example, strengthening buildings can lower the cost of building repairs and reduce the number of injuries and fatalities (Tamura *et al.*, 2000).

The costs associated with building alterations or reconstruction are significant, but so are the consequences of being unprepared. Decision theory techniques can help us make sound decisions about earthquake preparations. Tamura *et al.* (2000) use both an *expected utility model* and a *value function under risk model* to analyze the decision making process for a large earthquake. The value function under risk model is one that incorporates earthquake occurrence probabilities as an attribute of the evaluation. The authors classify large earthquakes as low probability/high consequence events and suggest that expected utility model is inadequate for this type of problem.

Tamura *et al.* (2000) formulate a multiattribute problem to compare their value function under

risk model to an expected utility model. This study shows that the expected utility model fails to show the impact of risk and is not adequate for analysis of low probability events, where probability that this event will occur during the next 10 years is $< 10^{-2}$ (Tamura *et al.*, 2000). On the other hand, the authors observed that the value function under risk model was useful for evaluating low probability/high consequence events such as large earthquakes because it could represent the impact of risk that the utility model could not (Tamura *et al.*, 2000).

This thesis models a generic disaster process that is quite different than a large earthquake. A large earthquake is characterized as a low probability event that occurs with little to no warning. The generic disaster process in this thesis provides ample warning to the decision maker and often has a moderate occurrence probability ($> 1\%$). Similar to Tamura *et al.* (2000), this thesis incorporates the disaster occurrence probability into the evaluation.

2.3 Probabilistic Evacuation Decision Model for River Floods

Flooding is a significant problem in the Netherlands. Much of that country's land lies below sea level or is susceptible to river flooding. About 65% of the Netherlands would experience either temporary or permanent flooding without the use of flood protection (Frieser, 2004). Experts believe that additional measures need to be taken to protect populated areas from future flooding. To make matters worse, predictions show a potential rise in flooding hazards due to climatic changes and land subsidence (Frieser, 2004).

In 1995, 240,000 people were ordered to evacuate the Nijmegen region of the Netherlands in response to potential flooding (Frieser, 2004). The evacuation was conducted successfully, but no flooding occurred in the evacuated region. This raised concerns about the necessity of the evacuation and the evacuation decision process itself. The current decision process for flood evacuations is based on a deterministic criterion (current water level) and expert advice about the condition of flood protection devices (Frieser, 2004).

Frieser (2004) presents a probabilistic evacuation decision model that determines an optimal decision for any point in time. This probabilistic model determines a optimal decision (evacuation, do not evacuate, or delay decision) by incorporating the cost of evacuation, potential flood damage, and the probability of flooding (Frieser, 2004). They test the probabilistic model using data from the 1995 flood, and the resulting decision is the same. Frieser (2004) determines that the uncertainty in water level predictions have the largest effect on the model; thus, water level predictions need to be improved in order to use this model in practice. Nevertheless, Frieser

(2004) asserts that this framework provides a useful tool for decision-makers.

Similar to Freiser's work, this thesis uses a dynamic model to determine an optimal decision for each point in time. In this thesis, the decision maker can make one of two decisions (evacuate or wait) and the optimal decision is determined using evacuation costs, potential disaster damage, and the disaster state.

2.4 Dynamically Modeling Hurricane Evacuation Decisions

Tropical cyclones affect the eastern portion of the United States and the Caribbean each year. These storms form as tropical depressions over the warm water in the Atlantic Ocean and proceed westward. As these storms move, they often strengthen into a tropical storm or a hurricane before making landfall. The National Hurricane Center (NHC) tracks tropical cyclones and issues forecast information every six hours. Forecasts are issued from the time that a tropical depression is formed until the storm dissipates (Czajkowski, 2007). These periodic forecasts are released daily at 5:00 am, 11:00 am, 5:00 pm, and 11:00 pm (Czajkowski, 2007). Potential evacuees monitor these forecasts and ultimately make the decision to either stay at home or to evacuate. Government officials also use this information to determine when to recommend or order evacuations in the areas under their control. When these evacuations occur, evacuation routes can quickly become congested, emergency shelters often reach capacity, and unfortunately some people become trapped and remain vulnerable to the effects of the storm.

There is a continued need to study the behavior of evacuees during a natural disaster such as a hurricane. In response to this need, Czajkowski (2007) creates a dynamic model to model hurricane evacuation behavior. This model is a multi-period model where each period corresponds to a NHC hurricane forecast (Czajkowski, 2007). This model incorporates actual hurricane forecast data from the Gulf of Mexico region and evacuation cost data from the the entire Southeast region of the US. Following some initial test runs, this model is calibrated using actual hurricane evacuation data. The results from this dynamic model provide a deeper understanding of evacuation timing data and support the analysis of numerous policy issues regarding the timing of evacuations (Czajkowski, 2007). Czajkowski (2007, 104) also shows a need for accurate modeling of evacuation behavior in a more precise and comprehensive way.

In a similar manner, this thesis models a generic disaster process using a multi-period dynamic model. Whereas Czajkowski (2007) uses his model to explain evacuation behavior in a specific region, this thesis model uses dynamic programming to examine optimal evacuation decision

policies. This thesis model does not attempt to model one specific type of disaster, so no outside data is used to calibrate the results.

2.5 Evacuation Decisions and Hurricane Track Uncertainty

The coastal population continues to grow in the United States, and so do the costs associated with hurricanes. Estimates show that the long term average cost of hurricane damage is about \$10-11 billion per year (Regnier, 2008). Hurricane evacuation costs are also significant. The average cost of hurricane evacuations is estimated at \$1.5 billion per year, of which \$1 billion is due to false alarm evacuations (Regnier, 2008).

False alarm evacuations are quite common since only 25% of the regions that come under a hurricane warning actually experience hurricane force winds (Regnier, 2008). Evacuating a region takes a significant amount of time, so evacuations must be ordered with considerable lead time. Regnier (2008) shows that the quality of forecast information declines as evacuation lead time increases. This means that public officials are faced with the difficult and dynamic problem of balancing the risk of delaying an evacuation decision with the improving accuracy of forecast data (Regnier, 2008).

Regnier (2008) makes the connection between the timing of the evacuation decision and the time profile of forecast uncertainty using a Markov model of storm motion, developed by Regnier & Harr (2006). Regnier (2008) quantifies the track uncertainty for storms. Four target cities are used to show how the trend in hurricane strike probabilities depends on geographic location. Regnier (2008) sets a framework for quantifying hurricane uncertainty as a function of lead time and geographic location, which are relevant terms to the decision maker.

Regnier (2008) considers the tradeoff between delaying an evacuation decision and the risk associated with waiting for more information (a better forecast). This thesis incorporates a similar consideration through two separate methods. The first method involves a rising evacuation cost, and the second involves varying levels of disaster uncertainty.

2.6 Thesis Contributions

This thesis supports a larger study that seeks to determine how people make decisions when faced with an impending disaster. The larger study will involve behavioral experiments in which human participants must make decisions about a simulated disaster. This thesis describes a disaster model that is under consideration for these behavioral experiments. This model was

selected because it is simple and intuitive, yet it reflects many of the tensions and uncertainties encountered in actual disaster situations. In support of the larger study, this thesis provides an analytical study of this generic disaster model and calculates optimal decision policies for a decision maker.

We model the movement of a generic disaster process with a stochastic “particle” that occupies a single discrete state at each point in time. It then uses dynamic programming techniques to determine optimal evacuation decision policies. Chapter 3 describes the stochastic particle model framework. Two “volatility” terms describe the uncertainty in particle (disaster) movement. This thesis uses simulation to generate histograms that show the distribution of final particle states. These histograms are used to illustrate how changes in the initial particle state and volatility affect the distribution of final particle states. This study also analytically derives an equivalence between disaster volatility and time to strike. Two-dimensional “probability heatmaps” are created to visualize the probability that the disaster will strike a given location (target) given the disaster’s current location. Chapter 3 ends with a description of how these heatmaps are affected by changes in particle volatility.

Chapter 4 describes how optimal evacuation decision policies are determined with a dynamic programming algorithm. These decision policies give the optimal decision (evacuate or stay) for all achievable particle states and calculate the expected cost associated with each state.

Model parameters are varied to determine their effects on these decision policies. These variations include rising evacuation costs and changes in particle volatility. Rising evacuation costs imply that decision makers generally face a penalty when the decision to evacuate is delayed. Particle volatility reflects the degree of uncertainty in when and if the disaster could strike. These model embellishments are used to add reality to the dynamic model and to provide additional insights into the evacuation decision process.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 3: MODELING THE DISASTER PROCESS

This chapter describes a model of a generic disaster process that is used to explore specific tensions and tradeoffs in the decision to evacuate. Figure 3.1 gives a schematic representation of this model, which includes a *particle* that moves within a bounded, two-dimensional (2D) *state space* toward a stationary *target*. In this representation, the particle moves from left to right according to a random walk until it reaches the rightmost boundary, at which point all motion ceases. The position of the particle within this 2D space is its *state*, and the movement of the particle is its *trajectory*. The upper and lower boundaries of the state space are *reflecting*, in the sense that they keep the trajectory of the particle within the bounded state space. Thus, a single trajectory for the particle always starts with the particle at the left side of the state space and ends with the particle reaching the rightmost boundary. If the ending point of the trajectory coincides with the target, we say that “the particle has *hit* the target,” otherwise we say that “the particle has *missed* the target.”

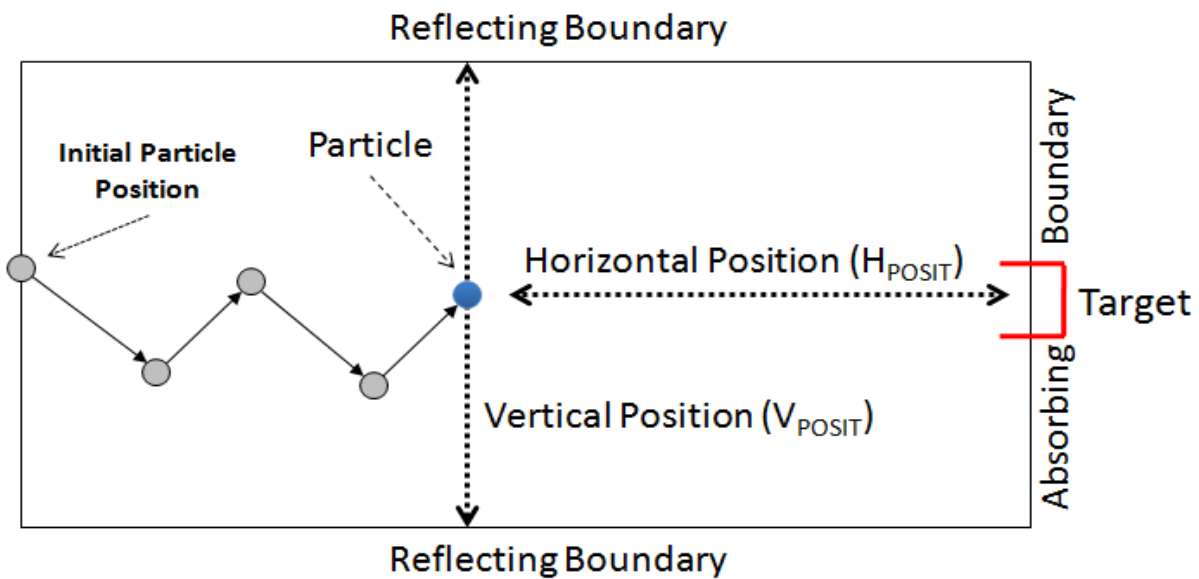


Figure 3.1: A schematic of the disaster process model. The trajectory of the particle represents the behavior of an impending disaster. If the particle hits the target, the disaster has affected the entity (e.g., the hurricane has hit the town). If the particle misses the target, this has not occurred.

The occurrence of a disaster is associated with the particle hitting the target. This association can be interpreted in two different ways. First, the particle might represent the disaster itself (e.g., a hurricane), and the target might represent an entity that could be affected by the disaster (e.g., a town). In this interpretation, if the particle hits the target, then the disaster has affected the entity (e.g., the hurricane has hit the town). If the particle misses the target, this has not occurred.

A second interpretation of the particle is that of a general stochastic process that results in one of two outcomes: either the disaster happens, or it does not.

At each discrete point in time t , let $H_{posit}(t)$ represent the particle's horizontal distance from the rightmost boundary of the state space, and let $V_{posit}(t)$ represent the particle's vertical distance from the center of the state space. Thus, the state of the particle is described by the pair $(H_{posit}(t), V_{posit}(t))$.

