Methods Development for Optimal Defense against Adaptive Adversaries – Quantification of Uncertain Preferences and Development of Computational Approaches

by

Chen Wang

A dissertation submitted in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

(Industrial Engineering)

at the

UNIVERSITY OF WISCONSIN–MADISON

2013

Date of final oral examination: 05/14/2013

The dissertation is approved by the following members of the Final Oral Committee:

Vicki M. Bier, Professor, Industrial and Systems Engineering

Michael C. Ferris, Professor, Computer Science

Leyuan Shi, Professor, Industrial and Systems Engineering

Michael A. Newton, Professor, Statistics

Peter Z.G. Qian, Associate Professor, Statistics

To My Parents: Ying Chen and Jianjun Wang.

Acknowledgments

First of all, I would like to express my greatest appreciation and gratitude to my advisor Professor Vicki Bier. Over the past six years I have received unwavering support and encouragement from her. She has been a mentor, colleague, and friend. Her guidance has fundamentally improved the quality of my graduate study, and made this a thoughtful and rewarding journey. I would also like to extend my sincere appreciation to my committee members Professors Michael Ferris, Leyuan Shi, Michael Newton, and Peter Qian, for sharing their astounding expertise with me through the years.

My appreciation also goes to Professors Stephen Robinson and Jeff Linderoth, who have provided invaluable support during my graduate coursework study and the process of seeking a faculty job. There are also many others in the Department of Industrial and Systems Engineering and the Center for Human Performance and Risk Analysis that have been helpful to me in many different ways. It is difficult to list all of these individuals. Instead, I would like to say thank you to the department and the center as a whole for making my staying here an enjoyable and memorable experience. In addition, I would like to thank Dr. Yuri Ermoliev at the International Institute for Applied Systems Analysis and Dr. James

Peerenboom at Argonne National Laboratory, for the wonderful summers I spent with them and their colleagues.

Of course, I would also like to thank my dearest parents, Ying Chen and Jianjun Wang. Without their love and encouragement, I would never have been able to grow into an enthusiastic and positive person. Therefore, I would like to dedicate this dissertation to them. I also feel very fortunate to have my husband, Yang Sun, accompany and support me during this precious time. The past six years spent in Madison will be forever remembered and truly treasured in my heart. Looking forward, I would like to quote my favorite Chinese poem: "The journey is long. I will search up and down." (Qu Yuan, 343-278 BC)

This research was supported by the United States Department of Homeland Security through the Center for Risk and Economic Analysis of Terrorism Events (CREATE) under grant numbers 2007-ST-061-000001 and 2010-ST-061-RE0001. However, any opinions, findings, and conclusions or recommendations in this dissertation are those of the author and do not necessarily reflect views of the United States Department of Homeland Security. I thank this financial support for allowing me the time to concentrate on my dissertation.

Contents

List of Tables

List of Figures

Abstract

This dissertation extends game-theoretic models for homeland security in three different ways, all motivated by the desire to make game theory ready for use in real-world security decisions. First, we introduce a simple elicitation process where subject-matter experts can provide only ordinal judgments of the attractiveness of potential targets, and the adversary preferences among targets are assumed to involve multiple attributes such as fatalities, property loss, and symbolic value. Probabilistic inversion or Bayesian density estimation is then used to mathematically derive probability distributions representing both defender uncertainty about adversary weights on the various attributes, and also defender ignorance about unobserved attributes that may be important to the adversary, but have not yet been identified by the defender.

Although the motivation for this work was the need for methods of estimating adversary preferences from ordinal data, this area of our work also makes methodological contributions to the field of expert elicitation in general, especially through the use of unobserved attributes to ensure the existence of feasible solutions in probabilistic inversion, and by elucidating the relationship between probabilistic inversion and Bayesian density estimation when applied to preference rankings. Moreover, our proposed elicitation process reduces the burden of

time-consuming orientation and training associated with traditional methods of attribute weight elicitation, and explicitly captures the existing uncertainty and disagreement among experts, rather than attempting to achieve consensus by eliminating them. We present both hypothetical and real-world case studies on elicitation of adversary preferences to illustrate the applicability of the proposed elicitation process.

Next, this dissertation attempts to fill a gap in the literature of game-theoretic models for homeland security by explicitly considering adversary capabilities in addition to just intent, since intelligence experts generally believe that adversary capabilities are at least as important as intent. In particular, we propose a Bayesian Stackelberg game capable of analyzing the joint effects of both attacker intent and capabilities on optimal defensive strategies. The novel feature of our model is the use of contest success functions from economics to explicitly capture the extent to which the success of an attack is attributable to the adversary's capability (as well as the level of defensive investment), rather than pure luck. Moreover, our model allows the effectiveness of attacker capabilities to differ across targets (e.g., civilian versus military targets) and attack modes (e.g., attacks using improvised explosive devices versus nuclear weapons).

The results of this phase of our work suggest that uncertainty about adversary capabilities can play a critical role in making defensive decisions. Our model is one of the first studies to include such uncertainty, and thus paves the way for homeland-security decision makers to base defensive strategies on a more realistic and comprehensive set of adversary characteristics. Results also show that precise assessment of attacker intent will generally not be necessary if the defender is highly uncertain about attacker capabilities. Therefore,

our model can provide useful guidance on how to prioritize intelligence collection about adversary capabilities versus intent.

Finally, we identify and evaluate rigorous and efficient computational tools to solve for equilibrium (optimal) defensive strategies in problems of realistic size and complexity. The Bayesian Stackelberg game with defender uncertainty about adversary characteristics is formulated as a two-stage stochastic programming problem with binary recourse, and solved using solution approaches based on sample-average approximation. In particular, we present conditions under which the general convergence properties of the sample-average approximation method can apply to our case, and investigate two categories of state-of-the-art optimization algorithms (one based on mixed-integer nonlinear programming, and the other based on derivative-free global optimization techniques) for solving games complex enough to be suitable for use in real-world decisions.

Moreover, the defender-adversary game of incomplete information we address is an example of a broader class of "stochastic selection problems." Thus, the techniques that we have investigated for this case are also applicable to other stochastic-selection problems (e.g., to obtain optimal budget allocations for marketing initiatives, in the face of uncertainty about the importance of different product features to customers).

Chapter 1

Introduction

The development of game-theoretic models for homeland security has advanced quite rapidly, and the related theories are by now relatively mature. Our game-theoretic model of defensive resource allocation is illustrated by the influence diagram in Figure 1.1. In particular, the defender moves first to allocate her limited defensive resources among a collection of potential targets or attack strategies. The attacker then observes that allocation, and makes decisions about whether to launch an attack, and if so, which target(s) to attack (or which attack strategies to use). Two main factors are used to describe the attacker's characteristics: intent; and capability. In reality, of course, the attacker's goals and motivations and levels of capabilities may not be fully known to the defender (as shown by the dashed arrows in Figure 1.1).

There are challenges in applying this type of model to real-world problems in a manageable way. For example, there is a need to develop methods for the quantification of uncertain attacker intent using subject-matter experts, due to the lack of empirical data to

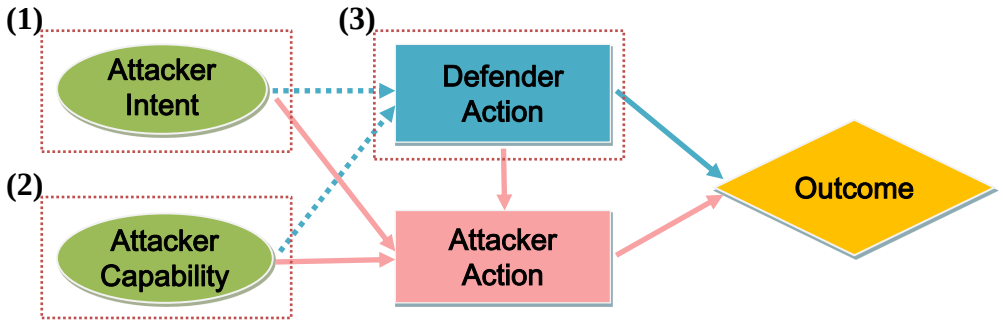

Figure 1.1: Influence Diagram Showing the Defender and Attacker Decisions

support purely statistical analysis. Moreover, there is a significant gap in the literature of game-theoretic models in homeland security to account in a comprehensive manner for attacker capabilities in addition to just intent, even though intelligence experts generally believe that attacker capabilities are at least as important as intent in defensive decisions. Finally, in order to use the proposed defender-attacker game as a strategic tool for problems of realistic size and complexity, sufficiently powerful computational tools are also required. Therefore, the goal of this dissertation is to meet these challenges by methods development in the three areas highlighted in Figure 1.1.

We first explore (1) rigorous and practical approaches for quantifying **attacker intent** using ordinal judgments from domain experts. We then extend the existing game-theoretic framework for modeling adaptive adversaries to account for (2) **attacker capabilities** rather than just intent. Finally, we investigate (3) sufficiently powerful computational tools for identifying optimal (equilibrium) **defender actions** in games complex enough to be suitable for use in real-world defensive decisions.

In this work, we adopt the Bayesian Stackelberg game developed by Bier et al. (2007) as the basic model. Their model assumes that the defender first decides on how to allocate her defensive resources; an attacker then observes the defensive allocation, and chooses attack actions according to his target valuations, which are in general not fully known to the defender. Following Wang and Bier (2011) and Bier et al. (2012), we represent the attacker's target preferences by a multiattribute utility function, and assume that the defender is uncertain about both the attacker weights on the various attributes and any unobserved attribute that may be important to the attacker, but have not been identified by the defender. The defender's choice is how much to spend on defense of each target, while the attacker makes a discrete choice of which target to attack (if any).

The first challenge in using this model for real-world homeland-security decisions is how to quantify the defender's uncertainty about attacker preferences. Eliciting information on the attacker's utility functions from individuals knowledgeable about terrorism is difficult, since many intelligence experts are not quantitatively trained, and may not be familiar with multiattribute utility theory. Moreover, many intelligence experts place a great deal of importance on achieving consensus, which is not conducive to accurately characterizing the level of uncertainty that may exist. As a result, there is a need for indirect elicitation methodologies that can "translate" information of the type that intelligence experts feel comfortable providing into risk-assessment "language."

Thus, in this dissertation, we show how to use ordinal judgments from subject-matter experts (e.g., rankings of the attractiveness of a subset of potential attacker targets or strategies) as a basis to infer probability distributions over the uncertain parameters in the attacker's multiattribute utility function. In particular, we discuss two existing methods,

probabilistic inversion and Bayesian density estimation, recently adapted by Wang and Bier (2013a) to apply to preference rankings.

A novel feature of our treatment of defender uncertainty is the use of unobserved attributes – which might be important to the adversary but are not known to the defender. In addition to providing for realistic levels of uncertainty and defensive "hedging," the amount of weight assigned to the unobserved attributes in a given analysis can also provide indication of whether the attributes that have already been identified are adequate to model adversary preferences. Our work also makes methodological contributions to the broad field of expert elicitation, by using the unobserved attribute to ensure the existence of feasible solutions in probabilistic inversion, and by elucidating relationships between probabilistic inversion and Bayesian density estimation when applied to preference rankings.

In addition, we have made several efforts to advance the applicability of the proposed elicitation process. For example, we explore the use of negative attribute weights, to automatically account for the possibility that experts may disagree on whether a larger value of an attribute corresponds to higher or lower attacker utility. By simulation-based sensitivity analysis, we also show that the breadth of the elicited probability distributions is generally a good indicator of the level of uncertainty and disagreement among experts, and probabilistic inversion is more likely to yield multimodal distributions than Bayesian density estimation in the face of expert disagreement. Moreover, we show that even partial rank orderings of the attacker targets or strategies can yield reliable results, which would further ease the elicitation burden. For illustrative purposes, we also provide a hypothetical case study and a realistic application (CREATE 2011) on elicitation of adversary preferences among

major U.S. urban areas and among different attack scenarios (e.g., chemical, biological, radiological, and nuclear attacks), using ordinal data of preference rankings.

Second, Bier et al. (2007), Wang and Bier (2011), and Bier et al. (2012) all focus on defender uncertainty about attacker goals and motivations without accounting for attacker capabilities. To fill this significant gap, this dissertation presents a comprehensive game-theoretic framework (following Wang and Bier 2013b) to explicitly capture the effects of attacker capabilities in addition to just intent. The novel feature of our model is the use of contest success functions from economics. A "decisiveness" parameter is introduced to reflect the extent to which the success of an attack can be attributed to the relative effort of attack over defense rather than pure luck. For example, an attack on a soft civilian facility using improvised explosive devices (IED) is assumed to be associated with a low level of decisiveness, since even an attacker with low capability may get lucky and succeed with such an attack. By contrast, only highly capable attackers who possess abundant resources are likely to be able to conduct a successful attack on a hardened military target with high decisiveness, while diligent efforts by the defender would be required to foil such an attack.

Results of our model suggest that uncertainty about attacker capabilities can play a critical role in making defensive decisions, especially when attack targets or strategies have widely differing levels of decisiveness. In fact, precise assessment of attacker intent might not even be necessary if the defender is highly uncertain about the attacker's level of capability. Therefore, it may be reasonable to prioritize intelligence collection about attacker capabilities over intelligence collection about intent.

We also extend our model to allow for multiple types of attacker capabilities such as capital, labor, technological sophistication, and flexibility of movement, by the use of a

generalized contest success function. This representation separates the task of estimating each individual type of attacker capability from the task of identifying the relative effects of the various types of capabilities, and thus provides a realistic way to assess the effects of multi-dimensional attacker capabilities. To validate the usability of the proposed game that accounts for both attacker intent and capabilities, we present an illustrative case study on protection against terrorist groups with different levels of capabilities (e.g., Northern Ireland's Real Irish Republican Army, the Palestinian group Hamas, and al Qaeda), using data from a recent project (CREATE 2011) as well as open sources such as Falkenrath et al. (1998) and Cragin and Daly (2004).

A third challenge of applying game-theoretic models to real-world homeland-security decisions is to develop sufficiently powerful computational tools to obtain optimal defensive resource allocations for games of realistic size and complexity. The sequential game with defender uncertainty about attacker characteristics can be computationally challenging to solve for large numbers of potential attack targets or strategies, and large numbers of target attributes (resulting in high-dimensional probability distributions over attacker attribute weights). It is therefore necessary to develop rigorous and efficient algorithms for determining optimal defensive resource allocations to an acceptable level of accuracy, in order to facilitate strategic planning decisions.

In response to this need, we investigate a broad range of state-of-the-art optimization algorithms that are applicable to our proposed game. In this area, we first generalize the defender-attacker game as a "stochastic-selection" problem where the decision maker's objective is to allocate a limited budget to a collection of entities against stochastic selections by an opponent. We then model it as a two-stage stochastic programming problem with

binary recourse, and solve it by solution approaches based on sample-average approximation (SAA). In particular, we present conditions under which the general convergence properties of the SAA method can apply to our case, and discuss two categories of computational algorithms (one based on mixed-integer nonlinear programming, and the other based on derivative-free optimization). Moreover, the techniques that we have investigated for this case are also applicable to other stochastic-selection problems (e.g., to obtain optimal budget allocations for marketing initiatives, in the face of uncertainty about the importance of different product features to customers).

By addressing the above three challenges, we believe that this work will provide homeland-security decision makers with ways of effectively and efficiently quantifying and solving games of realistic size and complexity, against uncertain and adaptive adversaries.

The next chapter provides a literature review. In Chapter 3, we present the basic game accounting for defender uncertainty about only attacker intent, as represented by a multiattribute utility function involving unobserved attributes. To quantify the uncertain attacker intent from simple ordinal preference rankings, Chapter 4 explores two mathematical approaches, probabilistic inversion and Bayesian density estimation. Next, Chapter 5 extends the basic game to capture the effects of attacker capabilities on optimal defensive decisions, and allows the defender to be uncertain about both attacker intent and capabilities. The defender-attacker game is then generalized as a stochastic-selection problem in Chapter 6, followed by an investigation of state-of-the-art optimization approaches for solving such types of problems. Finally, Chapter 7 concludes the dissertation and discusses directions for future work.

Chapter 2

Literature Review

In this chapter, we first review the development of game-theoretic models against adaptive adversaries, especially those accounting for defender uncertainty about adversary intent (Section 2.1). Then, we investigate existing models to represent and elicit uncertain adversary intent, some of which are from the general expert elicitation literature and beyond the context of homeland security (Section 2.2). Finally, we discuss computational tools that have been adopted for solving defender-adversary games of incomplete information (Section 2.3).

2.1 Game-Theoretic Models against Adaptive Adversaries

2.1.1 Adaptive Adversaries and Security Externalities

Both the Bush and Obama administrations agreed that homeland security decisions should be risk-based or at least risk-informed, since we cannot protect everything from every threat

(Bellavita 2009). The well-known risk assessment framework adopted by the Department of Homeland Security (DHS) decomposes risk into three multiplicative components; i.e.,

$$\text{Risk} = \text{Threat (T)} \times \text{Vulnerability (V)} \times \text{Consequence (C)}$$

According to the DHS Risk Lexicon (DHS 2008), "threat" is generally the likelihood of an attack being attempted by an adversary, while "vulnerability" can be interpreted as the likelihood that an attack is successful given that it is attempted. Finally, "consequence" is generally the expected magnitude of damage from a successful attack, accounting for both immediate and secondary effects.

Of these three, threat has been the most difficult to quantify (Willis et al. 2005; Parnell et al. 2008; Cox 2008; Brown and Cox 2011). Probabilistic risk assessment (PRA) has been widely used as a powerful tool to address risks in both natural and engineered systems, and has also been usefully applied to estimate terrorist threats (e.g., Ezell et al. 2010). However, the static risk estimates resulting from PRA have been criticized for not reflecting the dynamics of defense against intelligent, adaptive adversaries. For example, the National Research Council (2010) recommends that the U.S. Department of Homeland Security should use models that "explicitly recognize terrorists as intelligent adversaries who observe U.S. defensive preparations and seek to maximize achievement of their own objectives." Likewise, Parnell et al. (2008) argue that terrorists are "goal-oriented, resourceful adversaries who will, given the constraints they perceive, select the best agent and target to achieve their objectives."

In particular, PRA does not explicitly account for the negative externalities associated with defensive investments; namely, the fact that protecting one target may shift risks to other targets. Powell (2007) notes that unlike natural disasters, strategic adversaries "try to

strike where the defense is weak;" investing in one target can therefore make other targets more likely to be attacked. In fact, as noted by Bier et al. (2007) and Cox (2008, 2009), protecting targets without accounting for threat externalities can actually increase overall risk, instead of reducing it.

Game theory is a natural approach to consider if we want to account for attacker adaptivity and the corresponding security externalities. In a game-theoretic model, each player (e.g., defender or attacker) is assumed to make rational decisions based on available resources and knowledge. For example, Bier et al. (2007) model a fully strategic defender and attacker, deriving the attacker's behavior (and the externalities of defensive investments) directly from the underlying optimization problem. In that model, the defender may purposely leave a low-valued location unprotected, in order to lead the adversary away from more valuable targets. Shan and Zhuang (2013) investigate models involving a partially strategic attacker, but find that assuming a fully strategic attacker incurs less potential for loss than allowing attackers to be non-strategic. In other words, even if the assumptions of game theory are unrealistic in some cases, they at least generally yield conservative results.

Since early explorations by Sandler and Lapan (1988), game-theoretic models have been increasingly used in homeland security (especially after the tragedy of September 11th, 2001), providing a rigorous mathematical way to account for the actions of intelligent adversaries. Comprehensive reviews of game-theoretic models against adaptive adversaries can be found in Sandler and Siqueira (2009) and Guikema (2009). Recent applications of game-theoretic models to homeland security cover a wide range, including shipping-container inspection (Wein et al. 2006; Bier and Haphuriwat 2011; Haphuriwat et al. 2011), randomized inspection at airports (Pita et al. 2009; Paruchuri et al. 2008), choice of

bioterrorism countermeasures (Parnell et al. 2008), and risk assessment of radiological and nuclear terrorism (Streetman 2010).

Many of these models rely on the assumption of common knowledge, in which all players share the same understanding of the system, and have complete knowledge of the other players' objectives. However, as noted by Bier (2007), one reason for the existence of intelligence services is precisely because the defender may not know the attacker's preferences (and vice versa). The importance of defender uncertainty about attacker characteristics is also discussed by Farrow (2008). Therefore, Bayesian games of incomplete information are needed to reflect such uncertainties. We will review recent models involving games of complete and incomplete information, respectively, in more detail in the following two sections.

2.1.2 Games of Complete Information

In the literature, the defender's decision is in general to optimally allocate her defensive resources among a collection of potential targets or countermeasures, while the attacker's decisions may include selecting a target, choosing an attack mode, choosing a level of attack effort, or allocating personnel and other resources. Although some models allow continuous choices for both players (e.g., Major 2002; Siqueira and Sandler 2006; Zhuang and Bier 2007), Sandler and Sigueira (2009) note that to maintain model tractability, the number of continuous variables is generally limited. Therefore, to simplify the analysis and emphasize the defender side of the resource-allocation problem, researchers often assume continuous defender decisions but discrete attacker choices (e.g., Bier et al. 2007, 2008).

Game-theoretic models can be either simultaneous or sequential. In a simultaneous game, neither player can observe its rival's action (Bier et al. 2005; Patcha and Park 2006; Zhuang and Bier 2007; Hausken et al. 2009). Such models are particularly useful for allocating real-time defensive resources (e.g., police patrols), as opposed to long-term capital investments. By contrast, sequential games may be more appropriate for decisions on infrastructure protection, since in practice, many target-hardening measures are likely to be observable by attackers. A Stackelberg game, where the defender moves first to allocate her defensive resources and the attacker then launches an attack, is often used for this purpose. This type of game usually involves a first-mover advantage, in which the defender benefits from the ability to allocate her defensive investments in such a way as to deter attacks or deflect the attacker to less valuable targets. The standard Stackelberg game assumes complete information; i.e., both defender and attacker characteristics are assumed to be common knowledge. In this case, the defender's optimal decision is to allocate her defensive resources in such a way that all defended targets are equally attractive to the attacker (Major 2002; Powell 2007). Additional variants of the Stackelberg game include defender-attacker-defender games (Parnell et al. 2010), and games with multiple attacker types (Hausken and Bier 2011).

However, in reality, the defender may not be fully knowledgeable about the attacker's characteristics, which may compromise the accuracy of threat estimates if not explicitly taken into account. Sources of uncertainty include the sparsity of empirical data, the time-evolving nature of the threat, and the reliance on subjective judgment by experts with varying levels of expertise and background (Willis 2008). In fact, uncertainty about attacker goals and capabilities is the essential reason for intelligence collection and analysis, and is critical for

achieving realistic hedging in defensive resource allocation (Bier 2007; Baker et al. 2009). Therefore, it is important to incorporate defender uncertainty about adversary characteristics in game-theoretic models. We therefore review games of incomplete information in the next section.

2.1.3 Games of Incomplete Information

Games of incomplete information are widely used to capture defender uncertainty about attacker goals and motivations. For example, Powell (2007) discusses how to allocate a limited defensive budget against intentional attacks, but considers only two types of adversaries (e.g., military and political terrorists). The model by Bier et al. (2007) allows for defender uncertainty about a continuous range of adversary intent, represented by probability distributions over possible attacker valuations of the various targets or attack strategies.

Many of these models do not explicitly explore how the optimal defenses vary with the extent of defender uncertainty. However, some exploratory efforts have been made in this regard. For example, Bier et al. (2008) assume that attacker target valuations follow a Rayleigh distribution, and take the coefficient of variation of that distribution to reflect the extent of defender uncertainty. Similarly, Jenelius et al. (2008) assume a Gumbel distribution for the attacker target valuations, and use the scale parameter of that distribution (proportional to its standard deviation) as a measure of defender uncertainty. Interestingly, the results in Jenelius et al. (2008) show a significant impact of uncertainty on the optimal defensive resource allocations, while the results in Bier et al. (2008) show little effect of uncertainty. However, the results by Bier et al. do indicate that different measures of target attractiveness can yield significantly different optimal budget allocations, suggesting that a

single univariate measure of attractiveness may be inadequate to capture the full effects of defender uncertainty.

To account for the different attributes of attractiveness in a manageable way, Bier et al. (2012) assume that the attacker chooses targets or attack strategies according to a multiattribute utility function with random weights on the various attributes. Wang and Bier (2011) further allow for defender ignorance about any unobserved attributes that may be important to the attacker, but have not been identified by the defender. These models provide a flexible representation of uncertain adversary intent, and are generally able to achieve realistic levels of defensive hedging. Results of these models also indicate that as the defender becomes more uncertain about the adversary goals and motivations (i.e., intent), the optimal defensive allocations need to be based more heavily on the defender's own target valuations.

Of course, the nature and extent of defender uncertainty about the adversary intent may both change over time. Such dynamics can be captured using Bayesian methods. For example, Sticha et al. (2005) and Jha (2009) use Bayesian networks to update defender beliefs about the attacker's propensity to attack, based on the attacker's publicized intent (e.g., statements made in public media) and historical choices. Similarly, Wang and Bier (2011) allow the defender to update her knowledge about attacker target preferences, also from observing actual (or attempted) attacks in the past.

Even though there is an abundant literature modeling uncertain adversary intent, models that explore the effects of adversary capabilities have so far been quite limited. Two representative examples are Brown et al. (2006) and Zhuang and Bier (2007). In particular, Brown et al. represent the adversary's capability by the number of targets that can be

attacked simultaneously, while Zhuang and Bier consider a more sophisticated model that allows for continuous levels of attacker effort. However, both assume that levels of attacker capability or effort are common knowledge. By contrast, Lapan and Sandler (1993) allow the defender to be uncertain about the adversary's level of resources, but consider only two defensive options – to capitulate or resist the adversary.

Moreover, none of the above models explicitly reflects the multi-dimensional nature of adversary capabilities (e.g., money, personnel, technological sophistication, and flexibility of movement), nor do they account for the fact that a given type of adversary capability may differ in its effectiveness against different types of targets (e.g., military bases versus civilian buildings) or for different attack modes (e.g., nuclear versus IED attacks). This is a significant gap, since intelligence experts generally believe that adversary capabilities are at least as important as intent (Cragin and Daly 2004).

In this dissertation, we attempt to bridge this gap in the literature by accounting for defender uncertainty about both adversary intent and capabilities, so that homeland-security decision makers can base defensive strategies on a more realistic and comprehensive set of adversary characteristics. We start with the Bayesian Stackelberg game developed by Bier et al. (2007). In their model, the defender first allocates her limited resources among a collection of targets without full knowledge of the adversary's goals and motivations (i.e., intent); the adversary subsequently chooses the target with the highest expected payoff to attack. We extend that model to include adversary capability.

Note that adversary capabilities will in general affect the vulnerability of potential targets, since for example, attacks on hardened targets can be expected to have a higher success probability when the attacker is highly capable. The vulnerability of a target can thus be

viewed as being jointly determined by defender preparedness and attacker capability or level of effort (Farrow 2007; Willis 2007; Zhuang and Bier 2007; Keeney and von Winterfeldt 2011). To account for this, the concept of a "contest success function" is borrowed from economics (Skaperdas 1996; Hausken and Bier 2011) to describe how the defender and the adversary compete for a target by their investments in defense and attack, respectively. In particular, the contest success function we choose, following Wang and Bier (2013b), uses a "decisiveness" parameter to measure the extent to which vulnerability is affected by adversary and defender investments, rather than pure luck. This approach makes it possible to analytically explore the effects of defender uncertainty about adversary capabilities on optimal defensive strategies.

The vulnerability of potential targets is in turn expected to affect the adversary's target choices. Typically, adversaries will prefer attacks with higher success probabilities, at least among those with similar consequences for a successful attack. By explicitly capturing the effects of adversary capabilities on target vulnerabilities, our model is then able to endogenously determine the effects of adversary capabilities on threat probabilities.

Researchers have also studied the effects of attacker uncertainty. For example, Powell (2007) allows defenders to have private information about the vulnerability of potential targets, while the attacker observes only a probability distribution for the vulnerability of each target. Similarly, Zhuang et al. (2010) assume different defender types (e.g., with differing defense effectiveness and/or target valuations) that are not fully known to the attacker, and allow the defender to implement deceptive signals to mislead the attacker. In addition, Zhuang et al. (2010) also model two types of attacker learning, in which the attacker can observe both the defender's choice of security investments and the actual

results of past attacks. However, in this dissertation, we focus only on modeling defender uncertainties about attacker characteristics.

2.2 Quantification of Uncertain Adversary Intent

2.2.1 Representation of Adversary Intent

An adversary's objective could be a univariate function, such as maximizing the dollar-equivalent damage to the defender (Bier and Abhichandani 2003; Bier et al. 2005; Dillon et al. 2009). However, more complex multivariate measures of target attractiveness will in general be more realistic (Paté-Cornell and Guikema 2002; Beitel et al. 2004; Rosoff and John 2009). For example, in addition to evoking terror, terrorists may be interested in generating publicity, exacting revenge, achieving specific concessions, causing disorder, or provoking repression (Richardson 2006). The propaganda value or symbolic value of a target may also be important (Rubin and Rubin 2002).

Moreover, Woo (2002) recognizes that terrorists may consider not only psychological impact (e.g., evoking fear), but also execution difficulty (e.g., planning time, required personnel, technical difficulty, and consumption of financial and material resources), when choosing an attack strategy. Similarly, Rosoff and John (2009) argue that terrorists consider not only numbers of fatalities and injuries, terror, economic impact, and the symbolic value of their attacks, but also the time required to plan an attack, human resources required, and the costs of attack strategies. Beitel et al. (2004) likewise consider multiple attacker attributes, including both the resources required for attacks on particular targets, as well as the "returns" on those investments.

Several studies have discussed how to represent defender uncertainties about adversary intent, typically using probability distributions over their unknown target valuations (e.g., Rosoff and von Winterfeldt 2007; Barrett 2010). Researchers have also explored different ways of incorporating uncertainties into multiattribute adversary objective functions. For example, Bier et al. (2012) allow for random weights in multiattribute attacker utility functions. Inspired by Jenelius et al. (2008), Wang and Bier (2011) further account for the possibility of unobserved attributes that are important to the attacker but not known to the defender; and Rosoff and John (2009) propose a random utility model that allows uncertainty about not just the attribute weights, but also the attribute values.

Obtaining quantitative estimates for the parameters of interest in the adversary's multiattribute utility functions (e.g., attribute weights) from historical data is a difficult task. Although there are a number of well-maintained databases of international terrorist events (e.g., the Global Terrorism Database; International Terrorism: Attributes of Terrorist Events; and the Worldwide Incidents Tracking System), not much statistical analysis of these data has been done, perhaps because terrorist attacks have occurred relative rarely (at least for the severe attacks that are of most interest). Some exploratory work exists, including Enders and Sandler (2000), Barros and Proença (2005), and Mohtadi and Murshid (2006). However, none of these studies explicitly attempts to infer adversary preferences from historical data. Thus, pure statistical analysis of this issue is far from ready for use in practice.

It is therefore necessary to turn to subject-matter experts (intelligence experts, policy makers, security observers, terrorist proxies, etc.) for quantitative estimation of adversary preferences among potential targets or attack strategies. We now discuss existing methods

for elicitation of multiattribute preferences from subject-matter experts (not constrained to the elicitation of adversary preferences).

2.2.2 Expert Elicitation of Multiattribute Preferences

Multiattribute utility theory (MAUT; Keeney and Raiffa 1976) is one of the best-known and rigorous methods for trading off multiple (possibly conflicting) objectives; see recent advances in MAUT reviewed by Edwards et al. (2007). The simple additive multiattribute utility model has been shown to give satisfactory results in a wide variety of circumstances (Dawes and Corrigan 1974; Stewart 1996). One key task in the development of an additive multiattribute utility function is to obtain estimates for the weights (i.e., relative importance) of the various attributes.

Traditional elicitation methods including the ratio method (Edwards 1977), the swing-weighting method (von Winterfeldt and Edwards 1986), and the trade-off and pricing-out methods (Keeney and Raiffa 1976) involve asking stakeholders (e.g., experts) for cardinal assessments of attribute weights. However, these methods can be quite time-consuming to use, and may impose unwarranted precision (Schoemaker and Waid 1982; Borcherding et al. 1991). Moreover, assessing uncertainty over attribute weights would require the estimation of subjective probability distributions. While this approach has a long history of successful application (Edwards 1961; Cooke 1991; Hora and Jensen 2002), it generally requires extensive training and orientation, especially for elicitees with relatively non-quantitative backgrounds (Rosoff and John 2009).

Providing ordinal information rather than precise cardinal assessments is widely believed to be easier and more reliable. In fact, making pairwise comparisons has been suggested

to be the first step in the process of evaluating alternatives (Watson and Buede 1987). Moreover, ordinal judgments are generally easy to understand and interpret, avoiding the need for extensive orientation of experts prior to the elicitation process. Finally, while it is unrealistic to expect agreement on cardinal assessments in group decision process, consensus on rankings is often attainable (Kirkwood and Sarin 1985). There have also been a broad range of applications on eliciting multiattribute preferences from ordinal information, for example, in the fields of marketing research (Green and Srinivasan 1978, 1990; Green et al. 2001), health economics (McCabe et al. 2006; Ali and Ronaldson 2012), and adversary preference modeling (CREATE 2011).

Typical ordinal data for constructing multiattribute utility functions include rank orderings (or pairwise comparisons) of attribute weights, alternatives, or utility differences between paired alternatives (Horsky and Rao 1984). Some methods provide estimates for the overall utility values of the various alternatives from their rank orders directly. For example, Abbas (2004, 2006) defines "utility densities" analogous to probability densities, and applies the principle of maximum entropy to assign utilities to the alternatives directly according to their ranks. However, this method does not provide explicit estimates of attribute weights, and therefore cannot be applied to additional alternatives that have not been ranked.

Other elicitation methods utilize rank orders of attribute weights. For example, to avoid the difficulties associated with exact weight assessment, some researchers recommend instead simply rank ordering the importance of the various attributes (Dawes and Corrigan 1974; Stillwell et al. 1981). Among the existing formulae that convert rank orderings of attribute importance into cardinal estimates of attribute weights, SMARTER (Edwards and Barron 1994) has been shown to be the most accurate (Barron and Barrett 1996; Sarabando

and Dias 2006). (Note, however, that SMARTER has not been applied to cases with unobserved attributes, perhaps because it seems unrealistic to specify the rank of an attribute that we do not even know.)

Note that SMARTER generates only point estimates of attribute weights rather than probability distributions. Therefore, it may not be appropriate for situations in which uncertainty about preferences is a crucial consideration. However, Rao and Sobel (1980) have derived a marginal distribution for the kth largest weight, using the principle of maximum entropy (such that no more information is reflected by the distribution than that given by the ranks only), in which the marginal means of the attribute weights correspond exactly to the SMARTER weights (Barron and Barrett 1996).

In this dissertation, however, we focus our attention on indirect methods that yield estimates for attribute weights starting from rank orders of the overall attractiveness of alternatives (rather than rank orders of attribute importance). Conjoint analysis in marketing is one method for doing this (Green and Srinivasan 1978). Here, surveyed customers rank products in a factorial design; ordinal regression is then used to estimate the relative importance of each product attribute (Shocker and Srinivasan 1979; Green and Srinivasan 1990). A representative approach to conjoint analysis, LINMAP (Shocker and Srinivasan 1973), uses a linear program to find weight estimates that yield the smallest sum of pairwise "violations" from a single expert's ranking of the alternatives. However, like SMARTER, LINMAP provides only point estimates for the attribute weights. In addition, the optimal attribute weights may not be unique, and there is no agreement in the literature on how to handle multiple optima. Finally, it is an open question how to aggregate attribute weights obtained from multiple experts using LINMAP.

Another version of conjoint analysis uses logistic regression to obtain weight estimates that best fit a stakeholder's pairwise comparisons of alternatives (McFadden 1977). However, this approach is based on the unrealistic assumption that different pairwise comparisons of the alternatives (derived from a single rank ordering of all alternatives given by an expert) are independent. Moreover, it also yields only point estimates of attribute weights, which may not be adequate to capture heterogeneity among stakeholders. Mixed logit models can be applied to obtain probability distributions for the attribute weights, but often require that the distribution of attribute weights be multivariate normal (Revelt and Train 1998; McFadden and Train 2000; Sándor and Wedel 2002).

The ability to generate probability distributions (rather than only point estimates) is if anything even more important for an elicitation method that deals with ordinal data than for one that deals with cardinal data. In general, ordinal judgments are relatively weak, so it is important for an ordinal method to reflect the uncertainty that arises from the use of ordinal data. Therefore, probabilistic inversion (PI) and Bayesian density estimation (BDE) have recently been adapted to apply to ordinal preference rankings (Neslo et al. 2011; Wang and Bier 2013a). Both PI and BDE can generate probability distributions over attribute weights, explicitly capturing the existing uncertainty, without parametric assumptions such as normality.

In particular, probabilistic inversion (Cooke 1994; Bedford and Cooke 2001; Kraan and Bedford 2005; Kurowicka et al. 2010) aims to find a probability distribution over the attribute weights that can reproduce the experts' stated (marginal) rank orderings of alternatives. Neslo et al. (2011) was the first to apply PI to obtain distributions over multiple attribute weights that best fit the ordinal judgments given by experts. Recently, Wang and

Bier (2013a) extend that model to include an unobserved attribute that may be important but has not yet been identified, which ensures the existence of feasible solutions in PI at least in theory.

It is worth noting that the logic of PI is analogous to earlier work by Kadane et al. (1980), who elicited subjective conjugate distributions for the covariate coefficients in a multiple linear regression model using quantile estimates of the response variable. However, we believe that PI can be applied to a broader range of problems, since it does not require the use of conjugate priors.

Bayesian density estimation (Ferguson 1973, 1974, 1983; Escobar and West 1995) updates a decision maker's (possibly non-informative) prior distribution over the attribute weights by treating each expert's rank orderings as observations. BDE also allows the decision maker to specify a level of reliance on his or her own judgment, as opposed to the expert judgments. Erkanli et al. (1993) was the first to apply BDE to estimation using ordinal inputs. However, they first convert the rank orderings to cardinal values (a process that may introduce additional information and biases), and then treat those cardinal values as if they were independent (even though they could not be, since the underlying ordinal rankings clearly could not be independent). Wang and Bier (2013a) avoid these pitfalls by treating the entire set of rank orderings from a given stakeholder as a single observation (thus inherently accounting for the lack of independence between rank orderings given by the same stakeholder), and using the rank orderings directly (rather than converting them into cardinal values first).

The two basic elicitation methods, PI and BDE, provide practical yet methodologically rigorous ways for quantifying multiattribute adversary preferences using only ordinal

preference rankings, thus simplifying the task of quantifying threat probabilities for intelligence experts. The application of these methods may therefore increase the acceptance of quantitative approaches to risk estimation in the intelligence community.

2.3 Solving Defender-Adversary Games of Incomplete Information

Solving for equilibrium (optimal) defensive allocations in a defender-adversary game of incomplete information is generally difficult, because the computational complexity of the game may grow rapidly with both the number of uncertain parameters (e.g., attribute weights) and the number of potential targets or attack strategies. Moreover, a realistic model should allow the defender and attacker to have different objectives, so non-zero-sum games are generally needed, resulting in added computational complexity.

To our knowledge, only a few studies discuss algorithms for efficiently obtaining optimal defensive resource allocations against uncertain and adaptive adversaries. For example, Paruchuri et al. (2008) and Kiekintveld et al. (2010) develop a solver based on mixed-integer linear programming to randomize police checkpoints on the roadways of an airport, using a Bayesian Stackelberg game where the defender is uncertain about the discrete attacker types. Kiekintveld et al. (2011) further extend that model to represent defender uncertainty about attacker payoffs by continuous probability distributions, and explore several promising computational techniques for use in large-scale problems. However, these algorithms generally work well only with ”nice” objective functions; i.e., convex functions if minimizing, or concave functions if maximizing.

In spite of the sparse literature on how to solve games with defender uncertainty about attacker intent or capabilities, there is a rich body of algorithms on a related class of infrastructure protection problems – namely, network-interdiction problems (see for example Wood 1993; Washburn and Wood 1995). Research on the stochastic version of network interdiction problems is also well-established. In particular, it models a Bayesian Stackelberg game with two players, an interdictor and an evader (Cormican et al. 1998; Morton et al. 2007). The interdictor attempts to remove vertices or arcs from a network through which the evader then optimizes a given objective (e.g., maximizing the network flow, minimizing the traverse time, or finding the most reliable path from origin to destination). Stochastic variation comes from the fact that the interdictor may be uncertain about such factors as interdiction success, arc capacities, and the evader's characteristics. Applications of the stochastic interdiction model include disruption of the flow of enemy troops or terrorists, interception of illegal items (e.g., drugs, weapons, etc.), infectious disease control, and cyber security.

Stochastic network interdiction problems are often solved by sampling the objective function values. There are two sampling methods – i.e., the Monte Carlo simulation method and the decomposition method, depending on whether random samples are generated before or during the solution process. For example, the sample-average approximation (SAA) method is based on Monte Carlo simulation, and solves multiple independent replications of the sampling-based approximate problem (Shapiro and Homem-de-Mello 2000). This method is simple to implement, and especially suitable for parallel computing, which would significantly reduce the time required for large-scale problems (Janjarassuk and Linderoth 2008). However, it is reasonable to expect a waste of information in the

SAA method, since connections between samples are largely ignored. By contrast, the decomposition methods progressively reveal more information about the sample space, based on information obtained from previous steps of the solution process. For example, we may either sequentially increase the sample size (Higle and Sen 1991; Ermoliev 1983; Bayraksan and Morton 2011), or sequentially refine the partitions of the sample space (Kall et al. 1998).

A nice feature of stochastic interdiction problems is that the second-stage game is typically a well-structured network model (such as maximum flow or shortest path). Moreover, the interdictor and the evader are assumed to interact in a zero-sum game in most cases. Convexity of the recourse problem in the second stage can thus be exploited using duality, making the interdiction problem amenable to efficient computation methods such as the L-shaped method of Van Slyke and Wets (1969), or the generalized Benders decomposition of Geoffrion (1972). Moreover, convexity also permits pleasant derivation of statistical bounds for the true optimal objective value (see for example Cormican et al. 1998).