The particle updates its position in time according to the following rule:

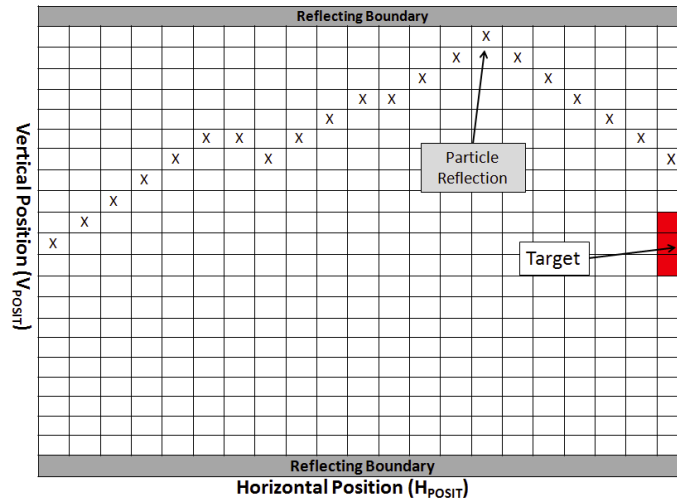
$$\begin{aligned} H_{posit}(t+1) &= H_{posit}(t) + \Delta H \\ V_{posit}(t+1) &= V_{posit}(t) + \Delta V \end{aligned}$$

where $\Delta H > 0$ is the horizontal step (to the right) and ΔV is the vertical step. These steps are discrete in nature and can either be deterministic or stochastic. When steps are stochastic, a new value for ΔH or ΔV is realized at each time step.

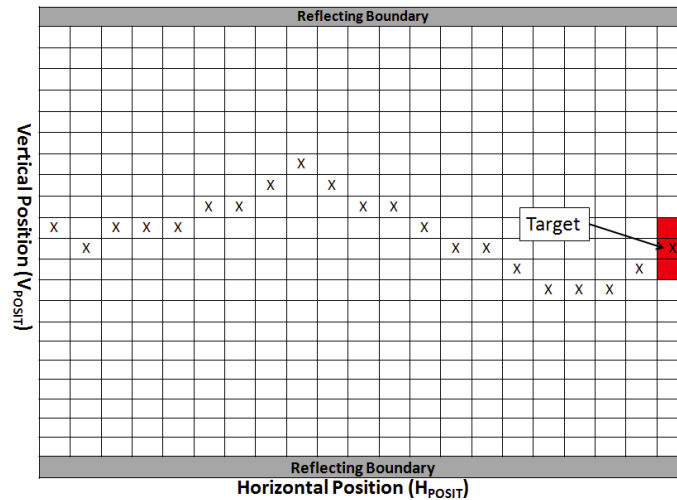
The particle's movement is constrained by the boundaries of the state space. As mentioned before, the upper and lower boundaries restricting the particle's vertical movement are *reflecting boundaries*. If the particle hits an upper or lower boundary, it "reflects" off that boundary and moves toward the center of the state space. For example, if the particle attempts to move n steps above the upper boundary, it is reflected to a new position n steps below the upper boundary. In contrast, the boundary to the right is an *absorbing boundary*.

Figure 3.2 shows two different particle trajectories. These two trajectories illustrate some key points about particle movement. The first trajectory, shown in Figure 3.2a, displays a particle that reflects off the upper boundary and misses the target. Figure 3.2b shows another trajectory in which the particle hits the target but does not make contact with either of the reflecting

boundaries. These figures include a grid to illustrate the discrete nature of the 2D state space. For the remainder of this thesis particle trajectories are displayed without the grid, and adjacent particle states are connected with a line. This is done to make it easier for the reader to observe a particle's movement.



(a) Sample particle trajectory in which the particle reflects off the upper boundary and misses the target.



(b) Sample particle trajectory in which the particle never encounters a reflecting boundary and hits the target.

Figure 3.2: Two sample particle trajectories that display some key characteristics of particle movement. Subfigure (a) shows a particle that reflects off the top boundary and misses the target. Subfigure (b) shows a particle hits the target but never encounters either of the reflecting boundaries.

3.1 Special Case: Deterministic Horizontal Motion ($\Delta H = 1$)

This section describes a special case of the disaster model that considers a deterministic horizontal step size ($\Delta H = 1$) and a stochastic vertical step size (ΔV). This special case is shown schematically in Figure 3.3. In this thesis, *volatility* describes a limit on stochastic particle movement. The magnitude of the particle's step in the vertical direction at time t is given by a discrete random variable ΔV that is uniformly distributed between $-VV$ and VV , where VV is an integer-valued parameter called the *vertical volatility* of the disaster process. Vertical volatility (VV) is defined as the maximum number of vertical steps that the particle can move in a single time step.

Horizontal movement is deterministic in this special case, so there is a 1-to-1 mapping between the particle's horizontal position and the time until the particle will strike the rightmost boundary. The time until the particle strikes the rightmost boundary is denoted as the *Time-To-Go* (TTG).

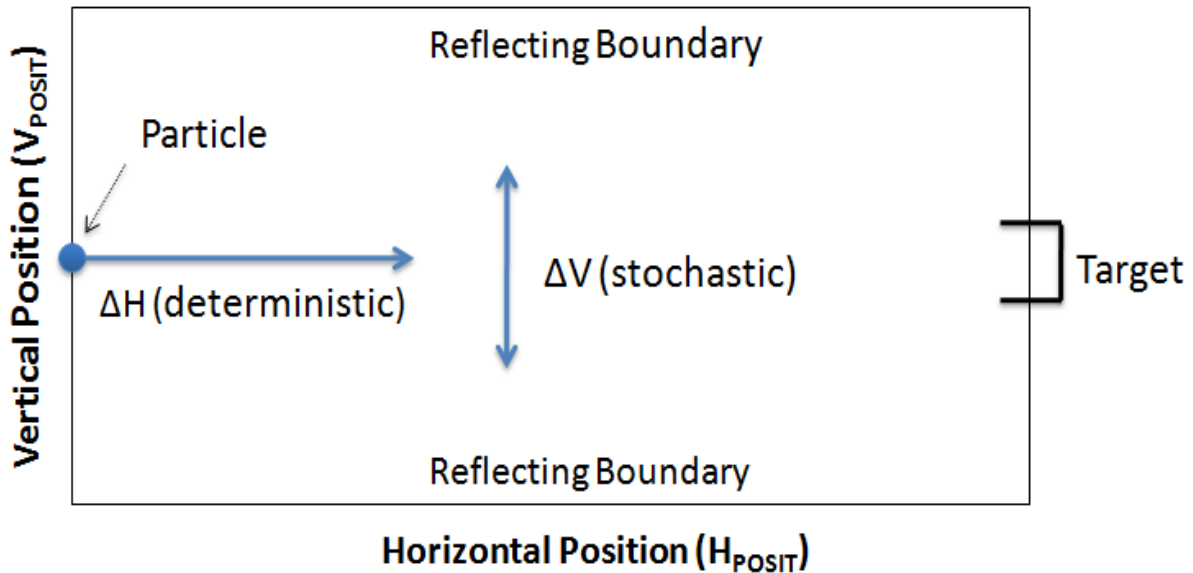


Figure 3.3: Special Case: Deterministic Horizontal Motion ($\Delta H = 1$). A random walk particle model that simulates a disaster process approaching an entity. The particle represents the disaster process, while the target represents the entity that could be affected by the disaster. The particle's state is given by its geographic position relative to the target. Each time the particle moves it takes one deterministic step toward the target along with some stochastic vertical movement. The particle's movement is limited by reflecting upper and lower boundaries.

3.1.1 Impact of Vertical Volatility on Final Particle State

One goal of this thesis is to characterize the probability that the particle will ultimately strike the target, given its horizontal and vertical position. We denote this probability as the *hit probability*, P_{hit} . In order to calculate P_{hit} , it is necessary to consider the way in which the particle can move. This subsection focuses on vertical movement of the particle, since horizontal movement is deterministic in this special case.

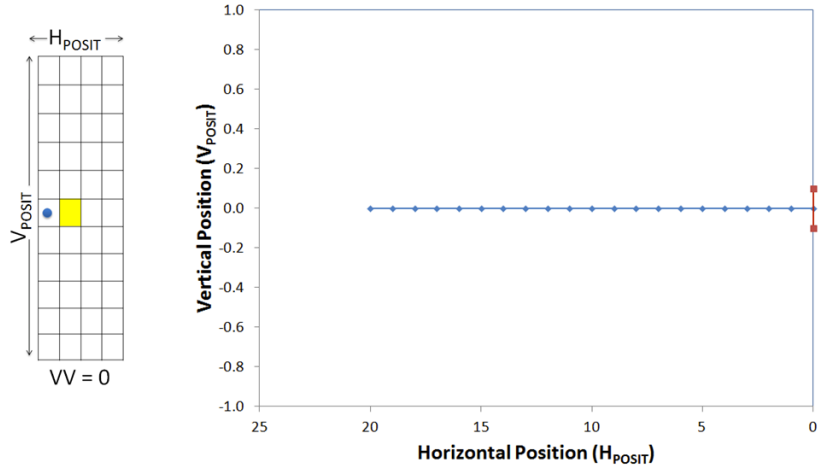
Figure 3.4 shows three particles with different vertical volatilities VV : $VV = 0$, $VV = 1$, and $VV = 5$. A particle with $VV = 0$ must remain in the same V_{posit} on the next time step; thus, its vertical movement is deterministic. A particle with $VV = 1$ can move to one of three vertical positions in the next time step. It can remain in the same position, move up one step, or move down one step. A particle with $VV = 5$ can move to one of eleven vertical positions in the next time step. For any VV , the particle can move to one of $2VV + 1$ vertical positions in a single time step. In all cases, the probability that the particle will proceed to one particular future vertical position is equal to $\frac{1}{2VV+1}$.

A particle's VV places a hard limit on how far the particle can travel vertically over a given number of steps. This limit implies some relationship between a particle's initial state, its vertical volatility, and its set of possible final states. We will now explore this relationship empirically; Appendix A explores it analytically.


Sample Histograms

This section describes the results of six computational experiments that are designed to show the relationship between a particle's initial state $(H_{posit}(0), V_{posit}(0))$, its vertical volatility VV , and its final state. The following is conducted for each experiment:

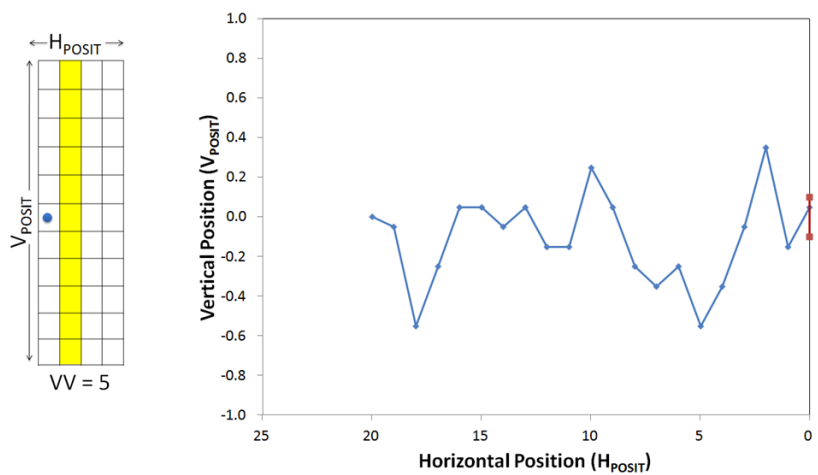
1. A particle is placed in some initial state and assigned a VV .
2. The size and position of the target are established.
3. The particle's trajectory is simulated until the particle strikes the right boundary. Under this special case, this occurs at time $T_f = H_{posit}(0)$.
4. The final position of the particle is recorded.



(a) Sample trajectory for a particle with $VV = 0$.



(b) Sample trajectory for a particle with $VV = 1$.



(c) Sample trajectory for a particle with $VV = 5$.

Figure 3.4: Three sample trajectories for particles with different levels of VV . For each trajectory, horizontal movement is deterministic, and vertical movement is stochastic. The leftmost figures show the range of stochastic vertical movement for each particle.

This process is repeated one million times for each set of initial conditions. A vector of final positions $V_{posit}(T_f)$ is created from the repeated experiments, and this vector is used to create a histogram showing the distribution of final particle states. From these final positions one can calculate P_{hit} .

Figure 3.5 explains the conditions for the six experiments. The experiments are labeled with the letters (a) through (f). These labels include the parameters that are changed between each experiment (H_{posit} and VV). Highlighted cells indicate the initial state of the particle for each experiment. The rightmost column of the grid contains twenty one numbered bins representing possible final particle states.

The histograms in Figures 3.6 show the distribution of final positions (final states). For each histogram, a probability value is displayed for each bin. These bins correspond to those shown in the rightmost column of Figure 3.5. A bin value represents the probability that a particle, with the given initial conditions, will land in that bin. These empirical probabilities are determined by dividing the number of particles that landed in that bin by the total number of replications. For example, assume that a bin has the probability 0.14 that resulted from 1,000,000 replications. This means that approximately 140,000 of the 1,000,000 particles landed in this bin.

Let us proceed with a discussion of the results shown in Figure 3.6. Figure 3.6a (Experiment (a)) shows a distribution of final positions for the following initial conditions: $V_{posit}(0) = 10$, $H_{posit}(0)$, and $VV = 1$. Experiment (a) is considered the “base case” for the remainder of discussion. Under these initial conditions, at least one particle lands in each of the bins, and the target is hit 42% of the time. In general, a majority of the particles land near the middle of the rightmost wall.

In Experiment (b), $H_{posit}(0) = 20$ and all other parameters remain the same. The resulting distribution shifts outward toward the reflecting boundaries, and P_{hit} drops to 31%. For Experiment (c), $H_{posit}(0) = 40$. Even more of the distribution shifts toward the boundaries, and P_{hit} lowers to 22%.

Experiment (d) implements an increase in VV from the base case. A comparison between Figures 3.6a and 3.6d shows a large shift in the distribution toward the boundaries. P_{hit} drops nearly in half for this change in VV . Similar observations can be made by comparing Experiments (b) with (e) and (c) with (f).

Experiment (f) shows a nearly uniform distribution of final positions. Any further increase in

VV or $H_{posit}(0)$ would yield similar results to those shown in Figure 3.6f.

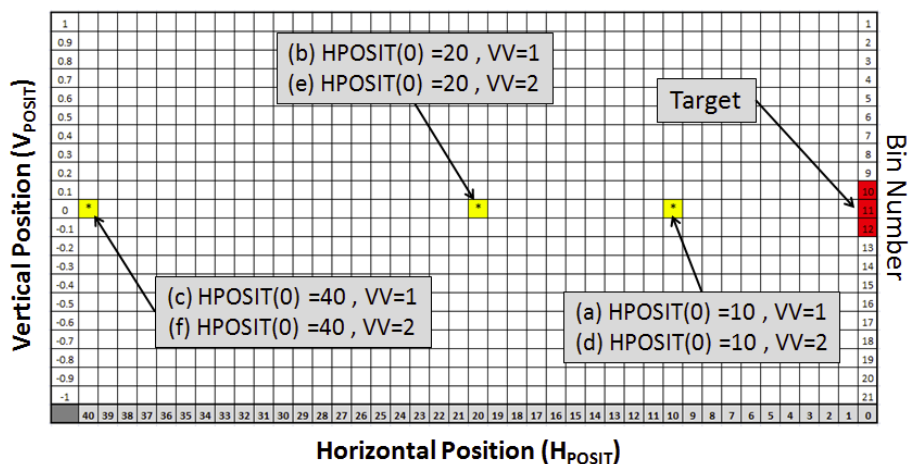


Figure 3.5: A gridded representation of the particle state space. This figure illustrates the initial particle positions for three separate experiments. The far right column contains numbered bins that represent the final position of a particle trajectory. Bins 10, 11 and 12 represent the target and are highlighted in red. The annotation: “(a) $H_{posit}(0) = 10$, $VV = 0$ ” means that Experiment (a) used the initial conditions $H_{posit}(0) = 10$, $VV = 0$

These empirical results indicate a clear relationship between a particle’s initial state, its vertical volatility, and its distribution of possible final states. In particular, the Shannon entropy (Klir, 2006, p. 69) of the distribution of final states can be increased via either an increase in vertical volatility or an increase in $H_{posit}(0)$. Appendix A explores the relative impact of these two factors analytically.

3.2 General Disaster Model

Figure 3.7 illustrates a general disaster model that includes stochastic horizontal motion. In this model, the particle can take random steps in both the vertical and horizontal directions. Figure 3.8 shows the trajectory of a particle that takes stochastic horizontal and vertical steps. The particle’s H_{posit} can no longer be directly mapped to its TTG ; however, TTG can be estimated using information about the random distribution that is used to determine ΔH .

3.2.1 Horizontal Volatility

Horizontal volatility (HV) describes the maximum number of horizontal steps that the particle can take in a single time step. Horizontal volatility (HV) is defined differently from VV . For horizontal movement, $HV = 0$ indicates deterministic movement; however, the particle is

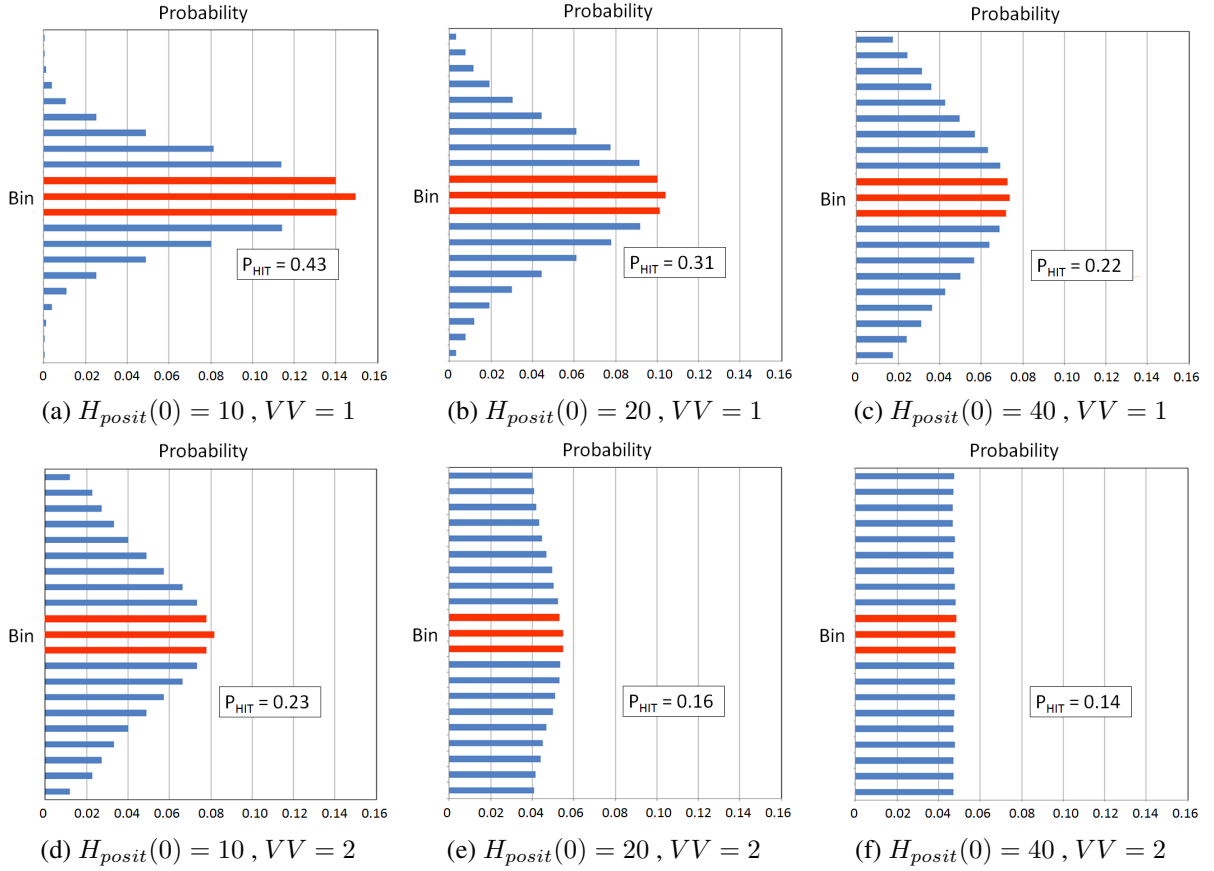


Figure 3.6: These histograms illustrate the effect of VV and the initial particle position ($H_{posit}(0)$) on the distribution of final particle positions. H_{posit} values are constant for each column of histograms and increase from left to right. VV values are constant for each row and increase from top to bottom. As H_{posit} increases, the distribution flattens as a larger proportion of the particles move away from the center of the state space. The distribution also flattens as VV is increased.

modeled as moving exactly one step to the right each time that it moves. When $HV = k$, the particle moves between one and $k + 1$ steps to the right.

Figure 3.9 shows four different particles with different values for VV and HV . This figure illustrates how far a particle can move in a single time step. A particle with $VV = 0$ and $HV = 0$ moves one step to the right but must remain at the same V_{posit} . A particle with $VV = 1$ and $HV = 0$ must take one deterministic step to the right but can move to one of three vertical positions. If a particle's $VV > 0$ and $HV > 0$, the particle can move stochastically in both the horizontal and vertical directions. A particle with $VV = 1$ and $HV = 1$ can move

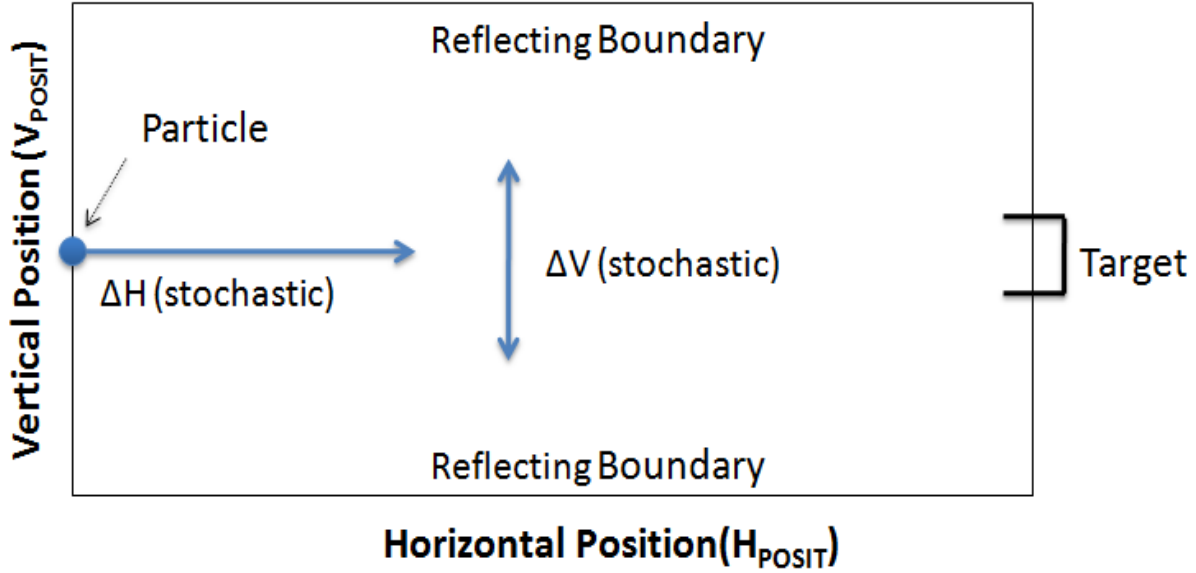


Figure 3.7: A generalization of the random walk particle model shown in Figure 3.3. This generalization introduces uncertainty to the time until the disaster could strike. Each time the particle moves, it takes one stochastic step towards the target along with some stochastic vertical movement.

to one of six positions, while a particle with $VV = 1$ and $HV = 2$ can move to one of nine positions.

When the particle's horizontal movement is stochastic, the particle can occupy one of $(2VV + 1)(HV + 1)$ possible positions in the next time step. Again, the particle moves to each possible position with probability $\frac{1}{(2VV+1)(HV+1)}$.

3.2.2 Expected TTG

Although there is no longer a one-to-one mapping between horizontal position at time-to-go (TTG) when horizontal movement is stochastic, TTG can be estimated using the particle's current horizontal position and information about the random variable that controls the horizontal movement.

A simple experiment is used to illustrate the difference between deterministic and stochastic horizontal movement. This experiment involves two particles with the same initial state ($H_{posit}(0) = 20, V_{posit}(0) = 0$), the same VV ($VV = 0$), but different HV . The results from this experiment are shown in Figure 3.10. $VV = 0$ for both cases, so the particles remains

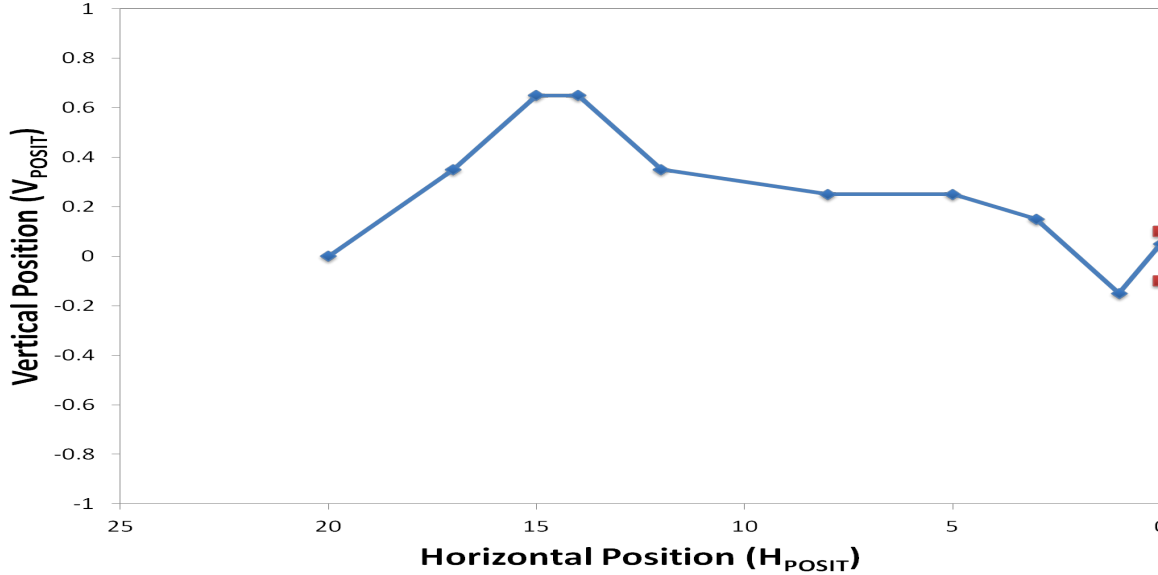
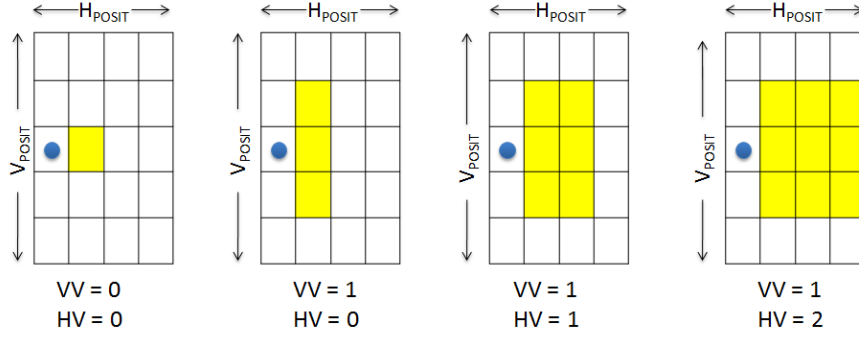