The defender-adversary game of incomplete information that we discuss in this dissertation is in part easier than the stochastic interdiction problems, since there is no complicated network structure. However, our problem generally forms a non-zero-sum game where the defender and attacker have different objectives, and involves binary recourse variables and nonconvex objective functions and constraints. All these features lead to added difficulties. Therefore, we conduct a comprehensive investigation of state-of-the-art optimization techniques with the purpose of solving our proposed defender-adversary game. In particular, we focus on the widely used Monte Carlo simulation-based sample-average approximation (SAA) method, and explore two categories of optimization algorithms, one based on

mixed-integer nonlinear programming and the other based on derivative-free optimization. We are mainly interested in identifying conditions and solvers with which our proposed defender-adversary game can be solved rigorously and efficiently within acceptable time constraints and levels of accuracy.

2.4 Summary

The development of game-theoretic models for homeland security has advanced to the point where those models are almost ready for use in real-world decisions. The goal of this dissertation is to provide ways of effectively and efficiently quantifying and solving game-theoretic models of realistic size and complexity.

We attempt to fill a gap in homeland-security literature by the explicit consideration of adversary capabilities in addition to just intent, and to address two of the most significant hurdles to making game theory applicable in practice; namely, the need to quantify uncertain adversary intent using subject-matter experts, and the need for powerful computational tools to solve for optimal defensive strategies. We hope this effort will pave the way for homeland-security decision makers to adopt game theory as a readily usable tool.

Chapter 3

Basic Game for Optimal Defenses against Adaptive Adversaries with Uncertain Intent

In this chapter, we introduce the Bayesian Stackelberg game to account for defender uncertainty about only adversary intent, following Bier et al. (2007) and Wang and Bier (2011). This model has two players: a defender; and an attacker. The defender first decides how to allocate her defensive resources, and the attacker then chooses a target to attack according to his target preferences (represented by a multiattribute utility function, which is not fully known to the defender). We present the basic game in Section 3.1, and discuss how to account for defender uncertainty about the multiattribute attacker utility function in Section 3.2. Then, Section 3.3 discusses briefly the behavior of the basic model. The

contents in this chapter are drawn heavily from Wang and Bier (2011), with some minor modifications.

3.1 Basic Model

We assume that the defender moves first to decide on how to allocate her defensive resources among a heterogeneous collection of potential targets, where x is a vector representing the resources allocated to the various targets (see Figure 3.1 for the corresponding game tree). The defender has only partial knowledge about the attacker type ω, represented by a subjective probability distribution Q over the space of all possible attacker characteristics Ω. Nature then randomly chooses an attacker type $\omega \in \Omega$ according to Q. Finally, the attacker observes both his type ω and the defender's resource allocation x, and chooses one target with the highest expected payoff to attack. For now, the attacker type ω is described only by his intent. In Chapter 5, we extend our discussion to consider both attacker intent and capabilities.

We quantify the outcome of the game (to the defender) using the TVC risk equation "Risk = Threat $\times$ Vulnerability $\times$ Consequence." In particular, we assume that the consequences suffered by the defender from a successful attack on a given target are independent of her protective strategy. However, spending more defensive investment on a given target will reduce its vulnerability (i.e., the success probability of an attack on that target). Moreover, for any given defensive allocation, the likelihood of a given target being attacked (i.e., threat) can be derived endogenously based on the defender's subjective assessment of the attacker's intent.

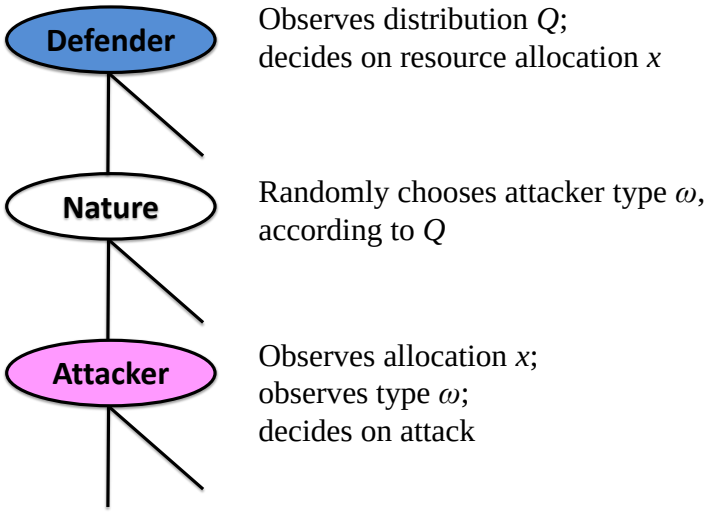

Figure 3.1: Extensive-Form Game of Incomplete Information

For simplicity, we assume that the defender and attacker are both risk neutral. The defender's objective is to allocate a fixed budget B to a set of N potential targets in such a way as to minimize her expected loss, as given by

$$\min_{x_1+\ldots+x_N \leq B} \int \sum_{n=1}^{N} p_n(x, U(\omega)) s_n(x_n) v_n dQ_U(U(\omega)) \tag{3.1}$$

where

N = number of targets

x_n = defensive resource allocated to target n; and $x = (x_1, \ldots, x_N)$

B = defensive budget

v_n = defender's valuation of target n; and $v = (v_1, \ldots, v_N)$

$s_n(x_n)$ = success probability of an attack on target n

$U_n(\omega)$ = attacker's valuation of target n; and $U(\omega) = (U_1(\omega), \ldots, U_N(\omega))$

$Q_U(\cdot)$ = defender's subjective joint distribution of attacker's target valuations $U(\omega)$

$p_n(x, U(\omega))$ = likelihood of an attack against target n given defensive allocation x and attacker target valuations $U(\omega)$

As in Bier et al. (2008) and Wang and Bier (2011, 2012), we assume that the success probability of an attack on target n is an exponential function of the defender's investment in that target, $s_n(x_n) = e^{-\lambda x_n}$, where λ is the cost effectiveness of defensive investment. For example, at $\lambda = 0.02$, if the x_n are measured in millions of dollars, then every million dollars of defensive investment will reduce the success probability of an attack by about 2%.

Given defensive resource allocation x, we assume that an attacker of type ω will choose to attack a single target with the highest expected payoff, and will choose randomly among multiple targets that all share the same maximum expected payoff. The likelihood of an attack on target n (for $n = 1, \ldots, N$) is then given by

$$p_n(x, U(\omega)) = \begin{cases} \frac{1}{Z} & \text{if } n \in \arg\max_i s_i(x_i) U_i(\omega) \\ 0 & \text{otherwise} \end{cases} \tag{3.2}$$

where Z is the cardinality of the set $\{i : \arg\max_i s_i(x_i) U_i(\omega)\}$.

Note that $U_n(\omega)$ is equal to the attacker's utility of a successful attack on target n in the absence of defense. To account for defender uncertainty about attacker intent (i.e., target preferences), we model the values of $U_n(\omega)$ as random variables. Note also that we have relaxed the zero-sum assumption in most existing game-theoretic models for homeland security and allow the defender and attacker to have different preferences among targets. This relaxation is realistic (Bier 2007; Wang and Bier 2011; Bier et al. 2012), since for

example, the defender and attacker may differ in how much weight they place on different attributes of attack consequences (such as fatalities, property loss, psychological impact, etc.).

3.2 Uncertain Multiattribute Adversary Utility Function

The attacker's target valuation $U_n(\omega)$ of target n is assumed to be a multiattribute utility function, and may in particular take into account the effects of unobserved attributes that are important to the attacker, but have not been identified by the defender. For simplicity, we assume that the attacker's utility is additive and linear in each of the various attacker attributes (including the unobserved attributes). Note that this choice is often adopted as a reasonable approximation of the true preference function; see for example Gigerenzer and Todd (1999). Of course, we could also consider utility dependence (Keeney and Raiffa 1976; Abbas and Bell 2011). However, in this dissertation, we focus on the simple case of additive independence among attributes, and adopt the form

$$U_n(\omega) = \sum_{m=1}^{M} W_m(\omega) a_{nm} + W_{M+1}(\omega) Y_n(\omega) \tag{3.3}$$

where

M = number of known attacker attributes

a_{nm} = attacker utility of target n on attribute m ($m = 1, \ldots, M$), where $a_{nm} \in [0, 1]$

$W_m(\omega)$ = attacker weight on attribute m, where $W_m(\omega) \geq 0$ for $k = 1, \ldots, M+1$, and $\sum_{m=1}^{M+1} W_m(\omega) = 1$

$Y_n(\omega)$ = utility of the unobserved attribute for target n

The values of a_{nm} represent the attacker's single-attribute utilities (not the actual attribute values) over the M attributes that are known to the defender, taking on values in [0, 1] with 1 representing the best possible value of the mth attribute and 0 representing the worst possible value. The attribute weights, $W(\omega) = (W_1(\omega), \ldots, W_{M+1}(\omega))$, and the utilities of the unobserved attribute for the various targets, $Y(\omega) = (Y_1(\omega), \ldots, Y_N(\omega))$, are assumed to be uncertain (as they will be to the defender). The values of the attribute weights $W(\omega)$ and utilities of the unobserved attributes $Y(\omega)$ are then used to describe the uncertain attacker type ω. Hereafter, we drop the type ω for convenience of notation; i.e., $(W(\omega), Y(\omega))$ = (W, Y).

To start with, we assume that the attribute weights W can take on only non-negative values. In Chapter 4, we allow the weights to be negative, to account for the possibility of conflicting views on whether a larger value of an attribute corresponds to higher or lower attacker utility. On the other hand, utilities of the unobserved attribute Y are introduced to represent the effects of any attributes that are unobserved by the defender, but could nonetheless be important to the attacker. The Y_n also take on values in [0, 1]. Since we generally do not know what the unobserved attributes are, we assume a priori that the Y_n are independent and identically distributed.

In principle, we could also allow the utilities of the known attributes a_{nm} to be random (as in Rosoff and John 2009), if putting probability distributions on W and Y is not adequate to capture the full effects of defender uncertainty about attacker preferences. However, for simplicity of the model, we limit ourselves here to the case where the utilities of the known attributes a_{nm} are constants.

Another way of incorporating uncertainty into multiattribute attacker utility functions would be to assess probability distributions directly over the target utilities U_n. However, the attractiveness of different targets may be correlated. For example, adversaries who would like to launch an attack on Chicago may also be inclined to attack other cities with major hub airports, such as Atlanta. It may be extremely difficult to assess a suitable joint probability distribution over such correlated utilities (Clemen and Reilly 1999). Incorporating uncertainty in the attribute weights seems to capture such correlations in a more natural way; see Bier et al. (2012).

3.3 Behavior of the Basic Model

In this section, we present three propositions and two examples to illustrate the behavior of the basic game. Proofs of the propositions are given in Appendix A.

Proposition 3.1. *The probability of an attack on any given target is non-increasing in the level of defensive resources allocated to it, if expenditures on other targets are held constant.*

In general, the more the defender invests in protecting a target, the less attractive it is to the attacker, and the less likely it is to be attacked, if defensive investments on other targets remain unchanged.

Proposition 3.2. *If the defender has a sufficiently large budget, she can ensure that any target she chooses has a probability arbitrarily close to one of being attacked, by defending the other targets.*

Thus, for example, with an adequate budget, the defender can effectively deter attacks against more valuable targets by deflecting attacks to the least valuable target.

Proposition 3.3. *When the effects of unobserved attributes are included, and the joint probability distribution of the weights W puts non-zero mass on every point in the set $\{w \in \mathbb{R}_{+}^{M+1} | \sum_{m=1}^{M+1} w_m = 1\}$ and the probability distribution of the utility of each target Y for the unobserved attribute puts non-zero mass on every value in $[0, 1]$, then all targets will have non-zero probabilities of being attacked. Moreover, this will remain true after any finite defensive investment.*

This suggests that the model with unobserved attributes may perform well in the event of "surprise" (e.g., an attack on a target that would appear to be of low attractiveness based on the known attributes). By contrast, a model without unobserved attributes could fail in the event of an attack on an unattractive target, especially one that was previously predicted to have a zero probability of being attacked.

Unfortunately, determining optimal resource allocations analytically is difficult. Therefore, we use the following example to illustrate the impact of defender uncertainty on the nature of the optimal defensive resource allocation. However, this particular case is not adequate to illustrate the complicated effects of defender uncertainty on her optimal expected utility, so we present another example to show those effects later. In both examples, we use simple hypothetical probability distributions for the uncertain attacker parameters, instead of distributions elicited from expert judgments.

Example 3.1. *Consider a two-target, two-attribute case with one known attribute and one unobserved attribute. For simplicity, we set the defender valuation of target 1 to*

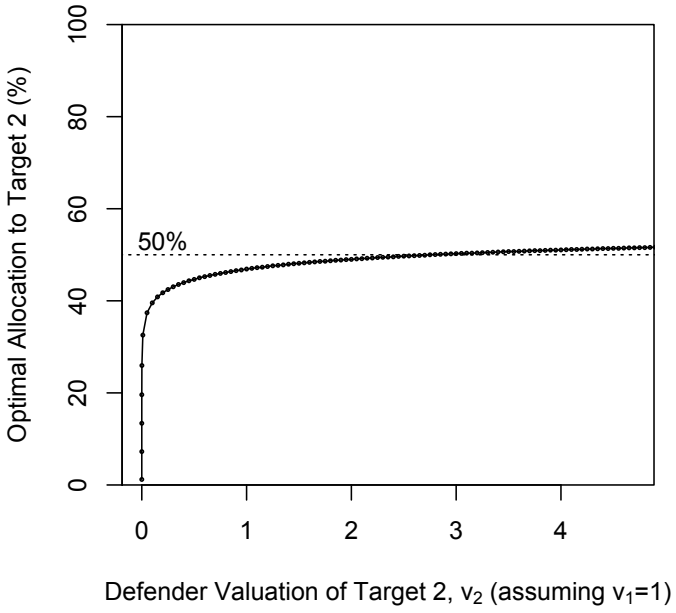

Figure 3.2: Optimal Defensive Resource Allocation for Example 3.1

$v_1 = 1$. *Suppose that target 1 is more attractive to the attacker on the known attribute* $(a_{11} > a_{21})$*; specifically, we assume that* $a_{11} = 1$ *and* $a_{21} = 0$. *Moreover, let the weight for the known attribute* (W_1) *and the utilities of the unobserved attribute* $(Y_1$ *and* $Y_2)$ *all be independently uniformly distributed in* $[0, 1]$. *Let the total defensive budget* $B = 100$*, and the cost effectiveness of defensive investment be given by* $\lambda = 0.1$. *For this case, the probability of each target being attacked can be derived in closed form (see proof in Appendix A.4), which helps simplify the determination of the optimal resource allocations. Figure 3.2 shows the optimal allocation of defensive resource to target 2 (*$\frac{x_2^*}{B}$*) as a function of the defender's valuation of target 2 (*v_2*).*

As we might expect, the optimal allocation to target 2 is increasing with the defender valuation of target 2, while the defender valuation of target 1 is fixed (e.g., at $v_1 = 1$). Furthermore, although target 2 is not attractive at all to the attacker on the known attribute, the optimal level of defensive investment allocated to this target can still be moderately high (even exceeding 50%) if the target is sufficiently important to the defender, because of "hedging" in the face of uncertainty about the unobserved attribute.

Powell (2007) shows that with perfect information about attacker intent, the defender would optimally allocate her resources in such a way as to equalize the attacker's expected payoffs across all defended targets. Otherwise, any defensive resources spent in protecting a less attractive target would be wasted. In that case, the optimal allocations will depend only on the attacker intent.

However, Example 3.1 shows that when the defender has only partial knowledge about the attacker's intent, she needs to assign more weight to her own target preferences. This example also suggests that in the presence of unobserved attributes, even a target that is believed to be of little value to the attacker on the known attributes may still receive significant defensive investment, to protect against the possibility that it may be attractive to the attacker on some unobserved attribute(s).

Example 3.2. *Consider a two-target, two-attribute case with one known attribute and one unobserved attribute, similar to Example 3.1. Again, target 1 is assumed to be more attractive to the attacker on the known attribute ($a_{11} = 1$ and $a_{21} = 0$), and the utilities of the unobserved attribute (Y_1 and Y_2) are assumed to be independently uniformly distributed on $[0, 1]$. However, we now let the weight on the known attribute (W_1) be a chosen parameter instead of a random variable for computational ease. In particular, a smaller value for W_1*

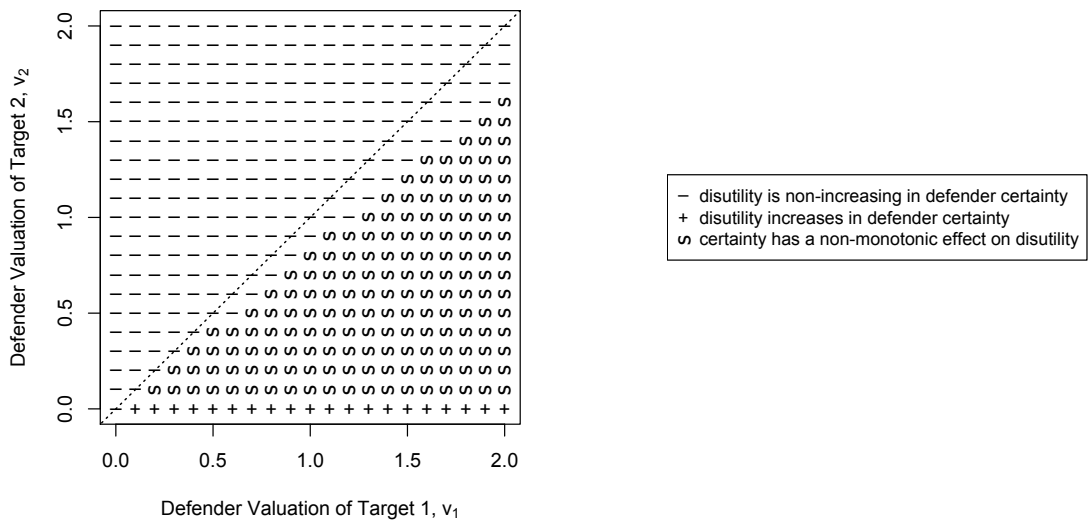

Figure 3.3: Effect of Defender Certainty on Equilibrium Expected Loss for Example 3.2

implies higher weight on the unobserved attribute, and thus less defender certainty about the attacker's target preferences. Figure 3.3 shows the relationship between the defender's optimal expected loss and the extent of her certainty (which is monotonically increasing in W_1), for different combinations of the defender's target valuations (v_1 and v_2).

The impact of defender uncertainty about attacker intent on the defender's equilibrium (optimal) expected loss is two-fold. On the one hand, high levels of uncertainty create the likelihood that some defensive resources will be wasted on the protection of targets that are not highly attractive to the attacker. On the other hand, high levels of certainty may not always lead to favorable outcomes, if the defender expects that the attacker will choose a target that is of high value to her.

When target 1 is less valuable to the defender than target 2 ($v_1 \leq v_2$, above the diagonal dashed line in Figure 3.3), the defender's optimal expected loss will be non-increasing as the defender becomes more confident that target 1 is more attractive to the attacker (since attacks on target 1 are less damaging to the defender, and the higher level of certainty allows the defender to concentrate her defensive resources on the attacker's preferred target). By contrast, when the defender values target 1 more than target 2 ($v_1 > v_2$, below the diagonal), these two effects work in opposite directions. In this case, higher certainty still allows the defender to concentrate her defensive resources, but now also increases the chance of an attack on the more damaging target. Depending on how much target 1 is preferred by the defender when $v_1 > v_2$, the defender's optimal expected loss can be non-increasing in her certainty level, non-monotonically related to her level of certainty, or even increasing in her certainty level (when the ratio of v_1 to v_2 is sufficiently large).

The basic game we have developed so far is capable of representing the effects of defender uncertainty about attacker intent, by using multiattribute attacker utility functions and taking into account factors that are important to the attacker but have not been identified by the defender. Moreover, the model also allows us to explicitly explore the impact of the defender's uncertainty on her optimal decisions. In the next chapter, we propose a rigorous yet practical elicitation process to obtain estimates of attacker intent using ordinal judgments provided by subject-matter experts. In Chapter 5, we extend the basic game to account for the effects of attacker capabilities in addition to just intent, and allow the defender to be uncertain about both attacker intent and capabilities. Chapter 6 discusses computational tools for solving these defender-adversary games within acceptable time constraints and levels of accuracy.

Chapter 4

Expert Elicitation of Adversary Intent Using Ordinal Judgments

In this chapter, we introduce a simple elicitation process to estimate attacker intent by asking subject-matter experts to provide only (partial) rank orderings of the attractiveness of potential targets. In particular, we investigate two mathematical approaches, probabilistic inversion (PI) and Bayesian density estimation (BDE) in Sections 4.1 and 4.2, respectively, followed by a discussion of the relationship between the two methods in Section 4.3. Next, Section 4.4 exhibits how PI and BDE handle unobserved attributes. Section 4.5 allows for negative attribute weights to automatically account for possible conflicting views on the direction of effect for each attribute. In Sections 4.6 and 4.7, respectively, we present a hypothetical case study and a realistic application on elicitation of adversary intent. Finally, we provide sensitivity analysis on how PI and BDE behave in the face of expert consensus or disagreement, and on the reliability of results using partial rankings in Section 4.8. The

contents in this chapter are drawn heavily from Wang and Bier (2013a), with newly added analyses in Sections 4.5, 4.6, and 4.8.3.

As in Section 3.2, we assume that the attacker's target valuations are represented by a multiattribute utility function, which may in particular include unobserved attributes that are important to the attacker, but have not been identified by the defender. For simplicity, we assume that the attacker's utility is linear in each of the various attributes, and these attributes are additively independent of each other. In particular, the attacker's valuation U_n of target n (n = 1, …, N) is given by

$$U_n(W, Y) = \sum_{m=1}^{M} W_m a_{nm} + W_{M+1} Y_n \tag{4.1}$$

where

N = number of potential targets (or attack strategies)

M = number of known attacker attributes

a_{nm} = attacker utility of target n on attribute m (m = 1, …, M), where a_{nm} takes on values in [0, 1], with 1 representing the best possible value of the mth attribute and 0 representing the worst possible value

W_m = weight on attribute m, where $W_m \geq 0$ for m = 1, …, $M + 1$, and $\sum_{m=1}^{M+1} W_m = 1$

Y_n = utility of the unobserved attribute for target n, also taking on values in [0, 1]

Note that the attribute weights, W = (W_1, …, W_{M+1}), and the utilities of the unobserved attribute for the various targets, Y = (Y_1, …, Y_N), are assumed to be uncertain (as they will be to the defender). We use lower-case letters w and y for realizations of the vector random variables W and Y, respectively. Let Ω be the space of all possible values of (W, Y), as

given by

$$\Omega = \Delta_{M+1}(1) \times [0,1]^N \tag{4.2}$$

where $\Delta_{M+1}(1)$ is the simplex defined by $\{w \in \mathbb{R}_+^{M+1} \mid \sum_{m=1}^{M+1} w_m = 1\}$. The task of expert elicitation is then to mathematically derive probability distributions over Ω that can match the rank orderings of attacker target valuations U_n provided by subject-matter experts.

In the following, we apply PI and BDE to preference rankings based on the additive multiattribute utility function in (4.1). However, our methods can also be extended in a straightforward way to accommodate multilinear utility functions (Keeney and Raiffa 1976).

4.1 Probabilistic Inversion

In this section, we first formalize the mathematics of probabilistic inversion, as applied by Neslo et al. (2011) to convert ordinal judgments of target attractiveness into cardinal estimates of attribute weights, and extend their model to include unobserved attributes. We then present a Monte Carlo-based approximation of the PI problem, and discuss two possible solution approaches for solving it.

4.1.1 Mathematical Basis of PI

Suppose that we ask K experts to rank the top R out of N targets based on their attractiveness to the attacker, no ties of targets being allowed. Note that when $R = N$, experts are asked to give a complete rank ordering of all targets. (The methods presented in this chapter can be extended in a straightforward manner to the case where the experts also provide rankings for some number R' of the least attractive targets.) We then specify an R-by-N empirical

distribution matrix of expert rankings, P, where element P_{rn} represents the probability that target n is ranked at the rth place by a randomly chosen one of the K experts, following Neslo et al. (2011). For example, suppose that three experts are asked to compare two targets. One of these experts thinks that target 1 is more attractive to the attacker, whereas the other two experts both think that target 2 is more attractive. In this case, the empirical distribution matrix for the three experts is $P = \begin{bmatrix} 1/3 & 2/3 \\ 2/3 & 1/3 \end{bmatrix}$. Note that since ties are not allowed, each expert can rank exactly one target in the rth place, so the row sums of P satisfy $\sum_{n=1}^{N} P_{rn} = 1$ for all r = 1, …, R.

To illustrate the PI approach, we consider the uncertain attacker parameters (W, Y) as "input," and treat the expert rank orderings as an "output" that depends on the values of (W, Y). PI aims to find the distribution Q over Ω, the space of all possible values of (W, Y), that can match the empirical distribution matrix of expert rankings P, and has the smallest Kullback-Leibler (K-L) divergence (Kullback and Leibler 1951) to a pre-determined (e.g., non-informative) starting probability measure, Q_0. In particular, the optimization problem is given by

$$\min_{Q} \int_{\Omega} \frac{dQ}{dQ_0} ln(\frac{dQ}{dQ_0}) dQ_0 \tag{4.3}$$

$$\text{s.t. } \int_{\Omega} J_{rn}(w, y) dQ = P_{rn} \quad \text{for } r = 1, \ldots, R;\ n = 1, \ldots, N$$

where $J(w, y)$ is an R-by-N indicator matrix. For a given set of attribute weights w and utilities y of the unobserved attribute, $J_{rn}(w, y)$ equals 1 if target n ranks in the rth place and 0 otherwise. In addition, $\frac{dQ}{dQ_0}$ is the Radon-Nikodym derivative of the probability measure Q with respect to the starting measure Q_0 (Seppäläinen 2010). Note that to have finite K-L divergence in (4.3), we must assume that Q is absolutely continuous with respect

to Q_0. Moreover, we use the convention that if $\frac{dQ}{dQ_0} = 0$, then $\frac{dQ}{dQ_0} ln(\frac{dQ}{dQ_0}) = 0$, following Csiszár (1975).

The use of K-L divergence to measure the closeness of two probability distributions in PI is recommended by Cooke (1994), mainly because it is closely related to the concept of entropy in information theory, and is associated with a readily usable computational tool – iterative proportional fitting (Csiszár 1975). Of course, there are several other ways to measure the closeness of two probability distributions. For example, we can use Hellinger distance (Beran 1977), total variation distance (Tierney 1996), or the more general f-divergence (Csiszár 1967); see Gibbs and Su (2002) for a comprehensive review.

We could choose the starting measure Q_0 to be non-informative, if the defender had little or no prior knowledge about the attacker preferences before any expert judgment became available. An easy choice for Q_0 in that case would be to adopt a "flat" starting measure; i.e., to assign equal probability to every possible value in the attacker parameter space Ω. In particular, the attribute weights W would be assumed to follow the Dirichlet distribution with all parameters equal to one; moreover, the utilities of the unobserved attribute for the various targets Y would be independently uniformly distributed in $[0, 1]$, and also independent of the attribute weights W. If desired, of course, we could also consider other types of non-informative starting measures (e.g., U-shaped instead of uniform), or an informative starting measure.

4.1.2 Monte Carlo-Based Approximation

The probability distribution Q^* that solves equation (4.3) has the density satisfying

$$\frac{dQ^*}{dQ_0}(w, y) = c \cdot \exp\{-\sum_{r,n} \lambda_{rn} J_{rn}(w, y)\} \tag{4.4}$$

where the λ_{rn} are Lagrange multipliers for the constraints in (4.3), and c is a normalizing constant. However, it is difficult to obtain the Lagrange multipliers λ_{rn} analytically in the important case of multiple experts. Therefore, we instead resort to Monte Carlo simulation. In particular, we randomly generate S independent samples for the attacker parameters (W, Y) from the starting measure Q_0, and let $\bar{\Omega}$ be the set of all simulated values ($w^{(s)}$, $y^{(s)}$), $s = 1, \ldots, S$. We construct the discretized starting measure $\bar{Q}_0$ by placing equal mass on every element ($w^{(s)}$, $y^{(s)}$) of $\bar{\Omega}$. The approximate PI problem is then to find a discrete distribution $q = (q_1, \ldots, q_S)$ over the elements of $\bar{\Omega}$ that yields the smallest K-L divergence from the discretized starting measure $\bar{Q}_0$, given that the mapping of q to the space of target rank orderings matches the empirical distribution matrix P. In particular, the approximate PI problem is given by

$$\min_{q \in \Delta_S(1)} \sum_{s=1}^{S} q_s ln(S q_s) \tag{4.5}$$

$$\text{s.t. } \sum_{s=1}^{S} q_s J_{rn}(w^{(s)}, y^{(s)}) = P_{rn} \text{ for } r = 1, \ldots, R;\ n = 1, \ldots, N$$

where $\Delta_S(1)$ is the simplex defined by $\{q \in \mathbb{R}_+^S \mid \sum_{s=1}^{S} q_s = 1\}$. Proposition 4.1 provides a sufficient condition to guarantee the existence of feasible solutions to the approximate PI problem in (4.5). Proofs of the propositions in this chapter can be found in Appendix B.

Proposition 4.1. *Suppose that a set of independent random samples* $(w^{(s)}, y^{(s)})$, $s = 1, \ldots, S$, *has been drawn from the starting measure* Q_0. *If for each expert,* $\exists s$ *such that* $(w^{(s)}, y^{(s)})$ *yields the rank ordering of targets specified by that expert, then the optimization program in (4.5) is feasible and has a unique optimal solution.*

When the approximate PI problem in (4.5) is feasible, we may employ the iterative proportional fitting (Csiszár 1975), which is a general algorithm to find the smallest K-L divergence between two discrete probability measures subject to linear constraints. In particular, the procedure begins with $q_s = \frac{1}{S}$ for $s = 1, \ldots, S$, and iteratively adjusts q to satisfy exactly one of the linear equations in (4.5) at a time. (See Appendix C.1 for the algorithm, adapted to our problem.) If we have a sufficiently large number of samples S, then for each expert, we are ensured to have at least one sample $(w^{(s)}, y^{(s)})$ that can match that expert's rank ordering of targets. However, when it takes too many samples to ensure the feasibility of (4.5) (e.g., for large numbers of targets), we could also use the iterative PARFUM algorithm developed by Du et al. (2006) to get a probability distribution for (W, Y) corresponding to a marginal ranking distribution that is "closest" to the empirical expert ranking matrix P.

4.2 Bayesian Density Estimation

We now discuss another elicitation method, Bayesian density estimation. First, we explain how BDE can be applied to our elicitation process. We then describe how Gibbs sampling could be used to obtain the (expected) joint posterior distribution for parameters in the

multiattribute attacker utility, using ordinal judgments of target attractiveness provided by experts.

4.2.1 Mathematical Basis of BDE

BDE allows the defender to specify a prior distribution Q_p over Ω, the space of attacker parameters (W, Y), and then treat expert judgments as observations to update that prior, leading to a posterior distribution Q. In particular, we assume that the defender's prior Q_p is randomly chosen in accordance to a Dirichlet process. This simplifies the Bayesian updating, since the posterior Q will still be a Dirichlet process, but with different parameters. Moreover, the defender can specify a self-trust degree α ($\alpha > 0$) to reflect the level of reliance on her own knowledge about the attacker preferences, as opposed to the expert judgments.

The definition of a Dirichlet process is given by Ferguson (1973), as follows.

Definition 4.2. *(Ferguson 1973) Let μ be a non-negative and finitely additive measure*[1] *on $(\Omega, \mathscr{A})$, where Ω is a space and $\mathscr{A}$ is a σ-field of subsets of Ω. Then a stochastic process Q_p indexed by elements A of $\mathscr{A}$ is a Dirichlet process on $(\Omega, \mathscr{A})$ with parameter μ if for all positive integers L and all measurable partitions $(A_1, \ldots, A_L)$ of Ω, the joint probability distribution of the random vector $(Q_p(A_1), \ldots, Q_p(A_L))$ is Dirichlet with parameters $(\mu(A_1), \ldots, \mu(A_L))$.*

The process Q_p may be considered a random probability measure on $(\Omega, \mathscr{A})$, since $Q_p(A)$ (which can be interpreted as the probability of falling into the set $A \in \mathscr{A}$) equals 1

[1] Definition of a finitely additive measure can also be found in Ferguson (1973).

when $A = \Omega$. Note that the scalar function $\mu(\cdot)$ maps $\mathscr{A}$ (the σ-field of subsets of Ω) to the real line $\mathbb{R}$, and is itself a non-negative measure (but not necessarily a probability measure). Let $Q_0(A)$ be the probability of being in region $A \in \mathscr{A}$ under the starting probability measure Q_0. Then we choose μ according to

$$\mu(A) = \alpha Q_0(A) \quad \text{for any } A \in \mathscr{A}$$

where $\alpha > 0$ is the defender's self-trust degree.

By choosing the Dirichlet parameter μ in this way, the expectation of the random prior distribution $\mathbf{E}[Q_p]$ equals Q_0, so the starting probability measure Q_0 is the defender's best guess for the uncertain attacker parameters. For given choices of α and Q_0, the random vector $(Q_p(A_1), \ldots, Q_p(A_L))$ for any partition of the attacker parameter space Ω would follow the Dirichlet distribution with parameters $(\alpha Q_0(A_1), \ldots, \alpha Q_0(A_L))$. Therefore, the expected values of the various elements in the random vector are given by

$$\mathbf{E}[Q_p(A_l)] = \frac{\alpha Q_0(A_l)}{\sum_{l'=1}^{L} \alpha Q_0(A_{l'})} = \frac{Q_0(A_l)}{\sum_{l'=1}^{L} Q_0(A_{l'})} \quad \text{for } l = 1, \ldots, L$$

Increasing the value of α will shrink the variance of the Dirichlet distribution, while keeping the expected values $\mathbf{E}[Q_p(A_l)]$ unchanged (Bier and Yi 1995). For example, when $\alpha \to \infty$, the random vector $(Q_p(A_1), \ldots, Q_p(A_L))$ will almost surely equal the deterministic vector $(Q_0(A_1), \ldots, Q_0(A_L))$. Thus, the value of α controls the dispersion of the random prior distribution Q_p around the defender's best guess Q_0. In fact, the larger α is, the closer Q_p is likely to be to Q_0.

Suppose again that we ask K experts to rank the top R out of N targets, no ties among targets being allowed. For convenience, we denote the rank ordering of expert k (for $k = 1, \ldots, K$) by an ordered set of N distinct target indices $RO^{(k)} = \{n_1^{(k)}, \ldots, n_N^{(k)}\}$, where $n_r^{(k)}$ is the index of the rth most preferred target for $r = 1, \ldots, R$; for $r = R{+}1, \ldots, N$, $n_r^{(k)}$ is the index of one of the $N - R$ unranked targets.

We then associate the (partial) rank ordering $RO^{(k)}$ provided by expert k with a subset of Ω by excluding all values of (w, y) that are inconsistent with that expert's judgment, called the "active region" $AR^{(k)}$ for expert k, as given by

$$\begin{aligned} AR^{(k)} = \Big\{ & (w,y) \in \Delta_{M+1}(1) \times [0,1]^N, \\ & \text{s.t. } U_{n_i^{(k)}}(w,y) \geq U_{n_j^{(k)}}(w,y) \quad \text{for } 1 \leq i \leq R \text{ and } i < j \leq N \Big\} \end{aligned} \tag{4.6}$$

where $U_n(w, y)$ is the multiattribute attacker utility as given in equation (4.1).

We randomly sample an observation $O^{(k)} \in AR^{(k)} \subseteq \Omega$ under the starting measure Q_0 to represent expert k's rank ordering $RO^{(k)}$. We then condition on $O^{(k)}$ instead of $RO^{(k)}$; in other words, we treat the rank ordering $RO^{(k)}$ as if it were equivalent to a single random point $O^{(k)} \in \Omega$ that generates rank ordering $RO^{(k)}$. The distribution $Q_0^{(k)}$ for random point $O^{(k)}$ is therefore proportional to

$$Q_0 \cdot \mathbf{1}_{\{(w,y) \in AR^{(k)}\}} \tag{4.7}$$

where $\mathbf{1}_{\{(w,y) \in AR^{(k)}\}}$ equals 1 if $(w, y) \in AR^{(k)}$ and 0 otherwise; i.e., $Q_0^{(k)}$ is a truncated version of Q_0 that puts non-zero mass on only those values (w, y) that are in the active

region $AR^{(k)}$. In what follows, we use the lower-case letter $o^{(k)}$ to denote a realization of $O^{(k)}$.

Suppose that the random observations $O^{(k)}$ (k = 1, …, K) for the various experts are independent of each other, and also independent of the prior Dirichlet process $Q_p \sim \mathscr{D}\{\alpha Q_0\}$. Then the posterior distribution Q conditional on the random observations $O^{(k)}$ (k = 1, …, K) is a mixture of Dirichlet processes (Antoniak 1974), as given by

$$Q \mid O^{(1)}, \ldots, O^{(K)} \sim \int \ldots \int \mathscr{D}\{\alpha Q_0 + \sum_{k=1}^{K} \delta(o^{(k)})\} dQ_0^{(1)} \ldots dQ_0^{(k)}$$

where $\delta(o^{(k)})$ is the probability distribution giving unit mass to the point $o^{(k)}$.

Moreover, the expectation $\mathbf{E}[Q]$ of the random posterior distribution is a weighted sum of probability distributions as given by

$$\mathbf{E}[Q] = \frac{\alpha}{\alpha + K} Q_0 + \frac{1}{\alpha + K} \sum_{k=1}^{K} Q_0^{(k)} \tag{4.8}$$

In traditional BDE, each of the $Q_0^{(k)}$ in (4.8) would reduce to a unit probability mass at a single point; however, in our application, $Q_0^{(k)}$ is instead a truncated version of the starting measure, maintaining non-zero mass over that portion of the domain Ω consistent with the rank orderings given by expert k. Note that larger values of α correspond to higher emphasis on the defender's prior guess Q_0, and lower trust in the expert judgments. In particular, one can interpret this as the defender weighting his or her judgment as equivalent to the judgments of some number α of experts.

4.2.2 Gibbs Sampling for BDE

Clearly, (4.8) is a linear pool of K + 1 probability distributions; i.e,

$$(W, Y) \sim \begin{cases} Q_0 & \text{with probability } \dfrac{\alpha}{\alpha + K} \\ Q_0^{(k)} & \text{with probability } \dfrac{1}{\alpha + K} \quad \text{for } k = 1, \ldots, K \end{cases}$$

which can be simulated by drawing random samples from the starting measure Q_0 with probability $\dfrac{\alpha}{\alpha + K}$, and from each of its truncated variants $Q_0^{(k)}$ with probability $\dfrac{1}{\alpha + K}$.

Note that for a given vector of utilities y of the unobserved attribute, the active region $AR^{(k)}$ for expert k is generally a polyhedron of the attribute weights w (and vice versa). Therefore, we propose to use Gibbs sampling, which is simple to implement, and is popular in the field of Bayesian analysis with constrained parameters (Geman and Geman 1984; Gelfand and Smith 1990; Gelfand et al. 1992). (See Appendix C.2 for the implementation of Gibbs sampling in our case.)

In particular, Gibbs sampling generates random samples for the attacker parameters $(W_1, \ldots, W_{M+1}, Y_1, \ldots, Y_N)$ cyclically, from the univariate conditional distribution for one parameter at a time, while keeping the values of other parameters fixed. It is generally much easier to sample from the univariate conditionals than from the joint distribution. For example, take the utility Y_n of the unobserved attribute for target n. With the values of all other parameters fixed, the possible values for Y_n consistent with the active region $AR^{(k)}$ are constrained to a bounded interval. Sampling of the attribute weights W_m is similar, except that we need to consider the fact that $\sum_{m=1}^{M+1} W_m = 1$.

Gibbs sampling is shown to produce random samples that converge in distribution to the target joint probability distribution (Tierney 1994). However, the convergence rate can be

slow, and it is not always clear when to stop the procedure in practice. In this dissertation, we simply predefine a large number of iterations (10^6), and remove the first 10% of the samples in order to get rid of the influence of the starting point. Another choice might be to set up a measure by which to observe convergence; e.g., by plotting a sequence of marginal posterior distributions in order to judge stability (Gelfand et al. 1992).

4.3 Relationship between PI and BDE

In this section, we explore the relationship between PI and BDE when applied to ordinal preference rankings. Most of our results here (except for Example 4.2) assume zero weight ($\alpha \to 0$) on the defender's judgment. This is done only for comparison purposes, since PI does not allow for an equivalent self-trust parameter. However, this does not mean that we would advocate putting zero weight on the defender's judgment in practice.

Hereafter, we choose the starting measure Q_0 to guarantee a feasible solution to the original PI problem (4.3).

Proposition 4.3. *Assume that the defender's self-trust degree $\alpha \to 0$. If there is only one expert, or the experts all give the same rank ordering, then PI and BDE will yield the same probability distributions for all attacker parameters.*

However, if the experts give different rank orderings, PI and BDE generally produce different results, even if we assume that the defender's self-trust degree $\alpha \to 0$. In order to investigate this discrepancy, we define the "composition of expert rank orderings" as a vector of the proportion of experts giving each possible (partial or full) rank ordering. For example, if one expert thinks that the top three out of five targets are targets 1, 2, and

3 (in that order), while three other experts all rank targets 1, 4, and 5 as the top three (in that order), then the composition of expert rank orderings is 25% for the rank ordering $1 > 2 > 3 > \{4,\ 5\}$, and 75% for the rank ordering $1 > 4 > 5 > \{2,\ 3\}$.

It is trivial to see that if multiple compositions of expert rank orderings yield the same empirical distribution matrix P, then PI will give the same results for all such compositions. By contrast, BDE can generate different results for different compositions of expert rank orderings, even when they correspond to the same empirical distribution matrix P. We show this by an example.