Figure 3.8: A trajectory that shows a particle ($VV = 3$, $HV = 3$) approaching and hitting the target. Movement in the both the horizontal and vertical directions is stochastic. Note that the initial particle state is 20 horizontal units away from the rightmost boundary, but the particle hits the target after taking 9 steps.

at the same vertical position throughout their trajectories. Figure 3.10a shows a trajectory for a particle with $HV = 0$. Horizontal movement is deterministic, so a one-to-one relationship exists between the particle's horizontal position and TTG . This particle strikes the target after taking exactly 20 horizontal steps. Figure 3.10b shows a trajectory for a particle with $HV = 1$. This particle can move either one or two horizontal units each time it moves. For this example, the particle travels twenty horizontal units in only fifteen steps. In any particular experiment, the particle could take more or fewer steps before striking the righthand wall. Average values for the TTG and the particle step-size (ΔH) are calculated to describe the behavior of this stochastic process.

The expected time-to-go ($\mathbb{E}[TTG]$) is the average time until a particle could strike the target. $\mathbb{E}[TTG]$ can also be thought of as the average number of horizontal movements that are required for the particle to move from its current state to the rightmost boundary. The expected step-size ($\mathbb{E}[\Delta H]$) is the average horizontal step-size that a particle is expected to take throughout a single trajectory. For a particle with $HV = 1$, the $\mathbb{E}[\Delta H] = 1/2(1) + 1/2(2) = 1.5$. When this particle is 30 horizontal units away from the target, the $\mathbb{E}[TTG] = 30/1.5 = 20$.



VV = Vertical Volatility **HV** = Horizontal Volatility

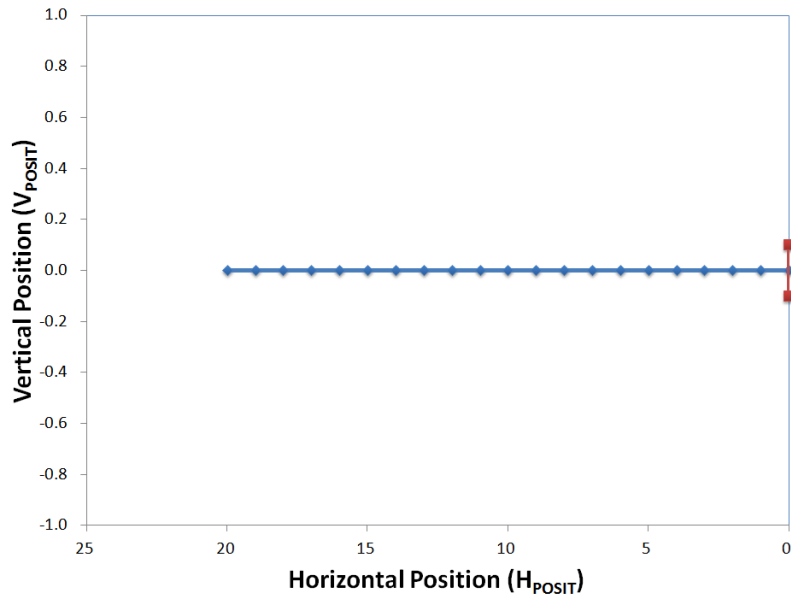
Figure 3.9: Four grids that illustrate particle volatility. A particle with $VV = 0$ will remain in the same V_{posit} throughout its trajectory. A particle with $VV = 1$ can either remain in the same V_{posit} , move up one step, or move down one step. A $HV = 0$ means that the particle will progress deterministically by moving one step to the right. If the $HV = 1$, the particle can move either one or two spaces to the right.

Just as there is a relationship between a particle’s initial state, its VV , and the distribution of possible final states, there is also a relationship between a particle’s initial state, its HV , and the distribution of possible final states. Figure 3.11a shows the distribution of final states for a particle with $H_{posit}(0) = 30$, $V_{posit}(0) = 0$, $VV = 1$, and $HV = 1$. $\mathbb{E}[TTG] = 20$ for a particle with these initial conditions. Figure 3.11b shows a similar distribution for a particle with $H_{posit}(0) = 20$, $V_{posit}(0) = 0$, $VV = 1$, and $HV = 0$. This distribution is similar because in both cases, the particles have the same VV and take about the same number of horizontal steps before reaching the righthand boundary.

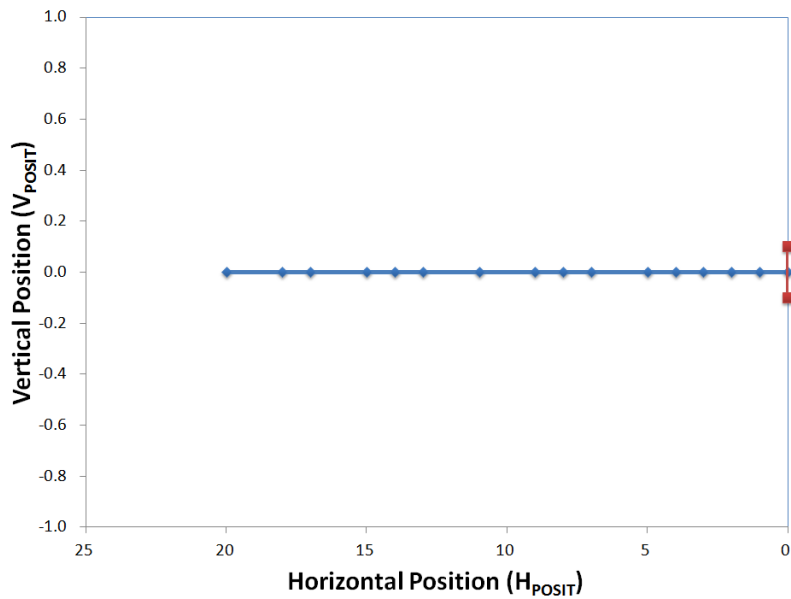
3.3 Calculating the Hit Probability (P_{hit})

The previous section used histograms to provide empirical probability data about final particle states. Each histogram provides information based on one initial particle state. It would be cumbersome to create numerous histograms to determine probability data for the entire state space. Fortunately, analytical methods provide a more efficient way to calculate this probability data. This section focuses on the probability that the particle will strike the target, denoted as P_{hit} . The goal is to map every possible particle state to a P_{hit} value. These probabilities can be calculated using either Markov transition matrices or a backward induction method. This thesis uses a backward induction method that is illustrated in Figure 3.12.

Figure 3.12 contains a small matrix of P_{hit} values. The target is located in rightmost column and




(a) Sample trajectory for a particle with $VV=0$ and $HV=0$. The particle strikes the target after making exactly 20 horizontal steps to the right.

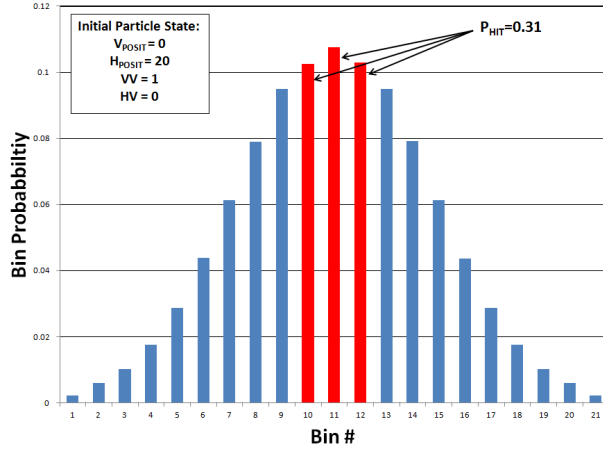


(b) Sample trajectory for a particle with $VV = 0$ and $HV = 1$. The particle strikes the target after making 12 horizontal steps to the right.

Figure 3.10: Two trajectories for particles with different levels of HV . The VV remains constant and is set to zero for both examples. For each of the two cases, the initial particle state is 20 time units away from the target with a V_{posit} of 0. In Figure 3.10a, the particle strikes the target after 20 horizontal steps while the particle in Figure 3.10b strikes after only 15 steps.



(a) Histogram that displays the results from million experiments using the following initial conditions: $V_{posit} = 0$, $H_{posit} = 30$, $VV = 1$, $HV = 1$.



(b) Histogram that displays the results from million experiments using the following initial conditions: $V_{posit} = 0$, $H_{posit} = 20$, $VV = 1$, $HV = 0$.

Figure 3.11: Similar histograms that display the results of two separate experiments with different initial conditions. The similarity between these two histograms shows the relationship between the initial H_{posit} and HV , when all other initial conditions remain identical.

corresponds to the cells where $P_{hit} = 1$. If the particle lands in one of these states, it has struck the target. $P_{hit} = 0$ for the remaining cells in this column. If the particle lands in one of these states, it cannot hit the target. The left column consists of possible particle states (geographic locations). Each row in this column contains a P_{hit} value that is calculated using backward induction. These specific values are calculated for a particle with $VV = 1$ and $HV = 0$. A particle that occupies the shaded box has a $P_{hit} = 2/3$, since this particle can occupy one of

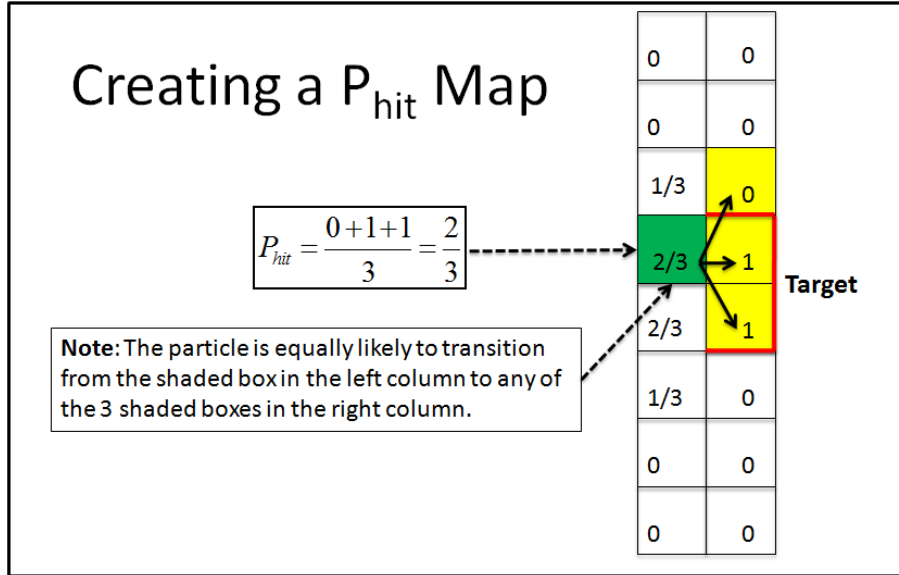


Figure 3.12: A diagram that show how to create a probability map. The probability maps shows the probability that a particle will strike the target given its current state and volatility. For a particle that is located in the green box with $VV = 1$ and $HV = 0$ has a $P_{hit} = 2/3$. This particle can move to one of three future positions, and two of these future states represent the target.

three future states, two of which fall within the target. Thus, its P_{hit} value is calculated as $\frac{0+1+1}{3}$. This backward induction technique can be repeated to add columns to the left side of this matrix. Figure 3.13 shows a larger version of the probability matrix (map) that is explained in Figure 3.12. This figure includes shading to help the reader differentiate between regions with high and low P_{hit} values.

Once a probability matrix is calculated using backward induction, one can refer to it to find P_{hit} values for each state in a particle's trajectory. Figure 3.14 shows a particle trajectory and an associated P_{hit} graph. This figure displays how the P_{hit} value changes as the particle approaches the righthand wall.

Initially, VV is set to zero. After each set of 1,000,000 replications, the average P_{hit} value is stored in a vector and VV is increased by 1.

Figure 3.15 displays the results for an experiment using the following initial conditions: $V_{posit}(0) = 0$, $H_{posit}(0) = 20$, $VV = 0$. This Figure shows the P_{hit} value for different levels of VV . For this experiment, the target size is 10% of the rightmost wall. The initial V_{posit} is aligned with the target. Thus, P_{hit} values are high when VV is low. As VV rises, the distribution of final states spreads outward and thus P_{hit} decreases. This spreading occurs because as VV increases, the particle can travel farther away from its original V_{posit} . Eventually, VV rises to a point where the distribution of final particle states is uniform over the righthand wall and $P_{hit} \approx \text{Target Size}$, which approximately equals 10%. Thus, the above process is repeated until the average P_{hit} is approximately equal to the Target Size.

3.5 Chapter Summary

This chapter described two variations of a simple disaster model. The disaster model simulates the movement of a particle (disaster) through a 2D state space. The particle moves from left to right within the state space and threatens to strike a target. This target represents a concerned entity, and P_{hit} is the probability that the particle will strike this entity. Both model variations include stochastic vertical movement which is characterized by the vertical volatility (VV). The difference between the two variations is that one implements deterministic horizontal movement while the other implements stochastic horizontal movement which is characterized by the horizontal volatility (HV). This chapter included a discussion of how VV and HV affect particle movement and the final distribution of system states. A backward induction algorithm was used to create a matrix (map) of P_{hit} values for every achievable particle state. The chapter ended with a discussion of how these P_{hit} values are affected by changes in VV .

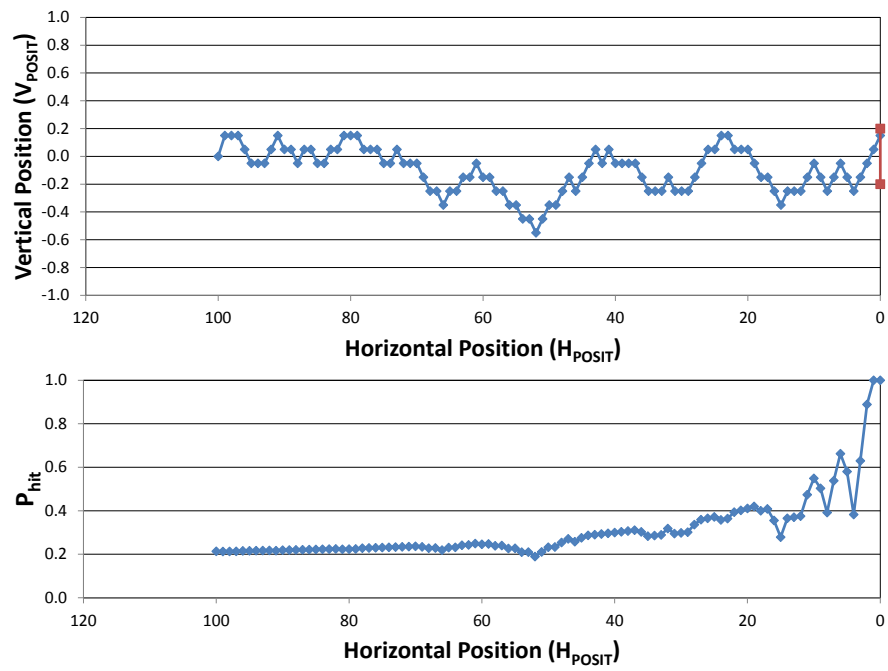


Figure 3.14: A particle trajectory with P_{hit} displayed data for each individual particle state. This figure shows how the P_{hit} value changes as the particle approaches the rightmost wall.

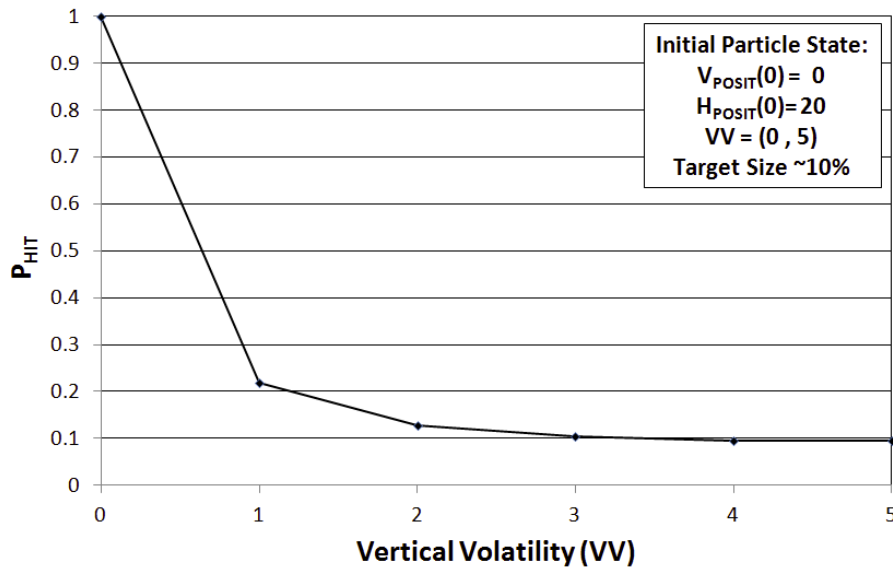


Figure 3.15: This graph shows the relationship between P_{hit} and VV . The initial V_{posit} is aligned with the target, so P_{hit} values are high when VV is low. As VV rises, the distribution of final particle states spreads outward away from the target, so P_{hit} decreases. Eventually, VV rises to a point where the distribution of final particle states is uniform over the righthand wall and $P_{hit} \approx$ Target Size, which approximately equals 10% of the righthand wall.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 4:

DECISION MAKING IN DISASTER SITUATIONS

This chapter focuses on the decision of an entity who risks being hit by an impending but uncertain disaster and must decide at each time t whether to “stay” or to “evacuate.” The decision to evacuate is irreversible, meaning that the entity cannot change this decision once it has been made. The decision to stay is reversible in the sense that the entity can decide to evacuate at a later time. The entity incurs a *cost* that measures the overall consequences of the outcome. This cost depends on whether or not the entity evacuated and whether or not the disaster hit. The entity’s objective is to choose the course of action that minimizes the overall expected cost.

In the context of this disaster model, the entity’s decision must balance several things. The entity observes the state $(H_{posit}(t), V_{posit}(t))$ of the particle along with the probability of hit (P_{hit}) at each point in time. The entity also knows the particle’s HV and VV . Let $C_e(H_{posit}(t))$ represent *the cost to evacuate* (C_e) when the disaster has horizontal position $H_{posit}(t)$. This cost is incurred only if the entity decides to evacuate. In the event the entity never decides to evacuate, the entity incurs a cost that depends on the disaster outcome. No cost is incurred if the entity stays and the disaster misses the target. On the other hand, the entity faces a significant cost if the disaster strikes and the entity chose to stay. Table 4.1 displays a matrix of possible costs.

This study assumes that the entity seeks to minimize the expected cost that will be incurred once the final disaster outcome is known. Therefore, the entity must understand the expected cost associated with either decision for all achievable particle states. It would be quite complicated to estimate the cost associated with both decisions over the entire state space with a single calculation. However, this complicated problem can be broken down into smaller subproblems to simplify these calculations. One process for breaking a large, complicated problem into subproblems is *dynamic programming*. This thesis uses a “top-down” version of dynamic programming where the results of certain calculations are stored and used in subsequent calculations (Wagner, 1995, p. 45). A complete description and tutorial of dynamic programming is given in the textbook *Dynamic Programming and Optimal Control* by Dimitri P. Bertsekas (Bertsekas, 2000). In this chapter, dynamic programming (DP) is used to create an optimal evacuation policy for the entity.

The DP algorithm uses a Cost-To-Go function (J^*) to determine an optimal decision policy. The J^* value for a given system state is the cost that the entity can expect to incur, over all future time, given that the system begins in that state and the entity makes optimal decisions for all time going forward. There are four primitive parts of a DP model that are necessary to calculate J^* . These parts are: the system state, the entity's possible actions, a cost function, and a state transition model. The system state describes both the particle (disaster) state and the entity state. The entity's actions are the decisions that it can make (evacuate or stay). The cost function defines what the entity pays based on both the entity's decisions and the final outcome. Lastly, the state transition model defines the movement of the particle through the state space and the transition of the entity state. As discussed in Chapter 3, HV and VV define the state transition properties of the disaster model.

4.1 Cost-To-Go Function

The Cost-To-Go function (J^*) is used to determine an optimal decision policy for any achievable state in the state space. In Chapter 3, the disaster state is defined as a geographic location (H_{posit}, V_{posit}) . An *evacuation indicator* (E) must be included in the system state in order to implement a DP algorithm. The evacuation indicator describes the entity's state. This indicator is included to ensure that the system state contains all the information that is needed to realize the value of the current state and to predict the next state. When $E = 0$, the entity has not evacuated. $E = 1$ indicates that the entity has made the irreversible decision to evacuate. Let $Y(t)$ be the system state at time t ; i.e., $Y(t) = (H_{posit}(t), V_{posit}(t), E)$.

Assuming that the entity has not evacuated, each time the particle moves, the entity has a choice to evacuate or to stay until the next step. An evacuation cost ($C_e(H_{posit}(t))$) is incurred if the entity decides to evacuate at time t . This cost is incurred when the decision to evacuate is made. No immediate cost is incurred if the entity decides to stay, but the state could become less favorable. The expected cost-to-go if the entity stays is $\mathbb{E}(J^*(Y(t+1)))$ where $\mathbb{E}(x)$ is the expected value of x . Thus, the cost-to-go at time t is equal to $J^*(Y(t)) = \min[C_e(H_{posit}(t)), \mathbb{E}(J^*(Y(t+1)))]$. If the entity decides to stay, the particle proceeds to the next state and this decision making process is repeated.

4.2 Creating Decision Cost Functions

In order to make an optimal decision, the entity must consider the cost that is associated with each option. The evacuation cost function is used to determine the cost of evacuating as a

function of the H_{posit} of the particle at the time the evacuation decision is made.

One typically expects that evacuation costs will rise over time or as the disaster approaches. Rising evacuation costs reflect the fact that, as the evacuation decision is delayed, the entity would likely experience increased road congestion, less robustness to failures and uncertainties, and an increased travel distance to find an emergency shelter with vacancies. Figure 4.1 displays four different cost functions $C_e(H_{posit}(t))$. Note that in these functions, cost can either remain constant or rise according to a concave, linear, or convex function.

Table 4.1: Costs incurred in the evacuation decision problem as determined by the disaster outcome and the entity's decision.

	Evacuated	Stayed
Disaster Strike Entity	$C_e(H_{posit}(t))$	10,000
Disaster Misses Entity	$C_e(H_{posit}(t))$	0

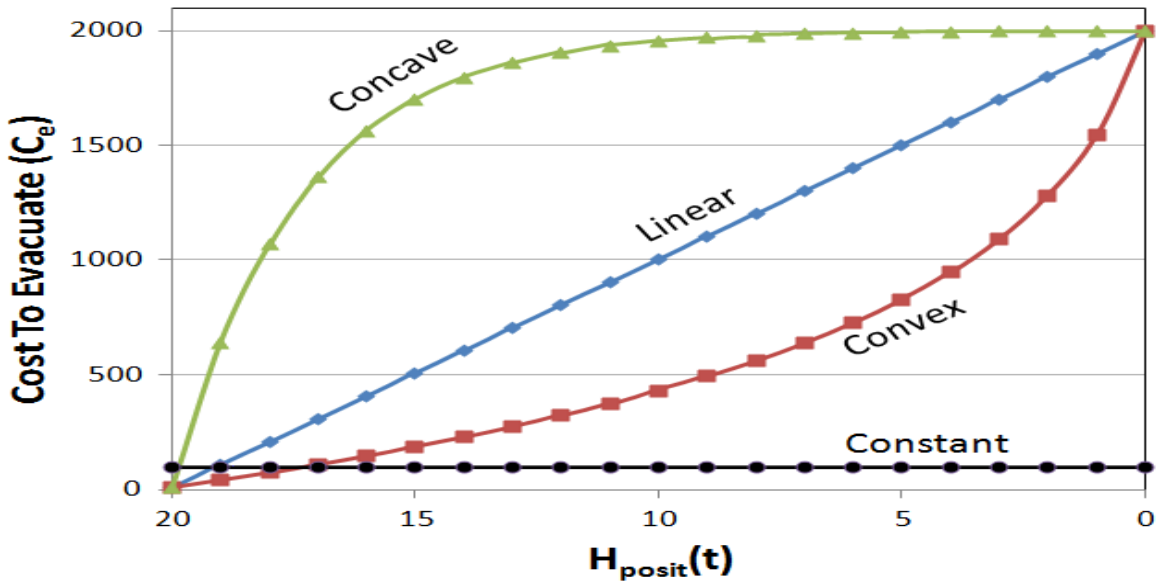


Figure 4.1: Cost-To-Evacuate Functions

If the entity never evacuates, the cost associated with that decision depends only on the final disaster state. Using the values from Table 4.1, this cost is 10,000 units if the particle hits the target and zero units if the particle misses. A backward induction method can determine the

4.3 How Evacuation Costs Affect Decision Policies

The DP model calculates the optimal decision (stay or evacuate) for all achievable particle states as well as the expected cost (J^*) associated with each state. Figure 4.2 shows the optimal policy for a constant C_e . The shaded cells on the right wall show the target location. The remainder of the state space is divided into two distinct regions, shaded and unshaded. The white, or unshaded, region shows where the entity should stay until the next step. The yellow shaded region indicates where the optimal decision is to evacuate.