Example 4.1. *Suppose that two groups of experts are asked to give full rank orderings of three targets described by two known attacker attributes. Assumed attribute values and hypothetical expert rank orderings are given in Tables 4.1 and 4.2, respectively. We assume that the starting measure Q_0 is flat. For comparison purposes, we also let the defender's self-trust degree $\alpha \to 0$.*

Table 4.1: Values of Adversary Attributes for Example 4.1

	Attribute 1	Attribute 2
Target 1	1	0
Target 2	0	1
Target 3	0.5	0.5

Figure 4.1 presents the resulting probability distributions for the first known attribute weight W_1 for both methods (PI and BDE) and both expert groups (A and B), along with the corresponding mean values. PI gives identical distributions for both groups (as shown in the upper panel of Figure 4.1), because they have identical empirical distribution matrices;

Table 4.2: Hypothetical Expert Rank Orderings for Example 4.1

Group	Expert	Rank Ordering of Targets
Group A	1	$1 > 2 > 3$
	2	$2 > 3 > 1$
	3	$3 > 1 > 2$
Group B	1	$1 > 3 > 2$
	2	$2 > 1 > 3$
	3	$3 > 2 > 1$

i.e., $P_A = P_B = \begin{bmatrix} 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 1/3 \end{bmatrix}$. *By contrast, BDE provides different results for the two groups (as shown in the lower panel of Figure 4.1). Note that the first expert in group B gives a rank ordering of targets that is perfectly consistent with their values on attribute 1, reflecting a high weight on that attribute, and thus resulting in a broader distribution over* W_1 *for group B than for group A. The following proposition describes the fundamental relationship between PI and BDE.*

Proposition 4.4. *Assume that the defender's self-trust degree* $\alpha \to 0$. *Consider all possible compositions of expert rank orderings that can yield the given empirical distribution matrix* P. *If BDE gives multiple probability distributions for those compositions, then among those BDE distributions, the one that has the smallest K-L divergence from the starting measure* Q_0 *coincides with the result given by PI using the same* Q_0.

The expected posterior distribution of BDE essentially uses a linear opinion pool based on the chosen starting measure to aggregate the probability distributions that are elicited from the individual experts. By contrast, PI is more complicated. If we find every possible

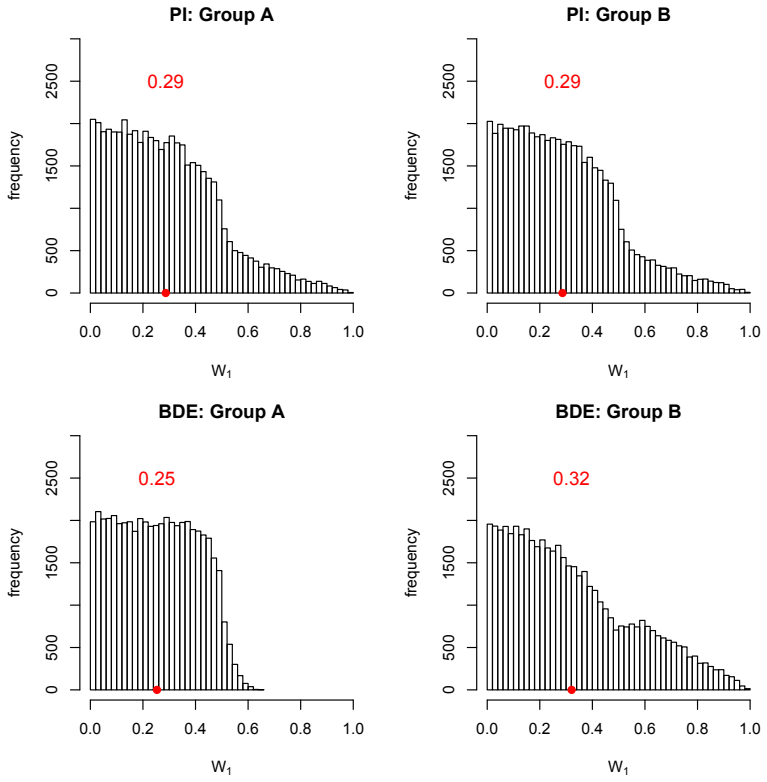

Figure 4.1: Elicited Probability Distributions for W_1 in Example 4.1

composition of expert rank orderings that satisfies the given empirical distribution matrix P (of which there can be infinitely many!) and apply BDE to all of them, then PI will pick the resulting BDE probability distribution that is closest to the chosen starting measure over the attacker parameters (e.g., by choosing the maximum entropy distribution if we adopt a flat starting measure). In Example 4.1, the distribution for W_1 generated by PI (see the upper panel of Figure 4.1) in fact coincides with the BDE result for a composition that matches the given P, but is different from that of either group A or B, with equal proportions of experts

giving all $3! = 6$ possible rank orderings of three targets. (In other words, with six experts, BDE and PI will give the same results if each expert gives a different rank ordering.)

Thus, we can see that PI exploits only marginal information about target rankings, and fails to take into account correlations among subgroups of experts (e.g., if those experts ranking target 1 higher than target 2 may also rank target 3 higher than target 4). By contrast, BDE is able to utilize that correlational information. However, one should note that using more information may not necessarily make BDE perform more sensibly, if for example the results seem overly sensitive to minor changes in expert rank orderings.

We would ideally like an elicitation method that is sensitive to the absolute amount of information provided by experts. For example, with only a small number of experts, we may want a method that yields a flatter distribution than when a large number of experts is available. Unfortunately, neither PI nor BDE is able to explicitly capture that idea (at least when we just use the expected posterior (4.8) and set the defender's self-trust degree $\alpha \to 0$ in BDE). However, when applying BDE, the defender could in fact assign a relatively high self-trust degree (i.e., large α), to at least qualitatively account for the lack of reliability in expert judgments when only a small number of experts is available. We use the following example to illustrate this.

Example 4.2. *Suppose that two experts are asked to give rank orderings of a collection of targets described by just one known attacker attribute. Suppose also that the elicited densities for the known attribute weight* (W_1) *using judgments of the two individual experts are given by Beta*$(20, 2)$ *and Beta*$(2, 20)$*, respectively. Figure 4.2 then shows the BDE densities under a flat starting measure, considering various levels of self-trust* α *for the*

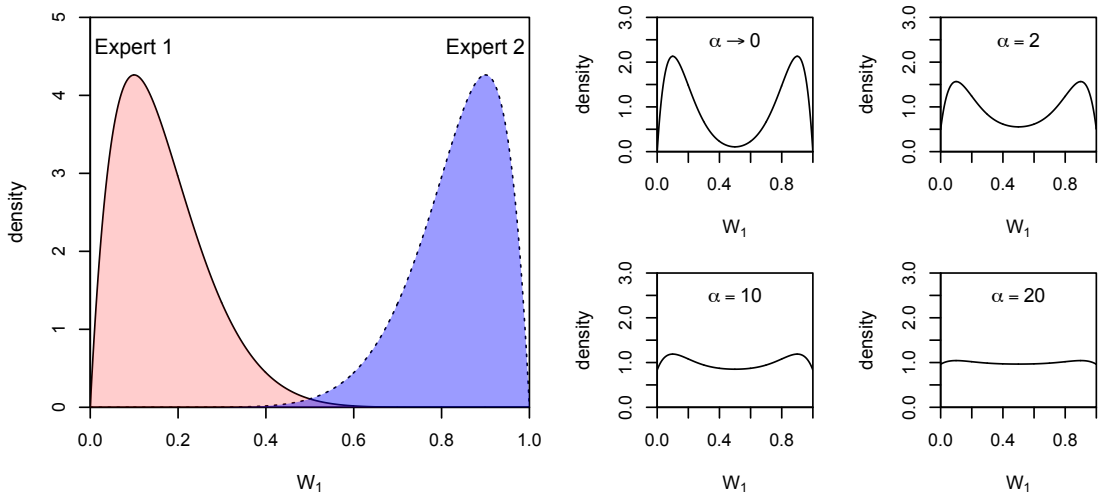

Figure 4.2: BDE Densities of W_1 for Different Levels of Defender Self-Trust

defender. As α increases, the aggregated probability density places less reliance on the expert judgments, and becomes less informative.

The computational complexity of BDE based on Gibbs sampling grows only linearly with the number of experts, the number of targets, and the number of uncertain attacker parameters. Therefore, one advantage of using BDE is that we can always anticipate obtaining a good solution within a controllable time constraint.

As for PI, if the experts do not deviate too much from the available set of known attacker attributes when giving their rank orderings, then the run time will be roughly equal for any number of experts. This favorable feature constitutes an advantage for PI when handling large numbers of experts. However, it sometimes requires too many samples from Q_0 to ensure a perfect match between the expert judgments and the distribution of rankings produced by the PI method. This is especially likely to occur if some expert judgments

cannot possibly be explained by the known attributes. This makes the computational behavior of PI somewhat difficult to analyze, since it may vary from case to case.

4.4 Treatment of Unobserved Attributes

The elicited weight for the unobserved attribute (from either PI or BDE) can be used as a measure of how well the given expert judgments can be explained by the assumed attacker attributes. In particular, for a given set of targets and their attribute values, the larger the weight we get for the unobserved attribute, the less capable the known attributes are of matching the expert opinions, and the more we need to investigate the nature of any possible unobserved attributes.

One caveat of our model is that even if the rank orderings of targets are perfectly consistent with their values on the known attributes, we could still get a non-zero weight for the unobserved attribute. Nonetheless, the mean of that weight generally decays as the number of targets N grows, and may become arbitrarily small when N gets sufficiently large. To illustrate this, consider the following example.

Example 4.3. *Suppose that an expert is asked to rank N targets described by just one known attacker attribute, where we assume a flat starting measure. Note that when there is only one expert, PI and BDE will yield the same result (if we set the defender's self-trust degree $\alpha \to 0$). We also assume that the target values on the known attribute $(a_{\cdot 1})$ form an arithmetic series with maximum 1 and minimum 0. For example, when we have N = 4 targets, the sorted attribute values are $(a_{11}, a_{21}, a_{31}, a_{41}) = (1, \frac{2}{3}, \frac{1}{3}, 0)$. If the expert gives a rank ordering of targets that is perfectly consistent with their values on the known attribute,*

then the mean elicited weight $\mathbf{E}[W_2]$ *for the unobserved attribute will be strictly decreasing in the number of targets* N *(at least for* $N \leq 1000$*). Moreover,* $\mathbf{E}[W_2]$ *gets arbitrarily close to zero as* $N \to \infty$ *(see proof in Appendix B.4).*

However, it is important to know how quickly the weight for the unobserved attribute declines in practice. To investigate this, we now randomly generate 500 sets of values for the known attribute for each given number of targets N. Figure 4.3 reports the 90% confidence interval of $\mathbf{E}[W_2]$ as a function of N assuming that the expert bases his judgment entirely on the simulated attribute values. As a rule of thumb, we may regard the weight for the unobserved attribute as being negligible in the case of perfect consistency when the number of targets is sufficiently large (e.g., $N \geq 15$). (Of course, if the rank orderings of target attractiveness provided by experts cannot be well explained by the known attacker attributes, no matter whether the judgments are based on a small or large number of targets, we will get a high weight for the unobserved attribute.)

Note that assigning a moderate weight to the unobserved attribute in the case of perfect consistency for small N may not actually be a flaw of our method. In fact, a given set of known attributes that can well explain the rank orderings of ten targets should essentially be more reliable than if those same attributes can explain the relative attractiveness of only two targets. In particular, when there are two targets to compare, the known attributes could easily give a perfect match just by coincidence, something that is less likely to happen for larger N. Therefore, the results of our method are conservative in the sense that they avoid placing too much weight on the known attributes when there are only a small number of targets whose rankings are being explained.

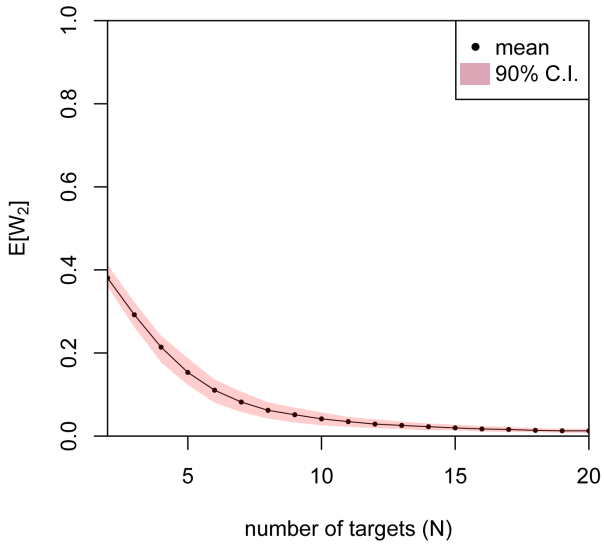

Figure 4.3: Confidence Interval for the Mean Elicited Weight of the Unobserved Attribute in the Case of Perfect Consistency

4.5 Allowing for Negative Attribute Weights

In using the methods developed so far, the defender must begin with a predefined set of attributes, as well as scales on which the attributes are measured. In particular, we need to specify the worst and best values of each attacker attribute. However, experts may differ on whether a larger value of a particular attribute corresponds to higher or lower attacker utility. For example, some experts may believe that a particular adversary group favors smaller rather than larger numbers of fatalities, since attacks with large number of fatalities could lead to reduced support for the terrorist's cause, and/or massive U.S. retaliation. However,

others may think that the adversary would prefer attacks that would cause more fatalities in order to evoke more fear.

To model this, CREATE (2011) proposes extending the standard additive multiattribute utility function in (4.1) to allow for negative weights, while restricting the sum of the absolute values of the various attribute weights to equal one (i.e., $\sum_{m=1}^{M+1} |w_m| = 1$); see also Wang and Bier (2013c). This choice of attribute weights allows enough flexibility to deal with expert disagreement on the direction of effect of the various attacker attributes, but still constrains the magnitudes of these weights.

Some attributes that seem totally unimportant when using only non-negative weights may turn out to be surprisingly important after we drop the non-negativity restriction. For example, if targets that were dominated by others on the known attributes using their original scales are judged to be quite attractive, the original model could account for this only by large weights on the unobserved attribute, whereas the new model may be able to account for that by negative attribute weights, reducing the weight on the unobserved attribute. Allowing for negative weights handles possible conflicting views on the direction of effect for each attribute in an automatic and realistic way, which is especially important for elicitation tasks with large numbers of experts (e.g., using on-line surveys).

In the next section, we use hypothetical expert judgments to illustrate the application of PI and BDE, with a focus on the inclusion of unobserved attributes and negative attribute weights.

4.6 Hypothetical Case Study

4.6.1 Sample Data

We now present a case study on elicitation of adversary preferences about major U.S. urban areas using PI and BDE. In particular, we consider 20 major U.S. urban areas with the highest expected damage from terrorism (according to Willis et al. 2005), including: New York City (NYC); Chicago; San Francisco; Washington, DC; Los Angeles (LA); Philadelphia; Boston; Houston; Newark; Seattle; Jersey City; Detroit; Las Vegas; Oakland; Orange County; Cleveland; San Diego; Miami; Minneapolis; and Denver. We then use hypothetical expert ordinal judgments based on two known attacker attributes: expected yearly property loss from terrorism and population density (Willis et al., 2005). The original values g_{nm} of each attacker attribute m (m=1, 2) for each city n ($n = 1, \ldots, 20$) are presented in Table 4.3.

Following Wang and Bier (2012), we assume that the attacker's single-attribute utility a_{nm} of attribute m for city n is proportional to $\ln \frac{g_{nm}}{\underline{g}_m}$ (where $\underline{g}_m = \min_n g_{nm}$), or in other words how much more attractive city n is than the least desirable city on attribute m. The utilities are normalized so that a_{nm} = 1 when city n is the most attractive target on attribute m. This choice of attacker utility is also consistent with Fechner's law, which states that human perceptions are typically logarithmic in the magnitude of the original stimuli (Fechner 1860). For example, if the expected property loss in Chicago is doubled from 115 to 230 million dollars, the attacker's single-attribute utility increases by only an additive increment proportional to $\ln(2)$.

We choose the same starting probability measure Q_0 for both PI and BDE in the following way. The three attacker attribute weights (including the weight on the unobserved attribute)

Table 4.3: Attribute Values for U.S. Cities with Highest Expected Terrorism Losses

	Property Loss ($ million), $g_{\cdot 1}$	Population Density (per sq mile), $g_{\cdot 2}$
NYC	413	8,159
Chicago	115	1,634
San Francisco	57	1,705
DC	36	756
LA	34	2,344
Philadelphia	21	1,323
Boston	18	1,685
Houston	11	706
Newark	7.3	1,289
Seattle	6.7	546
Jersey City	4.4	13,044
Detroit	4.2	1,140
Las Vegas	4.1	40
Oakland	4	1,642
Orange County	3.7	3,606
Cleveland	3	832
San Diego	2.8	670
Miami	2.7	1,158
Minneapolis	2.7	490
Denver	2.5	561

are assumed to follow the Dirichlet distribution $(1, 1, 1)$, and the utilities of the unobserved attribute for the various targets are independently uniformly distributed over $[0, 1]$, and are also independent of the attribute weights. Again, we set the defender's self-trust $\alpha \to 0$.

We first consider two groups of hypothetical experts who give partial rank orderings to these urban areas (as shown in Table 4.4). In particular, we assume that both groups I and II have five experts, each of whom ranks the top ten cities without ties, reflecting their knowledge about the attacker intent.

Table 4.4: Hypothetical Expert Rank Orderings (Groups I and II)

Expert	Top Ten Cities (in That Order)
	Group I
1	NYC, Chicago, San Francisco, LA, DC, Jersey, Boston, Phil., Houston, Newark
2	NYC, Chicago, LA, San Francisco, Jersey, Phil., Boston, DC, Newark, Orange
3	NYC, Jersey, Chicago, San Francisco, Phil. DC, LA, Boston, Orange, Houston
4	NYC, Jersey, Chicago, San Francisco, LA, Phil., DC, Orange, Newark, Boston
5	NYC, Jersey, Chicago, LA, San Francisco, Boston, Phil., DC, Orange, Newark
	Group II
1	NYC, Chicago, San Francisco, LA, DC, Phil., Boston, Houston, Newark, Jersey
2	NYC, Chicago, LA, San Francisco, DC, Phil., Boston, Houston, Newark, Jersey
3	NYC, Chicago, LA, San Francisco, DC, Phil., Boston, Houston, Las Vegas, Newark
4	NYC, Chicago, LA, San Francisco, DC, Phil., Boston, Houston, Newark, Las Vegas
5	NYC, Chicago, LA, San Francisco, DC, Phil., Boston, Houston, Newark, Seattle

4.6.2 Elicited Probability Distributions for Attacker Attribute Weights

We now apply PI and BDE to get probability distributions for the various attacker attribute weights (including the weight for the unobserved attribute) using hypothetical judgments of groups I and II. Results are shown in Figure 4.4.

Experts in group I are generally assumed to think that the attacker would put similar weight on property loss and population density. Therefore, the mean elicited weights on the two known attributes are shown to be close to each other for both PI and BDE (about 0.42 for property loss versus about 0.47 for population density, as shown in the top panel in Figure 4.4).

Note that expert judgments in group I are also correlated, in the sense that experts rank Jersey City higher than Chicago also tend to rank Orange County higher than Newark (and

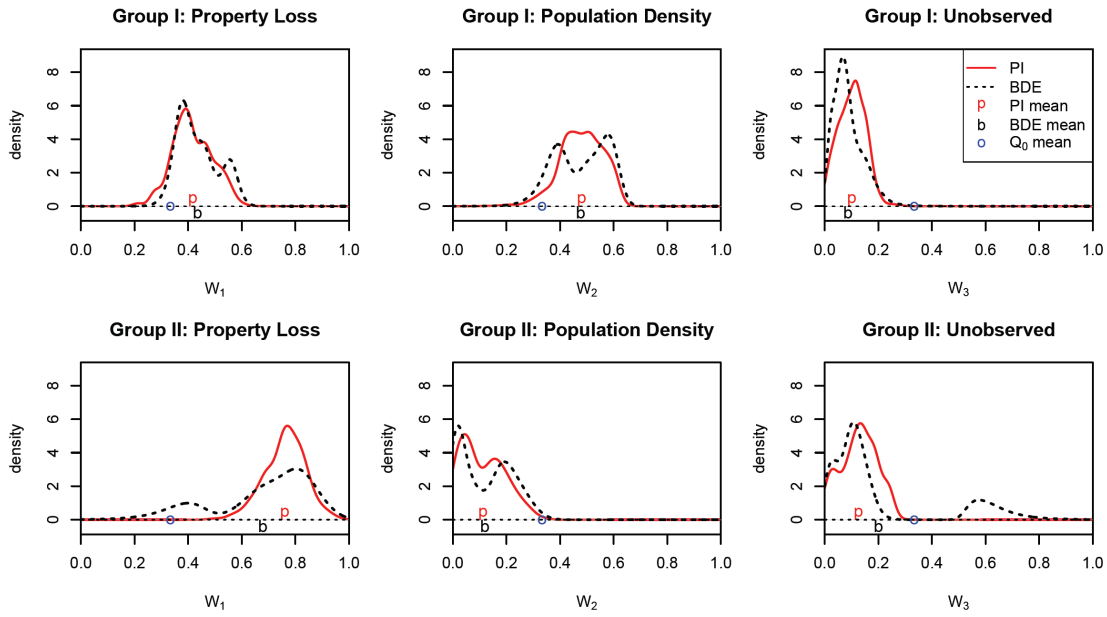

Figure 4.4: Elicited Marginal Probability Densities over Attribute Weights (Groups I and II)

vice versa). According to the attribute values given in Table 4.3, Jersey City and Orange County are rated distantly lower on property loss whereas distantly higher on population density than Chicago and Newark, respectively. Therefore, an expert who ranks Jersey City higher than Chicago (and also Orange County higher than Newark) sends a relatively clear signal of placing high weight on population density. On the other hand, if an expert thinks that Chicago is more attractive than Jersey City (and also Newark is more attractive than Orange County), it suggests that the particular expert thinks the adversary would give less weight to population density than to property loss. Hence, disagreement on the relative attractiveness of Chicago versus Jersey City (and also Newark versus Orange County) can be seen as evidence of disagreement on the weight assigned to population density.

BDE visualizes the two schools of thought by a bi-modal probability distribution for the weight on population density. However, PI fails to capture such disagreements (represented by correlations among subgroups of experts), resulting in a probability distribution that concentrates its mass in the middle between the differing expert views.

Experts in group II are assumed to rank targets mainly according to their ratings on property loss. Therefore, the mean elicited weight on property loss is shown to be pretty high (0.74 for PI, and 0.68 for BDE, as shown in the lower panel of Figure 4.4). However, these experts do not agree perfectly with each other, especially on the importance of population density and the presence of any unobserved attribute(s). For example, some experts think that Jersey City needs to be included in the top ten most attractive cities, indicating that population density of a city would affect the attacker's choice. However, most other experts rank Las Vegas in the top ten, suggesting the need to include some unobserved factors (other than property loss and population density) that can make Las Vegas more attractive, because Las Vegas is among the least attractive cities on both of the known attributes.

Results show that both PI and BDE give higher mean weight on the unobserved attribute for group II than for group I. Moreover, BDE further identifies some experts' strong feeling about the importance of some unobserved attribute (by generating a separate peak corresponding to high weight on the unobserved attribute, and low weight on property loss, accordingly). Note that in this case, PI also gives multi-modal distributions in the face of disagreement between experts. In Section 4.8, we further conduct simulation-based sensitivity analysis to explore the general tendency of PI or BDE to generate multi-modal distributions in the face of expert disagreement.

Note that results by using these indirect elicitation methods may be better than direct weight elicitation. For example, experts who are asked to provide attribute weights directly may place inappropriately high weight on some attribute (e.g., population density), without realizing that Jersey City has a population density much higher than that of larger cities like New York, or Los Angeles. Experts who are given the opportunity to rank the attractiveness of targets directly would be unlikely to make that mistake, since they presumably might not want to rank Jersey City as a top terrorist target; therefore, mathematically inferring attribute weights from expert judgments of target attractiveness might give more realistic results. By extrapolating, the elicited attribute weights could then be used to assess the attractiveness of a much larger number of targets than had been directly evaluated by the experts.

4.6.3 Interpretation of Unobserved Attributes

We have shown that the mean weight on the unobserved attribute is higher for group II than for group I, suggesting that the two known attacker attributes (property losses and population density) are less suited to capturing the rank orderings provided by experts in group II. With the inclusion of unobserved attributes and the explicit probability distributions over the utilities of the unobserved attribute for the various cities, we can further speculate about the nature of possible unobserved attribute(s).

Figure 4.5 shows the elicited probability densities over the utilities of the unobserved attribute for the 20 U.S. major cities. These utilities were initially assumed to be represented by independent uniform distributions on $[0, 1]$, before we observe any expert input. For each city whose utility on the unobserved attribute is significantly increased (or decreased) by

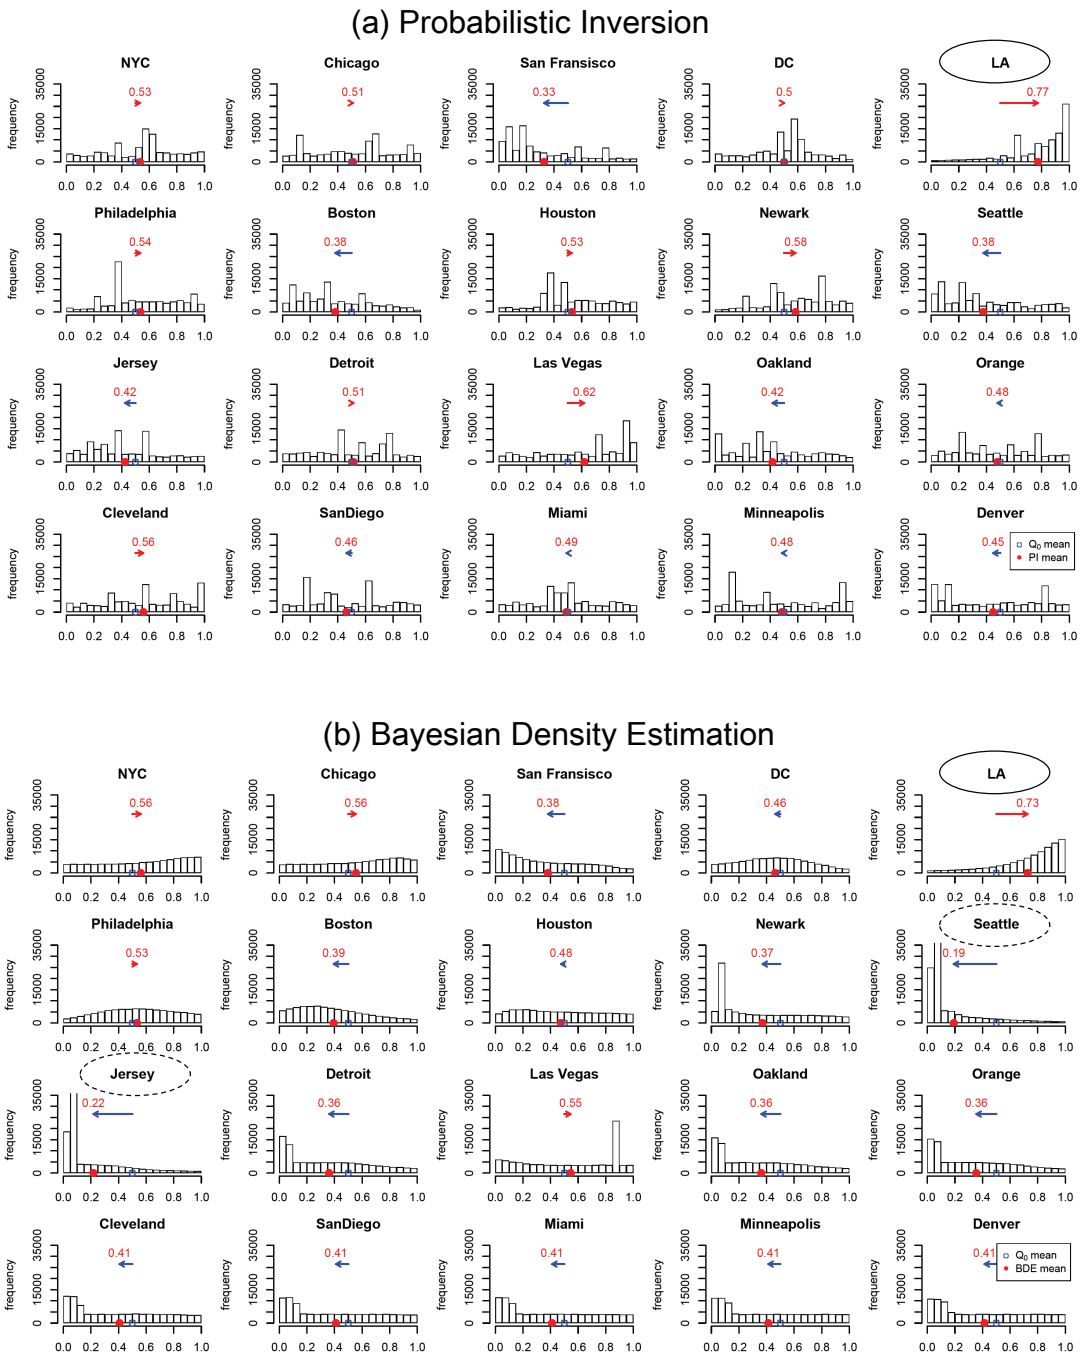

Figure 4.5: Elicited Marginal Probability Densities over Utilities of the Unobserved Attribute (Group II)

either PI or BDE, we use an arrow (starting from 0.5, and ending at the PI or BDE means) to show how the utility of that city on the unobserved attribute was affected by the expert opinions of group II.

Both PI and BDE suggest that LA is rated higher by experts in group II than their known attribute values would suggest. BDE further shows that Seattle and Jersey City are expected to have a low utility on the unobserved attribute. This indicates the desirability of including an additional attacker attribute on which LA is rated high while Seattle and Jersey City are rated low. For example, the presence of a large entertainment industry might be a good candidate. We could then include such an attribute in our analysis, and run PI or BDE again to see whether the weight on the unobserved attribute can be reduced or not. It would be ideal if the consideration of entertainment industry is able to make the updated set of attacker attributes sufficiently capable of explaining the expert judgments. Otherwise, there is a need to seek other possible unobserved attribute(s).

4.6.4 Allowing for Negative Attribute Weights

We now consider two additional groups of hypothetical experts (groups III and IV) to illustrate how the incorporation of negative attribute weights can automatically reflect conflicting views on the direction of effect for each attacker attribute. Table 4.5 present the hypothetical judgments of the two groups. Results of applying PI and BDE are shown in Figure 4.6.

In group III, all experts rank the four cities with the highest property loss as the most attractive cities (e.g., NYC, Chicago, San Francisco, and DC). This indicates a high attacker weight on the attribute of property loss. Note that Las Vegas is also ranked in the top ten

Table 4.5: Hypothetical Expert Rank Orderings (Groups III and IV)

Expert	Top Ten Cities (in That Order)
	Group III
1	NYC, Chicago, San Francisco, DC, LA, Las Vegas, Phil., Houston, Boston, Newark
2	NYC, Chicago, San Francisco, DC, LA, Phil., Boston, Houston, Seattle, Las Vegas
3	NYC, Chicago, San Francisco, DC, Phil., LA, Boston, Houston, Seattle, Las Vegas
4	NYC, Chicago, San Francisco, DC, Boston, LA, Las Vegas, Phil., Houston, Seattle
5	NYC, Chicago, DC, San Francisco, Las Vegas, LA, Boston, Phil., Houston, Seattle
	Group IV
1 – 5	Same as experts in group II
6 – 10	Same as experts in group III

by all experts in group III. When negative weights are not allowed, its low ratings on both property loss and population density suggest that Las Vegas would never be given such a high rank. The elicitation method then must place high weight on the unobserved attribute, and further seek utilities for the unobserved attribute for each target that make Las Vegas more attractive.

However, it is possible that some experts may regard a particular adversary to be averse to attack a densely populated city – e.g., if the attacker wishes to cause significant economic damage and make a big splash but without large numbers of fatalities. By allowing for negative attribute weights, we get a high probability of a weight on population density using the judgments of group III (with probability 1 for PI, and 0.86 for BDE; upper panel of Figure 4.6). As a result, if we reverse the scale of population density (so that Las Vegas can be highly attractive), then using only the known attributes is able to explain the expert judgments in group III well, reducing the demand for any unobserved attributes.

Group IV is constructed by putting together all the experts from group II and group III. Since some experts in group II were shown earlier to agree with the original scale

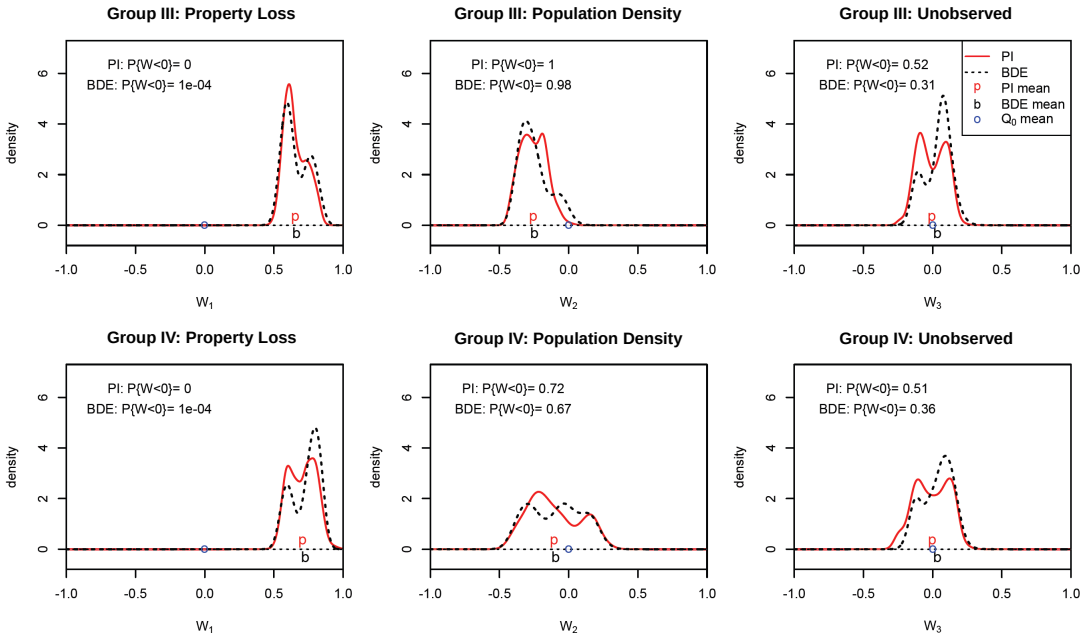

Figure 4.6: Elicited Marginal Probability Densities over Attribute Weights (Groups III and IV)

of population density in the previous analysis, and are now included in group IV, there are conflicting views on the scale of population density within group IV. Therefore, the elicited weight on population density for group IV has a broad distribution encompassing both negative and positive values (lower panel of Figure 4.6). The fact that the weight on population density takes on such a broad range of values demonstrates the disagreement within group IV on whether a larger population density means higher or lower attractiveness to the adversary.

4.7 Realistic Application

We have successfully applied probabilistic inversion to adversary preference elicitation in CREATE (2011). In that project, "proxy" experts (graduate students knowledgeable about terrorism, from countries where support for terrorism is relatively common) were asked to rank eight attack scenarios based on their attractiveness to the adversary, where attack scenarios are characterized by seven known attributes plus an unknown attribute. Figure 4.7 shows how the expected utilities of the various scenarios differ depending on whether a proxy's judgments were elicited using PI, or by direct elicitation using the random utility method of Rossof and John (2009). Both methods identified the same three least attractive scenarios (Pneumonic Plague, Dirty Bomb, and Blister Agent), and assigned relatively high utilities to another three scenarios (Chlorine Tank Explosion, Improvised Explosive Device, and Food Contamination). The only discrepancies are in Nerve Agent and Aerosol Anthrax (which were rated high using PI, but much lower using direct elicitation). However, we could take advantage of such discrepancies as input for convergent validation – e.g., by asking the proxy expert whether he puts more credence in his scenario rankings or his assessed attribute weights.

Moreover, the results in CREATE (2011) also suggest that applying PI to partial rather than full rank orderings can give reasonably reliable results. Figure 4.8 compares the results of PI using the complete set of eight scenario rankings given by one proxy expert, versus only four rankings (the top three most attractive scenarios and the least attractive scenario). The two sets of expected utilities are quite close. Similar results are obtained in the simulation-based sensitivity analysis in Section 4.8, supporting the idea that attribute

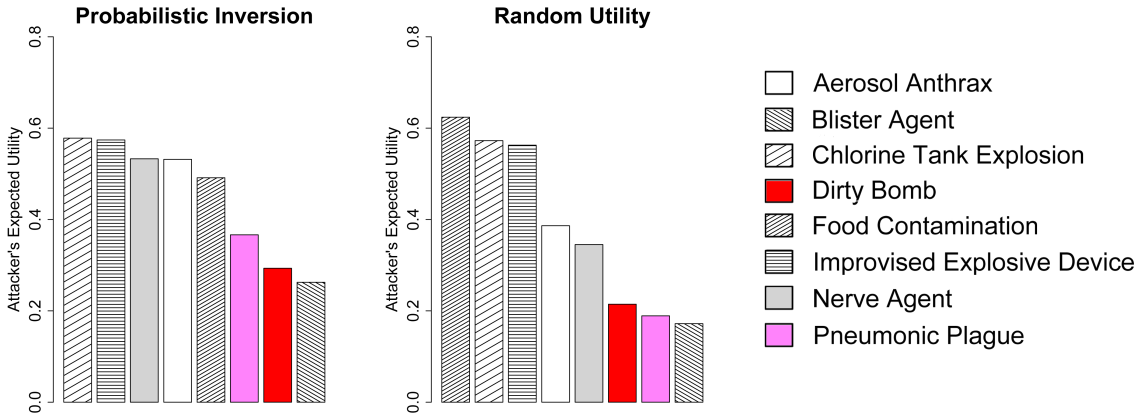

Figure 4.7: Expected Utilities of Attack Scenarios Obtained by PI vs. Direct Elicitation

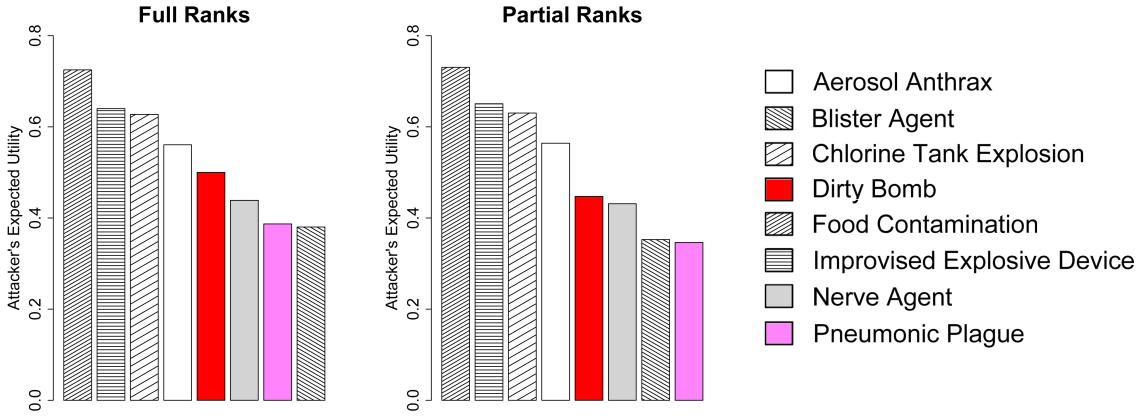

Figure 4.8: Expected Utilities Based on Full vs. Partial Ranks

weights can be estimated by asking experts to rank only a modest subset of alternatives, with no need to rank all alternatives.

4.8 Sensitivity Analysis

In this section, we conduct simulation-based sensitivity analysis to explore the behavior of both PI and BDE in the face of expert consensus or disagreement, and to also investigate the reliability of results obtained by using partial rather than full rank orderings. In particular, we attempt to answer the following three questions: (1) If there are different "schools of thought" between experts (i.e., subgroups of experts who hold similar views), do the elicitation methods tend to generate multimodal probability distributions, or do they generate distributions that assign most of their mass in the middle between the differing expert views? (2) Do the elicited probability distributions given by these methods adequately reflect the level of consensus or disagreement among the rank orderings given by the experts? (3) Can partial rank orderings of targets generate reliable results, compared to full rank orderings?

4.8.1 Tendency to Generate Multimodal Distributions

Hypothetically, one might speculate that whether PI and BDE will tend to produce multimodal distributions for the adversary parameters would depend on how far apart the differing expert views are from each other. Moreover, PI and BDE are anticipated to behave most differently when subsets of the ordinal judgments provided by the experts are correlated (e.g., experts ranking target 1 higher than target 2 would also rank target 3 higher than target 4, and vice versa), since in that case, PI will not be able to capture such correlations. We now conduct a Monte Carlo-based sensitivity analysis to test these hypotheses.

Consider a case where four experts are asked to rank six targets described by two known adversary attributes (with values of the known attributes given in Table 4.6). This choice of

Table 4.6: Values of Adversary Attributes (a_{nm}) for the Sensitivity Analysis

	Attribute 1	Attribute 2
Target 1	1	0
Target 2	0	1
Target 3	0.75	0.25
Target 4	0.25	0.75
Target 5	0.55	0.45
Target 6	0.45	0.55

problem scale is complicated enough to illustrate the question of interest reasonably well, and yet computationally inexpensive. In addition, we choose a flat starting measure for both PI and BDE.