The evacuation cost function (C_e) affects the size and shape of these two regions. This section discusses how and why the shapes of these regions depend on the both the magnitude and shape of the evacuation functions given in Figure 4.1. For clarity a small (20x20) particle state space is used. This is done so that the reader can see the J^* values associated with each state. However, the shape of these evacuation regions remains roughly the same for an arbitrary state space as long as the ratio $(2VV + 1)/(\#Bins)$ remains constant. Here $\#Bins$ is the number of vertical positions in the discrete state space. Appendix B discusses this relationship in more detail.

4.3.1 Constant Evacuation Cost

We now discuss the policy shown in Figure 4.2 in more detail. For this example, $C_e(H_{posit}(t)) = 100$ for all $H_{posit}(t)$. Notice that the optimal policy is to stay until $H_{posit}(t) = 2$ and to evacuate only if the particle is located near the center of the state space when $H_{posit}(t) \leq 2$.

A closer inspection reveals that **when C_e is constant, the entity can always wait until the last time step to evacuate.** The entity should evacuate at $H_{posit}(2)$ if and only if the entity would evacuate in the next time step anyway. To understand why this is so, note that C_e is small and remains constant throughout the entire state space. Thus, the entity can safely wait until $H_{posit} = 1$, and at that point the entity should evacuate if $10000 * P_{hit} > C_e$. Note that at $H_{posit} = 1$, the value $10000 * P_{hit}$ is the expected cost to stay.

The results for the constant C_e case are neither interesting nor realistic. In a disaster scenario, waiting to evacuate usually carries some sort of penalty. Therefore, it is reasonable to consider a C_e that rises as the particle draws closer to the target.

4.3.2 Increasing Evacuation Cost

Consider the cost function $C_e(H_{posit}(t)) = k*(m*H_{posit}(t)+b)$. Notice that $k*(m*H_{posit}(t)+b)$ describes a linearly increasing cost with slope m and intercept b , which is then scaled by a factor

k . Figure 4.3 shows decision policies for three such cost functions, of which each corresponds to a different value of k . Each decision policy has two distinct evacuation regions. In the left most region, the entity takes advantage of low C_e and evacuates. As H_{posit} increases, this region diminishes as C_e rises. The rightmost evacuation region begins when P_{hit} is large enough to make evacuation the optimal decision despite a high C_e .

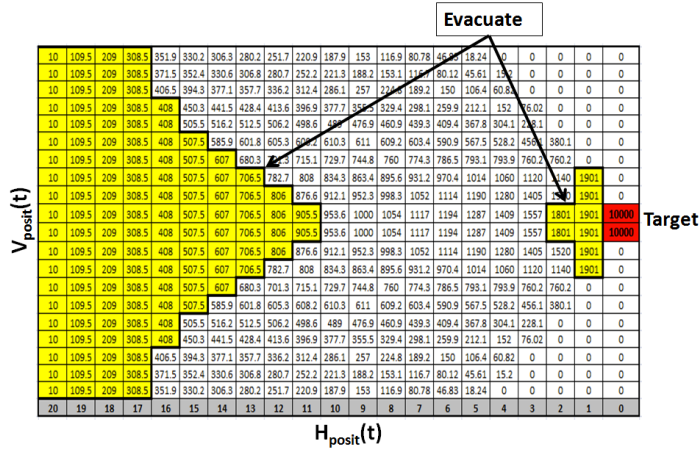
Observe the differences between Figures 4.3a, 4.3b, and 4.3c. As k increases, the evacuation regions shrink. This occurs because both the starting point and the slope rise with k , so the evacuation cost starts at a higher level and increases at a faster rate. This reduces the size of both evacuation regions.

To explore this phenomenon further, consider varying k from 1 to 10. Figure 4.4 shows the evacuation cost function and the corresponding decision policies for each of these values of k . This “heatmap” display allows us to see the shape of the evacuation region(s) and how this shape changes with k . For $k = 1$, Figure 4.4b shows two distinct evacuation regions. As k increases, the evacuation region shrinks. The rightmost region disappears when $k > 2$ and continues to recede to the left until $k = 10$. At $k = 10$, the evacuation region has diminished to the two leftmost columns.

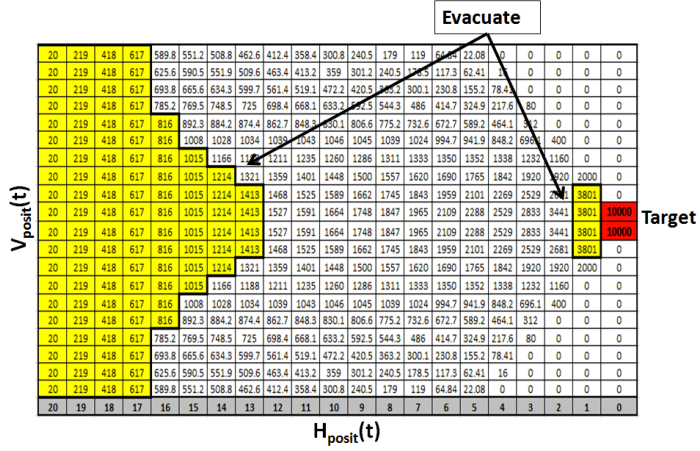
Figure 4.5 illustrates the evacuation region for a convex cost function $C_e(H_{posit}(t)) = k * (b - n * \ln(H_{posit}(t) + 1))$, where b and n are constants. Notice that there is one continuous evacuation region when $k = 1$. For this evacuation function, $C_e(H_{posit}(t))$ remains lower than the corresponding linear case throughout most of the state space. Within this region, $C_e(H_{posit}(t)) \leq \mathbb{E}_{t+1}(J^*(Y(t+1)))$ for all t and $k \leq 2$. This region becomes disconnected when $k \geq 3$. No evacuation region is shown for $k = 2$ because it is identical to the region for $k = 1$.

Next, let us discuss the results for a concave evacuation function where $C_e(H_{posit}(t)) = k * (b - (\exp(H_{posit}(t)/d) - 1))$. Figure 4.6 shows the evacuation functions and regions for this case. The evacuation regions are disconnected and the smaller than the linear case. This is expected since C_e is larger than the linear case across most of the evacuation region. Notice that no regions are labeled for $3 \leq k \leq 9$. This is because the regions for those k values are identical to the region labeled $k = 10$.


Clearly the shape of the evacuation cost function has a impact on the shape and size of the evacuation region. The results for the concave evacuation cost function are quite similar to the case involving a constant evacuation cost in that for most of the state space the optimal



(a) Linear Increasing Evacuation Cost and $k = 1$

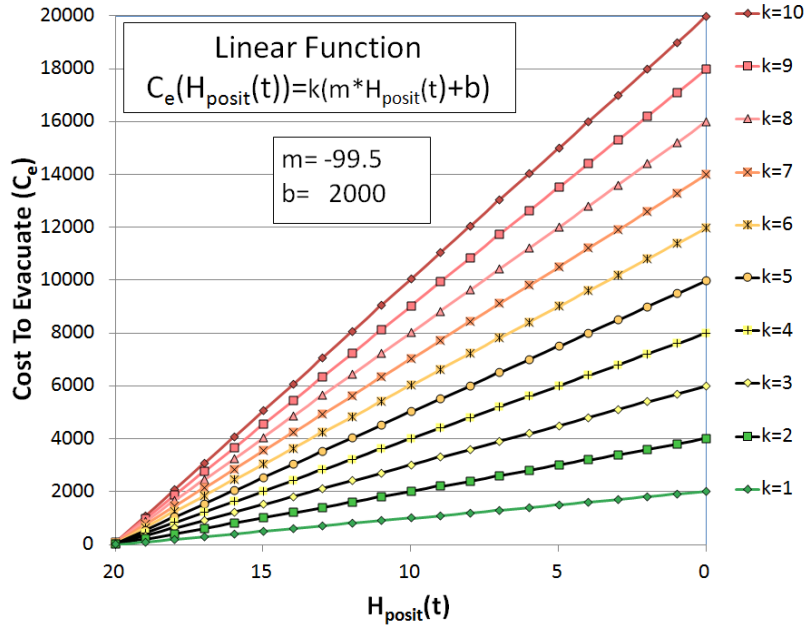


(b) Linear Increasing Evacuation Cost and $k = 2$

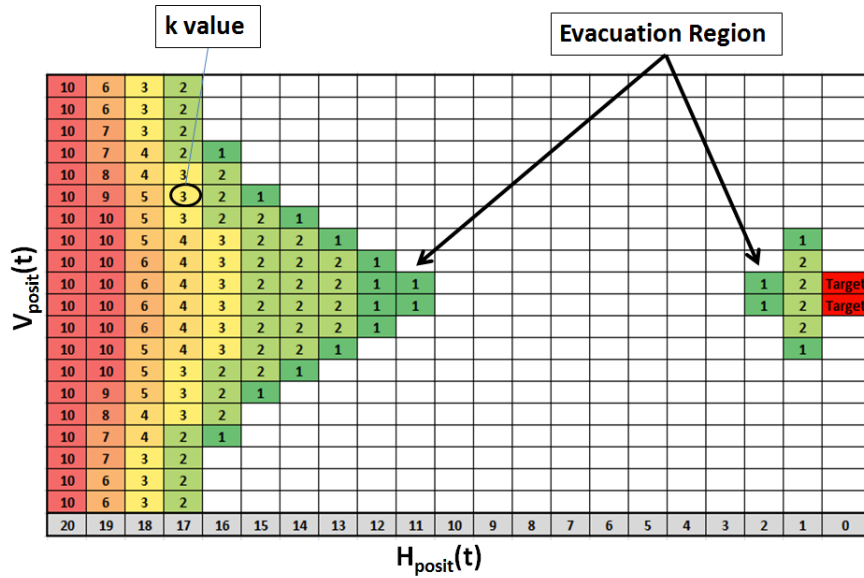


(c) Linear Increasing Evacuation Cost and $k = 3$

Figure 4.3: Decision matrices for linearly increasing evacuation cost and changing k . As evacuation cost increases with parameter k , the evacuation region shrinks.



(a) Linear Increasing Evacuation Functions ($k = 1$ to $k = 10$)



(b) Heatmap (Decision Matrix) for Linear Increasing Evacuation functions ($k = 1$ to $k = 10$)

Figure 4.4: A heatmap that shows ten overlapping evacuation regions for a linearly increasing evacuation cost. Each numbered region in Figure 4.4b corresponds to an evacuation function shown in Figure 4.4a. Notice that the evacuation region is disconnected in that there is a region in the center of the state space where the entity should not choose to evacuate.

policy is to wait until the last step to evacuate. Recall that the linear evacuation cost function gives two distinct evacuation regions for $k \leq 2$. In contrast, the convex cost function results in a connected evacuation region for $k \leq 2$. A comparison of these results suggests that there is some function that lies between the convex and linear cases used in this analysis, for which the evacuation region becomes disconnected for a given k .

To further explore this hypothesis, let us make C_e a convex combination of the linear and convex evacuation functions. We can define C_e as follows: $C_e(H_{posit}(t)) = \lambda * (k * (b - n * \ln(H_{posit}(t) + 1)) + (1 - \lambda) * (k * (m * H_{posit}(20) + b)))$. Figure 4.7a displays C_e for $0 \leq \lambda \leq 1$. When $\lambda = 0$, the function is identical to the linear $k = 1$ case in Figure 4.4a. When $\lambda = 1$, the function is identical to the convex $k = 1$ case in Figure 4.5a.

Figure 4.7b shows how the evacuation region changes with λ . Let us start with $\lambda = 0$ (linear), for which there are two disconnected evacuation regions. As λ increases, the two evacuation regions grow toward each other. The two regions connect to form one continuous region when $\lambda = 0.5$. The region remains connected but continues to grow as λ is increased to 1. When $\lambda = 1$ the evacuation region is identical to the $k = 1$ case in Figure 4.5b.