We now randomly simulate expert rank orderings as input for our analysis. We first introduce a set of random variables u_{kn} ranging between 0 and 1 to reflect the utility of target n according to expert k (k = 1, …, 4; n = 1, …, 6), and then derive rank orderings from the randomly generated target utilities u_{kn}. In this way, we can induce a dependency structure among the u_{kn} by adopting the Gaussian copula with desired levels of pairwise correlations (Bier and Yi 1995; Clemen and Reilly 1999; Hora 2010). In particular, we assume that experts 1 and 2 and experts 3 and 4 form two different schools of thought, and set the Pearson correlation coefficients according to

$$\mathbf{cor}[u_{1n}, u_{2n}] = \mathbf{cor}[u_{3n}, u_{4n}] = |\gamma| \text{ ; and}$$

$$\mathbf{cor}[u_{1n}, u_{3n}] = \mathbf{cor}[u_{1n}, u_{4n}] = \mathbf{cor}[u_{2n}, u_{3n}] = \mathbf{cor}[u_{2n}, u_{4n}] = \gamma, \text{ for } n = 1, \ldots, 6$$

where $\gamma \in (-1, 1)$ controls the "similarity" of the two differing schools of expert judgments. We also set the Pearson correlation coefficients between the utilities of targets 1 and 3 and

targets 2 and 4 for a given expert to both equal $\rho \in (0, 1)$; i.e.,

$$\mathbf{cor}[u_{k1}, u_{k3}] = \mathbf{cor}[u_{k2}, u_{k4}] = \rho, \text{ for } k = 1, \ldots, 4$$

where higher ρ means that experts who rank target 1 higher than 2 are likely to rank target 3 higher than 4, and vice versa. We can then construct a valid correlation matrix (i.e., symmetric and positive definite with all diagonal elements equal to one) for the target utilities u_{kn} that satisfies both of the above conditions (correlations between experts and between targets).

For each level of expert similarity γ and target correlation ρ, we randomly generate 500 sets of values for the u_{kn}, and derive the rank orderings accordingly. Applying either PI or BDE, we obtain 500 elicited distributions for the first known attribute weight W_1, and count the occurrence of multimodal distributions. Figure 4.9 shows the proportions of multimodal distributions (out of 500) for different choices of γ and ρ.

As we expected, the elicited distributions are more likely to have multiple peaks when the two subgroups of expert judgments are farther apart (corresponding to more negative values of γ), using either PI and BDE on a flat starting measure. However, there is a less than 30% chance of multimodal distributions using either PI or BDE, even when the two differing schools of expert views are almost opposite of each other (γ = -0.95).

Moreover, for a fixed level of expert similarity γ, more highly correlated rank orderings (corresponding to larger values of ρ) generally lead to more frequent occurrence of multimodal distributions. This effect is more pronounced for BDE than for PI, as expected. In general, BDE tends to give more multimodal distributions than PI. This tendency is especially significant when target rankings given by each expert are highly correlated (i.e., ρ = 0.95).

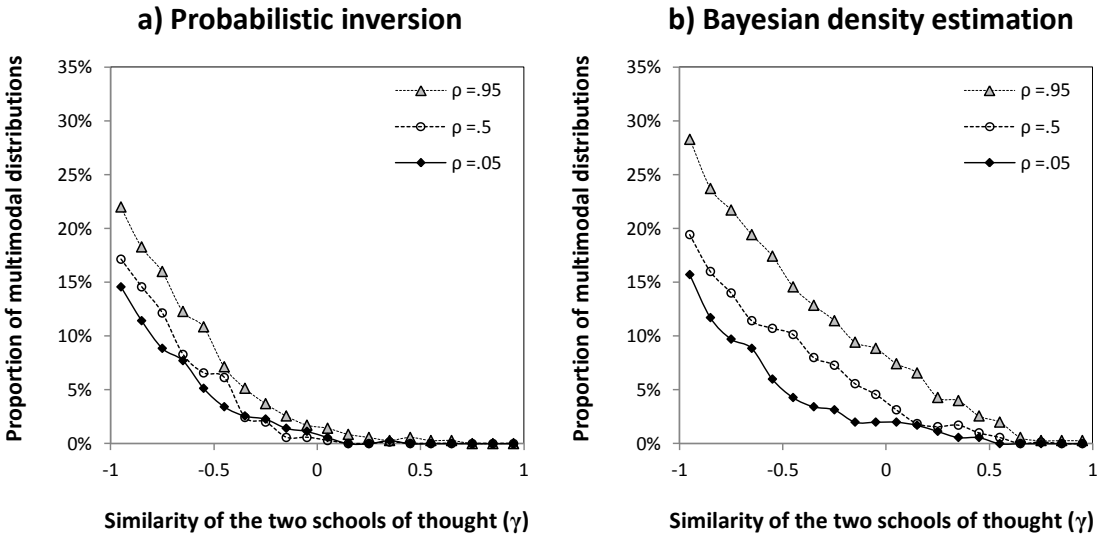

Figure 4.9: Proportions of Multimodal Distributions Resulting from PI vs. BDE

Whether we actually want to see a multimodal distribution from elicitation may depend on the problem under investigation. For example, we may prefer multimodal distributions when a large number of experts seem to form differing subgroups (in which similar views are held), because we might be fairly confident that any new expert would then give judgments that fall into some existing school of thought. By contrast, if a small number of experts disagree with each other, we may not want the elicited distributions to be too sensitive to differences between their judgments.

4.8.2 Expert Disagreement and Dispersion of Elicited Distributions

We now discuss another important issue. Ideally, probability distributions provided by a good elicitation method should adequately reflect the level of consensus or disagreement among

the rank orderings given by experts. We therefore conduct another Monte Carlo-based sensitivity analysis to explore whether higher levels of disagreement between the experts lead to broader probability distributions for the attribute weights.

In particular, we consider a case where two experts are asked to give full rank orderings to ten targets described by one known adversary attribute. We assume that the target values on the known attribute ($a_{\cdot 1}$) form an arithmetic series with maximum 1 and minimum 0. Target utilities u_{kn} according to the two experts (k = 1, 2; n = 1, …, 10) are then randomly generated, from which rank orderings are derived. Dependency between judgments of the two experts is induced by setting the Pearson correlation coefficients between the target utilities as

$$\mathbf{cor}[u_{1n}, u_{2n}] = \gamma, \text{ for } n = 1, \ldots, 10$$

where $\gamma \in (-1, 1)$ again controls the level of agreement between the two experts. The unrestricted correlations are then properly set to ensure that the correlation matrix for u_{kn} is valid.

We use the normalized variance to measure dispersion of the elicited distributions over the adversary attribute weights. Note that the normalized variance for a random variable $X \in [0, 1]$ is defined as $\mathbf{NV}[X] = \dfrac{\mathbf{Var}[X]}{\mathbf{E}[X](1 - \mathbf{E}[X])}$ (Bier and Yi 1995). In particular, we have $\mathbf{Var}[X] \leq \mathbf{E}[X](1 - \mathbf{E}[X])$ for $X \in [0, 1]$, with equality achieved when $\mathbf{E}[X]$ = 0 or 1, so $\mathbf{NV}[X]$ gives variance as a fraction of its maximum value.

We randomly generate 200 sets of target utilities u_{kn} that satisfy the correlation requirements, with expert rank orderings derived accordingly. This number of simulation runs seems reasonable, since the simulation errors for the quantity of interest (e.g., the normalized variance of the known attribute weight $\mathbf{NV}[W_1]$) are always less than $\pm 5\%$.

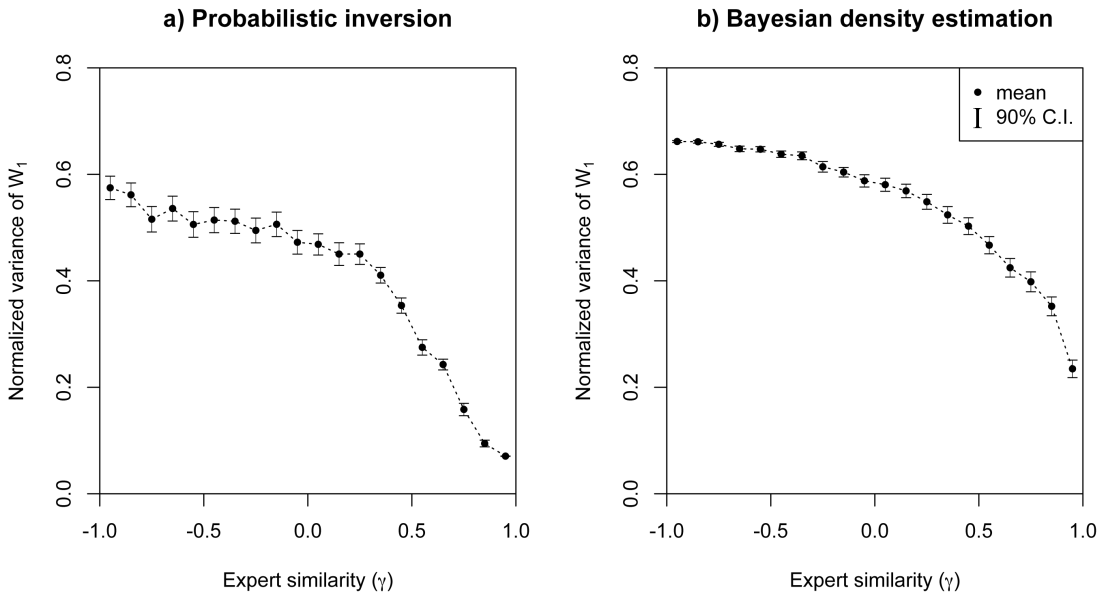

Figure 4.10: Average and Confidence Interval of $\mathbf{NV}[W_1]$ Resulting from PI vs. BDE

Figure 4.10 shows the average and the 90% confidence interval of the normalized variance $\mathbf{NV}[W_1]$ from both PI and BDE for varying levels of expert similarity γ, again adopting a flat starting measure.

In general, both PI and BDE conform to the predicted trend of generating probability distributions with higher normalized variance for higher levels of disagreement between experts. However, such a trend is less significant when experts strongly disagree with each other (i.e., $\gamma < 0$) than when they give similar judgments (i.e., $\gamma > 0$). In other words, the breadth of the elicited probability distributions does decrease with expert similarity, but is less sensitive to the precise degree of similarity when experts are highly likely to give opposite judgments.

4.8.3 Reliability of Partial Rankings

The elicitation methods discussed here are practical only if experts can provide partial rank orderings of the potential targets, rather than a full ranking of all targets, which may become cumbersome. In fact, there would be little reason to even obtain estimates of attribute weights, if experts had to provide direct rankings of all possible targets in order to generate those estimates. Therefore, we conduct Monte Carlo simulation-based sensitivity analysis to investigate whether elicitation results from partial rank orderings are sufficiently reliable for use in practice.

For simplicity, we apply PI and BDE to the case of a single expert, without allowing for negative attribute weights. We also choose a flat starting distribution, and set the self-trust degree $\alpha \to 0$. (Note again that when there is only one expert, PI and BDE give the same results.)

We consider three main factors that might affect the reliability of partial rank orderings. First, for a fixed number of targets, we expect that ranking more of the targets (i.e., larger R) would generate more reliable results than ranking fewer targets. However, it is important to know how quickly results become reliable as the number of ranked targets increases. Second, the total number of targets (N) also seems to be an important factor. However, we have no prior hypothesis on whether results would be more reliable for a larger or smaller total number of targets. Finally, we are interested in how the number of known attributes (M) affects the reliability of partial rank orderings. Partial rankings should hypothetically be less reliable for estimating large numbers of attribute weights, since in that case there are simply more parameters to estimate. Moreover, as the number of attributes grows, the various attributes would be more likely to be correlated, possibly leading to less stable

elicitation results. Therefore, when the total number of attributes is greater, we expect that results derived from partial rank orderings would deviate more from those based on full rank orderings.

In particular, we consider three levels for the total number of targets (N = 10, 20, or 50), and three levels for the number of known attributes (M = 2, 4, or 7). For each number of targets N, we compare different levels of partial ranking R against full ranking of all N targets. When N = 10, we consider one partial rank ordering identifying the top R = 5 targets; when N = 20, we consider two partial rank orderings, the top R = 5 or 10 targets; when N = 50, we consider three cases, the top R = 5, 10, or 20 targets.

Here, reliability is measured by the Euclidean distance between the (mean) attribute weights elicited using full rank orderings and those obtained from partial rank orderings. For ease of interpretation, we normalize this distance by dividing by the square root of 2, which is the largest possible Euclidean distance between two non-negative weights that sum up to unity, to yield a maximum normalized distance of 1. (Of course, the ideal distance is 0.)

For each of the above cases, we randomly generate 400 sets of attribute values and rank orderings of targets. 400 runs is sufficient to ensure that the simulation error is within $\pm 5\%$ of the quantity of interest. Figure 4.11 presents the estimated normalized distances between the attribute weights elicited using partial versus full rank orderings in our simulations.

As we expected, larger subsets of partial rankings (i.e., larger R) lead to more reliable estimates for the attribute weights (smaller normalized distances). In general, PI and BDE are practical to use, since ranking roughly 50% of the targets (e.g., the top 5 of 10, the top 10 of 20, or the top 20 of 50) can yield moderately reliable results (with a normalized distance

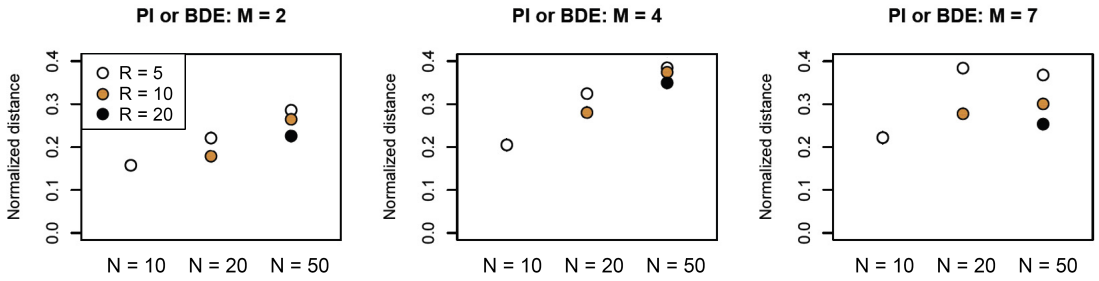

Figure 4.11: Normalized Distance between Estimated Weights Using Full Rankings vs. Partial Rankings

of less than 0.3) in most cases. However, if we fix the proportion of ranked targets at around 50%, the reliability of results using PI or BDE would be reduced as the total number of targets increases.

In addition, PI and BDE seem to work better for smaller numbers of attributes (i.e., smaller M). For example, when there are only two known attributes ($M = 2$), asking the experts to provide only the top 5 out of 50 targets (i.e., the top 10%) would generally produce quite reliable estimates for the attacker attribute weights.

4.9 Summary

In this chapter, we develop a simple expert elicitation process to generate probability distributions for uncertain adversary intent using only partial rank orderings of target attractiveness provided by experts. To accomplish this task, we propose to use two recently adapted mathematical methods, probabilistic inversion and Bayesian density estimation.

Rank orderings are believed to be easy to provide and interpret, so these methods may be quite useful in eliciting judgments of experts who do not have highly quantitative backgrounds. Moreover, they can serve not only as alternatives to direct elicitation, but also as a source of inputs for convergent validation, a common technique in decision analysis. For example, we could ask the experts to comment on any observed discrepancies – whether they put more credence in their rankings of targets or in their directly assessed attribute weights.

The simplicity and computational ease of the proposed elicitation process also makes it promising for large-scale elicitation tasks that involve many elicitees with non-quantitative backgrounds. For example, it is suitable for use in quantifying customer preferences in marketing, using on-line surveys.

Chapter 5

Extension to Account for Adversary Capabilities

This chapter extends the basic game in Chapter 3 to account for adversary capabilities, in addition to just intent. In Section 5.1, we start from a game of complete information to introduce an analytical representation of the effects of attacker capabilities, through the use of a contest success function from economics. Next, Section 5.2 extends that game to incorporate defender uncertainty about both attacker intent and capabilities. Sections 5.3 further allows the attacker to have multiple types of capabilities (such as money, personnel, technological sophistication, and flexibility of movement). In Section 5.4, the attacker is not only allowed to attack multiple targets simultaneously, but also to vary his levels of effort among those targets. Section 5.5 presents a case study to illustrate the applicability of our model, using data from a recent project (CREATE 2011) as well as other open-source information on terrorist capabilities.

5.1 Game of Complete Information

In addition to attacker intent (i.e., target valuations U_n), we also use level of capability A to describe the attacker's characteristics. Figure 5.1 illustrates the assumed decision processes of the defender and the attacker in the face of an intentional attack.

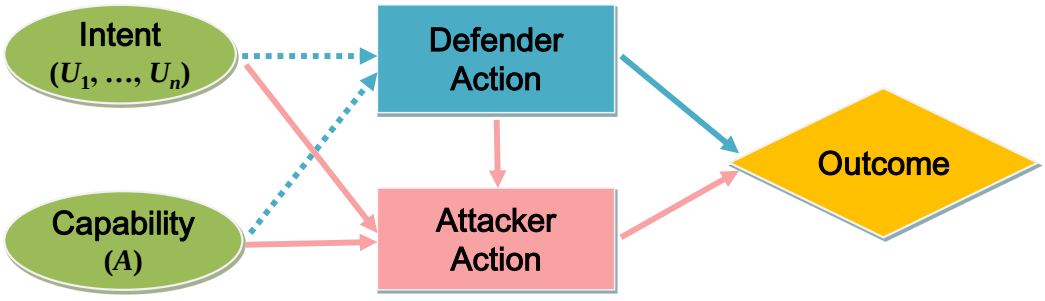

Figure 5.1: Influence Diagram Showing the Defender and Attacker Decisions

We first consider a game of complete information where the defender is fully knowledgeable about the attacker's intent and capabilities. Of course, in reality, those characteristics may not be fully known to the defender (as shown by the dashed arrows in the influence diagram of Figure 5.1); models that take into account such uncertainties are explored Section 5.2.

5.1.1 Modeling Adversary Capabilities

We assume that the defender moves first and allocates a fixed budget B among the N potential targets in such a way as to minimize her expected loss, as given by

$$\min_{x_1+\ldots+x_N \leq B} \sum_{n=1}^{N} p_n(x, A, U) s_n(x_n, A) v_n \tag{5.1}$$

where

N = number of targets

x_n = defensive resource allocated to target n; and $x = (x_1, \ldots, x_N)$

B = defensive budget

A = attacker capability

v_n = defender's valuation of target n (consequence); and $v = (v_1, \ldots, v_N)$

$s_n(x_n, A)$ = success probability of an attack on target n (vulnerability)

U_n = attacker's valuation of target n; and $U = (U_1, \ldots, U_N)$

$p_n(x, A, U)$ = likelihood of an attack against target n (threat)

Values of B, A, v, and U, and the functional forms of $s_n(\cdot, \cdot)$ are all assumed to be common information. For now, we assume that the two agents interact in a zero-sum game; i.e., the attacker's payoff always equals the defender's loss (with $U_n = v_n$ for all n). We also assume that the attacker can attack only a single target. (Later, Section 5.2 allows U_n and v_n to differ, and Section 5.4 allows for multiple simultaneous attacks.)

For a given attacker capability A and intent U (equal to v), as well as the defensive resource allocation x, the likelihood of an attack on target n (for $n = 1, \ldots, N$) is then

given by

$$p_n(x, A, U) = p_n(x, A, v) = \begin{cases} \frac{1}{Z} & \text{if } n \in \arg\max_i s_i(x_i, A)v_i \\ 0 & \text{otherwise} \end{cases} \tag{5.2}$$

where Z is the cardinality of the set $\{i : \arg\max_i s_i(x_i, A)v_i\}$.

The novel aspect of this model is the explicit consideration of the attacker's capability A in determining terrorism risks. We assume that the vulnerability of target n is determined not only by the amount of defensive investment spent on that target x_n, but also by the attacker's capability A. In particular, we borrow the idea of a contest success function from the rent-seeking model in economics (Skaperdas 1996; Hausken and Bier 2011), and define the success probability of an attack on target n by

$$s_n(x_n, A) = \frac{A^{\beta_n}}{A^{\beta_n} + (x_n)^{\beta_n}} = 1 - \frac{(x_n)^{\beta_n}}{A^{\beta_n} + (x_n)^{\beta_n}} \tag{5.3}$$

Thus, given a fixed level of defensive investment x_n in target n, a more capable attacker (with a larger capability A) will be able to launch a more effective attack on target n, causing more expected loss to the defender.

Note that the attacker capability A and the defensive investment B need to be measured in the same units. For example, if the defensive budget B is measured in millions of dollars, then the attacker capability A needs to also be represented in monetary form and measured in millions of dollars. However, we could of course introduce a positive scalar γ to achieve equivalence between different types of defense or attack resources; e.g., by specifying that every additional trained attacker is equivalent to a $\$\gamma$ increase in defensive funding. More

generally, Section 5.3 will discuss a generalized contest success function that accounts for multiattribute attacker capabilities.

Note also that the contest success function in (5.3) is "homogeneous," in the sense that if the attacker capability A and the defender investment x_n are multiplied by the same factor, then the success probability $s_n(x_n, A)$ of an attack on target n will remain unchanged (following Skaperdas 1996). Therefore, it is the ratio of the attacker and defender efforts $\frac{A}{x_n}$ that matters in determining the vulnerability of each target.

In addition, the contest success function in (5.3) uses a "decisiveness" parameter $\beta_n > 0$ to reflect the extent to which the success of an attack can be attributed to the relative effort of attack over defense rather than luck. Figure 5.2 illustrates how the level of decisiveness β_n affects the success probability of an attack on target n for different values of the relative effort $\frac{A}{x_n}$. As for the shape of the contest success function, it is easy to show that $s_n(x_n, A)$ is strictly concave in the attacker effort A and strictly convex in the defensive investment x_n if $0 < \beta_n \leq 1$, and is S-shaped in A and reverse S-shaped in x_n if $\beta_n > 1$. As β_n goes to infinity, $s_n(x_n, A)$ becomes a step function of both A and x_n.

For $\beta_n \to 0$, the success of an attack is purely attributable to luck. Both the attacker and defender have an equal chance to win the contest, regardless of the efforts expended. For moderate values of β_n (e.g., $\beta_n = 0.6$), the attacker will still have a good chance of destroying target n even if he does not have a capability advantage over the defender; for example, even an attacker with low capability may get lucky and succeed with an IED attack on a soft civilian target. By contrast, for large values of β_n (e.g., $\beta_n = 8$), whichever agent (attacker or defender) devotes more resources to target n is highly likely to succeed at destroying or protecting that target. For example, only highly capable attackers who possess

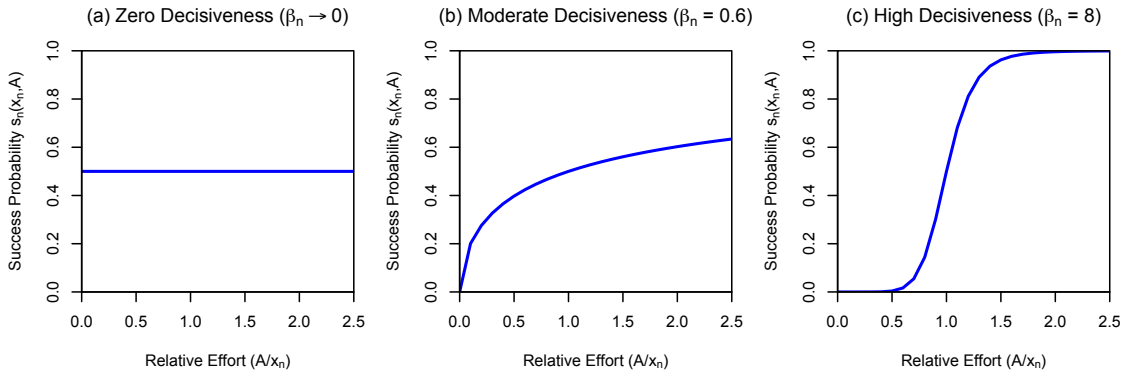

Figure 5.2: Effect of Decisiveness on Target Vulnerability

abundant resources are likely to be able to conduct a successful attack on a hardened military target, while diligent efforts by the defender would be required to foil such an attack.

5.1.2 Defender's Optimal Strategy

We are now ready to derive the optimal defensive strategy when the defender has complete information about the attacker's capability (A) and intent ($U_n = v_n$), and the decisiveness of an attack (β_n) on each of the targets. Proofs of the lemmas, propositions and corollaries in this chapter are provided in Appendix D.

Proposition 5.1. *If the N targets are rank ordered such that $v_1 \geq v_2 \ldots \geq v_N > 0$, then the necessary and sufficient condition for $(x_1^*, \ldots, x_N^*)$ to be a solution to the optimization problem defined by (5.1), (5.2), and (5.3) is that there exists a target n such that $1 \leq n \leq N$ and*

(i) $x_1^* + \ldots + x_n^* = B$;

(ii) $\frac{A^{\beta_1}}{A^{\beta_1} + (x_1^*)^{\beta_1}} v_1 = \ldots = \frac{A^{\beta_n}}{A^{\beta_n} + (x_n^*)^{\beta_n}} v_n \geq v_i$ *for all* $i > n$.

The solution given by Proposition 5.1 is guaranteed to exist, since for each target n, the expected loss from an attack, $\frac{A^{\beta_n}}{A^{\beta_n}+(x_n)^{\beta_n}}v_n$, is strictly decreasing in the defensive resource x_n allocated to that target. Proposition 5.1 suggests that the defender exhaust her entire budget B to reduce (and equalize) the attractiveness of attacks on as many high-valued targets as possible. If B is adequately large (or the attacker is sufficiently weak, with small A), then the defender will be able to protect all N targets, at least to some degree. However, if the defense budget B is small (or the attacker is strong, with large A), then the defender will need to leave inferior targets unprotected. The following lemma and corollary formalize these observations.

Lemma 5.2. *Let the N targets be rank ordered such that $v_1 \geq v_2 \ldots \geq v_N > 0$, and let $B^{(n)}$ $(2 \leq n \leq N)$ be the minimal level of defender budget needed to equalize the attractiveness of the n highest-valued targets; i.e., such that $\frac{A^{\beta_1}}{A^{\beta_1}+(x_1)^{\beta_1}}v_1 = \ldots = \frac{A^{\beta_n}}{A^{\beta_n}+(x_n)^{\beta_n}}v_n$ for $x_1 + \ldots + x_n = B^{(n)}$. Then $B^{(i)} \leq B^{(j)}$ for $\forall 2 \leq i < j \leq N$.*

Corollary 5.3. *If the N targets are rank ordered such that $v_1 \geq v_2 \ldots \geq v_N > 0$, and $B^{(n)}$ $(n = 2, \ldots, N)$ are chosen as in Lemma 5.2, then the solution $(x_1^*, \ldots, x_N^*)$ to the optimization problem defined by (5.1), (5.2), and (5.3) satisfies one of the following cases:*

(i) If $B > B^{(N)}$, then $x_i^* > 0$ for $\forall i = 1, \ldots, N$;

(ii) If $B^{(n)} < B \leq B^{(n+1)}$ $(2 \leq n \leq N-1)$, then $x_i^* > 0$ for $\forall i \leq n$ and $x_i^* = 0$ for $\forall i > n$;

(iii) If $0 < B \leq B^{(2)}$, then $x_1^* > 0$ and $x_i^* = 0$ for $\forall i > 1$.

We illustrate Proposition 5.1 using the following two examples.

Example 5.1. *Consider a two-target example. We assume that attacks on the two targets share the same decisiveness ($\beta_1 = \beta_2 = \beta > 0$). (Example 5.2 allows for differing levels of decisiveness across targets.)*

If both targets are equally attractive to the attacker in the absence of defense (i.e., $v_1 = v_2 > 0$), then the optimal defensive strategy is always $x_1^* = x_2^* = \frac{B}{2}$, independent of the values of A and β. Without loss of generality, we focus on the case where target 1 is more attractive than target 2 in the absence of defense ($v_1 > v_2 > 0$). Denote by $r = \frac{x_1}{B}$ the proportion of defensive budget allocated to target 1, the higher-valued target. We look at how the attacker advantage (i.e., the ratio of the attacker capability over defensive budget $\frac{A}{B}$) affects the optimal defensive resource allocation r^*. If $\frac{A}{B} < \left(\frac{v_1}{v_2} - 1\right)^{-1/\beta}$, then $0.5 \leq r^* < 1$, such that

$$\frac{(\frac{A}{B})^\beta + (r^*)^\beta}{(\frac{A}{B})^\beta + (1 - r^*)^\beta} = \frac{v_1}{v_2}$$

from which we can derive that r^* is strictly increasing as the attacker advantage $\frac{A}{B}$ increases (see Appendix D.6). If $\frac{A}{B} \geq \left(\frac{v_1}{v_2} - 1\right)^{-1/\beta}$, then $r^* = 1$, and the entire budget is spent on the higher-valued target.

Figure 5.3 shows the effects of attacker advantage $\frac{A}{B}$ on the optimal proportion of defensive resources allocated to the higher-valued target $\frac{x_1^*}{B}$, considering three different levels of shared decisiveness: low ($\beta = 0.1$); moderate ($\beta = 0.6$); and high ($\beta = 8$). Here, we also consider two representative cases for target valuations: similar targets ($v_1 = 1.2v_2$); and targets that differ greatly ($v_1 = 1.8v_2$).

When β is close to zero, as in Figure 5.3(a), the success of an attack is due primarily to pure fortune, rather than to defender or attacker effort. In this case, defensive investment is not highly effective, so large levels of investment in the more valuable target 1 are needed

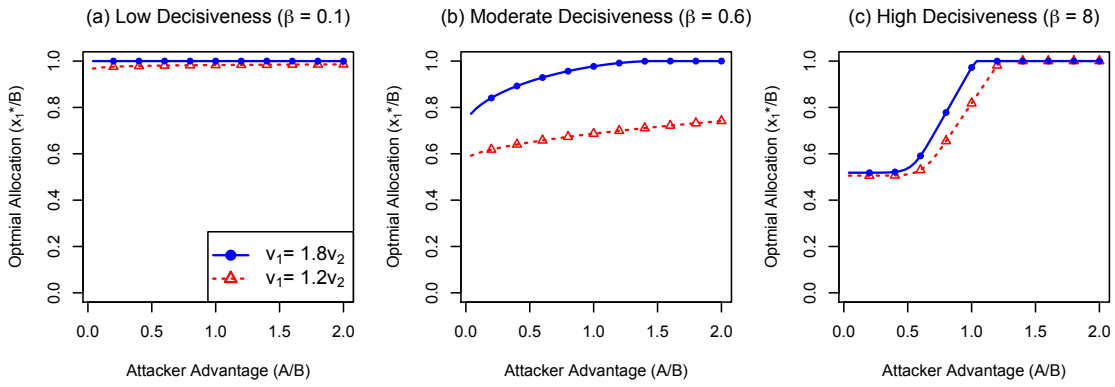

Figure 5.3: Effect of Attacker Advantage on Optimal Defensive Allocations in the Case of Equal Decisiveness

to make it no more attractive to the attacker than target 2. Therefore, the defender should devote the vast majority of her budget to the higher-valued target, regardless of the attacker's capability advantage ($\frac{A}{B}$).

When β is moderate, as in Figure 5.3(b), whether an attacker succeeds will depend non-trivially on the attacker capability relative to the defense level, but a less capable attacker may still get lucky in succeeding with an attack on either target. The optimal strategy in this case (even the attacker is believed to be quite strong with $\frac{A}{B} > 1$) is generally to hedge the allocation of defensive resources to provide at least some protection to both targets. Note also that the extent of hedging is much greater for similar targets ($v_1 = 1.2v_2$) than for targets that differ greatly in their valuations ($v_1 = 1.8v_2$).

As the decisiveness β gets sufficiently large, as in Figure 5.3(c), even a small advantage in attacker or defender resources is decisive. When faced with a weak attacker (e.g., $\frac{A}{B} < 0.5$), the defender is thus able to obtain an overwhelming advantage over the attacker by allocating resources almost evenly between the two targets. By contrast, if the attacker

is so capable that the defensive budget is inadequate to ensure effective protection of both targets (e.g., $\frac{A}{B} > 1$), then the optimal strategy is to put most of the defensive effort into defending the higher-valued target.

Example 5.2. *Consider two targets with widely different levels of decisiveness. In particular, target S is assumed to be a soft civilian facility with moderate decisiveness ($\beta_S = 0.6$), so even a weak attacker may get lucky with an attack on that target. By contrast, target H is assumed to be a hardened military base with high decisiveness ($\beta_H = 8$), so an attack on that target is likely to succeed only if the attacker's capability exceeds the defensive investment. Figure 5.4 shows the effects of the attacker advantage $\frac{A}{B}$ on the optimal proportion of defensive resources allocated to the soft target $\frac{x_S^*}{B}$, considering three cases for the target valuations: one case where the soft target is less attractive to the attacker than the hard target in the absence of defense ($v_S = 0.6v_H$); and two cases where the soft target is more attractive ($v_S = 1.2v_H$ and $v_S = 1.8v_H$). Although ordinarily we would expect the hard target to be more valuable than the soft target, there are cases where attacks on soft targets can cause severe consequences. For example, a large outdoor event that attracts a lot of people is a soft target associated with large numbers of casualties in case of IED attacks.*

If targets have unequal decisiveness, then knowing the level of attacker capability is critical to predicting the attacker's target choice. For example, a weak attacker (with small $\frac{A}{B}$) will generally prefer the soft civilian facility even if it is less valuable (e.g., when $v_S = 0.6v_H$), since he would not be likely to succeed in an attack on the hardened military base anyhow, but could get lucky in attacking the soft target. In this case, the defender is optimal to spend most of her resources on the soft target. By contrast, if the attacker is highly capable (with large $\frac{A}{B}$), then which target to defend depends on the relative valuations

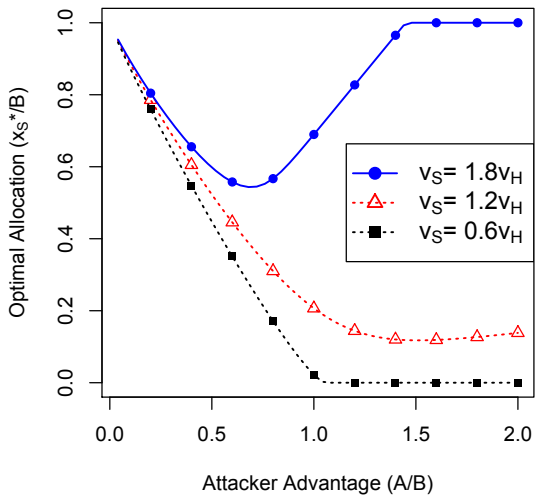

Figure 5.4: Effect of Attacker Advantage on Optimal Defensive Allocations in the Case of Unequal Decisiveness

of the soft and hard targets. The higher the valuation of the soft target, the more protection it merits.

With perfect knowledge about attacker intent and capabilities, it is reasonable for the defender to sometimes "put all her eggs in one basket," especially in the face of an extremely weak or strong attacker, and/or when one target is significantly more attractive than others. However, if the defender's beliefs about the attacker intent and/or capabilities are erroneous, then allocating budget in this way may lead to attacks with unexpectedly high consequences. In fact, real-world decision makers will generally want to hedge their defensive investments in case they have guessed wrong about the attacker characteristics (Bier 2007; Bier et al. 2013). In the next section, we therefore consider games of incomplete information where

the defender is uncertain about the attacker's capability and/or intent, in order to achieve realistic levels of hedging in homeland-security decisions.

5.2 Game of Incomplete Information

In this section, we consider a game of incomplete information to account for defender uncertainty about attacker characteristics. We first explore the effects of defender uncertainty about attacker capability (A) on the optimal defensive strategies, assuming that the two agents still interact in a zero-sum game; i.e., the defender's objective is to minimize the expected payoff to the attacker. Next, we allow the defender to be uncertain about the attacker intent (U_n), and to have her own target valuations (v_n) that may differ from the attacker's.

5.2.1 Defender Uncertainty about Attacker Capability

We here consider a zero-sum game (with $U_n = v_n$ for each target n), and assume that the defender is uncertain about only the attacker capability A; other model parameters (v_n and β_n) are all assumed to be common information. We use the cumulative distribution function $Q_A(\cdot)$ to describe the defender's imperfect knowledge about A. The defender's objective is then to allocate the budget B among the N potential targets in such a way as to minimize her expected loss from an attack, as given by

$$\min_{x_1+\ldots+x_N \leq B} \int \sum_{n=1}^{N} p_n(x, A, v) s_n(x_n, A) v_n dQ_A(A) \tag{5.4}$$

where $p_n(x, A, v)$ and $s_n(x_n, A)$ are the same as in equations (5.2) and (5.3), respectively. Note that given any defensive resource allocation $x = (x_1, \ldots, x_N)$, both the threat likelihood $p_n(x, A, v)$ and the success probability $s_n(x_n, A)$ are subject to random variations due to the stochastic nature of the attacker capability A. For computational convenience, we assume that for any given strategy of defense x, the probability of any two targets sharing the same highest level of attractiveness to the attacker is zero; i.e., Prob$\{Z > 1\} = 0$ where Z is the cardinality of the set $\{i : \arg\max_i s_i(x_i, A)v_i\}$. (This assumption can be satisfied if $Q_A(\cdot)$ is a continuous probability distribution, $\beta_n > 0$ for any n, and $v_i \neq v_j$ for any $i \neq j$.) Therefore, from the defender's perspective, the attacker is almost sure to choose only one target to attack, and the threat likelihood $p_n(x, A, v)$ takes on values of only 0 and 1 with probability one. Note that this assumption may not be applicable to multiple exchangeable targets such as warehouses and strip malls, but seems more suitable for iconic targets, which are unlikely to have exactly the same attractiveness.

Note that, after significant defensive investment has been made by the defender, multiple targets may have the same expected attractiveness under her subjective distribution of the attacker's target preferences. However, they will in general not be equally attractive to the attacker, if the defender has a continuous distribution over the attacker target preferences.

The following proposition gives the first-order condition for the defender's optimization problem (5.4).

Proposition 5.4. *The optimization problem (5.4) is equivalent to*

$$\min_{x_1+\ldots+x_N \leq B} \int \max_n \{s_n(x_n, A)v_n\} dQ_A(A) \tag{5.5}$$

Suppose that for any defensive allocation $x = (x_1, \ldots, x_N)$, *Prob*$\{Z > 1\} = 0$ *where* Z *is the cardinality of the set* $\{i : \arg\max_i s_i(x_i, A)v_i\}$. *Then the first-order necessary condition for an interior-point solution* x^* *to (5.5) is that* $\exists \mu$ *such that*

$$\int g_n(x^*, A) dQ_A(A) = \mu \text{ for } n = 1, \ldots, N$$

where $0 < x_n^* < B$ *for* $n = 1, \ldots, N$ *and* $x_1^* + \ldots + x_N^* = B$. *Here* $g_n(x_n, A)$ *is the marginal expected defender loss (equal to the marginal expected attacker payoff) of an attack on target* n *with an increase in the defense level* x_n, *as given by*

$$g_n(x_n, A) = \begin{cases} \frac{\partial}{\partial x_n} s_n(x_n, A)v_n = \frac{-\beta_n A^{\beta_n}(x_n)^{\beta_n - 1} v_n}{(A^{\beta_n} + (x_n)^{\beta_n})^2} & \text{if } s_n(x_n, A)v_n > s_i(x_i, A)v_i \text{ for } \forall i \neq n \\ 0 & \text{otherwise} \end{cases}$$

Proposition 5.4 indicates that an optimal hedging strategy (where all targets receive positive defense) must equalize the expectation of the marginal expected attacker payoffs (equal to the marginal expected defender losses) across all targets. This condition is not only necessary but also sufficient for an interior point to be optimal if the vulnerability functions $s_n(x_n, A)$ are strictly convex in x_n (e.g., $\beta_n \in (0, 1]$ for all n). However, if the vulnerability functions are not convex in x_n (e.g., $\beta_n > 1$ for some n), then the first-order condition in Proposition 5.4 is no longer sufficient for deriving the optimal defensive hedging strategies. Of course, corner solutions (where some targets are left unprotected) also exist, especially when there are greatly inferior targets. Take a two-target example. Target 1 will receive no defense at optimality, if $\frac{v_1}{v_2} < \frac{s_2(B,A)}{s_1(0,A)}$ for any value of A with a positive probability density.

For simple two-target cases, an exhaustive searching through all possible allocation plans will be efficient enough for finding solutions to the optimization problem (5.4), even if the objective function is not convex (e.g., when the vulnerability functions $s_n(x_n, A)$ are not convex in x_n for some n, or if we further consider a non-zero-sum game). However, for large numbers of targets, we will need more advanced global optimization techniques (as in Rios and Sahinidis 2012), or stochastic programming techniques that use binary variables representing the attacker's target choices. Detailed discussions on computation can be found in Chapter 6.

Example 5.3. *Consider two targets with unequal decisiveness. Again, S is a soft civilian facility with low decisiveness ($\beta_S = 0.6$), and H is a hardened military base with high decisiveness ($\beta_H = 8$). For convenience, we fix the defensive budget $B = 1$, and assume that the attacker capability A is bounded and can take on values only in $[0, 2]$. We then present the defender's imperfect knowledge of A as a uniform distribution with breadth α_A. For the case of a relatively weak attacker, we let the attacker capability A be uniformly distributed over $[0, \alpha_A]$. By contrast, we assume a uniform distribution over $[2 - \alpha_A, 2]$ for the capability of a relatively strong attacker. By restricting the value of α_A within $[0.05, 2)$, we can keep the expected value $\mathbf{E}[A] < 1$ for the weak attacker, while $\mathbf{E}[A] > 1$ for the strong attacker.*

Of course, the uniform distribution is just a convenient representation of defender uncertainties. Other approaches are also possible. For example, use of the Beta distribution would allow us to adjust the variance of A, while keeping its expected value unchanged. In fact, with sufficiently large variance, the Beta distribution becomes U-shaped, which is close to a discrete distribution (where the attacker capability can be either extremely large

or small, with little possibility of taking on intermediate values). Moreover, if the attacker capability is assumed to be unbounded, then other distributions could be considered, such as the Rayleigh distribution as adopted by Bier et al. (2008).

Figure 5.5 shows the optimal defensive allocation to the soft target (x_S^*) as a function of the extent of defender uncertainty about attacker capability (represented by the breadth of the uniform distribution α_A). We consider both a relatively weak attacker ($\mathbf{E}[A] < 1$) and a relatively strong attacker ($\mathbf{E}[A] > 1$). In addition, we also consider two cases for the target valuations: $v_S = 1.8v_H$; and $v_S = 0.6v_H$.

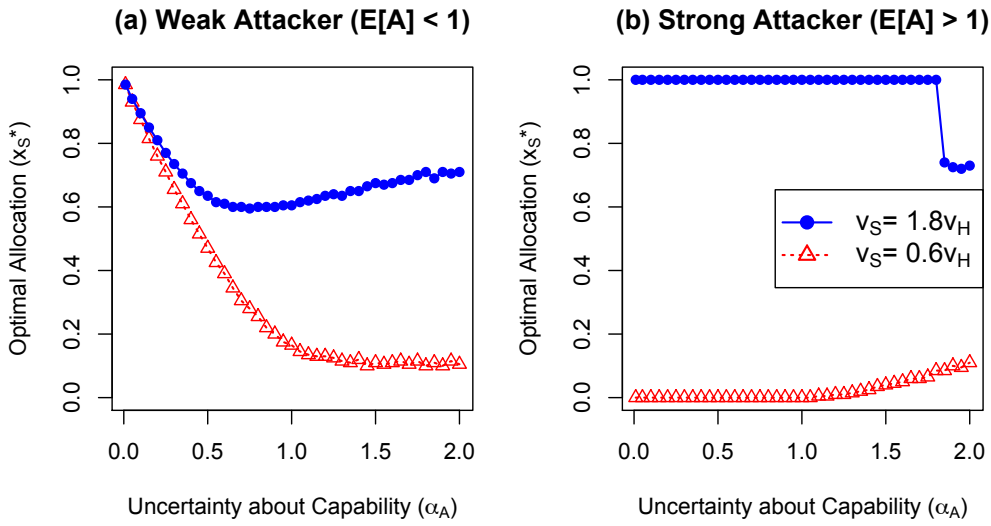

Figure 5.5: Effect of Defender Uncertainty about Attacker Capability

If the defender is highly confident that the attacker has low capability – i.e., for small values of α_A in Figure 5.5(a) – she would spend the vast majority of her investment on the soft target, since a weak attacker would generally prefer a less decisive attack. However, this strategy makes sense only if the attacker is highly likely to be weak. If the defender

believes that a stronger attacker is plausible (for large values of α_A), then she needs to protect both targets.

Similarly, it is reasonable for the defender to harden only the more valuable target if she is facing a sufficiently strong attacker who is highly likely to be able to damage any target – i.e., for small values of α_A in Figure 5.5(b). However, an ill-informed defender (with large α_A) must hedge her defense against the possibility that a less capable attacker may prefer low-valued targets with high chances of success.

In principle, more hedging is needed as the defender becomes more uncertain about attacker capability (i.e., as α_A increases). Note that in the limit for $\alpha_A = 2$, the optimal allocation of defensive resources to the various targets will depend only on the characteristics of those targets; i.e., their valuations and decisiveness. This can be seen by comparing the levels of defense at $\alpha_A = 2$ in Figures 5.5(a) and (b).

5.2.2 Defender Uncertainty about Attacker Intent and Capability

We now allow the defender to be uncertain about attacker intent, as well as capabilities. We here assume that the attacker's target valuations U_n can differ from the defender's valuations v_n, and may not be fully known to the defender. In particular, we represent the defender's imperfect knowledge of $U = (U_1, \ldots, U_N)$ by the joint cumulative distribution function $Q_U(\cdot)$. The defender's objective is then to allocate the budget B among the N potential targets to minimize her expected loss, as given by

$$\min_{x_1+\ldots+x_N \leq B} \int \int \sum_{n=1}^{N} p_n(x, A, U) s_n(x_n, A) v_n dQ_A(A) dQ_U(U) \tag{5.6}$$

where $p_n(x, A, U)$ and $s_n(x_n, A)$ are given by equations (5.2) and (5.3), respectively.

If we allow the two agents to differ in their target valuations, then there may exist targets that are attractive to the attacker but of relatively low value to the defender. In that case, an optimal strategy for the defender might be to purposely keep those targets unprotected or lightly protected (Bier 2007). However, if the defender's knowledge about the attacker's intent is wrong, then such a strategy is likely to cause unexpectedly high damages. The following example illustrates how defender uncertainties about both attacker intent and capabilities jointly affect the optimal defensive decisions.

Example 5.4. *Consider the two-target case with a soft civilian facility ($\beta_S = 0.6$) and a hardened military base ($\beta_H = 8$). The defender is assumed to be uncertain about both the attacker's capability (A) and the attacker's target valuations (U_S and U_H). For convenience, we set the defensive budget $B = 1$. The defender's imperfect knowledge of A is described in the same way as in Example 5.3. In addition, we fix the attacker's valuation of the hard target $U_H = 1$, and assume the attacker's valuation of the soft target U_S to be uncertain. For a relatively low-valued soft target, we set $Q_U(U_S) \sim$ Uniform$[0, \alpha_U]$ with $\mathbf{E}[U_S] < 1$; for a relatively high-valued soft target, we set $Q_U(U_S) \sim$ Uniform$[2 - \alpha_U, 2]$ with $\mathbf{E}[U_S] > 1$. The extent of defender uncertainty about the attacker's intent is again measured by the breadth of the uniform distribution, denoted as α_U with $\alpha_U \in [0, 2)$.*

Figure 5.6 shows the optimal allocation of defensive resource to the soft target (x_S^*) as a function of the extent of defender uncertainty about both attacker intent and capabilities. We consider two cases for the expected attacker capability (a weak attacker with $\mathbf{E}[A] < 1$ and a strong attacker with $\mathbf{E}[A] > 1$), two cases for the expected value of the soft target to the attacker ($\mathbf{E}[U_S] < 1$ and $\mathbf{E}[U_S] > 1$), and two cases for the defender's (deterministic) target valuation ($v_S = 0.6v_H$ and $v_S = 1.5v_H$).

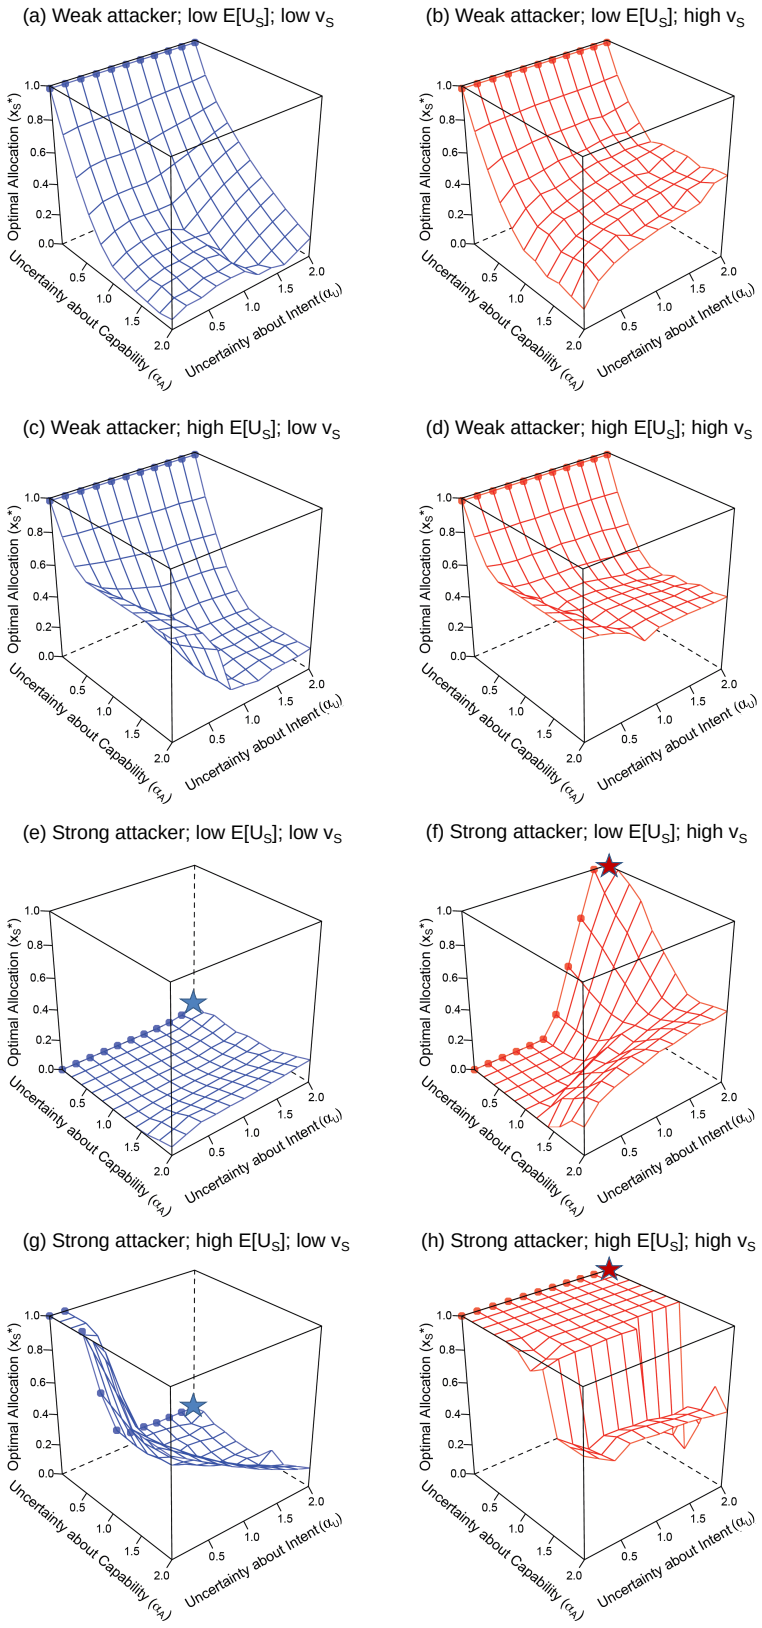

Figure 5.6: Effects of Defender Uncertainty about Attacker Intent and Capability

If the defender believes that the attacker is likely to be weak, then the defender should protect only the soft target, as shown by the large values of x_S^* for small values of α_A highlighted by the dark circles in Figures 5.6(a)-(d), regardless of the attacker's intent ($\mathbf{E}[U_S]$ and $\mathbf{E}[U_H]$), the defender's target valuations (v_S and v_H), and no matter how uncertain the defender is about the attacker's intent (α_U).

By contrast, if the defender is highly confident that the attacker is strong, then the defensive strategies should be based mainly on the attacker's intent, as shown by the dark circles for small values of α_A in Figures 5.6(e)-(h). In particular, the defender needs to spend more on the target that is more valuable to the attacker. This can be seen by the fact that the attacker spends more on the soft target (x_S^*) in Figure 5.6(g) than 5.6(e), and likewise more in 5.6(h) than 5.6(f).

Moreover, as the defender becomes more uncertain about the strong attacker's intent (for small values of α_A and large values of α_U in Figures 5.6(e)-(h)), she needs to hedge more against the possibility that her estimates of U_S and U_H might be erroneous, by assigning more weight to her own target valuations v_S and v_H. Thus, when v_S is small, as in Figures 5.6(e) and (g), the defender spends little on defense of the soft target, as shown by the stars. Conversely, when v_S is large, as in Figures 5.6(f) and (h), the defender's entire budget is spent on the soft target, as also shown by the stars.

However, when the defender is highly uncertain about the attacker's capability in addition to intent (for large values of α_A and α_U in Figures 5.6(a)-(h)), the extent of her uncertainty about attacker intent (α_U) has only a small impact on the optimal defensive strategies. That may be because significant hedging is already required when the defender is highly uncertain about the attacker's capability, as shown in Figure 5.5.

A summary of the optimal strategies for defensive resource allocations between a soft target S and a hard target H is presented in Table 5.1. If the attacker is known to have only limited resources, then the defender should protect only the soft target(s). Of course, if there are numerous soft targets, only the more valuable ones can be defended. By contrast, if the attacker is believed to be highly capable, then knowing his goals and motivations will be critical to achieving effective defense. Moreover, if the attacker's capability is unknown within the time frame of the defender's decision, then substantial hedging is needed and precise assessment of attacker intent might not be necessary. It is therefore reasonable to prioritize intelligence collection about the attacker's capabilities over intelligence collection about intent.

Table 5.1: Priority of Target Protection

	Weak attacker	Uncertain capability	Strong attacker
S more valuable to attacker	S	Both	S
Uncertain intent	S	Both	Target w/ higher defender value
H more valuable to attacker	S	Both	H

We have so far extended the game-theoretic model of adaptive adversary in Bier et al. (2007) to account for defender uncertainty about adversary capabilities in addition to just intent. In the next two sections, we further extend the game to handle multiattribute attacker capabilities and to allow for multiple simultaneous attacks.

5.3 Multiattribute Attacker Capability

According to Cragin and Daly (2004), an adversary group's capabilities can be divided into two categories: organizational capability (such as leadership, recruitment pools, and publicity); and operational capability (such as weapon sources, technical expertise, trained personnel, money, and flexibility of movement). In this section, we borrow the "generalized contest success function" developed by Rai and Sarin (2007) to capture the multiattribute nature of attacker capabilities. For illustrative purposes, we consider only two representative attributes, capital (C) and labor (L). However, the analysis can be extended to more than two attributes in a straightforward way.

Suppose that the attacker has amounts A_C of capital and A_L of labor to invest in a potential attack. Then for a fixed defensive investment x_n, the success probability of an attack on target n is given by the following generalized contest success function

$$s_n(x_n, A_C, A_L) = \frac{(A_C)^{\kappa_n \beta_n}(A_L)^{\lambda_n \beta_n}}{(A_C)^{\kappa_n \beta_n}(A_L)^{\lambda_n \beta_n} + (x_n)^{\beta_n}} \tag{5.7}$$

where $\beta_n, \kappa_n, \lambda_n > 0$, and $\kappa_n + \lambda_n = 1$ (following Rai and Sarin 2007). Note that the vulnerability function (5.7) also satisfies the condition that its value will remain unchanged if the attacker's capital and labor resources and the defender's investment are all multiplied by the same factor.

By defining the "effective attacker capability" as $A_n = (A_C)^{\kappa_n}(A_L)^{\lambda_n}$, we can reduce the vulnerability function in (5.7) to the single-attribute case in (5.3), except that the values of the effective capabilities A_n will differ across targets (because of the effects of κ_n and λ_n). Note that the functional form $A_n = (A_C)^{\kappa_n}(A_L)^{\lambda_n}$ coincides with the Cobb-Douglas production

function in economics (Douglas 1976), with κ_n and λ_n being the "output elasticities" of capital and labor, respectively. For example, if $\kappa_n = 0.9$, then a 1% increase in the attacker's capital would lead to approximately an 0.9% increase in his effective capability A_n for an attack on target n.

As before, the value of β_n in equation (5.7) represents the overall decisiveness of the effective attacker capability (or defensive investment) devoted to target n. We normalize the elasticities κ_n and λ_n to sum to one, so that they capture the relative decisiveness of capital versus labor for an attack on target n. For example, a high-technology attack (e.g., a nuclear attack) would be expected to have a higher capital elasticity κ_n, whereas a less sophisticated attack (e.g., an IED attack) may have roughly equal effectiveness of capital and labor ($\kappa_n \approx \lambda_n$). The following example illustrates the use of the generalized contest success function in modeling multiattribute attacker capabilities.

Example 5.5. *Consider two targets: a soft civilian facility S ($\beta_S = 0.6$, $\kappa_S = 0.5$, $\mu_S = 0.5$); and a hardened military base H ($\beta_H = 8$, $\kappa_H = 0.9$, $\lambda_H = 0.1$). For convenience, we assume that the defensive budget $B = 1$. Figure 5.7 presents contour plots of the optimal defensive resource allocation to the soft target (x_S^*) for different levels of attacker capital (A_C) and labor (A_L), assuming that there is no defender uncertainty about the attacker characteristics, and that the defender and attacker share the same target valuations ($U_S = v_S$ and $U_H = v_H$). Curves with higher labels correspond to more resources allocated to the soft target. Three cases for target valuations are considered: one where the soft target is less valuable than the hard target in the absence of defense ($v_S = 0.1v_H$); one where both targets are of equal value ($v_S = v_H$); and one where the soft target is more valuable ($v_S = 1.8v_H$).*

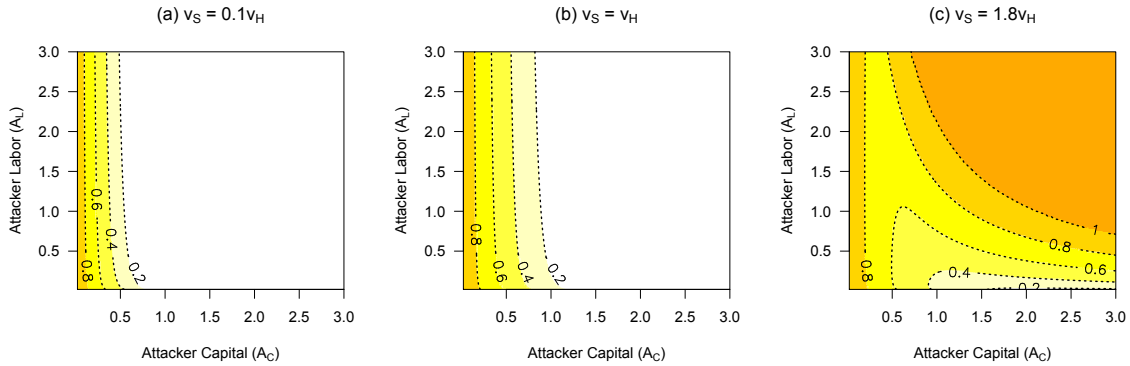

Figure 5.7: Contour Plots of Optimal Defensive Allocation (x_S^*)

When the soft target is less valuable than or equally valuable to the hard target (in the absence of defense), as in Figures 5.7(a) and (b), the attacker would prefer to attack the hard target unless his capital (A_C) is too small. The optimal defensive strategy is primarily driven by the attacker's capital level (A_C), since changes in the attacker's labor (A_L) have little effect on the relative attractiveness of the hard target. Therefore, the optimal defensive strategy is mainly determined by the attacker's capital, and is approximately independent of how many people the attacker has available. The more capital the attacker has, the more protection is needed by the hard target, and the less should be spent on the soft target.

By contrast, if the soft target is more valuable (in the absence of defense), as in Figure 5.7(c), then the defender's priority is to protect the soft target in most cases. However, an attacker with abundant capital (large A_C) but insufficient labor (small A_L) would still prefer a technically sophisticated attack on the hard target, to take advantage of the greater effectiveness of capital against such targets. In that case, the hard target requires more protection.

The generalized contest success function allows homeland-security decision makers to consider multiple attributes of adversary capability (such as capital, labor, technological sophistication, and flexibility of movement) in a tractable way. This representation separates the task of estimating each individual attribute of capability (e.g., availability of capital or labor) from the task of identifying the effects of the various attributes, and thus provides a realistic way for intelligence experts to estimate adversary capabilities. For example, by changing the output elasticities (e.g., the κ_n and λ_n), we can vary the relative importance of the various attributes. Moreover, defender uncertainty about attacker capability can then be described by probability distributions over the attacker's individual capability attributes (i.e., A_C and A_L), and/or the elasticities (e.g., κ_n and λ_n).

5.4 Multiple Simultaneous Attacks

Zhuang and Bier (2007) point out that an attacker may find attacking multiple targets simultaneously to be an attractive strategy even if any one target by itself may not have been sufficiently attractive individually (e.g., the four simultaneous airplane attacks involved in the 9/11 tragedy in 2001). In this section, we extend our model to allow the attacker capability to be divisible among targets, and provide an analytical framework to explore the optimal defensive strategy against multiple simultaneous attacks. For simplicity, we only consider single-attribute adversary capabilities.

In particular, we allow the attacker to allocate his total level of capability (or effort) A among the N targets, $a_1 + \ldots + a_N \leq A$, where $a_n \geq 0$ is the attacker effort spent on target n. The likelihood of target n being attacked is 1 if $a_n > 0$ and 0 if $a_n = 0$. For simplicity,

only the zero-sum game of complete information is considered; i.e., $U_n = v_n$ for all n. The defender's objective is then to minimize the expected loss from any attack(s) launched by the attacker, as given by

$$\min_{x_1+\ldots+x_N \leq B} \max_{a_1+\ldots+a_N \leq A} \sum_{n=1}^{N} s_n(x_n, a_n) v_n \tag{5.8}$$

where $s_n(x_n, a_n)$ is the vulnerability function as given by (5.3). The following proposition gives the necessary condition for the attacker's best response for given levels of defensive resource allocation.

Proposition 5.5. *For a given defense plan $x = (x_1, \ldots, x_N)$, the necessary condition for the attacker's best-response function $\hat{a}_n(x)$, to the optimization problem as defined by (5.8) and (5.3), is that for any $n \in \{i : \hat{a}_i > 0\}$*

$$h_n(x_n, \hat{a}_n) \begin{cases} = h_j(x_j, \hat{a}_j) & \textit{for} \quad \forall j \in \{i : \hat{a}_i > 0\} \\ \geq h_j(x_j, \hat{a}_j) & \textit{for} \quad \forall j \in \{i : \hat{a}_i = 0\} \end{cases}$$

where $\hat{a}_1 + \ldots + \hat{a}_N = A$ and $h_n(x_n, a_n)$ is the marginal expected payoff to the attacker of attacking target n at effort level a_n, as given by

$$h_n(x_n, a_n) = \frac{\partial}{\partial a_n} s_n(x_n, a_n) v_n = \frac{\beta_n (a_n)^{\beta_n - 1} (x_n)^{\beta_n} v_n}{((a_n)^{\beta_n} + (x_n)^{\beta_n})^2}$$

At optimality, any target the attacker chooses to attack with positive effort will yield the same marginal expected payoff, and any targets that are not chosen will cause (weakly) smaller marginal expected payoffs. If the vulnerability functions $s_n(x_n, a_n)$ are concave in attacker effort a_n for all n – i.e., if $\beta_n \in (0, 1]$ – then the condition in Proposition 5.5 is also sufficient for finding the attacker's best responses. However, if $\beta_n > 1$ for some n, then we

can write the attacker's objective function $\sum_{n=1}^{N} s_n(x_n, a_n)v_n$ as the sum of strictly convex and strictly concave functions, and apply the concave-convex procedure developed by Yuille and Rangarajan (2003) for determining the attacker's optimal strategies. To obtain the defender's optimal strategies is even more difficult, and more advanced global optimization techniques as in Rios and Sahinidis (2012) would be required.

A weak attacker (with small $\frac{A}{B}$) would generally prefer attacking only one target at a time, so the optimal defensive strategy will be similar to that in Section 5.1. On the other hand, a highly capable attacker (with large $\frac{A}{B}$) is more likely to conduct multiple simultaneous targets. We here focus on the case of a strong attacker (with $\frac{A}{B} = 2$) choosing between a soft target S and a hard target H to attack. Suppose that the defender shares equal target valuations with the attacker ($U_S = v_S$ and $U_H = v_H$), and is fully knowledgeable about the attacker's total level of effort A. Three cases for the target valuations are considered: $v_S = 0.2v_H$; $v_S = v_H$; and $v_S = 1.8v_H$.

Figure 5.8 illustrates the equilibrium of the sequential game in which the defender moves first. If the attacker can choose only a single target, the hollow triangle in each plot shows the defender's optimal resource allocation and the attacker's single target choice after observing that allocation. Note that multiple equilibria may exist, as shown by the two hollow triangles in Figure 5.8(b). By contrast, if the attacker is allowed to attack multiple targets simultaneously, the solid line in each plot represents the attacker's best response (i.e., the optimal allocation of his effort to the soft target S) to any given defensive strategy. The solid triangle then identifies the defender's optimal strategy considering the attacker's best responses.

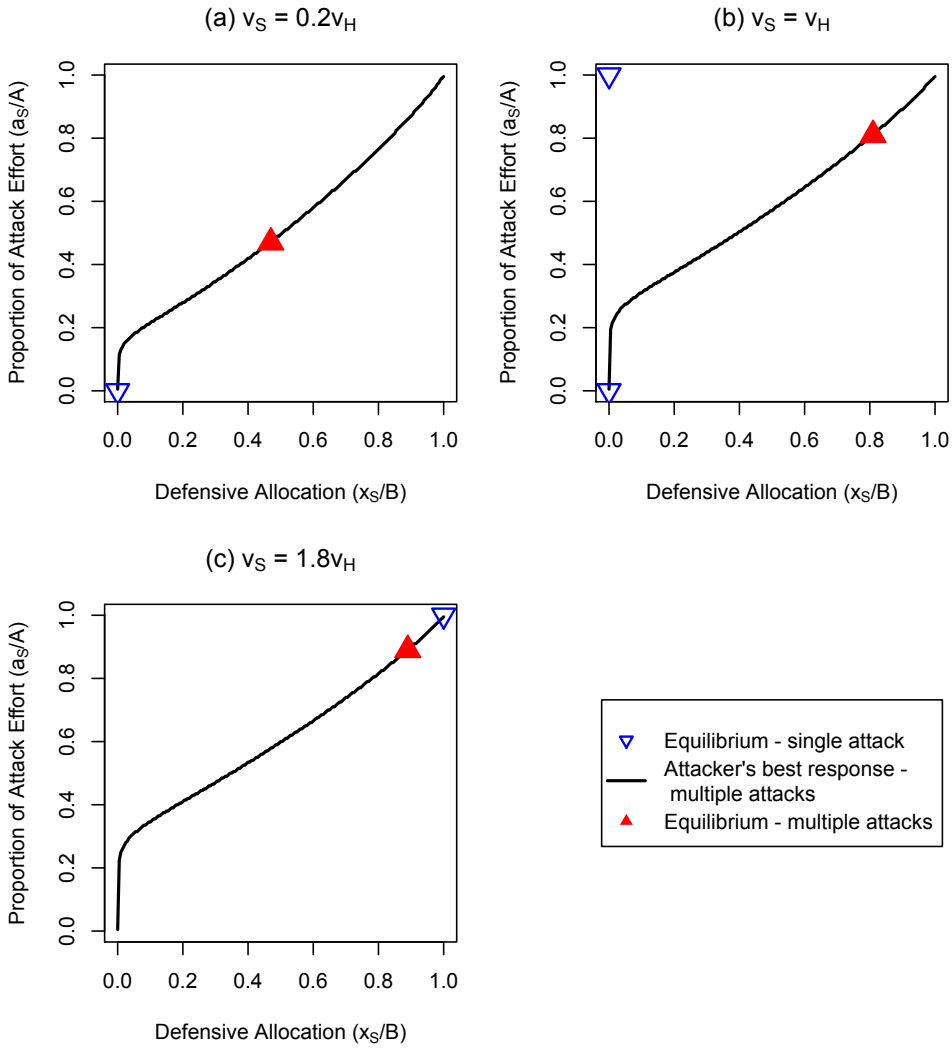

Figure 5.8: Optimal Defensive Resource Allocation against Multiple Simultaneous Attacks by a Strong Attacker

If a highly capable attacker can choose only a single target, as shown by the hollow triangles in Figure 5.8, he will choose the one with higher value, in order to cause more damage. Therefore, the optimal strategy of a poorly-equipped defender is to devote all her resources to the more valuable target – i.e., the hard target in Figure 5.8(a), or the soft target in Figure 5.8(c). If the two target values are similar, as in Figure 5.8(b), the optimal defensive strategy is to make the attacker indifferent between attacking either of the two targets, as shown by the existence of two hollow triangles, representing the possibility of an attack on either target.

However, if the highly capable attacker is allowed to attack both targets simultaneously, as shown by the solid triangles in Figure 5.8, he will generally prefer to do so, in order to avoid wasted effort and cause more damage. Therefore, a relatively weak defender will want to hedge her investment to protect both targets, at least to some extent. Additionally, the defender spends much more on the soft target in the case of multiple attacks, even if the two target values are similar, as in Figure 5.8(b). This is because protecting the hard target against a strong attacker would not be as effective as investing in protection of the less decisive soft target.

In general, our model can help to achieve more effective protection than the approaches used in previous studies. In particular, our model not only allows for simultaneous attacks on multiple targets, but also takes into account how much effort the attacker is likely to devote to the various targets based on both capability and intent.

5.5 Illustrative Case Study

In this section, we present a hypothetical case study to illustrate the usability of the proposed game-theoretic model. We consider three possible organizations: Northern Ireland's Real Irish Republican Army (RIRA)[1]; the Palestinian group Hamas; and al Qaeda. Cragin and Daly (2004) give notional estimates for the three organizations' capabilities using a 1-5 rating scale as shown in Table 5.2. Those estimates are based on both organizational and operational factors. The table also gives explanations for those notional values. For example, RIRA, with capability "1," is capable of killing or injuring at least 50 people in a single attack.

Table 5.2: Notional Estimates for Attacker Capabilities

	Capability	Explanation
RIRA	1	Kill or injure 50 or more people in a single attack
Hamas	2	Intentionally target unguarded foreign nationals
al Qaeda	5	Successfully coordinate multiple attacks

We then consider eight potential attack scenarios that are available to adversary groups, including three types of biological attacks (aerosolized anthrax, food contamination, and pneumonic plague), three types of chemical attacks (blister agent, chlorine tank explosion, and nerve agent), and attacks using improvised explosive devices (IED) or radiological dispersal devices (RDD). CREATE (2011) conducted an elicitation study where "proxy" experts (graduate students knowledgeable about terrorism, from countries where support for terrorism is relatively common) were asked to give estimates for the fatalities and economic impact of those attack scenarios (as shown in Table 5.3). Note that a chlorine

[1]The Real Irish Republican Army, otherwise known as the Real IRA or RIRA, is a successor to the original Irish Republican Army which existed from 1922 to 1969.

Table 5.3: Notional Estimates for the Various Attack Scenarios

	Fatalities	Economic Impact ($)	Decisiveness (β_n)
Aerosolized Anthrax	7.3×10^3	2.7×10^9	1.5
Food Contamination	220	2.7×10^6	1.5
Pneumonic Plague	1.5×10^3	2.7×10^6	1.5
Blister Agent	100	1.4×10^8	0.6
Chlorine Tank Explosion	1.5×10^4	2.7×10^6	0.6
Nerve Agent	4.4×10^3	8.1×10^7	0.6
IED	70	5.0×10^6	0.6
RDD	440	1.4×10^9	4

tank explosion is expected to cause the largest number of fatalities, while a biological attack using aerosolized anthrax is estimated to generate the highest economic loss.

Cragin and Daly (2004) also provide estimates of the hostility of the various adversary groups. In particular, al Qaeda is the most hostile group, and Hamas is moderately hostile. The intent of the RIRA is estimated to be relatively benign. Assuming that a more hostile attacker would place more weight on fatalities as opposed to economic impact, we assign valuations to the various attack scenarios (U_n) for each adversary group in the absence of defense (as shown in Table 5.4). In particular, we assign a weight of 0.9 to fatalities (and 0.1 to economic impact) for al Qaeda. We then set 0.5 and 0.1 as the weights of fatalities for Hamas and RIRA, respectively. As recommended by CREATE (2011), we use 6 million dollars as the value of a statistical life to both the defender and attacker.

We also give notional estimates for the decisiveness of each attack scenario (β_n), based mainly on the descriptive analysis in Falkenrath et al. (1998) on the difficulty of attackers obtaining required materials and the effectiveness of defensive countermeasures. For example, chemical weapons are relatively easy for many non-state actors to acquire,

Table 5.4: Attacker Valuations of the Various Attack Scenarios (U_n; \$)

	RIRA	Hamas	al Qaeda
Aerosolized Anthrax	2.87×10^{10}	3.54×10^{10}	4.21×10^{10}
Food Contamination	1.34×10^{8}	6.61×10^{8}	1.2×10^{9}
Pneumonic Plague	9.02×10^{8}	4.5×10^{9}	8.1×10^{9}
Blister Agent	1.86×10^{8}	3.70×10^{8}	5.54×10^{8}
Chlorine Tank Explosion	9.00×10^{9}	4.50×10^{10}	8.10×10^{10}
Nerve Agent	2.70×10^{9}	1.32×10^{10}	2.38×10^{10}
IED	4.70×10^{7}	2.13×10^{8}	3.79×10^{8}
RDD	1.50×10^{9}	2.00×10^{9}	2.50×10^{9}

since the precursor materials are commercially available and weaponization is not difficult. By contrast, biological attacks are believed to be more difficult, and hence show greater decisiveness. Even though the seed stocks for biological agents are accessible by some non-state actors, delivery devices are generally difficult to produce. Note that both chemical and biological attacks are difficult to identify and prevent in their early stages. However, radiological attacks may be deterred with high probability if we install adequate radioactive sensors.

For simplicity, we consider the zero-sum game of complete information where the defender's objective is to minimize the attacker's expected payoff, and the attacker can choose only a single attack scenario. Figure 5.9 presents the optimal defensive resource allocation in the face of each adversary group, as a function of the defensive budget (measured in multiples of the capability of the least capable adversary group RIRA).

RIRA is hypothesized to have a low capability, and weight economic impact more heavily than fatalities. Therefore, RIRA is likely to choose an anthrax attack, since it has a relatively low level of decisiveness and can cause large economic damage. As a result, it is optimal for the defender to focus her protection on aerosol anthrax attacks against RIRA, as

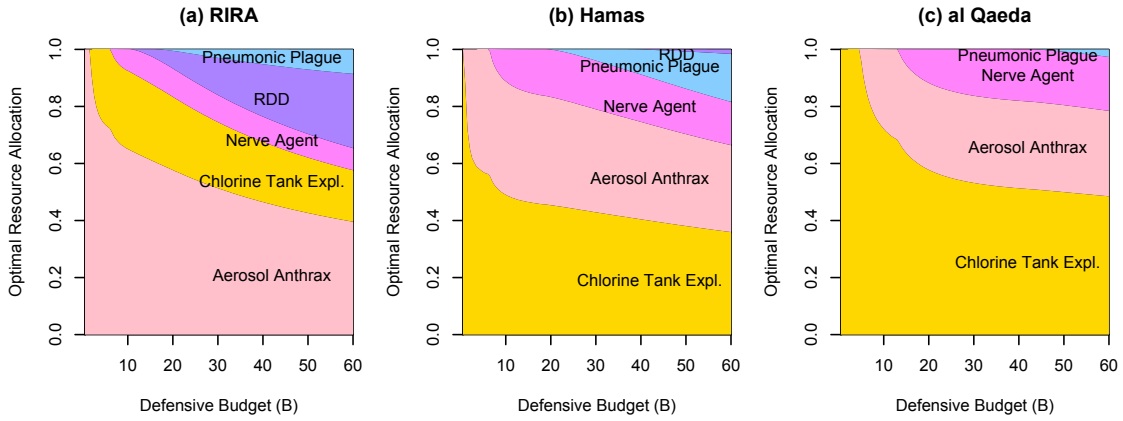

Figure 5.9: Optimal Defensive Resource Allocation in the Case Study

shown in Figure 5.9(a). Of course, as the defender's budget increases, she is able to also invest in protection against less economically damaging and more highly decisive attacks, such as RDD.

By contrast, the more hostile adversary groups, Hamas and al Qaeda, are predicted to prefer a chlorine tank explosion, because it can kill the most people. Accordingly, chlorine tank explosions are also the priority for protection, as shown in Figures 5.9(b) and (c). Moreover, a defender with a fixed budget is able to protect against more attack scenarios in total if the adversary group she is facing is the less capable Hamas, rather than al Qaeda.

5.6 Summary

Threat is widely regarded as the most difficult component to estimate in the TVC equation for quantifying terrorist risks. This chapter provides a novel game-theoretic framework for analyzing the effects of attacker intent and capabilities on both threat and optimal defensive

strategies. The novel feature of this model is the use of contest success functions to explicitly capture the extent to which the success of an attack is attributable to the attacker's capability (as well as the defensive investment) rather than pure luck. Moreover, our model allows the effectiveness of attacker capabilities to differ across targets (e.g., civilian versus military targets) and attack modes (e.g., IED versus nuclear attacks), and also allows for multiple types of attacker capability (such as capital versus personnel).

This work also extends the Bayesian Stackelberg game of Bier et al. (2007) to account for defender uncertainty about attacker capabilities, rather than just intent. Attackers with scarce resources are limited in their choices of targets or attack modes, since they are likely to succeed only when the targets or attack modes they choose have low decisiveness. By contrast, strong attackers are more capable of choosing targets according to their goals and motivations (i.e., intent), regardless of decisiveness. Therefore, uncertainty about attacker capabilities plays a critical role in making defensive decisions, especially when targets or attack modes have widely differing decisiveness.

Preliminary results suggest that precise assessment of attacker intent might not be necessary if the attacker's capability is highly uncertain. This result could provide practical guidance for determining the optimal trade-off between intelligence collection on attacker intent versus intelligence collection on capabilities.

Chapter 6

Computational Approaches for Determining Optimal Defenses

Solving the defender-adversary game of incomplete information is a challenging task. The defender's loss is subject to random variations because she is uncertain about the attacker's choice of target. In this chapter, we investigate state-of-the-art stochastic programming techniques for solving games of realistic size and complexity, to an acceptable level of accuracy within a reasonable time constraint.

We first generalize the defender-adversary game as a "stochastic-selection" problem, in which a decision maker is to allocate a limited budget among a collection of entities against stochastic selections by an opponent (Section 6.1). Next, we model the stochastic-selection problem as a two-stage stochastic programming problem with binary recourse, and investigate rigorous and efficient solution approaches based on sample-average approximation (SAA). In particular, the SAA method uses Monte-Carlo simulation to estimate the objective function

of the decision maker (e.g., defender) via sampling. We present conditions under which the general convergence properties of the SAA method can apply to our case (Section 6.2), and then identify two categories of computational algorithms for solving the simulation-based approximate problem, one based on mixed-integer nonlinear programming and the other based on derivative-free optimization (Section 6.3).

Finally, we use the basic defender-adversary game in Chapter 3 to evaluate the SAA method and the various computational algorithms (Section 6.4). Games of realistic size typically involve 50 to 100 decision variables (e.g., defensive resources allocated to the various targets) and a random sample space of dimension 10 to 50 (representing defender uncertainty about adversary characteristics). For the purpose of strategic planning where model parameters might need to be varied for comparative analysis, we would like to constrain the solution time of each optimization instance to be within six to ten CPU minutes. Best computation algorithms within that time limit are identified based on numerical experiments.

6.1 Stochastic-Selection Problem

Consider a sequential game with two players, *Blue* and *Red*. In the first stage of the game, *Blue* allocates a limited budget B among a finite collection of N targets. In the second stage, *Red* subsequently observes that budget allocation, and selects a target based on his characteristics (e.g., preferences and capabilities) which are not fully known to *Blue*. Selection of a particular target will bring *Red* returns while imposing costs on *Blue*. The interaction between the two agents can be modeled as a Bayesian Stackelberg game where

Blue is the leader and *Red* is the follower with private information. Moreover, we can model this game as a two-stage stochastic programming problem with recourse, as given by

$$\begin{aligned} &\min_{x \in X} f(x) = \mathbf{E}[F(x, \omega)] \\ &\text{where } F(x, \omega) = \sum_{i=1}^{N} z_i l_i(x) \\ &\text{s.t. } z_i = 0 \text{ for } \forall i \notin \arg\max_j R_j(x, \omega) \\ &\sum_{n=1}^{N} z_i(x, \omega) = 1 \\ &z \in \{0, 1\}^N \end{aligned} \tag{6.1}$$

where $X = \Delta_{N+1}(B) = \{x \in \mathbb{R}_+^N \mid \sum_{i=1}^N x_i \leq B\}$ is the set of *Blue*'s feasible allocations of budget B among N targets. We use $\omega \in \Omega$ to capture *Red*'s type (e.g., preferences and/or capabilities). Moreover, we assume that *Red* has full knowledge of ω, but *Blue* observes only a subjective distribution Q over Ω, the space of all possible *Red* characteristics. For a fixed budget allocation plan x of *Blue*, *Red* is then assumed to select a single target i with the highest reward $R_i(x, \omega)$ based on his type ω. We represent *Red*'s target choice by a vector of binary variables $z = (z_1, \ldots, z_N)$ with z_i equal to 1 if target i is selected and equal to 0 otherwise. If target i is selected, i.e., $z_i = 1$ and $z_j = 0$ for $\forall j \neq i$, then *Blue* will suffer a loss of $l_i(x)$ independent of *Red*'s type ω. From *Blue*'s perspective, her objective is to minimize the expected loss $\mathbf{E}[F(x, \omega)] = \mathbf{E}\left[\sum_{i=1}^N z_i l_i(x)\right]$ from *Red*'s stochastic target choice according to the distribution Q over Ω.

6.1.1 Example: Defender-Adversary Game

The defender-adversary game in this dissertation is a special case of the stochastic-selection problem (6.1). For illustrative purposes, we here focus on the basic game in Chapter 3 where only the effects of adversary intent are considered. The more complex game in Chapter 5 that further accounts for the effects of adversary capabilities can be handled in a similar manner.

In particular, the defender (i.e., *Blue*) needs to determine how to allocate the limited defensive resources B among N potential targets (or attack strategies) against an attacker (i.e., *Red*), with uncertainty about the attacker's preferences over those targets. For a given defensive allocation plan x, the payoff of attacking target i to an attacker of type ω is given by

$$R_i(x, \omega) = e^{-\lambda x_i} U_i(\omega)$$

where $U_i(\omega)$ is the inherent value of target i to the attacker in the absence of any defense. We assume that the attacker's target valuations $U_i(\omega)$ are not fully known to the defender. This uncertainty is described by a joint probability distribution Q over $(U_1(\omega), \ldots, U_N(\omega))$. In addition, the exponential function $e^{-\lambda x_i}$ specifies the relationship between how much investment has been spent in protecting a target and how likely an attack on that target will be successful. As in Section 3.1, the value λ reflects the cost effectiveness of defense. (For example, if we choose $\lambda = 0.02$, then every unit of additional defensive investment will reduce the success probability of an attack by about 2%.)

On the other hand, the defender's loss in case of an attack on target i is given by

$$l_i(x) = e^{-\lambda x_i} v_i$$

where v_i represents the defender's deterministic disutility of losing target i in the absence of defense. The cost effectiveness λ, defender valuations v_i, and functional forms of $R_i(x, \omega)$ and $l_i(x)$ are all assumed to be common knowledge, except for the attacker's type ω.

Of course, we can also apply stochastic-selection modeling to other fields. For example, a marketing manager (i.e., *Blue*) may wish to allocate resources to various types of marketing efforts (commercials, price promotion, etc.), under uncertainty about which product features may be most influential to the customers (i.e., *Red*); see for example Rust et al. (2004) and Lilien et al. (2007).

6.1.2 Existence of Optimal Solutions

We now present three assumptions which combined are sufficient to ensure the existence of optimal solutions to the stochastic-selection problem (6.1).