4.4 Impact of Particle Volatility on Decision Policies

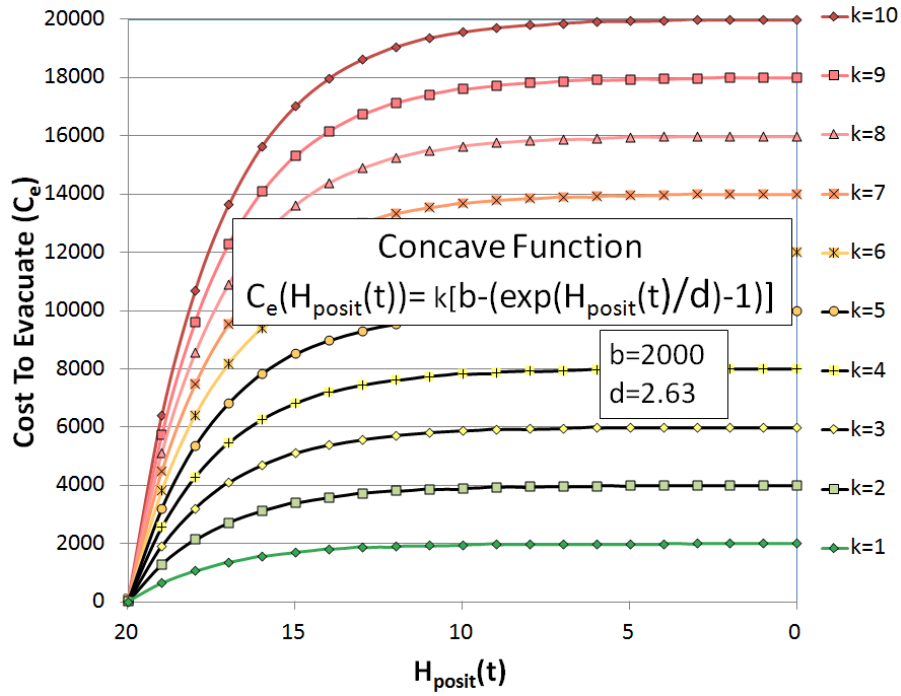
This section discusses how the parameters VV and HV affect the shape and size of the evacuation region. First, let us examine the effects of HV . Recall that horizontal movement is deterministic when $HV = 0$ and stochastic when $HV > 0$. Figure 4.8 shows decision policies (matrices) for three different levels of HV . In all cases, $VV = 2$ and $C_e(H_{posit}(t)) = 100$ for all time t . Comparing the evacuation regions in Figures 4.8a - 4.8c, we can see the evacuation region grows leftward as HV is increased. This occurs because as HV increases, the region of relatively large P_{hit} values expands to the left away from the target. A particle can take larger horizontal steps as HV increases, therefore the particle is more likely to take a large step ($\Delta H_{posit} > 1$) and strike the target.

Now let us consider a more general case in which both VV and HV are varied. Figure 4.9 shows a 3×3 matrix of decision policies. Each column represents a constant HV while each row represents a constant VV . If we compare the decision policies along a given row, we observe changes to the evacuation region similar to those shown in Figure 4.8. We see the effects of VV when we observe the changes along a column Figure 4.9. We notice that the evacuation region grows vertically with increasing VV . This vertical growth occurs because,

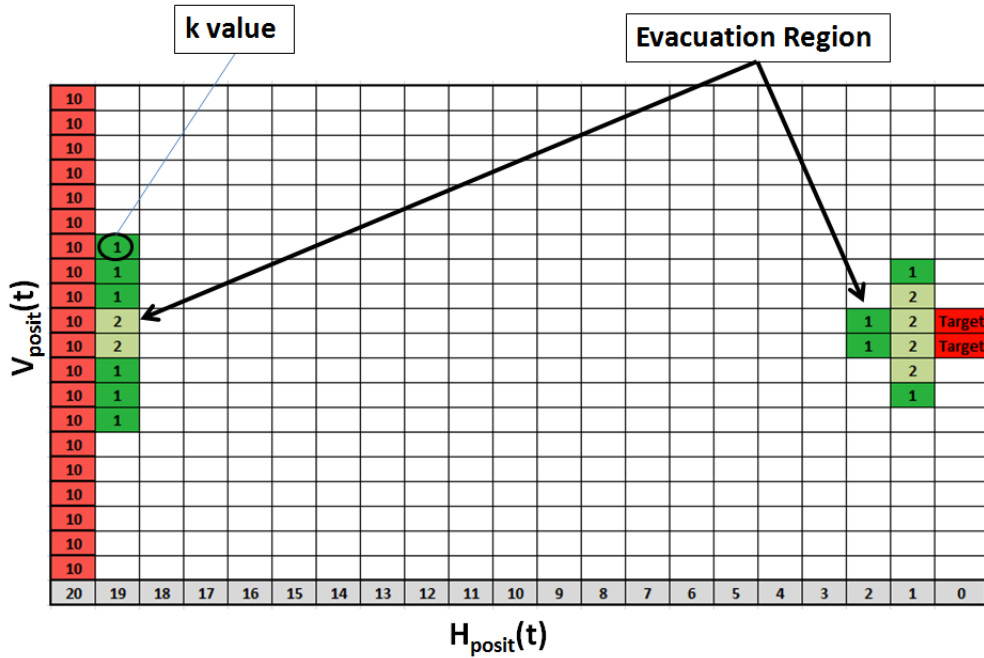
as VV increases, the particle can move farther in the vertical direction. So in general, the evacuation region grows leftward with increasing HV and vertically with increasing VV .

4.5 Chapter Summary

This chapter has considered an entity that is faced with a complex series of decisions regarding an uncertain disaster. A dynamic programming technique was introduced as a way to determine optimal decision policies for the entity. This technique was used to examine optimal decision policies for a number of cost functions and state transition models. Decision matrices were used to display the resulting policies. Each matrix showed the optimal decision for each state along with the expected costs associated with those decisions. These decision policies (matrices) had regions where the entity's optimal decision was to evacuate. We showed how the size and shape of these evacuation regions were affected by the magnitude and shape of the evacuation cost functions as well as the state transition parameters HV and VV .




(a) Concave Increasing Evacuation Functions ($k = 1$ to $k = 10$)

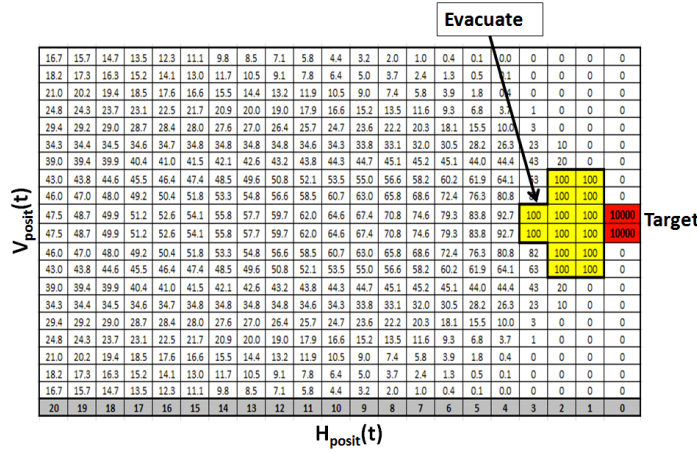


(b) Heatmap (Decision Matrix) for Concave Increasing Evacuation functions ($k = 1$ to $k = 10$)


Figure 4.6: A heatmap that shows ten overlapping evacuation regions for a concave increasing evacuation cost. Each numbered region in Figure 4.6b corresponds to an evacuation function shown in Figure 4.6a. Notice that the evacuation region is disconnected, like the linear case, but has a much smaller region in which the entity should evacuate.



(a) Constant Evacuation Cost ($VV = 2 HV = 0$)



(b) Constant Evacuation Cost ($VV = 2 HV = 1$)



(c) Constant Evacuation Cost ($VV = 2 HV = 5$)

Figure 4.8: Three decision matrices that illustrate the affect of changing horizontal volatility (HV). Increasing HV makes the evacuation region grow. This occurs because as HV grows, the time at which the disaster could strike the entity becomes more uncertain.

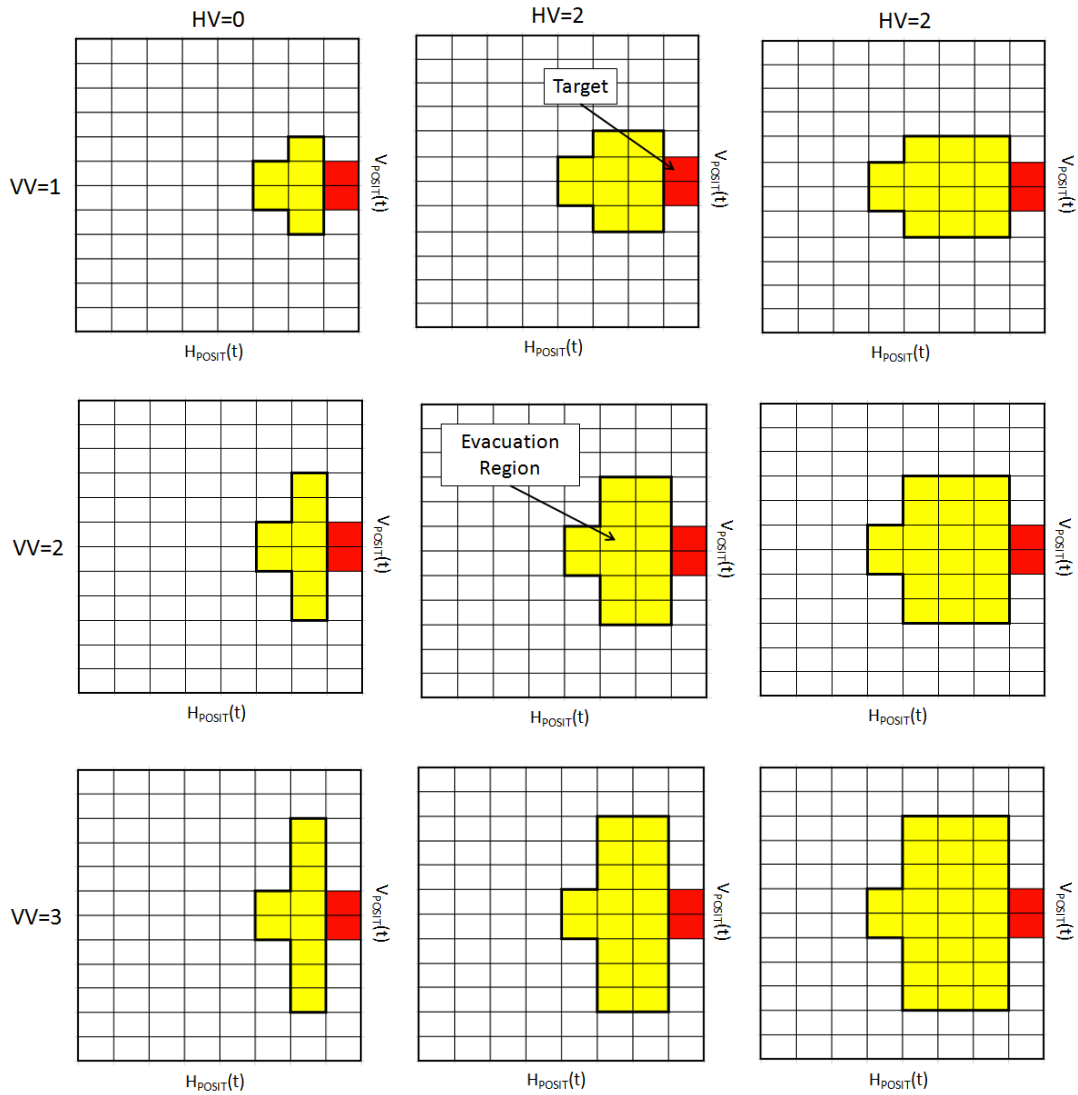


Figure 4.9: Grid of decision matrices that illustrates the effect of volatility (VV and HV) on the evacuation region. HV values are constant for each column and increase from left to right. VV values are constant for each row and increase from top to bottom. As HV increases, the evacuation region grows along the vertical axis. As VV increases, the evacuation region grows along the horizontal axis.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 5:

CONCLUSIONS

We conclude this thesis by summarizing our results and recommending some ideas for future research on the evacuation decision problem.

5.1 Summary

We developed a disaster model that idealizes natural disasters such as hurricanes, wildfires, and floods. This model uses a particle that simulates the movement of a generic disaster through a discrete and bounded 2D state space. We defined the position of the particle as its state and the movement of the particle as its trajectory. We then defined two volatility parameters (VV and HV) that control particle movement. A particle trajectory starts at the left side of the state space and ends with the particle reaching the rightmost boundary. If the ending point of the trajectory coincides with the target, we say that “the particle has hit the target,” otherwise we say “the particle has missed the target.”

We defined the probability of hit (P_{hit}) as the probability that the particle will ultimately strike the target given its state (H_{posit}, V_{posit}). Then we showed a clear relationship between a particle’s initial state, its volatility, and the distribution of final states. In general, the Shannon entropy of the distribution of final states can be increased via either an increase in vertical volatility (VV) or an increase in the $H_{posit}(0)$. More specifically, as VV , $H_{posit}(0)$, or both are increased, the distribution of final states become more uniformly distributed across the rightmost boundary.

We introduced time uncertainty into the disaster model with stochastic horizontal movement. Under this implementation, the particle can take randomly-sized horizontal steps toward the target. We defined the terms *Expected Step-Size* ($\mathbb{E}[\Delta H]$) and an *Expected Time-To-Go* ($\mathbb{E}(TTG)$), and showed that there is a relationship between a particle’s initial state, its HV , and the distribution of its possible final states. For a particle with a fixed initial state, the particle will strike the rightmost wall in fewer steps on average as HV is increased; therefore, the entropy of the distribution lowers.

The target represents an entity who risks being hit by the impending but uncertain disaster and must decide at each time t whether to “stay” or “evacuate.” We describe that the decision to evacuate is irreversible and the decision to stay is reversible. The consequences of the entity’s

decisions are expressed as a cost. The entity can evacuate and incur the cost $C_e(H_{posit}(t))$ at the time of this decision. The entity can delay evacuation until just before the disaster strikes; however, if the entity has not evacuated when the disaster strikes, the entity incurs a cost based on the outcome of the disaster. No cost is incurred if the entity stays and the disaster does not strike the target; however, the entity incurs a significant cost if it stays and the disaster strikes the target.

A dynamic programming (DP) algorithm was used to find optimal decision policies for the entity. We assumed that the entity seeks to minimize the expected cost that will be incurred once the disaster outcome is known. A top-down version of DP was used to calculate an optimal decision (stay or evacuate) for all achievable states as well as the expected cost (J^*) associated with each state. We observed that these decision policies can have as many as two distinct evacuation regions and that the size and shape of these regions is affected by the size and shape of the evacuation cost function ($C_e(H_{posit}(t))$).

With a constant $C_e(H_{posit}(t))$, the entity can wait until the last time step to make its decision. In the last time step, the entity can decide to either stay or evacuate. We showed that both the concave increasing $C_e(H_{posit}(t))$ and linearly increasing $C_e(H_{posit}(t))$ can result in two disconnected regions. For the leftmost region, the entity should evacuate to take advantage of a low C_e . The rightmost evacuation region forms where high P_{hit} and rising C_e make evacuation optimal. A convex $C_e(H_{posit}(t))$ yields a continuous evacuation region for low values of k . We compared the convex $C_e(H_{posit}(t))$ results with the linear $C_e(H_{posit}(t))$ results and used a convex combination of these two cost functions to determine a λ value where the evacuation region becomes disconnected.

Lastly, we described the impact of particle volatility on the size and shape of the evacuation region. For a constant $C_e(H_{posit}(t))$, the evacuation region expands leftward as HV increases and vertically with increasing VV .

5.2 Future Work

Natural disasters have varying levels of intensity, in that one disaster may have a larger potential for destruction than another. For example, a Category 5 hurricane would likely cause more damage to a given region than one with a lower classification. It can be difficult to predict hurricane intensity. Czajkowski (2007) shows that there is a significant level of uncertainty regarding hurricane intensity at 24 hours before landfall. Our disaster model does not account

for varying levels of disaster intensity. A model that incorporates varying intensity levels would provide more insight into the evacuation decision problem. We suggest that this extension should make the cost to stay and incur the disaster a function of disaster intensity, and that disaster intensity should be uncertain. We have argued that the entity should incur a penalty for delaying the decision to evacuate. This penalty was reflected in an increasing evacuation cost. One could also imagine that vacancies at emergency shelters might diminish as the disaster draws closer to a region. In this thesis, the entity maintains the ability to evacuate until the disaster strikes. We recommend a model extension in which the entity stands to lose the ability to evacuate due to a scarcity of shelter spaces. We believe this extension will provide additional insight into the evacuation decision by introducing another source of realism to the decision process.

5.3 Final Thoughts

Natural disasters will continue to threaten and affect people throughout the world. Unfortunately, there will always be limits on how well we can predict when and where these disasters will occur. We must continue to improve our decision policies regarding how and when to prepare for and evacuate from these disasters. This thesis has laid a foundation for future work by developing a disaster model and providing initial results on optimal decision policies and how these policies are impacted by uncertainties about the disaster process. We believe that future extensions of this model will provide a more detailed understanding of the decisions surrounding disaster evacuations.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

- Bertsekas, Dimitri P. 2000. *Dynamic Programming and Optimal Control*. 2nd ed. Vol. 1. Belmont, MA: Athena Scientific.
- Blake, Eric S., Rappaport, Edward N., & Landsea, Christopher W. 2007. *The Deadliest, Costliest, and Most Intense United States Tropical Cyclones from 1851 to 2006 (And Other Frequently Requested Hurricane Facts)*. Tech. rept. National Weather Service National Hurricane Center Miami, Florida.
- Czajkowski, Jeffery. 2007. *Is It Time To Go Yet? Dynamically Modeling Hurricane Evacuation Decisions*. Tech. rept. International Hurricane Research Center, Florida International University.
- Devore, Jay L. 2009. *Probability and Statistics for Engineering and the Sciences*. Belmont, CA: Brooks/Cole, Cengage Learning.
- Frieser, B. I. 2004. *Probabilistic Evacuation Decision Model For River Floods In The Netherlands*. M.Phil. thesis, TU Delft.
- Henry, Claude. 1974. Investment Decisions Under Uncertainty: The Irreversibility Effect. *The American Economic Review*, 64, 1006–1012.
- Klir, George J. 2006. *Uncertainty and Information Foundations of Generalized Information Theory*. Hoboken, NJ: John Wiley and Sons.
- Regnier, Eva. 2008. Public Evacuation Decisions and Hurricane Track Uncertainty. *Management Science*, 54, 16–28.
- Regnier, Eva, & Harr, Patrick. 2006. A Dynamic Decision Model Applied to Hurricane Landfall. *Weather and Forecasting*, 21, 764–780.
- Stewart, James. 2008. *Calculus Early Transcendentals*. Belmont, CA: Brooks/Cole, Cengage Learning.
- Tamura, Hiroyuki, Yamamoto, Kouji, Tomiyama, Shinji, & Hatano, Itsuo. 2000. Modeling and analysis of decision making problem for mitigating natural disaster risks. *European Journal of Operational Research*, 12, 461–468.
- Wagner, David B. 1995. Dynamic Programming. *The Mathematica Journal*, 5, 42–51.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX A: RELATIONSHIP BETWEEN VERTICAL VOLATILITY AND INITIAL HORIZONTAL POSITION

This appendix describes the relationship between a particle's initial state, its vertical volatility VV , and the distribution of its possible final states. Consider a particle that is placed at an initial position given by $H_{posit}(0)$ and $V_{posit}(0)$. Assume that the particle's horizontal volatility (HV) is equal to zero; thus, the particle will take exactly $H_{posit}(0)$ steps before striking the right wall. The particle's final vertical position depends on both the initial vertical position ($V_{posit}(0)$) and the sum of $H_{posit}(0)$ random vertical movements. As described in Chapter 3, these random vertical movements are defined by a symmetric distribution which is characterized by VV . For simplicity, this appendix assumes that the particle's initial vertical position is given by $V_{posit}(0) = 0$; thus, its final vertical position is given by the sum of its random vertical movements. This appendix also assumes that the reflecting walls are distant enough that the particle never encounters them.

By definition, each random vertical movement is an independent and identically distributed (i.i.d.) random variable; thus, the particle's final position is given by a sum of i.i.d. random variables. The distribution of the sum of i.i.d. random variables approximates a Gaussian distribution, which is fully characterized by its mean and variance (Devore, 2009). The mean of this distribution is zero due to our assumption that $V_{posit}(0) = 0$ and the fact that the distribution controlling vertical movement is symmetric about zero. The variance of this distribution depends on both VV and $H_{posit}(0)$. Intuitively, one would expect that low values of VV and $H_{posit}(0)$ should result in low variance, and that variance could be increased by increasing either VV or $H_{posit}(0)$. The following analysis describes the tradeoff between VV and $H_{posit}(0)$ in terms of the variance of the distribution of possible final V_{posit} values.

Consider a particle with $VV = n$, and recall that ΔV is equally likely to take on any value from the set $\{-n, -(n-1), \dots, -1, 0, 1, \dots, n-1, n\}$, i.e., any integer between $-n$ and n . Therefore, the variance of ΔV is

$$\begin{aligned}
Var[\Delta V] &= E[\Delta V^2] - (E[\Delta V])^2 \\
&= E[\Delta V^2] \\
&= \frac{2(1^2 + 2^2 + \dots + n^2)}{2n + 1}.
\end{aligned} \tag{A.1}$$

It has been shown (Stewart, 2008) that

$$\sum_{i=1}^n i^2 = \frac{n(n+1)(2n+1)}{6}.$$

Substituting this expression into Equation A.1., we obtain

$$\begin{aligned}
Var[\Delta V] &= \frac{n(n+1)(2n+1)}{3(2n+1)} \\
&= \frac{n(n+1)}{3}.
\end{aligned}$$

After k steps, the variance is

$$Var[V_{posit}(k)] = \frac{kn(n+1)}{3}.$$

Consider a particle with $VV = 1$ that takes k steps. Using the above formula, we see that the variance of its $V_{posit}(k)$ is equal to $\frac{2k}{3}$. Now consider the question; how many steps should a particle with volatility $VV = n$ take in order to obtain the same variance in its $V_{posit}(k)$?

To answer this question, we must find k' such that

$$\frac{k'n(n+1)}{3} = \frac{2k}{3}.$$

Solving for k' , we obtain an expression for the number of steps that must be taken by a particle with vertical volatility $VV = n$ in order to achieve roughly the same uncertainty in its final vertical position as particle with vertical volatility $VV = 1$ that takes k steps:


$$k' = \frac{2k}{n(n+1)} = \frac{k}{\frac{n(n+1)}{2}} = \frac{k}{1+2+\dots+n}.$$

Thus, for values of n up to 5, we have the following values for k' :


Table A.1: Number of Steps (k') with $VV = n$

n	k'
1	$\frac{k}{1}$
2	$\frac{k}{3}$
3	$\frac{k}{6}$
4	$\frac{k}{10}$
5	$\frac{k}{15}$

This analytic result can easily be confirmed experimentally. Figure A.1a shows distributions of possible final V_{posit} values for five different initial particle states. One distribution is for a particle with an initial $H_{posit} = 30$ and a $VV = 1$. The remaining distributions represent possible final states of particles with vertical volatilities of 2, 3, 4, and 5. The initial positions of these particles are chosen according to Table A.1; thus, they are 10, 5, 3, and 2, respectively. Note the similarity in the distributions; after only 2 steps, the particle with $VV = 5$ achieves nearly the same distribution of possible vertical positions as a particle with $VV = 1$ that takes 30 steps. The similarity in distributions is even more striking when the particles begin with greater H_{posit} values; see Figures A.1-A.2.




(a) Particle with $VV = 1$ takes 30 steps.




(b) Particle with $VV = 1$ takes 90 steps.

Figure A.1: Final V_{posit} distributions for five initial particle states, where each particle's initial H_{posit} is chosen so that its distribution of final states matches that of the particle with $VV = 1$.



(a) Particle with $VV = 1$ takes 150 steps.



(b) Particle with $VV = 1$ takes 270 steps.


Figure A.2: Final V_{posit} distributions for five initial particle states, where each particle's initial H_{posit} is chosen so that its distribution of final states matches that of the particle with $VV = 1$.

THIS PAGE INTENTIONALLY LEFT BLANK


APPENDIX B: SENSITIVITY OF OPTIMAL EVACUATION POLICIES TO STATE SPACE DISCRETIZATION

In this thesis, decision matrices are used to show the optimal decision (stay or evacuate) for all achievable particle states as well as the expected cost (J^*) associated with each state. For clarity, the figures in this thesis use a small (20x20) state space so that the reader can see the J^* values associated with each state. However, as this appendix demonstrates, the shapes of these evacuation regions remain roughly the same for an arbitrary state space as long as the ratio $\frac{2VV+1}{\#Bins}$ remains constant, where $\#Bins$ is the number of vertical positions in the discrete state space.


To illustrate this property, Figure B.1 displays three decision matrices for a linearly increasing evacuation cost. Each matrix contains a different number of vertical bins ($\#Bins$) and value for VV . The parameters for each scenario are chosen such that the ratio $\frac{2VV+1}{\#Bins} = 0.25$. Note that the size and shape of the evacuation regions remain roughly the same in each figure. Figure B.2 shows similar results for a convex increasing evacuation cost.



(a) # Vertical Bins = 20 and $VV = 2$




(b) # Vertical Bins = 36 and $VV = 4$




(c) # Vertical Bins = 76 and $VV = 9$

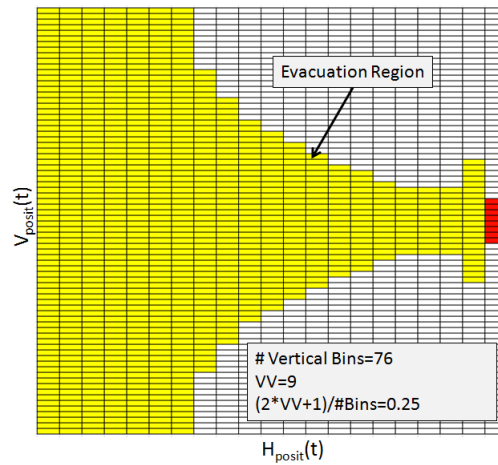
Figure B.1: Decision matrices for a linearly increasing evacuation cost with various state space discretizations. The sizes and shapes of the evacuation regions in these figures are similar because the ratio $\frac{2VV+1}{\#Bins} = 0.25$ remains constant.



(a) # Vertical Bins = 20 and $VV = 2$



(b) # Vertical Bins = 36 and $VV = 4$



(c) # Vertical Bins = 76 and $VV = 9$

Figure B.2: Decision matrices for a convex increasing evacuation cost with various state space discretizations. The sizes and shapes of the evacuation regions in these three figures are similar the ratio $\frac{2VV+1}{\#Bins} = 0.25$ remains constant.

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. Dr. Emily Craparo
Naval Postgraduate School Operations Research Department
Monterey, California
4. Dr. David Alderson
Naval Postgraduate School Operations Research Department
Monterey, California
5. Dr. Jean Carlson
University of California Physics Department
Santa Barbara, California
6. Dr. Thomas Otani
Naval Postgraduate School Computer Science Department
Monterey, California