Assumption 6.1. *Blue's loss if target i is selected, $l_i(x)$, is continuous on X for $i = 1, \ldots, N$.*

Remark. This assumption also guarantees that *Blue*'s random loss $F(x, \omega)$ as given in (6.1), also termed as the "recourse" function in the stochastic-programming literature, is finite on X and has well-defined expectation and higher moments.

Assumption 6.2. *Red's reward of selecting target i, $R_i(x, \omega)$, is continuous at x on X for almost every*[1] *$\omega \in \Omega$ with respect to the probability distribution Q for $i = 1, \ldots, N$.*

Assumption 6.3. *For every $x \in X$, the set $\{j : \arg\max_j R_j(x, \omega)\}$ is a singleton for almost every $\omega \in \Omega$ with respect to the probability distribution Q.*

[1] A property holds almost everywhere with respect to a probability measure if the set of elements for which the property does not hold is a set of probability zero.

Remark. The last assumption says that given any feasible budget allocation of *Blue*, the event that two targets would yield the same highest payoff to *Red* has probability zero. Note that in the context of counter terrorism this assumption may not be applicable to multiple exchangeable targets (such as warehouses or strip malls), but seems suitable for iconic targets, which are unlikely to have exactly the same attractiveness. Using Assumptions 6.2 and 6.3 together implies that the set of points at which *Red* would change his target selection in response to an infinitesimally small change in defensive allocation by *Blue* has probability zero. In other words, *Red* is unlikely to respond drastically to a negligible perturbation in *Blue*'s decision.

Note that, after significant defensive investment has been made by *Blue*, multiple targets may have the same expected attractiveness under *Blue*'s subjective distribution of *Red*'s preferences. However, they will in general not be equally attractive to *Red*, if *Blue* has a continuous distribution over *Red*'s preferences.

Stochastic-selection problems under these assumptions are capable of covering a wide variety of applications. Note that Assumptions 6.1 and 6.2 are trivial. To satisfy Assumption 6.3 is non-trivial but generally simple in practice. Take the defender-adversary game in Section 6.1.1 for example. The attacker's reward of attacking target i is given by $R_i(x, \omega) = e^{-\lambda x_i} U_i(\omega)$. If we assume continuous distributions for the attacker's target valuations $U_i(\omega)$, then the induced probability distributions for the corresponding reward functions $R_i(x, \omega)$ are also continuous for any fixed $x \in X$. It follows that the probability of $R_i(x, \omega) = R_j(x, \omega)$ is zero for any $i \neq j$. As a result, Assumption 6.3 holds.

The following proposition provides sufficient conditions for ensuring the existence of optimal solutions to the stochastic-selection problem (6.1). Proofs of the lemmas and propositions in this chapter are given in Appendix E.

Proposition 6.1. *Suppose that Assumptions 6.1, 6.2, and 6.3 hold. Then Blue's objective function $f(x)$ in (6.1) is continuous on X. Moreover, $\exists x \in X$ such that $f(x)$ attains its minimum value on X.*

Note that earlier work by Robinson (1993a, 1993b) discusses a similar two-agent, non-zero-sum, and non-cooperative game to model *Blue-Red* interactions in a combat field. However, in that model, *Blue* and *Red* are assumed to decide on their strategies simultaneously without knowing each other's action. Conditions for the existence of equilibrium in a simultaneous game are in fact stronger than those for our sequential game. Putting (6.1) into the context of Robinson (1993) where *Red* is not allowed to observe *Blue*'s budget allocation, we need to further require convexity of $l_i(x)$ to ensure the existence of optimal solutions, in addition to just continuity.

Since *Blue*'s objective function $f(x)$ in (6.1) cannot be expressed in closed form in most cases, we need to estimate it through random sampling. In this chapter, we focus on the widely used Monte Carlo simulation-based stochastic optimization technique, the sample-average approximation (SAA) method.

6.2 Solution Approach

6.2.1 Sample-Average Approximation

The sample-average approximation (SAA) method uses Monte Carlo simulation to estimate the objective function $f(x)$ in (6.1), i.e., the expectation of the random recourse function $F(x, \omega)$ with respect to ω under probability distribution Q. In particular, suppose that we have randomly generated a set of S scenarios for *Red*'s characteristics $\{\omega^{(1)}, \ldots, \omega^{(S)}\}$ from the sample space Ω according to *Blue*'s subjective probability distribution Q. An approximate version of the stochastic-selection problem (6.1) is then given by

$$\begin{aligned} &\min_{x \in X} \hat{f}_S(x) = \frac{1}{S} \sum_{s=1}^{S} F(x, \omega^{(s)}) \\ &\text{where } F(x, \omega^{(s)}) = \sum_{i=1}^{N} z_i^{(s)} l_i(x) \\ &\text{s.t. } z_i^{(s)} = 0 \text{ for } \forall i \notin \arg\max_j R_j(x, \omega^{(s)}), \forall s \\ &\quad \sum_{i=1}^{N} z_i^{(s)} = 1 \text{ for } \forall s \\ &\quad z^{(s)} \in \{0, 1\}^N \text{ for } \forall s \end{aligned} \tag{6.2}$$

Note that the set of S random samples $\{\omega^{(1)}, \ldots, \omega^{(S)}\}$ can be generated by either independent or dependent sampling methods. Studies on the convergence rate and statistical bounds of the SAA method are relatively mature with independent sampling (Shapiro 1993; Shapiro et al. 2009). By contrast, dependent sampling schemes (such as Latin hypercube sampling) have the potential of improving the performance of SAA compared to independent sampling, but limited theoretic work exists on how much improvement we could expect;

see for example Koivu (2005) and Drew and Homem-de-Mello (2012). Since the goal of this chapter is to adapt and assess state-of-the-art optimization techniques for solving stochastic-selection problems rather than theoretic development, we assume that the S scenarios are independent and identically-distributed (IID) samples.

For convenience, we denote by f^* and X^* the true optimal value and the set of true optimal solutions, respectively, to the original problem (6.1). We then denote by $\hat{f}_S^*$ and $\hat{X}_S^*$ the approximate optimal value and the set of approximate optimal solutions, respectively, to (6.2).

We now introduce an additional assumption to achieve semi-continuity of the approximate objective function $\hat{f}_S(x)$ in (6.2). This is useful for ensuring the existence of optimal solutions to the approximate problem (6.2) and for deriving its convergence properties.

Assumption 6.4. *For any given pair of $x \in X$ and $\omega \in \Omega$ such that the set $I(x,\omega) = \{j : \arg\max_j R_j(x,\omega)\}$ is not a singleton, we let the recourse function $F(x,\omega) = l_i(x)$ for some $i \in I(x,\omega)$ with $l_i(x) \leq l_j(x)$ for $\forall j \in I(x,\omega)$.*

Remark. If Assumption 6.3 in Section 6.1 holds and thus the set $I(x,\omega)$ is a singleton with probability one, then we can assign arbitrary values to the recourse function $F(x,\omega)$ for those x and ω that make $I(x,\omega)$ not a singleton. Doing this will not change the value of the true objective function $f(x)$ for any $x \in X$. In that case, the introduction of Assumption 6.4 has no effect on the original problem (6.1).

The following lemma provides sufficient conditions for achieving semi-continuity of the approximate objective $\hat{f}_S(x)$, and for the existence of optimal solutions to the approximate problem (6.2).

Lemma 6.2. *Suppose that Assumptions 6.1, 6.2, 6.3, and 6.4 hold. Then the approximate objective function $\hat{f}_S(x)$ in (6.2) is lower semi-continuous on X for any set of random samples $\{\omega^{(1)}, \ldots, \omega^{(S)}\} \subset \Omega$; i.e., the epigraph* $\operatorname{epi} \widehat{f}_S(\cdot)$ *is a closed set. Moreover, $\exists x \in X$ such that $\hat{f}_S(x)$ attains its minimum value on X for any set of random samples $\{\omega^{(1)}, \ldots, \omega^{(S)}\} \subset \Omega$.*

6.2.2 Consistency and Convergence Rate

We are now ready to discuss convergence properties of the approximate problem (6.2). We start with showing that the approximate objective function $\hat{f}_S(x)$ converges to the true objective function $f(x)$ as the sample size S gets sufficiently large.

Lemma 6.3. *Suppose that Assumptions 6.1, 6.2, and 6.3 hold. Then the approximate objective function $\hat{f}_S(x)$ in (6.2) converges to the true objective function $f(x)$ in (6.1) uniformly on X with probability one as $S \to \infty$.*

Remark. Since the true objective function $f(x)$ is continuous, as shown in Proposition 6.1, we expect more continuity of the approximate objective $\hat{f}_S(x)$ as the sample size S grows.

The following proposition establishes consistency of the approximate optimal value and solutions to (6.2), based on known results of the SAA method (Shapiro et al. 2009).

Proposition 6.4. *Suppose that Assumptions 6.1, 6.2, 6.3, and 6.4 hold. Then $\hat{f}_S^* \to f^*$ with probability one as $S \to \infty$. Moreover, for any $\hat{x}_S^* \in \hat{X}_S^*$, we have $\inf_{x^* \in X^*} |\hat{x}_S^* - x^*| \to 0$ with probability one as $S \to \infty$.*

Remark. Note that Assumptions 6.1 – 6.4 lead to Lemmas 6.2 and 6.3. According to Theorem 3.5 of Shapiro et al. (2009), Lemmas 6.2 and 6.3 combined are sufficient for reaching the conclusions of the above proposition. In fact, Proposition 6.4 is true no matter which sampling scheme is used (e.g., independent or dependent), as long as the sampling distribution converges in distribution to Q with some mild regularity conditions (e.g., Assumption 3.5 in Dupačová and Wets 1988).

Moreover, we are also interested in the rate of convergence in addition to just consistency. Kaniovski et al. (1995) show that if the true objective function $f(x)$ is lower semi-continuous and the approximate objective function $\hat{f}_S(x)$ is random lower semi-continuous, then the probability of $\hat{f}_S^*$ being no larger than f^* (with an ϵ-error) will approach one exponentially fast as S increases. We adapt this result to the case of our stochastic-selection problem in the following proposition.

Proposition 6.5. *Suppose that Assumptions 6.1, 6.2, 6.3, and 6.4 hold. Then* $\forall \epsilon > 0$, $\exists S_0 \in \mathbb{Z}_+$ *such that for any integer* $S \geq S_0$

$$P\{\hat{f}_S^* - f^* \leq \epsilon\} \geq 1 - \exp\{-S \sup_{x^* \in X^*} h_{x^*}(\epsilon)\} = p(S, \epsilon) \tag{6.3}$$

where $h_{x^*}(\theta) = \sup_{\eta \in \mathbb{R}}[\theta\eta - \log \phi_{x^*}(\eta)]$, *and* $\phi_{x^*}(\eta) = E[e^{\eta\{F(x^*,\omega) - f^*\}}]$ *is the moment generating function of* $F(x^*, \omega) - f^*$.

Remark. The probability bound $p(S, \epsilon)$ in the right-hand side of (6.3) is derived by Kaniovski et al. (1995) using the theory of large deviations. However, in practice, $p(S, \epsilon) = 1 - \exp\{-S \sup_{x^* \in X^*} h_{x^*}(\epsilon)\}$ is generally unknown since the set of true optimal solutions X^* is unavailable. We could instead use the set of approximate optimal solutions

$\hat{X}^*_{S'}$ to obtain a (typically biased) estimate $\hat{p}_{S'}(S,\epsilon) = 1 - \exp\{-S \sup_{x^*_{S'} \in X^*_{S'}} \hat{h}_{x^*_{S'}}(\epsilon)\}$, where $\hat{h}_{x^*_{S'}}(\theta) = \sup_{\eta\in\mathbb{R}}[\theta\eta - \log \hat{\phi}_{x^*_{S'}}(\eta)]$, and $\hat{\phi}_{x^*_{S'}}(\eta) = \hat{\mu}_{S''}[e^{\eta\{F(x^*_{S'},\omega) - \hat{f}_{S'}(x^*_{S'})\}}]$; $\hat{\mu}_{S''}$ is the sample mean based on data of size S''. Note that the sample sizes S, S', and S'' could be different.

Proposition 6.5 is also useful for deriving statistical lower bounds for the true objective value f^*. For completeness, we first briefly describe a simple way of generating the upper bound, and then show how to get the lower bound using the above proposition. Specifically, for a candidate solution $x' \in X$ and a reference IID sample of size S' (typically much larger than the sample size S for the approximate problem), we have

$$\sqrt{S'}\{\hat{f}_{S'}(x') - f(x')\} \xrightarrow{\mathscr{D}} \mathscr{N}(0,\ \sigma^2_F(x'))$$

by the central limit theorem, where $\sigma^2_F(x') = \mathbf{Var}[F(x',\omega)]$, and can be estimated by the sample variance $\hat{\sigma}^2_{S'}[F(x',\omega)] = \frac{1}{S'-1}\sum_{s=1}^{S'}\left(F(x',\omega^{(s)}) - \hat{f}_{S'}(x')\right)^2$. For a sufficiently large reference sample size (e.g., $S' \geq 10^6$), we are confident that

$$P\left\{f^* - \hat{f}_{S'}(x') - \frac{\Phi(0.99)\hat{\sigma}_{S'}[F(x',\omega)]}{\sqrt{S'}} \leq 0\right\} \geq 0.99 \tag{6.4}$$

where $\Phi(\cdot)$ is the inverse distribution function of the standard normal.

We now describe how to use Proposition 6.5 to get statistical lower bounds for the true objective value f^*. Suppose that we have solved T independent replications of the approximate problem (6.2), then we can use the maximum of these approximate objective values $\max_{t=1,\ldots,T} \hat{f}^*_{S,t}$ as the statistical lower bound for f^*. According to (6.3), we have

the following inequality

$$P\left\{f^* - \max_{t=1,\ldots,T} \hat{f}^*_{S,t} \geq -\epsilon\right\} \geq [p(S,\epsilon)]^T \tag{6.5}$$

For a given level of confidence $1-\alpha$, we can then deduce the required sample size S by setting $1-\alpha = [p(S,\epsilon)]^T$.

The nice feature of deriving lower bounds in this way is that we only need to solve a small number T of replications of the approximate problem (6.2). We could even set $T=1$ by solving only one approximate problem. However, it may possibly require a quite large sample size S for a practical confidence level (e.g., $1-\alpha \geq 95\%$), which would potentially be computationally burdensome. An alternative to obtain lower bounds for the true optimal value f^* is to solve a large number T of replications of the approximate problem (6.2), each with a relatively small sample size S; see details in the following section.

6.2.3 Replication-Based Lower Bound

Shapiro and Homem-de-Mello (2000) show that the expectation of the approximate optimal value $\mathbf{E}[\hat{f}^*_S]$ provides an exact lower bound for the true optimal value f^*; i.e., $\mathbf{E}[\hat{v}^*_S] \leq f^*$. This argument is true by requiring only the recourse function $F(x,\omega)$ to be finite almost surely, so Assumption 6.1 solely will suffice for our case of stochastic-selection problems.

Suppose that we have solved T independent replications of the approximate problem (6.2), each with the same sample size S. We could then estimate the expectation $E[\hat{f}^*_S]$ by the sample mean of the T approximate optimal values $\hat{\mu}_T[\hat{f}^*_S] = \frac{1}{T}\sum_{t=1}^{T} \hat{f}^*_{S,t}$. Since the T replications of the approximate problem are independently created, we can apply the central

limit theorem as follows

$$\sqrt{T}\{\hat{\mu}_T[\hat{f}_S^*] - \mathbf{E}[\hat{f}_S^*]\} \xrightarrow{\mathscr{D}} \mathscr{N}(0,\ \sigma_f^2)$$

where $\sigma_f^2 = \mathbf{Var}[\hat{f}_S^*]$, which can be estimated by $\hat{\sigma}_T^2[\hat{f}_S^*] = \frac{1}{T-1}\sum_{t=1}^{T}\left(\hat{f}_{S,t}^* - \hat{\mu}_T[\hat{f}_S^*]\right)^2$. Therefore, for a given confidence level $1-\alpha$, we can derive a statistical lower bound for the true objective value f^* in the following way

$$P\left\{f^* - \hat{\mu}_T[\hat{f}_S^*] - \frac{\Phi(1-\alpha)\hat{\sigma}_T[\hat{f}_S^*]}{\sqrt{T}} \geq 0\right\} \geq 1-\alpha \tag{6.6}$$

The idea of using large number of replications (T) rather than large sample size (S) is especially suitable for parallel computing (Janjarassuk and Linderoth 2008). This feature would significantly reduce the computing time required for large-scale problems. However, one caveat is that the lower bound derived from approximate problems with small S may not be close enough to the true optimal value f^*, no matter how many replications are implemented (i.e., how big T is), hindering such bounds from being extremely useful for proving optimality.

6.3 Algorithms and Software

In this section, we investigate state-of-the-art computational algorithms for solving the approximate problem (6.2). Note that the statistical lower bounds (6.5) and (6.6) are valid only if each replication of the approximate problem can be solved exactly. Therefore, we aim to seek for algorithms that can obtain global solutions with proved optimality. However,

we also discuss heuristic approaches that are able to provide near-optimal solutions within a reasonable time constraint. In particular, two classes of solution approaches are considered – namely, mixed-integer nonlinear programming (MINLP) and derivative-free optimization (DFO).

6.3.1 Mixed-Integer Nonlinear Programming

We first cast the approximate problem (6.2) as a mixed-integer nonlinear programming (MINLP) model. For convenience, since *Blue*'s loss functions $l_i(x)$ and *Red*'s reward functions $R_i(x, \omega)$ are both finite on X, we normalize them to take on values in $[0, 1]$. Then we have the following MINLP model (note that the notations y and w in this chapter have nothing to do with the multiattribute adversary utility in Chapters 3 and 4)

$$
\begin{aligned}
&\min_{x \in X, y_s, w_s, z^{(s)}} \frac{1}{S} \sum_{s=1}^{S} y_s \\
\text{s.t. } & y_s \geq l_i(x) + z_i^{(s)} - 1 && \forall i, s \\
& w_s - R_i(x,\ \omega^{(s)}) \geq 0 && \forall i, s \\
& w_s - R_i(x,\ \omega^{(s)}) \leq 1 - z_i^{(s)} && \forall i, s \\
& \sum_{i=1}^{N} z_i^{(s)} = 1 && \forall s \\
& z^{(s)} \in \{0,\ 1\}^N && \forall s
\end{aligned}
\tag{6.7}
$$

Blue's budget allocation $x \in X$ is determined in the first stage, independent of the scenarios $\omega^{(s)}$. Second-stage variables are determined after the scenarios $\omega^{(s)}$ are fully revealed, including *Red*'s binary choices $z^{(s)} \in \{0, 1\}^N$ and the nuisance variables y_s and w_s

for $s = 1, \ldots, S$. Note that the model (6.7) has a relatively complete recourse; i.e., for every feasible first-stage variable $x \in X$ and almost every set of scenarios $\omega^{(s)} \in \Omega$ $(s = 1, \ldots, S)$ with respect to the probability distribution Q, there exist second-stage variables $z^{(s)}$, y_s, and w_s that satisfy all constraints of (6.7).

We can use the general-purpose global optimization software BARON (Branch-And-Reduce Optimization Navigator; Sahinidis and Tawarmalani 2011) to solve the MINLP model (6.7). Global optimum can be both obtained and proved by the current version of BARON, if the MINLP set-up involves only specific forms of nonlinearity including multiplications and divisions, as well as exponential, logarithmic, and absolute-value functions.

In addition, a number of other MINLP solvers are also available, such as DICOPT (DIscrete and Continuous OPTimizer; Viswanathan and Grossmann 1990), SBB (Simple Branch and Bound; Bussieck and Drud 2001) and AlphaECP (Extended Cutting Plane; Westerlund and Pörn 2002). However, these solvers are capable of producing and proving global optimum only when certain conditions of convexity hold; e.g., if the continuous relaxation of (6.7) is convex (or pseudo-convex). In our case, such conditions are satisfied if ***Blue***'s loss functions $l_i(x)$ are convex in x, and ***Red***'s reward functions $R_i(x, \omega)$ are linear in x for almost every $\omega \in \Omega$ with respect to Q. It also suffices if $R_i(x, \omega)$ can be transformed to linear functions of x while preserving the ordinal relationships among the various targets for any $x \in X$ and $\omega \in \Omega$. Take the defender-adversary example in Section 6.1.1. If we have $R_i(x,\ \omega) = e^{-\lambda x_i} U_i(\omega)$ with $U_i(\omega) > 0$ almost surely, then we could take its logarithm and get a linear function $\overline{R}_i(x, \omega) = -\lambda x_i + \ln U_i(\omega)$.

6.3.2 Piecewise-Linear Approximation

If the loss and reward functions, $l_i(x)$ and $R_i(x, \omega)$ are separable in x, we can also apply the piecewise-linear approximation (PLA) technique to remove any non-linearity from the MINLP model (6.7). Take ***Blue***'s loss functions $l_i(x)$ for instance. If $l_i(x)$ is separable, then we can write $l_i(x) = \sum_{j=1}^{N} l_{ij}(x_j)$, and approximate each component $l_{ij}(x_j)$ by, for example, three line segments defined on $x_j \in [0, B]$, as given by

$$\begin{aligned} l_{ij}(x_j) \approx & c_1\gamma_1 + c_2\gamma_2 + c_3\gamma_3 + c_4\gamma_4 \\ x_j = & a_1\gamma_1 + a_2\gamma_2 + a_3\gamma_3 + a_4\gamma_4 \\ & \gamma_1 + \gamma_2 + \gamma_3 + \gamma_4 = 1 \\ & (\gamma_1, \gamma_2, \gamma_3, \gamma_4) \in \text{SOS2} \\ & \gamma_1, \gamma_2, \gamma_3, \gamma_4 \geq 0 \end{aligned} \tag{6.8}$$

where a_1, a_2, a_3, a_4 and c_1, c_2, c_3, c_4 are parameters defining the line segments. Note that γ_1, γ_2, γ_3, γ_4 are special-ordered set type 2 (SOS2) variables; i.e., at most two adjacent variables can be non-zero (Williams 1999). We can then solve an approximate version of (6.7) using mixed-integer linear programming (MILP) solvers such as the IBM Ilog CPLEX. Of course, when the component functions $l_{ij}(x_j)$ are convex, we can also approximate each of them as the maximum of four linear functions $\max_{q=1,2,3,4} l_{ijq}(x_j)$, without the need for SOS2 variables.

Although the use of PLA (either with or without SOS2 variables) is able to guarantee global optimality, the quality of solutions would heavily depend on accuracy of the involved linear approximations, and is somewhat intractable. Moreover, even though MILP solvers

are generally more efficient and reliable for large-scale problems than MINLP solvers, the problem size (e.g., numbers of variables and constraints) of model (6.7) will still explode as the sample size S increases.

6.3.3 Derivative-Free Optimization

In this section, we pursue a different avenue by reformulating the approximate problem (6.2) as a derivative-free optimization (DFO) model. We observe that for any given budget allocation $x \in X$, the approximate objective function $\hat{f}_S(x)$ in (6.2) is cheap to calculate even with a fairly large sample size (e.g., $S \geq 10^6$). Therefore, a derivative-free approach may be preferable. We now convert the approximate problem (6.2) into a DFO model with bounded constraints, as given by

$$\min_{x \in [0,B]^N} \hat{d}_S(x) = \frac{1}{S} \sum_{s=1}^{S} F(x, \omega^{(s)}) + M \max\{0, \sum_{i=1}^{N} x_i - B\} \tag{6.9}$$

where $F(x, \omega^{(s)})$ is the recourse function as given in (6.2), and $M > 0$ is a big penalty number. Under Assumptions 6.1–6.4, for a given set of samples $\omega^{(s)}$ for $s = 1, \ldots, S$, the objective function $\hat{d}_S(x)$ in (6.9) is lower semi-continuous but generally non-differentiable and non-convex.

State-of-the-art DFO algorithms are thoroughly investigated in Rios and Sahinidis (2012). In particular, they classify DFO algorithms as direct (i.e., search directions are determined by calculating the objective values directly) or model-based (i.e., local surrogate models are constructed to guide searches); local or global (depending on whether the algorithm is able to escape from local optima or not); deterministic or stochastic (depending on whether

random search steps are required or not). Rios and Sahinidis (2012) also observe that some algorithms that are originally designed for global search may perform better than local solvers even for convex objective functions.

In this chapter, we apply only a selective set of global DFO solvers that are shown by Rios and Sahinidis (2012) to outperform others on average. In particular, we consider deterministic solvers such as DIRECT (DIviding RECTangles; Jones 2001), MCS (Multilevel Coordinate Search; Huyer and Neumaier 1999), and stochastic solvers such as SNOBFIT (Stable Noisy Optimization by Branch and FIT; Huyer and Neumaier 2008) and MLSL (Multi-Level Single-Linkage; Rinnooy Kan and Timmer 1987). Among them SNOBFIT utilizes statistical surrogate models while others use the objective values directly to guide searches. Moreover, we also couple simulated annealing (Kirkpatrick et al. 1983) with local solvers such as BOBYQA (Bound Optimization BY Quadratic Approximation; Powell 2009), and Subplex (or SBPLX), a variant of Nelder-Mead (1965) by Rowan (1990). The last hybrid approach is selected because it showed good performance in both the literature (e.g., Martin and Otto 1993) and our preliminary studies on the stochastic-selection problem.

One caveat is that even if there is some sort of convexity with the objective $\hat{d}_S(x)$ in (6.9), none of these global DFO algorithms is able to provide measures for the quality of any found solution. Therefore, we treat these algorithms as heuristic approaches.

6.3.4 Summary of Selected Solvers

We consider four MINLP solvers for obtaining solutions to model (6.7), including BARON, DICOPT, SBB, and AlphaECP. By applying piecewise-linear approximation, the MINLP

model (6.7) then reduces to a MILP model, and is solved by CPLEX. All the above solvers are available to implement on the platform of GAMS (http://www.gams.com/).

As for the DFO solvers, we call DIRECT, BOBYQA, Subplex, and MLSL from the open-source MATLAB NLopt library 2.2.4 (Johnson 2011). Matlab codes for MCS and SNOBFIT are obtained separately from the link (http://www.mat.univie.ac.at/ neum/software/). The implementation of simulated annealing is coded from scratch in MATLAB (see Appendix E.6 for the algorithm adapted to our case).

The following table summarizes all solvers considered in this chapter, in terms of whether they converge to global optimum or not, and whether they can prove global optimality or not. Note again that the statistical bounds in Section 6.2 are valid only if global optimum to the approximate problem (6.2) can be both obtained and proved.

Table 6.1: Summary of Selected Solvers

<table>
<tr><th colspan="2">Model</th><th>Solver</th><th>Convergence to Global Optimum</th><th>Proof of Global Optimum</th></tr>
<tr><td rowspan="5">MINLP Model</td><td rowspan="4">MINLP</td><td>BARON</td><td>√</td><td>√</td></tr>
<tr><td>DICOPT</td><td rowspan="3">√
(only if relaxation is convex)</td><td rowspan="3">√
(only if relaxation is convex)</td></tr>
<tr><td>SBB</td></tr>
<tr><td>AlphaECP</td></tr>
<tr><td>Reduced to MILP</td><td>CPLEX</td><td>√</td><td>√</td></tr>
<tr><td rowspan="6">DFO Model</td><td rowspan="2">Deterministic</td><td>DIRECT</td><td rowspan="2">√</td><td rowspan="2">×</td></tr>
<tr><td>MCS</td></tr>
<tr><td rowspan="4">Stochastic</td><td>SNOBFIT</td><td rowspan="4">√</td><td rowspan="4">×</td></tr>
<tr><td>MLSL</td></tr>
<tr><td>SA-BOBYQA</td></tr>
<tr><td>SA-SBPLX</td></tr>
</table>

6.4 Numerical Results

We now use the defender-adversary example in Section 6.1.1 as a test problem to compare the performance of the various computational algorithms. Our first goal is to identify best algorithms for solving the approximate problem (6.2) with varying numbers of targets N and sample sizes S. In addition, we are also interested in finding suitable combinations of the sample size S and the number of replications T for the SAA method, in order to generate good optimality bounds for the true optimal value f^* to the original problem (6.1).

Practically we hope to solve a problem with 50 decision variables (corresponding to the defender's budget allocations to 50 targets), and a sample space of dimension 50 (representing defender uncertainty about the adversary's valuations of the 50 targets), within a reasonable time constraint. Specifically, assuming that parallel computing is possible, we would like each optimization instance (i.e., each replication of the approximate problem) to be solved within six to ten CPU minutes.

6.4.1 Test Problem

Following the example in Section 6.1.1, the attacker (i.e., *Red*) is assumed to get a reward of $R_i(x, \omega) = e^{-\lambda x_i} U_i(\omega)$ by attacking target i. We set the random component $U_i(\omega) = H_i(\omega) u_i$ for each $i = 1, \ldots, N$, where u_i is the average payoff of successfully destroying target i to the attacker, and $H_i(\omega)$ is independently and identically distributed as Fréchet (2, 1). The exponential function $e^{-\lambda x_i}$ describes the impact of defense on the success probability of an attack, where $\lambda = 0.02$ is the cost effectiveness.

On the other hand, we assume that the defender (i.e., *Blue*) will suffer a loss of $l_i(x) = e^{-\lambda x_i} v_i$ if target i is attacked, where v_i is the loss from an attack on target i in the absence of defense. Note that the $l_i(x)$ are convex in x and the $R_i(x, \omega)$ can be logarithmically transformed to linear functions of x as $\overline{R}_i(x, \omega) = -\lambda x_i + \ln\{H_i(\omega) u_i\}$. Therefore, the continuous relaxation of the MINLP model (6.7) is convex.

This test problem is particularly chosen also because the true objective function $f(x)$ can be expressed in closed form as given by

$$\min_{x \in X} f(x) = \sum_{i=1}^{N} \frac{e^{-0.06 x_i} u_i^2}{\sum_{j=1}^{N} e^{-0.04 x_j} u_j^2} v_i \tag{6.10}$$

The above problem can then be solved exactly using the global nonlinear programming (NLP) solver BARON, in order to evaluate the quality of statistical optimality bounds generated by the SAA method.

6.4.2 Comparison of Solvers

We first compare the performance of different solvers on the approximate optimization problem (6.2). In particular, we are interested in the execution time required by each solver to reach a solution x' that can yield an sample-average objective value $\hat{f}_S(x')$ within 1% of the approximate optimal value $\hat{f}_S^*$. The approximate optimal values $\hat{f}_S^*$ are obtained and proved by the fastest of the various MINLP solvers (e.g., BARON, DICOPT, SBB, or AlphaECP). In addition, the piece-wise linear approximations are solved and proved separately using CPLEX. All computations are performed on a 64-bit Linux workstation with eight 2.66Ghz processors and 16 Gb memory. MATLAB R2012a version 7.14.0.739

and GAMS version 24.0.1 are used. Default settings are kept for all solvers, except that we change the stopping criteria of DICOPT to "1" that can facilitate global convergence of convex MINLP problems (whose continuous relaxations are convex). Running times are reported in CPU seconds.

Figures 6.1 and 6.2 show the results for $N = 5$ and $N = 50$, respectively. In each case, we consider two sample sizes; e.g., $S = 20$ or 100 for $N = 5$, and $S = 10$ or 50 for $N = 50$. Twenty test instances (i.e., values of the defender's target valuations v_i and the attacker's expected target valuations u_i for $i = 1, \ldots, N$) are randomly generated for each design of parameter settings. Common random numbers are used across different solvers. A time limit is set at 360 CPU seconds for each approximate problem. Note that the DFO solvers DIRECT and MLSL are not shown in these figures, because they were not able to provide a 1%-optimal solution for any test instance within the 360-second time constraint.

When the number of targets and the sample size are both small (e.g., $N = 5$ and $S = 20$), as shown in Figure 6.1(a), all solvers but SBB are able to produce a good solution within 60 CPU seconds. However, as the sample size increases (e.g., $N = 5$ and $S = 100$), as shown in Figure 6.2, the DFO algorithms seem to outperform the mixed-integer solvers. In particular, efficiency of all MINLP or MILP solvers is significantly limited by the sample size S, while performance of the various DFO solvers is less sensitive to the sample size S.

Similar patterns are observed for the MINLP and MILP solvers when the number of targets is large (e.g., $N = 50$), as shown by comparing the left panels of Figures 6.2 (a) and (b). Larger sample sizes again lead to poorer performance for these mixed-integer solvers. By contrast, we observe an interesting phenomena associated with the DFO solvers. They seem to work better when the number of sample size S is larger (e.g., $S = 50$), as shown

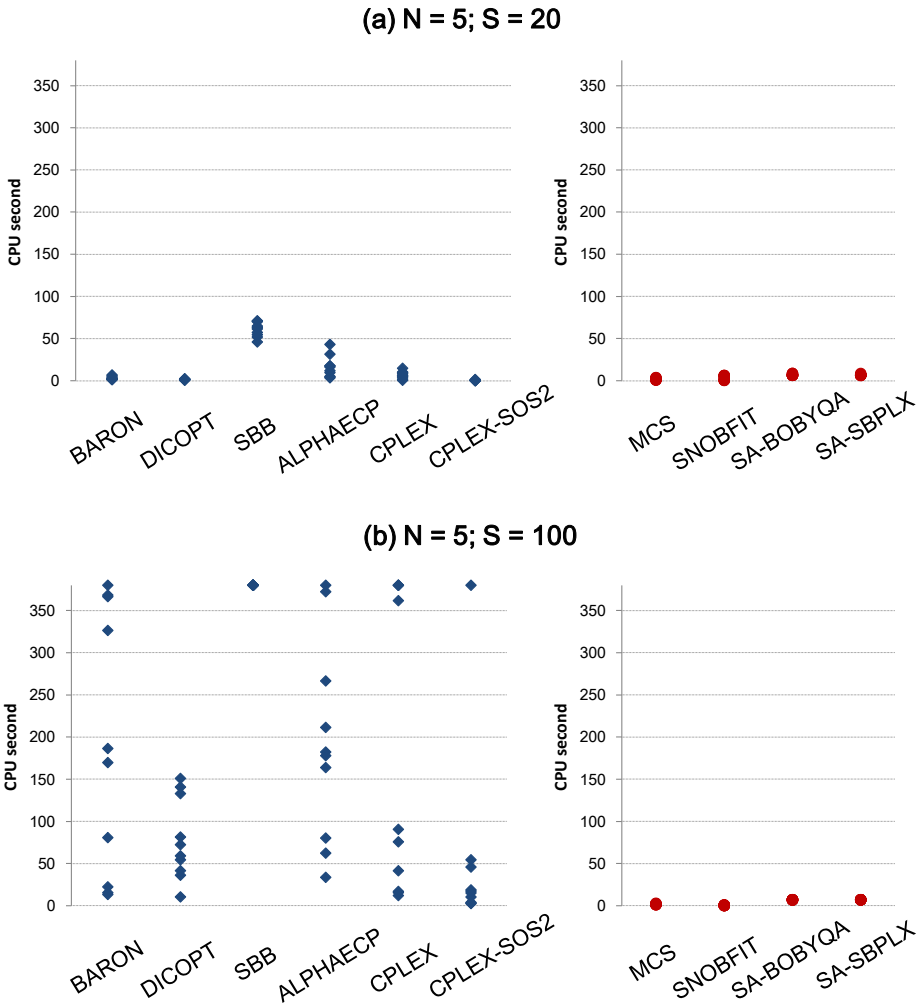

Figure 6.1: Execution Time for N = 5

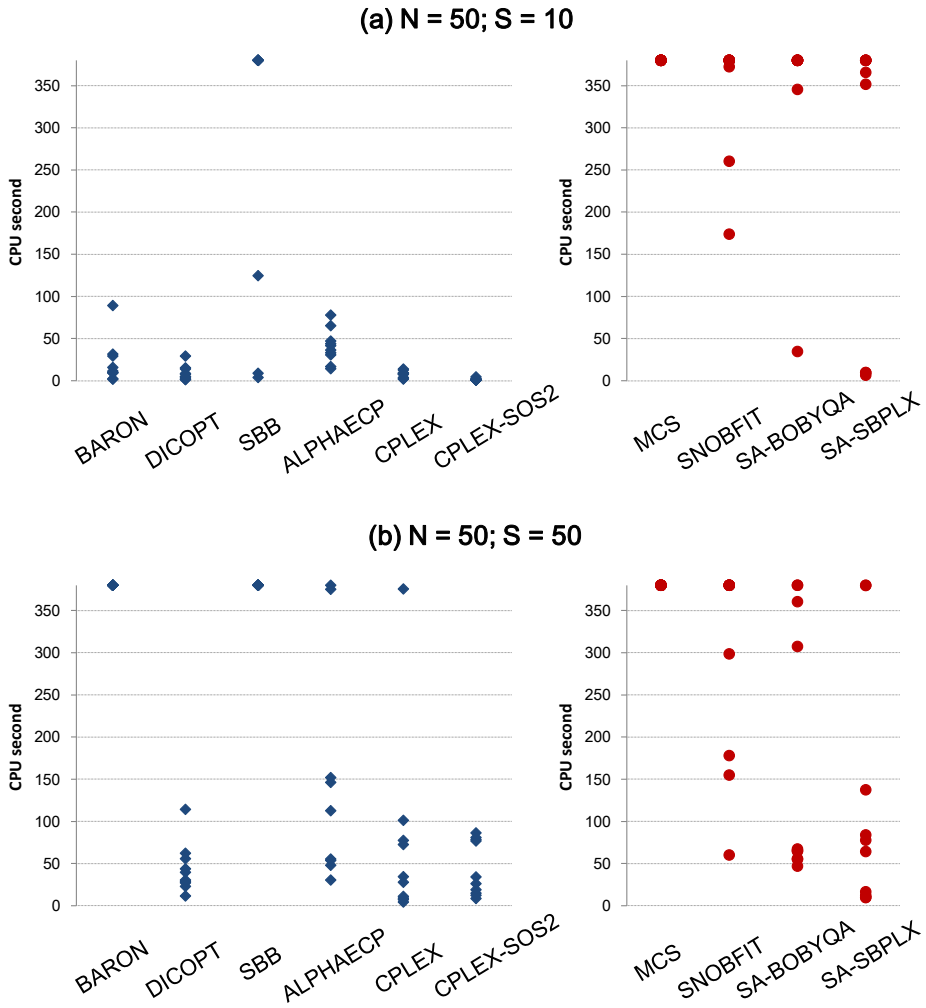

Figure 6.2: Execution Time for N = 50

in Figure 6.2(b), especially the SA-SBPLX solver. This is reasonable since DFO solvers may have improved performance if the objective function to optimize is more continuous or smoother, which is the case as the sample size S increases.

Unfortunately, we have not been able to obtain and prove global optimality for the important case of interest with $N = 50$ and a large sample size $S > 10^5$. In general, MINLP or MILP solvers become ineffective when the sample size S is greater than 10^3 even if we relax the time limit to 1200 CPU seconds. However, some of the DFO solvers are shown to yield 1%-optimal solutions within 360 CPU seconds in many cases with $N = 50$ and $S = 50$, as shown in Figure 6.2. We thus anticipate that the DFO solvers would be capable of solving problems with the same N but larger S at a similar or better performance level (since they seem to work better with larger S).

Therefore, we compare the best possible solutions each DFO solver can find within the same time constraint. In particular, we count the cumulative number of instances when the minimal approximate objective value achieved by a particular solver is within 5% of the best value among all DFO solvers. Figures 6.3 and 6.4 show the results for $N = 5$ and $N = 50$, respectively, both with a large sample size ($S = 10^4$ for $N = 5$, and $S = 10^5$ for $N = 50$). We set the termination time for the case of $N = 5$ at 60 CPU seconds, while that for the case of $N = 50$ at 360 CPU seconds.

When the number of targets is small (e.g., $N = 5$), as shown in Figure 6.3, SNOBFIT and MCS are the best solvers, followed by SA-SBPLX and SA-BOBYQA. By contrast, when the number of targets is large (e.g., $N = 50$), as shown in Figure 6.4, SA-SBPLX and SNOBFIT constantly perform better than others, while MCS is able to achieve the best performance in only limited cases. Surprisingly, our implementation of SA-BOBYQA does

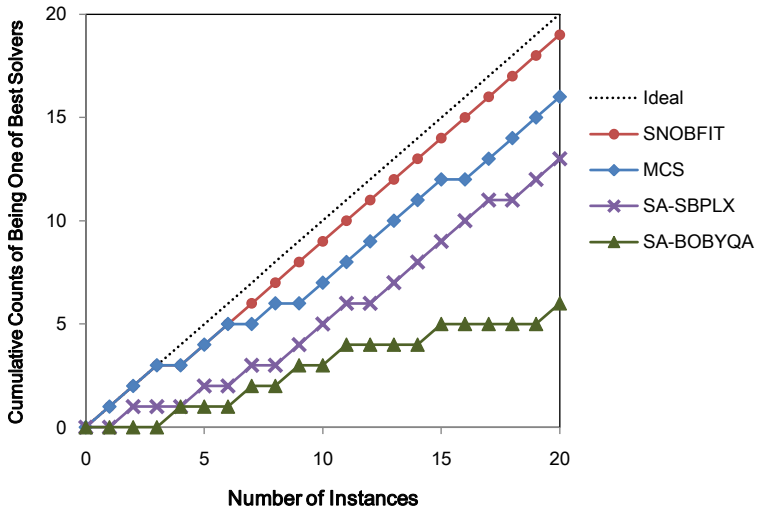

Figure 6.3: Cumulative Counts of Being One of the Best Solvers (N = 5; S = 10000)

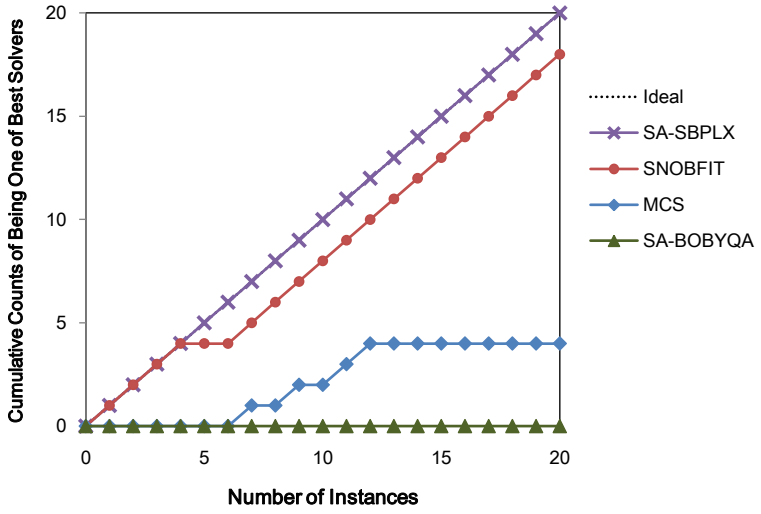

Figure 6.4: Cumulative Counts of Being One of the Best Solvers (N = 50; S = 100000)

Table 6.2: Best Solvers for Stochastic-Selection Problems with Convex MINLP Relaxation

N	S	Solvers of Choice
5	< 100	BARON, DICOPT, CPLEX-SOS2
5	100 - 1000	DICOPT, CPLEX-SOS2
5	> 1000	SNOBFIT, MCS
50	< 50	DICOPT, CPLEX-SOS2, CPLEX
50	50 - 100	DICOPT, CPLEX-SOS2
50	> 100	SA-SBPLX, SNOBFIT

not prevail when $N = 50$, although BOBYQA is believed to perform especially well in high-dimensional (despite local) optimization problems. We speculate that it is because the approximate objective function is not twice-differentiable, or perhaps $N = 50$ is not large enough for BOBYQA to demonstrate its strength.

Table 6.2 presents the best solvers for different cases of N and S. We select those algorithms based on the following two principles. First, when MINLP or MILP solvers and DFO solvers require similar time to generate near-optimal solutions, we prefer MINLP or MILP solvers, since they can also possibly provide an informative global optimality gap to assess the quality of currently found solutions. Second, when MINLP and MILP solvers are equally fast in yielding solutions of similar quality, we prefer MINLP solvers since no added intractability due to piece-wise linear approximation is involved.

In general, we should formulate the approximate stochastic-selection problem (6.2) as an MINLP model if the sample size S is small, and then solve it using DICOPT, or BARON if the number of targets N is also small. We could further apply piece-wise linear approximation to the MINLP model by using SOS2 variables, and reduce it to an MILP model that is solvable via CPLEX. By contrast, if we require a large sample size S, then

DFO solvers are better in finding good heuristic solutions. For example, we could use SNOBFIT and MCS for problems with small numbers of targets N, while SA-SBPLX and SNOBFIT for problems with large N.

6.4.3 Finding Suitable SAA Settings

We have two different ways of trading off the sample size S with the number of replications T when applying the SAA method. In particular, we could solve the approximate problem (6.2) with a large sample size S for only a small number of replications T. Alternatively, we could also solve the approximate problem with a moderate or small sample size S, but for a large number of replications T. We here propose a practical rule-of-thumb procedure to identify a suitable combination of S and T, and illustrate the procedure using a realistic defender-adversary game where the defender's objective is to protect 47 US urban areas against a possible terrorism attack.

We assume that the attacker's expected target valuations u_i equal to the expected property damages of the various urban areas from terrorism, and the defender's target valuations v_i equal to the population densities (see Table 6.3). These data are taken from Willis et al. (2005). The total defensive budget is \$700 million (which roughly equals the annual budget of the Urban Area Security Initiative program, DHS 2011), and the cost effectiveness of defense is $\lambda = 0.01$ per million dollars.

Since the defender's objective function $f(x)$ is available in closed form, we can use BARON to obtain a global optimal solution to the original problem (6.1). Specifically, the true optimal objective value is $f^* = 668.587$, and the true optimal allocations are $x_1^* = 133.375$ (New York), $x_2^* = 57.261$ (Chicago), $x_3^* = 47.942$ (San Francisco), $x_4^* = 13.499$

Table 6.3: Attacker and Defender Target Valuations of 47 Major U.S. Cities

Urban Area	Property Damage ($million) Attacker Value u_i	Population Density (per sq mile) Defender Value v_i
New York, NY	413	8,159
Chicago	115	1,634
San Francisco	57	1,705
Washington, DC-MD-VA-WV	36	756
Los Angeles-Long Beach	34	2,344
Philadelphia, PA-NJ	21	1,323
Boston, MA-NH	18	1,685
Houston	11	706
Newark, NJ	7.3	1,289
Seattle-Bellevue-Everett	6.7	546
Jersey City	4.4	13,044
Detroit	4.2	1,140
Las Vegas, NV-AZ	4.1	40
Oakland, CA	4	1,642
Orange County, CA	3.7	3,606
Houston	3	832
San Diego	2.8	670
Miami, FL	2.7	1,158
Minneapolis-St. Paul, MN-WI	2.7	490
Denver	2.5	561
Baltimore	2.4	979
Atlanta	2.3	672
Dallas	2.1	569
St. Louis, MO-IL	2.1	407
Portland-Vancouver, OR-WA	2	381
Phoenix-Mesa	1.9	223
San Jose	1.7	1,304
Charlotte-Gastonia-Rock Hill, NC-SC	1.1	444
Kansas City, MO-KS	1.1	329
Milwaukee-Waukesha, WI	1.1	1,028
New Haven-Meriden, CT	1.1	1,261
Buffalo-Niagara Falls	1	747
Pittsburgh	1	510
Cincinnati, OH-KY-IN	0.9	493
Tampa-St. Petersburg-Clearwater	0.9	938
New Orleans	0.8	394
Columbus, OH	0.7	490
Indianapolis	0.7	456
Sacramento	0.7	399
Louisville, KY-IN	0.6	495
Orlando	0.6	471
Memphis, TN-AR-MS	0.5	378
Albany-Schenectady-Troy	0.4	272
Richmond-Petersburg	0.4	338
San Antonio	0.4	479
Baton Rouge	0.2	380
Fresno	0.2	114

(Washington, DC), $x_5^* = 44.696$ (Los Angeles), $x_6^* = 18.404$ (Philadelphia), $x_7^* = 20.041$ (Boston), $x_{11}^* = 14.782$ (Jersey City), and $x_i^* = 0$ for all other cities.

We first apply the statistical lower bound (6.5) derived from the theory of large deviations. In particular, for a given confidence level $1 - \alpha$ and a given number of replications T, we have

$$P\left\{f^* - \max_{t=1,\ldots,T} \hat{f}_{S,t}^* \geq -\epsilon\right\} = [p(S,\epsilon)]^T = 1 - \alpha$$

where $p(S,\epsilon)$ is as given in (6.5). Suppose that we set $\alpha = 0.05$ and $T = 5$. The above equation then requires that $p(S,\epsilon) \geq 0.99$. By fixing $\epsilon = 0.01$, we can estimate $\hat{p}(S, 0.01) = 1 - \exp\{-S \sup_{x^* \in X^*} \hat{h}_{x^*}(0.01)\}$ as a function of the sample size S, where $\hat{h}_{x^*}(0.01)$ is as given in Proposition 6.5, and identify a suitable S to make $\hat{p}(S, 0.01) \geq 0.99$ hold.

In this particular test case, we have the true optimal solution x^* available. In practice, however, we will need to use the approximate solution $x_{S'}^*$ for some S' to replace x^*. Note also that there is no guarantee for x^* to be the unique optimal solution. For the purpose of finding optimality bounds only, we conveniently assume that the global optimum is unique. (This will in fact lead to a conservative lower bound, since the possible presence of other optimal solutions could produce a faster convergence rate of $p(S,\epsilon)$.) We can calculate that $\hat{h}_{x^*}(0.01) = 7.25 \times 10^{-4}$, and thus $\hat{p}(S, 0.01) = 1 - \exp\{-7.25 \times 10^{-4} S\}$. Figure 6.5 shows the values of $\hat{p}(S, 0.001)$ as a function of the sample size S.

It follows that to achieve a confidence level of 0.95 with $T = 5$ independent replications of the approximate problem, the required sample size of each replication is $S_p = 6352$. In other words, if we solve the approximate problem $T = 5$ times, each with the same sample

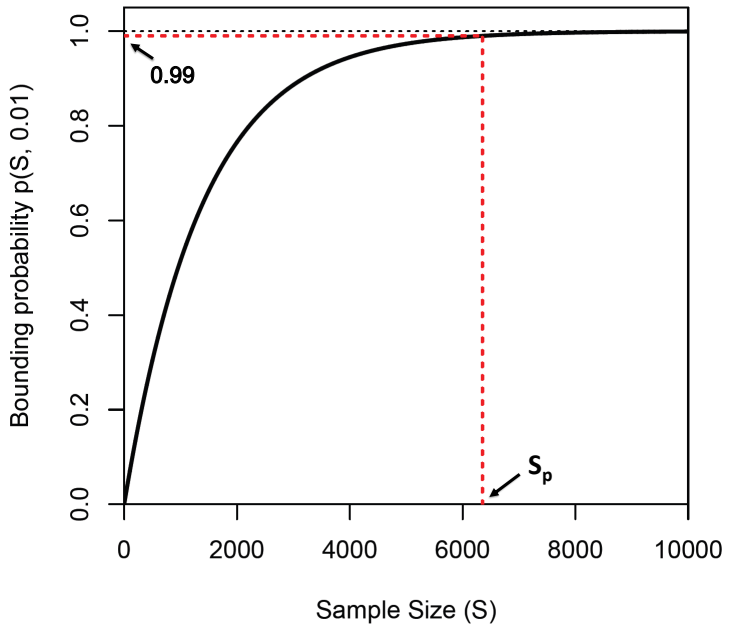

Figure 6.5: Convergence Rate of the Approximate Optimal Value

size $S \geq 6352$, then there is at least a 95% chance that the maximum of the five approximate optimal values $\max_{t=1,\ldots,5} \hat{f}^*_{S,t}$ should bound the true optimal value f^* from below.

Figure 6.6 represents the optimal objective values $\hat{f}^*_S$ for $T = 5$ independent replications of the approximate problems with sample size $S = 8000 > S_p$. They are solved by the DFO solver SA-SBPLX by setting a time limit at 1200 CPU seconds. Although SA-SBPLX is a heuristic optimization algorithm unable to provide exact global solutions, we anticipate it to yield near-optimal solutions at the level of 1% on average, according to the empirical experiences in Section 6.4.2.

All $T = 5$ approximate optimal values $\hat{f}^*_{S,t}$ are smaller than the true optimal objective value 668.587. The maximum of the approximate optimal values $\max_{t=1,\ldots,5} \hat{f}^*_{S,t}$ provides a

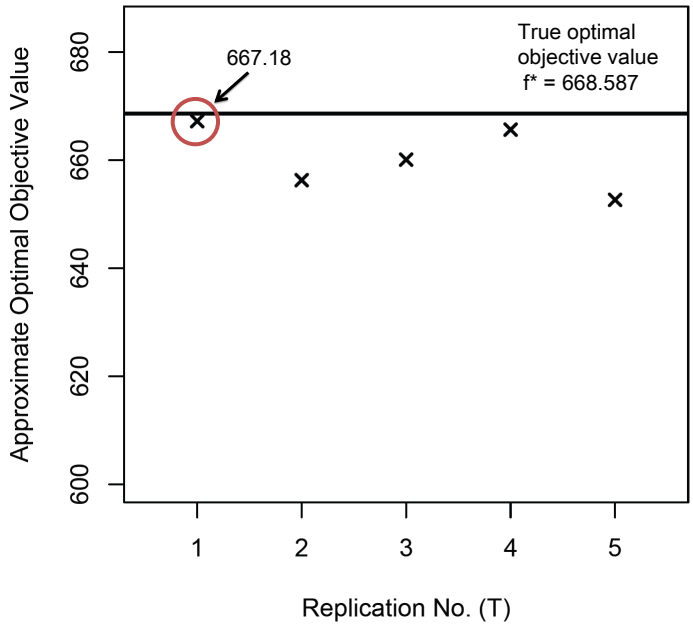

Figure 6.6: Statistical Lower Bound Derived from Theory of Large Deviations

quite tight statistical lower bound for the true optimal value, in spite of the fact that these values may be inflated by about 1% due to the heuristic optimization solver.

There are cases where solving the approximate problem with a large sample size S might be impractical, but computing resources are adequate for large numbers of replications T, for example, via a computing grid. In that case, the replication-based lower bound (6.6) should be used. As a rule of thumb, we propose to first try to set the sample size $S = S_p/4$, and increase the number of replications T until the bound as given in (6.6) gets stable. It is possible that this choice of S may be too small to produce a tight bound even if T is sufficiently large, so S needs to be increased. Otherwise, the value of S could be further reduced.

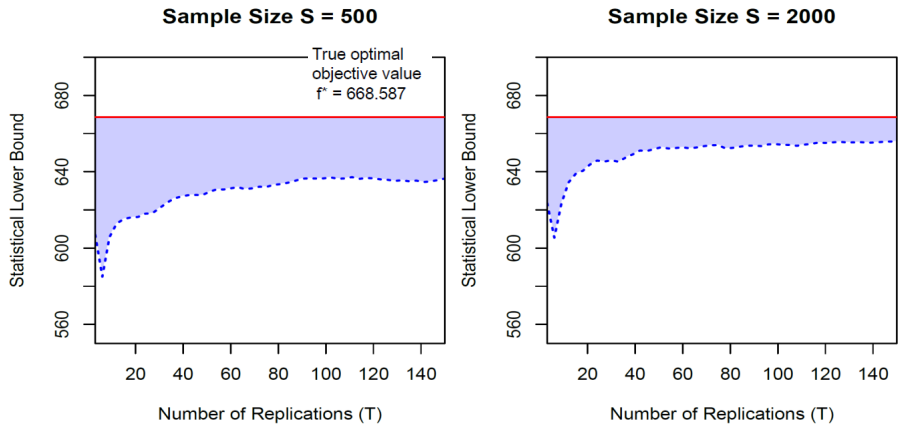

Figure 6.7: Replication-Based Lower Bound

We therefore present the replication-based statistical lower bound for two different sample sizes, $S = 500$ and $S = 2000$ in Figure 6.7. In this particular case, $T > 150$ is required for this bound to converge, but the converged value may be far below the true optimal value if the sample size S is too small (e.g., $S = 500$). By contrast, it appears that a larger sample size (e.g., $S = 2000$) can give a tighter statistical lower bound, and thus more capable of proving optimality.

6.5 Summary

This chapter investigates rigorous and efficient optimization techniques for solving the defender-adversary game of incomplete information. We first generalize the game as a stochastic-selection problem, and model it as a stochastic programming program with binary recourse. Solution approaches are then discussed in the framework of the sample-average

approximation method. In particular, we present conditions under which the general results on the convergence rate and statistical bounds for the SAA method can apply to the stochastic-selection problem. Moreover, we adapt and assess two categories of optimization algorithms to solve the sample-average approximate problem. Based on numerical experiments, we identify best solvers that can give satisfactory solutions within acceptable time constraints and levels of accuracy, and illustrate by example how to trade off the sample size and the number of replications when applying the SAA method to the defender-adversary games.

Chapter 7

Conclusions and Directions for Future Work

In this chapter, we first summarize our results in Section 7.1. Next, Section 7.2 discusses validation of the models that have been developed and investigated in this dissertation. Finally, we provide several directions for future work in Section 7.3.

7.1 Conclusions

In this dissertation, we extend game-theoretic models for homeland security in three different directions, all motivated by the desire to make game theory ready for use in real-world security decisions. We summarize our results in all three areas below.

7.1.1 Expert Elicitation of Adversary Intent Using Ordinal Judgments

First, we introduce a simple elicitation process to estimate uncertain adversary intent using only ordinal judgments from subject-matter experts. In particular, we assume that the adversary will choose targets or attack strategies according to a multiattribute utility function. Defender uncertainty is then represented by probability distributions over the uncertain attacker attribute weights, as well as the values of the various targets on an unobserved attribute that may be important to the adversary, but has not been identified by the defender. Our approach asks subject-matter experts to provide (partial) rank orderings of the attractiveness of the various targets, and then mathematically derives probability distributions for the uncertain attacker parameters using either probabilistic inversion or Bayesian density estimation.

Although the motivation for this aspect of our work was the need for methods of estimating adversary intent from ordinal data, this area of our work also makes methodological contributions to the field of expert elicitation in general, especially through the idea of using unobserved attributes (which ensures the existence of feasible solutions in probabilistic inversion), and by elucidating the relationship between probabilistic inversion and Bayesian density estimation when applied to preference rankings. Moreover, our proposed elicitation process dramatically reduces the burden of traditional methods of attribute weight elicitation, and explicitly captures the extent of uncertainty or disagreement among experts, rather than attempting to achieve consensus. We apply this elicitation process to a realistic elicitation of adversarial preferences over attack scenarios such as biological attacks, chemical attacks,

IED's or radiological attacks, to demonstrate the applicability of the proposed elicitation process.

7.1.2 Extension of Games to Account for Attacker Capabilities

Next, we fill a gap in the literature of game-theoretic models for homeland security by explicitly considering adversary capabilities in addition to just intent, since intelligence experts generally believe that adversary capabilities are at least as important as intent. In particular, we propose a Bayesian Stackelberg game capable of analyzing the joint effects of both attacker intent and capabilities on optimal defensive strategies. One novel feature of our model is the use of contest success functions from economics to explicitly capture the extent to which the success of an attack is attributable to the adversary's capabilities and the level of defensive investment, rather than pure luck. Moreover, our model also allows a given type of attacker capability to have differing effectiveness across targets (e.g., civilian versus military targets) and/or attack modes (e.g., attacks using IED's versus nuclear weapons).

Our model is one of the first studies to address uncertainty about attacker capabilities, and thus paves the way for homeland-security decision makers to base defensive strategies on a more realistic and comprehensive set of adversary characteristics. In particular, results show that precise assessment of attacker intent will generally not be necessary if the defender is highly uncertain about attacker capability. Therefore, our model provides useful guidance on how to prioritize intelligence collection about adversary capabilities versus intent.

7.1.3 Computational Approaches for Determining Optimal Defenses

Finally, we identify, evaluate and implement rigorous and efficient computational tools to solve for equilibrium (optimal) defensive strategies in problems of realistic size and complexity. We show how our Bayesian Stackelberg game (including defender uncertainty about adversary intent and capabilities) can be formulated as a two-stage stochastic programming problem with binary recourse, and solved using a variety of approaches based on sample-average approximation. In particular, we demonstrate that the sample-average approximation method can be applied to our case, and investigate two categories of state-of-the-art optimization algorithms (one based on mixed-integer nonlinear programming, and the other based on derivative-free global optimization techniques) for solving games of realistic size and complexity.

Moreover, the defender-adversary game of incomplete information we address is an example of a broader class of stochastic selection problems. Thus, the techniques that we investigated for this case are also applicable to other stochastic-selection problems (e.g., to obtain optimal budget allocations for marketing initiatives, in the face of uncertainty about the importance of different product features to customers).

In summary, this dissertation fills a significant gap in the homeland-security literature by explicitly considering adversary capabilities in game-theoretic modeling of adaptive adversaries, in addition to just intent. Moreover, we have also addressed two other significant hurdles to making game theory applicable in practice; namely, the need to quantify uncertain adversary intent using subject-matter experts, and the need for powerful computational tools to solve for optimal defensive strategies within acceptable time constraints and levels of

accuracy. We hope that the results of this effort will therefore increase the acceptance and applicability of game theory as a readily usable tool for homeland-security decision makers.

7.2 Model Validation

For our models to be useful in homeland security decision making, they should ideally be validated using either empirical or simulated data. In this dissertation, we have attempted to assess and demonstrate the validity and applicability of our models to real-world problems whenever possible. However, such models are obviously difficult to validate, due to the sparsity of historical data and the difficulty of accessing intelligence experts (March et al. 1991). We here summarize our efforts at model validation.

To advance the applicability of the expert elicitation process using ordinal preference rankings, we demonstrated that the elicited probability distributions using our method are able to capture the level of uncertainty and disagreement among experts, and that even partial rank orderings of the attacker targets or strategies can yield reliable results. For illustrative purposes, we analyzed both a hypothetical case study of major U.S. urban areas and a real-world application (CREATE 2011) on elicitation of adversary preferences among different attack scenarios (e.g., chemical, biological, radiological, and nuclear attacks), and obtained realistic results in both cases.

We also made several efforts to validate our analytic model of adversary capabilities. For example, we illustrated the reasonableness of the model by investigating how the attacker and defender target valuations and the degree of defender uncertainty affect the optimal choice of defensive strategies. Moreover, we present a realistic case study to illustrate the

applicability of our model using open-source information on terrorist capabilities. The ability to quantify our model of attacker capabilities using readily available data is a significant accomplishment, since quantification is often a weak point of game-theoretic models.

Finally, we demonstrated that state-of-the-art stochastic programming theories can be applied to solve our adversary-defender games. Moreover, we also evaluated the performance of a variety of computational algorithms using a single class of test problems with a limited number of test instances (due to the lack of access to powerful computing facilities). To validate these computation tools in a more systematic way, investigations with a much larger number of test problems and test instances would be needed.

7.3 Directions for Future Work

In this section, we present several possible directions for future work. Some of them are incremental extensions to methods already developed in this dissertation. However, this dissertation also motivates fundamental explorations in such field as the statistical properties of ordinal data, the calibration of subjective judgments, and more complex security games.

One immediate extension is to allow for tied rank orderings in the elicitation process using either probabilistic inversion or Bayesian density estimation, since in practice, experts may believe that multiple targets are equally attractive to the attacker. There are at least three possible interpretations for ties in a stochastic environment. For example, we could assume that experts who give tied rankings are totally uncertain about which of the tied targets is more attractive to the attacker. Another interpretation is to treat the tied targets as having exactly the same level of attractiveness to the attacker with probability one. We could

also allow the attractiveness of the tied targets to be merely close to each other; however, it is not clear what "close" should mean in practice. Future research is needed to investigate the performance of probabilistic inversion and Bayesian density estimation in the face of tied rank orderings using these different interpretations.

Our results have also shown that when the various attributes in the multiattribute attacker utility function are highly correlated, the elicitation results can be unstable in the face of small changes in the attribute values. Therefore, it is important to explore the relationship between the level of attribute collinearity and the variability of the elicited attribute weights, by introducing tiny perturbations in the attribute values. Among other factors, the number of attacker attributes could play an important role in the impact of collinearity. For example, when the number of attacker attributes is larger than the number of (ranked) targets, there will always exist some degree of collinearity. In that case, unstable elicitation results may be virtually inevitable. However, it is still necessary to investigate under what circumstances the effect of collinearity is most significant.

The notion of attribute collinearity here seems analogous to that in multiple regression. In particular, when predictor variables in a regression model are highly correlated, the regression coefficients may change erratically in response to small changes in the data. Inspired by this analogy, it may be possible to develop a "significance test" based on whether removing a known attribute from the model significantly changes the remaining attribute weights (and the importance of the unobserved attribute). For example, if multiple known attributes are highly linearly correlated, then removing any one of these attributes from the model should not significantly affect the capability of the remaining attributes to explain

the expert judgments, and therefore should not significantly increase the weight assigned to the unobserved attribute.

In order to assess adversary intent, adversary capabilities, and other parameters (such as target vulnerabilities, and countermeasure costs, etc.), expert judgments need to be well calibrated to minimize the effects of possible overconfidence and biases. In fact, studies show that experts tend to give overconfident judgments (e.g., overly narrow prediction intervals) – although the extent of overconfidence varies significantly across domains and experts (Tversky and Kahneman 1974; Lin and Bier 2008). In addition, expert judgments can also be biased (i.e., systematically greater or lower than the corresponding true values). For example, costs of new technologies tend to be overestimated, since the creation of a market generally leads to cheaper technology over time (Finkel 1996). By contrast, it is common to underestimate construction costs in the planning stage, because of possible time delays and other contingencies during the course of a construction project (Dillon et al. 2002).

One possible way to control overconfidence and bias is to ask subject-matter experts to give estimates (e.g., medians and prediction intervals) not only for the uncertain quantities of interest, but also for a set of "seed variables" whose values are verifiable using historical or anticipated near-future data as originally proposed by Cooke (1991); see also Cooke (2012) and Aspinall and Cooke (2013). Performance on the seed variables can then be used to empirically assess each expert's normative ability to provide accurate probabilistic judgments. A consensus probability distribution for the quantity of interest can then be derived by linearly pooling the various experts' individual predictions in a way that

minimizes the impact of any subjective biases or overconfidence as identified using the seed variables (Hora 2010).

Three mathematical approaches are available for obtaining such a linear opinion pool. The first is to weight the various experts according to their performance on the seed variables – e.g., by applying Cooke's classical method (1991), or the continuous ranked probability scoring rule of Gneiting and Raftery (2007). The second option is to generate expert weights to minimize the overall overconfidence and bias (or maximize the calibration) of the pooled predictions. Finally, we could adopt the approach of Apostolakis (1986), of broadening each expert's distribution for the quantity of interest to offset any overconfidence identified using the seed variables, and then pooling the adjusted distributions using equal weights. (In this approach, distributions could also be shifted to account for observed biases, in addition to broadening.)

Note that the last debiasing approach requires explicit estimation of each expert's tendency for location and precision biases. Earlier work by Shlyakhter et al. (1994) uses a compound normal distribution to jointly estimate those biases in a relatively simple way. However, we could adapt more modern methods of Bayesian quantile regression (e.g., Lancaster and Jun 2010) to account for three additional phenomena: (1) asymmetry of prediction intervals around the median (instead of assuming or requiring normality); (2) correlations between different quantiles (e.g., the fact that a large estimate for the median of a particular quantity will typically be associated with a large value for the 95th percentile as well); and (3) correlations between different experts representing the same "school of thought" (Merrick 2008).

However, some of the parameters in our model are not directly observable (such as the decisiveness of differing targets or attack modes, and the relative effectiveness of each type of attacker capability), and thus cannot pass the "clairvoyant test" of Howard (1988); i.e., whether a clairvoyant could give an unambiguous value for that quantity. Asking for estimates of unobservable model parameters is problematic, since experts may disagree with the particular risk model for which parameters are being assessed, or differ in their interpretation of the model parameters. If required, estimates for these unobservable quantities could instead be derived indirectly, using techniques such as probabilistic inversion (Cooke 1994; Kraan and Bedford 2005; Wang and Bier 2013a).

With regard to the game-theoretic portion of our model, it would be interesting to extend the consideration of adversary intent and capabilities to protection of a network of targets, such as power grids, transportation networks, or computer networks (see for example Morton 2011; Ertem and Bier 2013). It might also be interesting to model the nature of optimal defender efforts to interdict adversary networks, such as networks for weapon smuggling or terrorist financing (e.g., Biersteker and Eckert 2007).

Another worthwhile extension to the analytic model in this dissertation would be to consider repeated games with non-myopic players who care about long-term rather than immediate payoffs. For example, a non-myopic defender may make current resource-allocation decisions in a way that increases her ability to observe the attacker's preferences, if the cost of deceiving the attacker is not catastrophic. Moreover, in conditions of high uncertainty about terrorist preferences, the defender may want to save resources for use in later periods, when her investment may be more effective because of greater information about attacker characteristics.

Such a model could also be extended to distinguish between long-term investment and short-term expenses. Short-term expenses (such as salaries for police officers) would largely evaporate by the end of the current decision period, while the effect of long-term investments is generally cumulative, although it may deteriorate slowly over time (Zhuang et al. 2010). Our model could in principle be extended to treat long-term capital investments differently from short-term expenses, and potentially identify the optimal trade-off between them.

Finally, there are a couple of directions for potential explorations with regard to the development of computational tools. For example, one worthwhile extension might be to apply the generalized L-shaped method of Carøe and Tind (1998) to reduce the complexity of the optimization problem we address, and therefore the required computation time. We could also explore the use of variance-reduction techniques to enhance the quality of the sampling-based optimization approach used in this research (e.g., to facilitate faster convergence rate of the statistical optimization bounds), by using Latin hypercube sampling (Koivu 2005; Drew and Homem-de-Mello 2012) or other dependent sampling schemes such as those developed by Homem-de-Mello (2008), Qian and Wu (2009) and Qian (2009).

Appendix A

Proofs of Propositions and Derivation of Results of Examples in Chapter 3

A.1 Proof for Proposition 3.1

For a given defensive resource allocation $x = (x_1, \ldots, x_N)$ to the N targets, the attack probability on target 1 can be expressed as

$$\begin{aligned} h_1(x) =& P\{e^{-\lambda x_1} U_1 \geq e^{-\lambda x_k} U_k \ \forall k \neq 1\} \\ =& P\{e^{-\lambda x_1} U_1 \geq e^{-\lambda x_k} U_k \ \forall k \neq 1 | U_1 = 0\} P\{U_1 = 0\} \\ &+ P\{e^{-\lambda x_1} U_1 \geq e^{-\lambda x_k} U_k \ \forall k \neq 1 | 0 < U_1 \leq 1\} P\{0 < U_1 \leq 1\} \\ =& P\{\frac{U_k}{U_1} \leq e^{-\lambda(x_1 - x_k)} \ \forall k \neq 1 | 0 < U_1 \leq 1\} P\{0 < U_1 \leq 1\} \end{aligned}$$

While keeping $x_2, \ldots, x_N$ constant, we have $\frac{de^{\lambda(x_1-x_k)}}{dx_1} = -\lambda e^{-\lambda(x_1-x_k)} < 0$ for $\lambda > 0$, and thus $e^{-\lambda(x_1-x_k)}$ is strictly decreasing in x_1 for $\forall k \neq 1$. Therefore, the attack probability $h_1(x)$ is strictly decreasing in x_1.

A.2 Proof for Proposition 3.2

Without loss of generality, we suppose that the defender wants to make target 1 the most attractive target to the attacker out of the N targets. For a given defensive resource allocation $x = (x_1, \ldots, x_N)$ to the N targets, the attack probability on target 1 is given by

$$\begin{aligned} h_1(x) =& P\{e^{-\lambda x_1} U_1 \geq e^{-\lambda x_k} U_k \ \forall k \neq 1\} \\ =& P\{e^{-\lambda x_1} U_1 \geq \max_{k \neq 1} e^{-\lambda x_k} U_k\} \end{aligned}$$

Consider the case where the defender spends no investment on target 1, i.e.,

$$\begin{aligned} h_1(x) =& P\{U_1 \geq \max_{k \neq 1} e^{-\lambda x_k} U_k\} \\ \geq& P\{U_1 \geq \max_{k \neq 1} e^{-\lambda x_k}\} \end{aligned}$$

because the attacker target valuations U_k are assumed to take on values in $[0, 1]$. With a sufficiently large budget, the defender can make $\max_{k \neq 1} e^{-\lambda x_k}$ arbitrarily close to 0, and thus the attack probability $h_1(x)$ arbitrarily close to 1.

A.3 Proof for Proposition 3.3

Consider a two-target, two-attribute case with one known attribute and one unobserved attribute. According to Appendix A.1, it follows that for a fixed budget B and allocation x_1 to target 1, the attack probability on target 1 is given by

$$h_1(x) = P\left\{(e^{-\lambda x_1}a_{11} - e^{-\lambda(B-x_1)}a_{21})\frac{W_1}{1-W_1} \geq -e^{-\lambda x_1}Y_1 + e^{-\lambda(B-x_1)}Y_2\right\}$$

where W_1, Y_1 and Y_2 are independent random variables taking on values in $[0, 1]$. In this proposition, we assume that the probability distributions for W_1, Y_1 and Y_2 put positive mass on every value in $[0, 1]$.

If $e^{-\lambda x_1}a_{11} - e^{-\lambda(B-x_1)}a_{21} = 0$, then $h_1(x_1) = P\{0 \geq -e^{-\lambda x_1}Y_1 + e^{-\lambda(B-x_1)}Y_2\} > 0$, because the probability distribution of the right hand side $-e^{-\lambda x_1}Y_1 + e^{-\lambda(B-x_1)}Y_2$ puts positive mass on every value in $[-e^{-\lambda x_1}, e^{-\lambda(B-x_1)}]$.

If $e^{-\lambda x_1}a_{11} - e^{-\lambda(B-x_1)}a_{21} > 0$, then $h_1(x_1) = P\left\{\frac{W_1}{1-W_1} \geq \frac{-e^{-\lambda x_1}Y_1+e^{-\lambda(B-x_1)}Y_2}{e^{-\lambda x_1}a_{11}-e^{-\lambda(B-x_1)}a_{21}}\right\} > 0$, because the probability distribution of the left hand side $\frac{W_1}{1-W_1}$ puts positive mass on every value in $(0, +\infty)$, and the right hand side $\frac{-e^{-\lambda x_1}Y_1+e^{-\lambda(B-x_1)}Y_2}{e^{-\lambda x_1}a_{11}-e^{-\lambda(B-x_1)}a_{21}}$ is bounded above. Similarly, if $e^{-\lambda x_1}a_{11} - e^{-\lambda(B-x_1)}a_{21} < 0$, we also have $h_1(x_1) = P\left\{\frac{W_1}{1-W_1} \leq \frac{-e^{-\lambda x_1}Y_1+e^{-\lambda(B-x_1)}Y_2}{e^{-\lambda x_1}a_{11}-e^{-\lambda(B-x_1)}a_{21}}\right\} > 0$.

A.4 Derivation of Results in Example 3.1

In this example, we assume that the defender has only one known attribute, and the values of the attribute for each target are $a_{11} = 1$ and $a_{21} = 0$. For simplicity, we denote c as the

proportion of the total defensive resource allocated to target 1, instead of the actual amount of resource. For a given proportion c of the total resource allocated to target 1, the attack probability on target 1 is given by

$$\begin{aligned} h_1(c) =& P\{e^{-\lambda Bc}[W_1 + Y_1(1-W_1)] \geq e^{-\lambda B(1-c)}Y_2(1-W_1)\} \\ \overset{0<W_1\leq 1}{=\!=}& P\{e^{-\lambda Bc}\frac{W_1}{1-W_1} \geq e^{-\lambda B(1-c)}Y_2 - e^{-\lambda Bc}Y_1\} \end{aligned}$$

We let $a = e^{-\lambda Bc}$ and $b = e^{-\lambda B(1-c)}$, and introduce two new random variables $R = \frac{W_1}{1-W_1}$ and $Z = bY_2 - aY_1$. Then the attack probability can be simplified as $h_1(c) = P\{aR \geq Z\}$. Since W_1, Y_1 and Y_2 are independently uniformly distributed on [0, 1], it can be obtained that $P\{R \geq r\} = \frac{1}{1+r}$ for $r \geq 0$. In addition, if we assume $b > a$ (the case $b \leq a$ yields similar results), the PDF of Z is then given by

$$q(z;a,b) = \begin{cases} \frac{z}{ab} + \frac{1}{b} & -a \leq z < 0 \\ \frac{1}{b} & 0 \leq z \leq b-a \\ -\frac{z}{ab} + \frac{1}{a} & b-a < z \leq b \end{cases}$$

The attack probability on target 1 can then be calculated as follows

$$\begin{aligned} h_1(c) =& P\{aR \geq Z\} \\ =& \int_{-a}^{b} P\{R \geq \frac{z}{a} | Z = z\} q(z) dz \\ =& \int_{a}^{0} \frac{1}{1+z/a}(\frac{z}{ab} + \frac{1}{b})dz + \int_{0}^{b-a} \frac{1}{1+z/a}\frac{1}{b}dz + \int_{b-a}^{b} \frac{1}{1+z/a}(-\frac{z}{ab} + \frac{1}{a})dz \\ =& \frac{a}{b} + (\frac{a}{b}\ln b - \frac{a}{b}\ln a) + (-\frac{a}{b} + \frac{a+b}{b}\ln(a+b) - \frac{a+b}{b}\ln b) \\ =& -\frac{a}{b}\ln a + \frac{a+b}{b}\ln(a+b) - \ln b \end{aligned}$$

Appendix B

Proofs of Propositions and Derivation of Results of Examples in Chapter 4

B.1 Proof for Proposition 4.1

If we have a sufficiently large number of samples S, then for each expert, we are ensured to have at least one sample ($w^{(s)}$, $y^{(s)}$) that can match that expert's rank ordering of targets. This choice of q satisfies the constraints of (4.5), and thus feasibility is ensured. Moreover, the objective function in (4.5) is strictly quasi-convex and the constraints are linear, so the optimization problem has a unique solution.

B.2 Proof for Proposition 4.3

Note that the probability distribution $Q_0^{(1)}$ given by BDE in (4.7) is just a truncated version of the starting measure Q_0, reflecting the single rank ordering provided by the expert(s).

Note also that the PI problem in (4.3) has a unique solution. So, it suffices to show that $Q_0^{(1)}$ satisfies linear constraints of the form $\int_\Omega J_{rn}(w,y)dQ = P_{rn}$ (r = 1, …, R; n = 1, …, N), and can be derived from (4.4) for suitable λ_{rn} and c.

Then consider the truncated distribution $Q_0^{(1)}$. It would place zero probability on values (w, y) that are inconsistent with the given rank ordering; i.e., $\frac{dQ}{dQ_0}(w,y) = 0$ for $(w,y) \notin AR^{(1)}$, and would assign equal values for $\frac{dQ}{dQ_0}(w,y)$ for every $(w,y) \in AR^{(1)}$. Note that given any $(w,y) \in AR^{(1)}$, we have $J_{rn}(w,y) = 1$ if target n ranks in the rth place and 0 otherwise. Note also that the empirical distribution matrix of expert rankings P for the given single rank ordering has only binary elements (i.e., either zero or one), and relates directly to $AR^{(1)}$. Thus, $Q_0^{(1)}$ satisfies $\int_\Omega J_{rn}(w,y)dQ_0^{(1)} = P_{rn}$ for r = 1, …, R; n = 1, …, N.

Moreover, suppose that we set λ_{rn} = 1 if target n ranks in the rth place according to the given rank ordering, and $\lambda_{rn} \to \infty$ otherwise, and derive the corresponding normalization constant c. Then (4.4) yields exactly the truncated distribution $Q_0^{(1)}$. Therefore, we have $Q^* = Q_0^{(1)}$.

B.3 Proof for Proposition 4.4

We prove this proposition for the approximate PI model as given by (4.5). Suppose that a discretized sample space $\bar{\Omega}$ and the associated discretized starting measure $\bar{Q}_0$ are constructed by generating adequate random samples from Q_0 to ensure the feasibility of (4.5). It then suffices to show that the optimal solution q^* to (4.5) can be expressed as a weighted sum of

a collection of truncated variants of $\bar{Q}_0$, each corresponding to a particular rank ordering of targets.

Note that the iterative proportional fitting algorithm of Csiszár (1975) proceeds in such a way that if ($w^{(i)}$, $y^{(i)}$) and ($w^{(j)}$, $y^{(j)}$) yield the same rank ordering of targets, then the values of q_i and q_j (i, j = 1, …, S) will be equal throughout the procedure. In other words, the final convergent solution q^* must have equal probability over any subset of $\bar{\Omega}$ that is consistent with a particular rank ordering. Thus, each of these subsets is associated with a truncation of $\bar{Q}_0$, and q^* is equivalent to a weighted sum of those truncated measures for all expert rank orderings.

B.4 Derivation of Results in Example 4.3

For simplicity, we denote a_n (n = 1, …, N) as the value of target n on the single known adversary attribute. Without loss of generality, we set a_1 = 1, and $a_i - a_j = (N-1)^{-1}$ for any $j = i + 1$. Let Y_n (n = 1, …, N) be the utility of the unobserved attribute for target n, and W be the weight for the unobserved attribute. Moreover, we assume that Y_n and W are independent and uniformly distributed in [0, 1]. Then, for a rank ordering that is perfectly consistent with the values of a_n, we have the following relationship

$$a_n(1-W) + Y_nW \geq a_{n+1}(1-W) + Y_{n+1}W \text{ for } n = 1, \ldots, N-1$$

By aggregating and rearranging these inequalities, we get the following result

$$W \leq [1 + (N-1)X_N]^{-1}$$

where we set $X_N = \max_{n=1,\ldots,N-1}\{Y_{n+1} - Y_n\}$ which is bounded below and above by $-(N-1)^{-1} \leq X_N \leq 1$.

Let $F_N(x)$ be the cumulative distribution function of X_N such that $F_N(x) = P\{X_N \leq x\}$. Define also that $Z_N = [1+(N-1)X_N]^{-1}$ and let $G_N(z)$ be the cumulative distribution function such that $G_N(z) = P\{Z_N \leq z\}$. The mean weight for the unobserved attribute can then be derived as

$$
\begin{aligned}
\mathbf{E}[W] &= \int_0^1 \frac{1}{2} z dG_N(z) + \int_1^\infty \frac{1}{2} dG_N(z) \\
&= \frac{1}{2}\int_0^1 [1+(N-1)x]^{-1} dF_N(x) + \frac{1}{2}\int_{-(N-1)^{-1}}^0 dF_N(x) \\
&= \frac{1}{2}\{[1+(N-1)x]^{-1} F_N(x) \mid_0^1 - \int_0^1 F_N(x) d[1+(N-1)x]^{-1}\} + \frac{1}{2}F_N(0) - \frac{1}{2}F_N[-(N-1)^{-1}] \\
&= \frac{1}{2}\{\frac{1}{N} - F_N(0) - \int_0^1 F_N(x) d[1+(N-1)x]^{-1}\} + \frac{1}{2}F_N(0) \\
&= \frac{1}{2N} + \frac{1}{2}\int_0^1 -F_N(x) d[1+(N-1)x]^{-1} \\
&= \frac{1}{2N} + \frac{1}{2}\int_0^1 \frac{F_N(x)(N-1)}{[1+(N-1)x]^2} dx
\end{aligned}
\tag{B.1}
$$

Unfortunately, it is difficult to get the closed-form solution for (B.1). Instead, we calculate the numerical values of $\mathbf{E}[W]$ as a function of N using Monte Carlo simulation (with sample size 10^7), and find that $\mathbf{E}[W]$ is strictly decreasing in N, at least for N = 2, …, 1000.

We now look at the limiting behavior of (B.1) as $N \to \infty$. It is easy to show that the random variable X_N stochastically dominates the largest order statistic of N i.i.d. random

variables that follow the triangular distribution with parameters $(-1, 0, 1)$; i.e.,

$$F_N(x) = P\{X_N \leq x\} \leq [1 - \frac{1}{2}(1-x)^2]^{N-1}$$

In addition, if we set $\hat{x}_N = 1 - \sqrt{2 - 2(N-1)^{(1-N)^{-1}}}$, then for any $x \in [0, \hat{x}_N] \subseteq [0, 1]$, we have $F_N(x) \leq (N-1)^{-1}$. Therefore, (B.1) is bounded above by

$$\begin{aligned}
&\frac{1}{2N} + \frac{1}{2}\int_0^{\hat{x}_N} \frac{1}{[1+(N-1)x]^2} dx + \frac{1}{2}\int_{\hat{x}_N}^1 \frac{[1-\frac{1}{2}(1-x)^2]^{N-1}(N-1)}{[1+(N-1)x]^2} dx \\
&\leq \frac{1}{2N} + \frac{1}{2}\int_0^1 \frac{1}{[1+(N-1)x]^2} dx + \frac{1}{2}\int_{\hat{x}_N}^1 \frac{[1-\frac{1}{2}(1-x)^2]^{N-1}(N-1)}{[1+(N-1)x]^2} dx \\
&= \frac{1}{2N} + \frac{1}{2N} + \frac{1}{2}\int_{\hat{x}_N}^1 \frac{[1-\frac{1}{2}(1-x)^2]^{N-1}(N-1)}{[1+(N-1)x]^2} dx \\
&\leq \frac{1}{N} + \frac{1}{2}(1-\hat{x}_N)\frac{N-1}{[1+(N-1)\hat{x}_N]^2}
\end{aligned}$$

Note that the second term above converges to 0 as N gets sufficiently large, because $\hat{x}_N \to 1$ as $N \to \infty$. Therefore, the expression for $\mathbf{E}[W]$ in (B.1) gets arbitrarily close to zero as $N \to \infty$, since we must have $\mathbf{E}[W] \geq 0$.

Appendix C

Computation Algorithms in Chapter 4

C.1 Algorithm of Iterative Proportional Fitting for Probabilistic Inversion

1. Set a lower bound for the number of samples S_0, and a tolerance level $\epsilon > 0$;

2. Sequentially draw independent random samples from Q_0; stop when either the number of samples reaches S_0 or the condition in Proposition 4.1 is satisfied, whichever happens later. Let S be the total number of samples that have been drawn. Record the discretized parameter space $\bar{\Omega}$ as a collection of all the generated samples $(w^{(s)}, y^{(s)})$, $s = 1, \ldots, S$. Set the initial probability vector $q^0 = (\frac{1}{S}, \ldots, \frac{1}{S})$;

3. Find the optimal probability vector $q^* = (q_1^*, \ldots, q_S^*)$ to the optimization program (4.5) using the following procedure.

   $t' = 0$;

while $\| \sum_{s=1}^{S} q_s^t J(w^{(s)}, y^{(s)}) - P \|_\infty \geq \epsilon$

$t' = t' + 1$;

for $r = 1, \ldots, R$ and $n = 1, \ldots, N$

$t = t'RN + (r-1)N + n$; defining $S_{rn} \subset \{1, \ldots, S\}$ to be the set of all samples that rank target n in the rth place, we then update the probability vector q^t by the following steps:

a. Calculate $\bar{q}_{S_{rn}}^t = \sum_{s \in S_{rn}} q_s^{t-1}$;

b. For $s \in S_{rn}$ we solve q_s^t by

$$q_s^t \bar{q}_{S_{rn}}^t = q_s^{t-1} P_{rn}$$

For $s \notin \Sigma_{rn}$ we solve q_s^t by

$$q_s^t (1 - \bar{q}_{S_{rn}}^t) = q_s^{t-1} (1 - P_{rn})$$

end

end

4. Output the generated samples $(w^{(s)}, y^{(s)})$ for $s = 1, \ldots, S$, and the optimal probability vector $q^* = q^t$.

C.2 Gibbs Sampler to Generated the Truncated Starting Measure in Bayesian Density Estimation

This algorithm is to generate random samples from the truncated version of the starting measure $Q_0^{(k)}$ associated with the rank orderings provided by expert k.

1. Set the total number of samples as S;

2. Find a starting point $(w^{(0)}, y^{(0)}) = (w_1^{(0)}, \ldots, w_{M+1}^{(0)}, y_1^{(0)}, \ldots, y_N^{(0)})$ that belongs to the active region $AR^{(k)}$;

3. Obtain random samples $(w^{(s)}, y^{(s)})$ from $Q_0^{(k)}$ using the following procedure.

   $s = 0$;

   while $s < S$

   $s = s + 1$;

   for $m = 1, \ldots, M$

   $$W_m^{(s)} \sim Q_0^{(k)}\left(w_m \;\middle|\; \begin{array}{l} W_j = w_j^{(s)} \text{ for } j < m; W_j = w_j^{(s-1)} \text{ for } j > m+1; Y = y^{(s-1)}; \\ W_m + W_{m+1} = 1 - \sum_{j<m} w_j^{(s)} - \sum_{j>m+1} w_j^{(s-1)} \end{array}\right)$$

   end

$$W_{M+1}^{(s)} \sim$$
$$Q_0^{(k)}\left(w_{M+1} \;\middle|\; \begin{array}{l} W_j = w_j^{(s)} \text{ for } 1 < j < M; Y = y^{(s-1)}; \\ W_{M+1} + W_1 = 1 - \sum_{1<j<M} w_j^{(s)} \end{array}\right)$$

for $n = 1, \ldots, N$

$$Y_n^{(s)} \sim Q_0^{(k)}\left(y_n \middle| Y_i = y_i^{(s)} \text{ for } i < n; Y_i = y_i^{(s-1)} \text{ for } i > n; W = w^{(s)}\right)$$

end

Record the sth sample $(w^{(s)}, y^{(s)})$.

end

4. Output the generated samples $(w^{(s)}, y^{(s)})$ for $s = 1, \ldots, S$. To reduce the influence of the starting point $(w^{(0)}, y^{(0)})$, we may discard the first 10% of the S samples.

Appendix D

Proofs of Lemmas, Propositions and Corollaries and Derivation of Results of Examples in Chapter 5

D.1 Proof for Proposition 5.1

Suppose that the N targets are rank ordered such that the defender's (as well as the attacker's) valuations of these targets satisfy $v_1 \geq v_2, \ldots \geq v_N > 0$. We first give an intuitive explanation of the condition in Proposition 5.1, and then show that the condition is both sufficient and necessary for deriving the defender's optimal strategy.

In the absence of any defense, the attacker will attack only target 1 if $v_1 > v_n$ for $n > 1$, while randomly choose one of the first K targets to attack if they share the same highest value; i.e., $v_1 = v_n$ for $n \leq K$ and $v_1 > v_n$ for $n > K$. If there are K targets that share

the same highest value, then the defender needs to allocate her resources to make the K targets equally attractive to the attacker. There is only a unique way of doing this since the expected payoff of an attack on each target $s_n(x_n, A)v_n$ is strictly decreasing in the defensive investment x_n (assuming $\beta_n > 0$). Deviating from that strategy will make some one of the K targets less attractive than others, and defensive investments on those less attractive targets will be wasted.

However, as the defender spends more on the K highest-valued target(s), eventually they become equally attractive to the attacker as the $(K+1)$st target. Then additional resources should be spread over the first $K+1$ targets, as well as targets of the same value as v_{K+1}, to equalize and reduce their expected attacker payoff to be no smaller than the value of any unprotected target. If the defender continues this process until her budget is exhausted, then the condition in Proposition 5.1 is satisfied. By expanding the set of protected targets in this way, the defender is able to defend as many high-valued targets as possible with limited resources.

This condition is sufficient, because the defender is not able to achieve any lower expected loss with the same level of budget, by using a different defensive strategy other than the above process. On the other hand, the condition is also necessary. If there exists a target that receives positive protection but is not the most attractive to the attacker, then the defender can always gain profit by spending less on that target, and moving the saved effort to more attractive target(s). .

D.2 Proof for Lemma 5.2

Suppose that the N targets are rank ordered such that $v_1 \geq v_2 \ldots v_N > 0$. Let $C(i,j)$ $(2 \leq i < j \leq N)$ be the minimal level of defensive budget needed to equalize the attractiveness of the first i targets with the jth target; i.e., such that $\frac{A^{\beta_1} v_1}{A^{\beta_1} + (x_1)^{\beta_1}} = \ldots = \frac{A^{\beta_i} v_i}{A^{\beta_i} + (x_i)^{\beta_i}} = \frac{A^{\beta_j} v_j}{A^{\beta_j} + (x_j)^{\beta_j}}$ for $x_1 + \ldots + x_i + x_j = C(i,j)$. Then it is easy to show that $B^{(i)} \leq C(i-1,j) \leq B^{(j)}$ for $2 \leq i < j \leq N$.

D.3 Proof for Corollary 5.3

This corollary follows directly from the discussion in the proof for Proposition 5.1.

D.4 Proof for Proposition 5.4

Recall that for any given x and A, the threat likelihood $p_n(x, A, v)$ is given by

$$p_n(x, A, v) = \begin{cases} \frac{1}{Z} & \text{if } n \in \arg\max_j s_j(x_j, A) v_j \\ 0 & \text{otherwise} \end{cases}$$

where Z is the cardinality of the set $\{j : \arg\max_j s_j(x_j, A) v_j\}$. The attacker's expected payoff given x and A can then be calculated as given by

$$\begin{aligned} & \sum_{n=1}^{N} p_n(x, A, v) s_n(x_n, A) v_n \\ & = \textstyle\sum_{n \in \arg\max_j s_j(x_j, A) v_j} \frac{1}{Z} s_n(x_n, A) v_n \\ & = \max_j s_j(x_j, A) v_j \end{aligned}$$

Therefore, we have the following

$$\min_{x_1+\ldots+x_N\leq B}\int\sum_{n=1}^{N}p_n(x,A,v)s_n(x_n,A)v_n dF_A(A)$$
$$=\min_{x_1+\ldots+x_N\leq B}\int\max_j\{s_j(x_j,A)v_j\}dF_A(A)$$

We then prove the first-order necessary condition. For convenience, we denote $l(x,A)=\max_j s_j(x_j,A)v_j$ and $L(x)=\int l(x,A)dF_A(A)$. We are interested in

$$L(x+te^n)-L(x)=\int l(x+te^n,A)-l(x,A)dF_A(A) \tag{D.1}$$

where $t>0$ and e^n is the nth unit coordinate vector. Suppose that for any given defensive strategy x, the set $\{j:\arg\max_j s_j(x_j,A)v_j\}$ is a singleton with probability one under the distribution $F_A(\cdot)$. If we denote $n^*(x,A)=\arg\max_j s_j(x_j,A)v_j$, then with probability one the following inequality holds

$$s_{n^*}(x_{n^*},A)v_{n^*}>s_n(x_n,A)v_n \text{ for } \forall n\neq n^*$$

Note that $s_n(x_n,A)v_n$ is continuous and strictly increasing in x_n for all n. Therefore, $\exists\bar{t}>0$ such that for $\forall 0<t<\bar{t}$, we have

$$s_{n^*}(x_{n^*}+t,A)v_{n^*}>s_n(x_n,A)v_n \text{ and}$$
$$s_{n^*}(x_{n^*},A)v_{n^*}>s_n(x_n+t,A)v_n \text{ for } \forall n\neq n^*$$

So, $l(x+te^n,A)=s_{n^*}(x_{n^*}+t\cdot\mathbf{1}_{\{n=n^*\}},A)v_{n^*}$ for $n=1,\ldots,N$ with probability one, where $\mathbf{1}_{\{n=n^*\}}$ equals 1 if $n=n^*$ and 0 otherwise. Take such $\bar{t}$ and for $0<t<\bar{t}$ rewrite

(D.1) as given by

$$\begin{aligned} & L(x + te^n) - L(x) \\ = & \int s_{n^*}(x_{n^*} + t \cdot \mathbf{1}_{\{n=n^*\}}, A)v_{n^*} - s_{n^*}(x_{n^*}, A)v_{n^*} dF_A(A) \end{aligned} \tag{D.2}$$

By dividing both sides of (D.2) and taking the limit at $t \to 0$, we have the following equation

$$\begin{aligned} & \lim_{t\to 0} \frac{L(x + te^n) - L(x)}{t} \\ = & \int \lim_{t\to 0} \frac{s_{n^*}(x_{n^*} + t \cdot \mathbf{1}_{\{n=n^*\}}, A)v_{n^*} - s_{n^*}(x_{n^*}, A)v_{n^*}}{t} dF_A(A) \end{aligned}$$

Equivalently,

$$\frac{dL(x)}{dx_n} = \int g_n(x, A, v) dF_A(A)$$

where

$$g_n(x, A, v) = \begin{cases} \frac{\partial}{\partial x_{n^*}} s_{n^*}(x_{n^*}, A)v_{n^*} & \text{if } n = n^* \\ 0 & \text{if } i \neq n^* \end{cases}$$

Finally, since the defender's objective function $L(x)$ is strictly decreasing in x_n for all n, the total budget B needs to be exhausted at optimality. The Karush-Kuhn-Tucker (KKT) condition for an interior-point solution to the optimization problem $\min_{x_1+\ldots+x_N=B} L(x)$ is then given by $\frac{dL(x)}{dx_n} = \mu$ for all $n = 1, \ldots, N$, with $x_n > 0$ and $x_1 + \ldots + x_N = B$.

D.5 Proof for Proposition 5.5

For given defensive resource allocation $x = (x_1, \ldots, x_N)$ and the attacker's allocation of effort $a = (a_1, \ldots, a_N)$, we consider two cases. First, if $a_i > 0$ and $a_j = 0$ for $i \neq j$, and

the marginal expected attacker payoffs of attacking targets i and j satisfy

$$h_i(x_i, a_i) < h_j(x_j, a_j)$$

then the defender can gain additional benefits by reducing a small portion $\epsilon > 0$ from the effort devoted to target i, and spending it on target j. In particular, the difference between the attacker's total expected payoffs before and after the reallocation is given by

$$\begin{aligned}
&[s_i(x_i, a_i)v_i + s_j(x_j, a_j)v_j] - [s_i(x_i, a_i - \epsilon)v_i + s_j(x_j, a_j + \epsilon)v_j] \\
&= s_i(x_i, a_i)v_i - s_i(x_i, a_i - \epsilon)v_i + s_j(x_j, a_j)v_j - s_j(x_j, a_j + \epsilon)v_j \\
&= h_i(x_i, a_i)\epsilon + h_j(x_j, a_j)(-\epsilon) + o(\epsilon) \\
&= [h_i(x_i, a_i) - h_j(x_j, a_j)]\,\epsilon + o(\epsilon)
\end{aligned}$$

where $o(\epsilon)$ is the first-order Taylor remainder. Since $h_i(x_i, a_i) < h_j(x_j, a_j)$, the above equation is strictly smaller than zero for sufficiently small ϵ.

On the other hand, if $a_i > 0$ and $a_j > 0$ for $i \neq j$, and $h_i(x_i, a_i) < h_j(x_j, a_j)$, then the defender can also gain profit by reducing a small portion $\epsilon > 0$ from the effort devoted to target i and spending it on target j.

Therefore, the condition in Proposition 5.5 is necessary.

D.6 Derivation of Results in Example 5.1

We focus on the case where $\frac{A}{B} < \left(\frac{v_1}{v_2} - 1\right)^{-1/\beta}$. For convenience, we denote $R = \frac{A}{B}$ and $r = \frac{x_1}{B}$, and then investigate the relationship between R and r as given by

$$\frac{R^\beta + r^\beta}{R^\beta + (1-r)^\beta} = \frac{v_1}{v_2}$$

By rearranging the above equation, we get an implicit function of r and R as given by

$$z(r, R) = (1-r)^\beta v_1 - r^\beta v_2 - R^\beta (v_2 - v_1) = 0$$

We then apply the implicit function theorem, and obtain the first-order derivative of r with respect to R, as given by

$$\begin{aligned} \frac{dr}{dR} &= -\frac{\partial z/\partial R}{\partial z/\partial r} \\ &= \frac{R^{\beta-1}(v_1 - v_2)}{(1-r)^{\beta-1} + r^{\beta-1}} > 0 \end{aligned}$$

for $R > 0$ and $0 < r < 1$. Note that we assume $v_1 > v_2$, so the optimal proportion of defensive investment allocated to target 1, $r^* = \frac{x_1^*}{B}$, is strictly increasing with the attacker's capability advantage $R = \frac{A}{B}$.

Appendix E

Proofs of Propositions and Computation Algorithms in Chapter 6

E.1 Proof for Proposition 6.1

Assumptions 6.2 and 6.3 imply that for every $x \in X$, $\exists \delta > 0$ such that $\sup_{||x-x'||<\delta} ||z(x,\omega) - z(x',\omega)|| = 0$ for almost every $\omega \in \Omega$ with respect to Q. Together with Assumption 6.1, it is easy to show that $f(x) = \int_\Omega \sum_{i=1}^N z_i(x,\omega) l_i(x) dQ(\omega)$ is continuous. Moreover, $f(x)$ attains its minimum on X because X is a compact set and $f(x)$ is continuous on X.

E.2 Proof for Lemma 6.2

For every $\omega \in \Omega$, Assumption 6.1–6.4 imply that $F(x,\omega)$ is a lower semi-continuous at x on X. The sample-average objective function $\hat{f}_S(x)$ is a random lower semi-continuous function because it is the finite sum of S random lower semi-continuous functions.

It also follows that the solution set X^* to the true problem (6.1) is closed, and the solution set $\hat{X}_S^*$ to the SAA problem (6.2) is a random closed set (Kaniovski et al. 1995). Given the compactness of the feasible set X and the finite recourse function $F(x, \omega)$, it also implies the existence of an approximate optimal solution $\hat{x}_S^*$ w.p. 1; i.e., $P\{\hat{X}_S^* \neq \emptyset\} = 1$.

E.3 Proof for Lemma 6.3

Assumptions 6.1 – 6.3 imply that for any fixed $x \in X$, the recourse function $F(x, \omega)$ is finite and continuous at x for almost every $\omega \in \Omega$ with respect to Q. Then the approximate objective function $\hat{f}_S(x)$ converges to the true objective function $f(x)$ uniformly on X with probability one, according to Proposition 7 in Shapiro (2003).

E.4 Proof for Proposition 6.4

Assumptions 6.1 – 6.4 lead to Lemmas 6.2 and 6.3. According to Theorem 3.5 of Shapiro et al. (2009), Lemmas 6.2 and 6.3 combined are sufficient for reaching the conclusions of the above proposition.

E.5 Proof for Proposition 6.5

Assumptions 6.1 – 6.4 guarantee that the approximate objective functions $\hat{f}_S(x)$ are lower semi-continuous in x for any set of random samples $\{\omega^{(1)}, \ldots, \omega^{(S)}\} \subset \Omega$, and the true objective function $f(x)$ is continuous in x. Then Theorem 4.1 of Kaniovski et al. (1995) implies that this proposition holds.

E.6 Implementation of Hybrid Simulated Annealing

1. Consider a cooling schedule $T(t) = T_0/\ln(1+t)$, where $T(t)$ is the temperature at time t, and T_0 is the initial temperature based on the level of difficulty of escaping from local optima. Start from a feasible allocation plan x^0.

2. For each trial, randomly simulate a step size $\rho(t) \sim Beta(T_0/T(t),\ 1)$ and two targets $i,\ j \leq N$. Set $\overline{\rho} = \min\{\rho, x_i^t\}$ and find a trial point x' with $x'_i = x_i^t - \overline{\rho}$, $x'_j = x_j^t + \overline{\rho}$, and $x'_k = x_k^t$ for $k \neq i, j$. Do multiple trials and record the best trial point as $x_{[TR]}$.

3. If $\hat{d}_S(x_{[TR]}) < \hat{d}_S(x^t)$, then apply Subplex/Nelder-Mead or BOBYQA starting from $x_{[TR]}$ until a local stationary point $x_{[LC]}$ is reached. Set $x^{t+1} = x_{[LC]}$.

   Otherwise, set $x^{t+1} = x_{[TR]}$ with probability $p(t) = \exp\left\{-\frac{\hat{d}_S(x_{[TR]}) - \hat{d}_S(x^t)}{T(t)}\right\}$, and $x^{t+1} = x^t$ with probability $1 - p(t)$.

4. Stop when the temperature $T(t)$ is sufficiently low. Otherwise return to Step 2.

Bibliography

[1] Abbas, A. E. 2004. Entropy method for adaptive utility elicitation. *IEEE Transactions on Systems, Science and Cybernetics* **32**(2) 169–178.

[2] Abbas, A. E. 2006 Maximum entropy utility. *Operations Research* **52**(2) 277–290.

[3] Abbas, A. E., D. E. Bell. 2011. One-switch independence for multiattribute utility functions. *Operations Research* **59**(3) 764–771.

[4] Ali, S., S. Ronaldson. 2012. Ordinal preference elicitation methods in health economics and health services research: Using discrete choice experiments and ranking methods. *British Medical Bulletin* **103** 21–44.

[5] Antoniak, C. E. 1974. Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *The Annals of Statistics* **2** 1152–1174.

[6] Apostolakis, G. 1982. Data analysis in risk assessments. *Nuclear Engineering and Design* **71** 375.

[7] Aspinall, W.P., R.M. Cooke. 2013. Quantifying scientific uncertainty from expert judgement elicitation. In J. Rougier, S. Sparks, L. Hill, eds. *Risk and Uncertainty Assessment for Natural Hazards*. Cambridge University Press, New York, NY.

[8] Baker, J., M. Wool, A. Smith, J. Kahan, C. Ansel, P. Hammar, D. McGarvey, M. Phillips, R. Lark. 2009. *Risk Analysis and Intelligence Communities Collaborative Framework*. Report, Homeland Security Institute, Arlington,

VA. Retrieved June 10, 2013 http://www.homelanddefense.org/downloads/Risk-Intel%20Collaboration%20Final%20Report.pdf.

[9] Barrett, A. M. 2010. Cost effectiveness of on-site chlorine generation for chlorine truck attack prevention. *Decision Analysis* **7**(4) 366–377.

[10] Barron, F. H., B. E. Barrett. 1996. The efficacy of SMARTER – Simple multi-attribute rating technique extended to ranking. *Acta Psychologica* **93** 23–36.

[11] Barros, C. P., I. Proença. 2005. Mixed logit estimation of radical Islamic terrorism in Europe and North America: A comparative study. *Journal of Conflict Resolution* **49**(2) 298–314.

[12] Bayraksan, G., D. P. Morton. 2011. A sequential sampling procedure for stochastic programming. *Operations Research* **59**(4) 898–913.

[13] Bedford, T. J., R. M. Cooke. 2001. *Probabilistic Risk Analysis: Foundations and Methods.* Cambridge University Press, New York & Cambridge.

[14] Beitel, G. A., D. I. Gertman, M. M. Plum. 2004. Balanced scorecard method for predicting the probability of a terrorist attack. C. A. Brebbia, ed. *Risk Analysis IV.* WIT Press, Southampton, UK, 581–592.

[15] Bellavita C. 2010. Changing Homeland Security: Twelve Questions From 2009. *Homeland Security Affairs* **6**(1).

[16] Beran, R. 1977. Minimum Hellinger distance estimates for parametric models. *The Annals of Statistics* **5**(3) 445–63.

[17] Bier, V. M. 2004. Implications of the research on expert overconfidence and dependence. *Reliability Engineering and System Safety* **85** 321–329.

[18] Bier, V. M. 2007. Choosing what to protect. *Risk Analysis* **27**(3) 607–620.

[19] Bier, V. M., V. Abhichandani. 2003. Optimal allocation of resources for defense of simple series and parallel systems from determined adversaries. Y. Y. Haimes, D. A.

Moser, E. Z. Stakhiv, eds. *Risk-Based Decisionmaking in Water Resources X.* American Society of Civil Engineers, Santa Barbara, CA, 59–76.

[20] Bier, V. M., N. Haphuriwat, J. Menoyo, R. Zimmerman, A. Culpen. 2008. Optimal resource allocation for defense of targets based on differing measures of attractiveness. *Risk Analysis* **28**(3) 763–770.

[21] Bier, V. M., J. B. Menoyo, C. Wang. 2012. Achieving realistic levels of defensive hedging based on non-monotonic and multi-attribute terrorist utility functions. J. Herrmann, ed. *Handbook of Operations Research, Homeland Security and Emergency Preparedness.* Springer, New York, NY, 125 – 139.

[22] Bier, V. M., A. Nagaraj, V. Abhichandani. 2005. Optimal allocation of resources for defense of simple series and parallel systems from determined adversaries. *Reliability Engineering and System Safety* **87** 313–323.

[23] Bier, V. M., S. Oliveros, L. Samuelson. 2007. Choosing what to protect. *Journal of Public Economic Theory* **9**(4) 563-âŁ“587.

[24] Bier, V. M., W. Yi. 1995. A Bayesian method for analyzing dependencies in precursor data. *International Journal of Forecasting* **11**(1) 25–41.

[25] Biersteker, T. J., S. E. Eckert, eds. 2007. *Countering the Financing of Terrorism.* Routledge, Oxford & New York .

[26] Borcherding, K., T. Eppel, D. von Winterfeldt. 1991. Comparison of weighting judgments in multiattribute utility measurement. *Management Science* **37** 1603–1619.

[27] Brown, G., C. Matthew, J. Salmeron, K. Wood. 2006. Defending critical infrastructure. *Interfaces* **36**(6) 530–544.

[28] Brown, G. G., L. A. Cox, Jr. 2011. How Probabilistic Risk Assessment Can Mislead Terrorism Risk Analysts. *Risk Analysis* **31**(2) 196–204.

[29] Bussieck, M. R., A. Drud. 2001. SBB: A new solver for mixed integer nonlinear programming. *Recent advances in nonlinear mixed integer optimization*, INFORMS Fall, Invited talk.

[30] Carøe, C. C., J. Tind. 1998. L-shaped decomposition of two-stage stochastic programs with integer recourse. *Mathematical Programming* **83**, 451–464.

[31] Center for Risk and Economic Analysis of Terrorism Events. 2011. *Adaptive Adversary Modeling for Terrorism Risk Analysis: Final Report*, University of Southern California, Los Angeles, CA.

[32] Clemen, R. T., T. Reilly. 1999. Correlations and copulas for decision and risk analysis. *Management Science* **45**(2) 208–224.

[33] Cormican, K. J., D. P. Morton, R. K. Wood. 1988. Stochastic network interdiction. *Operations Research* **46**(2) 184–197.

[34] Cooke, R. M. 1991. *Experts in Uncertainty: Opinion and Subjective Probability in Science.* Oxford University Press, Oxford, UK.

[35] Cooke, R. M. 1994. Parameter fitting for uncertain models: Modeling uncertainty in small models. *Reliability and Engineering System Safety* **44** 89–102.

[36] Cooke, R.M. 2012. Uncertainty analysis comes to integrated assessment models for climate change and conversely. *Climate Change*, Special Issue on Improving the Assessment and Valuation of Climate Change Impacts for Policy and Regulatory Analysis.

[37] Cox, L. A., Jr. 2008. Some limitations of "Risk = Threat $\times$ Vulnerability $\times$ Consequence" for risk analysis of terrorist attacks. *Risk Analysis* **28**(6) 1749–1761.

[38] Cox, L. A., Jr. 2009. Game theory and risk analysis. *Risk Analysis* **29**(8) 1062–1068.

[39] Cragin, R. K., S. A. Daly. 2004. *The Dynamic Terrorist Threat: An Assessment of Group Motivations and Capabilities in a Changing World*. RAND Corporation, Santa Monica, CA.

[40] Csiszár, I. 1967. Information-type measures of difference of probability distributions and indirect observation. *Studia Scientiarum Mathematicarum Hungarica* **2** 229–318.

[41] Csiszár, I. 1975. I-divergence geometry of probability distributions and minimization problems. *Annals of Probability* **3** 146-158.

[42] Dawes, R. M., B. Corrigan. 1974. Linear models in decision making. *Psychological Bulletin* **81** 95–106.

[43] Department of Homeland Security, Risk Steering Committee. 2008. *DHS Risk Lexicon*. Department of Homeland Security, Washington, DC.

[44] Department of Homeland Security, Office of the Press Secretary. 2011. *DHS Announces Grant Guidance for Fiscal Year (FY) 2011 Preparedness Grants*, http://www.dhs.gov/news/2011/05/19/dhs-announces-grant-guidance-fiscal-year-fy-2011-preparedness-grants.

[45] Dillon, R.L., R. John and D. von Winterfeldt. 2002. Assessment of cost uncertainties for large technology projects: A methodology and an application. *Interfaces* **32**(4) 52–66.

[46] Dillon, R. L., R. M. Liebe, T. Bestafka. 2009. Risk-based decision making for terrorism applications. *Risk Analysis* **29**(3) 321–335.

[47] Douglas, P. H. 1976. The Cobb-Douglas production function once again: Its history, its testing, and some new empirical values. *Journal of Political Economy* **84**(5) 903–916.

[48] Drew, S., T. Homem-de-Mello. 2012. Some large deviations results for Latin hypercube sampling. *Methodology and Computing in Applied Probability* **14**(2) 203–232.

[49] Du, C., D. Kurowicka, R. M. Cooke. 2006. Techniques for generic probabilistic inversion. *Computational Statistics and Data Analysis* **50**(1) 1164–1187.

[50] Dupačová, J., R. J-B Wets. 1988. Asymptotic behavior of statistical estimators and of optimal solutions of stochastic optimization problems. *The Annals of Statistics* **16** 1517–1549.

[51] Edwards, W. 1961. Behavioral decision theory. *Annual Review of Psychology* **12** 473–498.

[52] Edwards, W. 1977. How to use multiattribute utility measurement for social decision making. *IEEE Transactions on Man, Systems, and Cybernetics* **7** 326–340.

[53] Edwards, W., F. H. Barron. 1994. SMARTS and SMARTER: Improved simple methods for multiattribute utility measurement. *Organizational Behavior and Human Decision Processes* **60** 306–325.

[54] Edwards, W., R. Miles, D. von Winterfeldt, eds. 2007. *Advances in Decision Analysis.* Cambridge University Press, Cambridge, New York, NY.

[55] Enders, W., T. Sandler. 2000. Is transnational terrorism becoming more threatening: A time-series investigation. *The Journal of Conflict Resolution* **44**(3) 307–332.

[56] Erkanli, A., D. K. Stangl, P. Muller. 1993. A Bayesian analysis of ordinal data. ISDS Discussion, Paper 93-A01. Duke University, Durham, NC.

[57] Ermoliev, Y. 1983. Stochastic quasi-gradient methods and their applications to systems optimization *Stochastics* **9** 1-36.

[58] Ertem, M., V. M. Bier. 2013. A stochastic network-interdiction model for cyber security. Under review at *European Journal of Operations Research*.

[59] Escobar, M. D., M. West. 1995. Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association* **90** 577–588.

[60] Ezell, B. C., S. P. Bennett, D. von Winterfeldt, J. Sokolowski, A. J. Collins. 2010. Probabilistic risk analysis and terrorism risk. *Risk Analysis* **30**(4) 575–589.

[61] Falkenrath, R. A., R. D. Newman, B. A. Thayer. 1998. *American's Achilles' Heel*. The MIT Press, Cambridge, MA.

[62] Farrow, S. 2007. The economics of homeland security expenditures: Foundational expected cost-effectiveness approaches. *Contemporary Economic Policy* **25**(1) 14–26.

[63] Farrow, S. 2008. The economics of homeland security expenditures: Foundational expected cost-effectiveness approaches. *Contemporary Economic Policy* **25**(1) 14–26.

[64] Ferguson, T. S. 1973. A Bayesian analysis of some nonparametric problems. *The Annals of Statistics* **1** 209–230.

[65] Ferguson, T. S. 1974. Prior distributions on spaces of probability measures. *The Annals of Statistics* **2** 615–629.

[66] Ferguson, T. S. 1983. Bayesian density estimation by mixtures of normal distributions. M. H. Rizvi, J. Rustagi, D. Siegmund, eds. *Recent Advances in Statistics.* Academic Press, New York, 287–409.

[67] Finkel, A.M. 1996. Who is exaggerating? *Discover*, May 1996, 48.

[68] Gelfand, A. E., A. F. M. Smith. 1990. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association* **85**(410) 398–409.

[69] Gelfand, A. E., A. F. M. Smith, T. M. Lee. 1992. Bayesian analysis of constrained parameter and truncated data problems using Gibbs sampling. *Journal of the American Statistical Association* **87** 523–32.

[70] Geman, S., D. Geman. 1984. Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **6** 721–741.

[71] Gneiting, T., A.F. Raftery. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* **102**(477) 359–378.

[72] Geoffrion, A. M. 1972. Generalized Benders decomposition. *Journal of Optimization Theory and Applications* **10**(4) 237–260.

[73] Gibbs, A. L., F. E. Su. 2002. On choosing and bounding probability metrics. *International Statistical Review* **70** 419–435.

[74] Gigerenzer, G., P. M. Todd, ABC Research Group. 1999. *Simple Heuristics That Make Us Smart.* Oxford University Press, New York, NY.

[75] Green, P. E., A. M. Krieger, Y. Wind. 2001. Thirty years of conjoint analysis: Reflections and prospects. *Interfaces* **31** (3) S56–S73.

[76] Green, P. E., V. Srinivasan. 1978. Conjoint analysis in consumer research: Issues and outlook. *Journal of Consumer Research* **5** (2) 103–123.

[77] Green, P. E., V. Srinivasan. 1990. Conjoint analysis in marketing: New developments with implications for research and practice. *Journal of Marketing* **54**(4) 3–19.

[78] Guikema, S. D. 2009. Modeling intelligent actors in reliability analysis: An overview of the state of the art. V. M. Bier, N. Azaiez, eds. *Combining Reliabilityand Game Theory*, Springer Series on Reliability Engineering.

[79] Haphuriwat, N., V. M. Bier, H. Willis. 2011. Deterring the smuggling of nuclear weapons in container freight through detection and retaliation. *Decision Analysis* **8**(2) 88–102.

[80] Haphuriwat, N., V. M. Bier. 2011. Trade-offs between target hardening and overarching protection. *European Journal of Operational Research* **213**(1) 320–328.

[81] Hausken, K., V. M. Bier. 2011. Defending against multiple different attackers. *European Journal of Operational Research* **211** 370–384.

[82] Hausken, K., V. M. Bier, J. Zhuang. 2009. Defending against terrorism, natural disaster, and all hazards. V. M. Bier and M. N. Azaiez, eds. *Game Theoretic Risk Analysis of Security Threats*. Springer, New York, NY, 65–97.

[83] Higle, J. L., S. Sen. 1991. Stochastic decomposition: An algorithm for two-stage linear programs with recourse. *Mathematics of Operations Research* **16** 650-669.

[84] Homem-de-Mello, T. 2008. On rates of convergence for stochastic optimization problems under non-independent and identically distributed sampling. *SIAM Journal on Optimization* **19**(2) 524–551.

[85] Hora, S. 2010. An analytic method for evaluating the performance of aggregation rules for probability densities. *Operations Research* **58**(5) 1440–1449.

[86] Hora, S., M. Jensen. 2002. *Expert Judgement Elicitation.* Swedish Radiation Protection Authority, Stockholm, SE.

[87] Horsky, D., M. R. Rao. 1984. Estimation of attribute weights from preference comparisons. *Management Science* **30** (7) 801–822.

[88] Howard, R. A. 1988. Decision analysis: Practice and promise. *Management Science* **34**(6) 679–695.

[89] Huyer, W., A. Neumaier. 1999. Global optimization by multilevel coordinate search. *Journal of Global Optimization* **14** 331–355.

[90] Huyer, W., A. Neumaier. 2008. SNOBFIT–Stable noisy optimization by branch and fit. *ACM Transactions on Mathematical Software* **35** 1–25.

[91] Janjarassuk, U., J. Linderoth. 2008. Reformulation and sampling to solve a stochastic network interdiction problem. *Networks* **52**(3) 120–132.

[92] Jenelius, E., J. Westin, Å. J. Holmgren. 2010. Critical infrastructure protection under imperfect attacker perception. *International Journal of Critical Infrastructure Protection* **3**(1) 16–26.

[93] Jha, M. K. 2009. Dynamic Bayesian network for predicting the likelihood of a terrorist attack at critical transportation infrastructure facilities. *Journal of Infrastructure Systems* **15**(1) 31–39.

[94] Johnson, S. G. 2011. The NLopt nonlinear-optimization package. http://ab-initio.mit.edu/nlopt

[95] Jones, D.R. 2001. The DIRECT global optimization algorithm. C. A. Floudas, P. M. Pardalos, eds. *Encyclopedia of Optimization*, vol. 1, Kluwer, Boston, MA, 431–440.

[96] Kadane, J. B., J. M. Dickey, R. L. Winkler, W. S. Smith, S. C. Peters. 1980. Interactive elicitation of opinion for a normal linear model. *Journal of the American Statistical Association* **75**(372) 845–854.

[97] Kall, P., A. Ruszczyński, K. Frauendorfer. 1988. Approximation techniques in stochastic programming. Y. Ermoliev, R. J-B. Wets, eds. *Numerical Techniques for Stochastic Optimization*. Springer-Verlag, Berlin, 33–64.

[98] Kaniovski, Y. M., A. J. King, R. J-B. Wets. 1995. Probabilistic bounds (via large deviations) for the solutions of stochastic programming problems. *Annals of Operations Research* **56** 189–208.

[99] Keeney, R., H. Raiffa. 1976. *Decisions with Multiple Objectives: Preferences and Value Trade-Offs.* John Wiley & Sons, Inc, New York, NY.

[100] Keeney, R. L., D. von Winterfeldt. 2011. A Value Model for Evaluating Homeland Security Decisions. *Risk Analysis* **31**(9) 1470–1487.

[101] Kiekintveldt, C., M. Tambe. 2010. Robust Bayesian methods for Stackelberg security games (extended abstract). *9th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2010)*, Toronto, Canada.

[102] Kirkpatrick, S., C. D. Gelatt Jr., M. P. Vecchi. 1983. Optimization by simulated annealing. *Science* **220**(4598) 671–680.

[103] Kirkwood, C. W., R. K. Sarin. 1985. Ranking with partial information: A method and an application. *Operations Research* **33** (1) 38–48.

[104] Koivu, M. 2005. Variance reduction in sample approximations of stochastic programs. *Mathematical Programming* **103**(3) 463–485.

[105] Kraan, B., T. Bedford. 2005. Probabilistic inversion of expert judgments in the quantification of model uncertainty. *Management Science* **51**(6) 995–1006.

[106] Kullback, S., R. A. Leibler. 1951. On information and sufficiency. *Annals of Mathematical Statistics* **22**(1) 79–86.

[107] Kurowicka, D., M. Nauta, K. Jozwiak, J. Katarzyna, R. M. Cooke. 2010. Updating parameters of the chicken processing line model. *Risk Analysis* **30**(6) 934–944.

[108] Lancaster, T., S.J. Jun. 2010. Bayesian quantile regression methods. *Journal of Applied Econometrics* **25**(2) 287–307.

[109] Lapan, H. E., T. Sandler. 1993. Terrorism and signaling. *European Journal of Political Economy* **9**(3) 383–397.

[110] Lilien, G. L., A. Rangaswamy, A. D. Bruyn. 2007. *Principles of Marketing Engineering*. Trafford Publishing, Bloomington, IN.

[111] Lin, S.W., V.M. Bier. 2008. A study of expert overconfidence. *Reliability Engineering and System Safety* **93** 711–721.

[112] Major, J. 2002. Advanced techniques for modeling terrorism risk. *Journal of Risk Finance* **4**(1) 15–24.

[113] March, J. G., L. S. Sproull, M. Tamuz. 1991. Learning from samples of one or fewer. *Organization Science* **2**(1) 1–13.

[114] Martin, O. C., S. W. Otto. 1993. Combining simulated annealing with local search heuristics. *Annals of Operations Research* **63**(1) 57–75.

[115] McCabe, C., J. Brazier, P. Gilks, A. Tsuchiya, J. Roberts, A. O'Hagan, K. Stevens. 2006. Using rank data to estimate health state utility. *Journal of Health Economics* **25** 418–431.

[116] McFadden, D. 1977. Quantal choice analysis: A survey. *Annals of Economic and Social Measurement* **5** 363–390.

[117] McFadden, D. 1994. Contingent valuation and social choice. *American Journal of Agricultural Economics* **76** 689–708.

[118] McFadden, D., K. Train. 2000. Mixed MNL models for discrete response. *Applied Econometrics* **15** 447–470.

[119] Merrick, J.R.W. 2008. Getting the right mix of experts. *Decision Analysis* **5**(1) 43–52.

[120] Mohtadi, H., A. Murshid. 2009. The risk of catastrophic terrorism: An extreme value approach. *Journal of Applied Econometrics* **24** 537–559.

[121] Morton, D. P. 2011. Stochastic network interdiction. Cochran, J. J., ed. *Wiley Encyclopedia of Operations Research and Management Science*. John Wiley & Sons, Hoboken, NJ.

[122] Morton, D. P., F. Pan, K. J. Saeger. 2007. Models for nuclear smuggling interdiction. *IIE Transactions* **39** 3–14.

[123] National Research Council, Committee to Review the Department of Homeland Security's Approach to Risk Analysis. 2010. *Review of the Department of Homeland Security's Approach to Risk Analysis*.

[124] Nelder, J.A., R. Mead. 1965. A simplex method for function minimization. *The Computer Journal* **7**(4) 308–313.

[125] Neslo, R., F. Micheli, C. V. Kappel, K. A. Selkoe, B. S. Halpern, R. M. Cooke. 2011. Modeling stakeholder preferences with probabilistic inversion: Application to prioritizing marine ecosystem vulnerabilities. I. Linkov, E. Ferguson, V. Magar, eds. *Real Time and Deliberative Decision Making: Application to Risk Assessment for Non-chemical Stressors.* Springer, Amsterdam, NL, 265–284.

[126] Parnell, G. S., L. L. Borio, G. G. Brown, D. Banks, A. G. Wilson. 2008. Scientists Urge DHS to Improve Bioterrorism Risk Assessment. *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science* **6**(4) 353–356.

[127] Paruchuri, P., J. Pearce, J. Marecki, M. Tambe, F. Ordonez, S. Kraus. 2008. Playing games with security: An efficient exact algorithm for solving Bayesian Stackelberg games. *AAMAS-2008 Conference*. Estoril, Portugal, 895–902.

[128] Patcha, A., J. M. Park. 2006. A game theoretic formulation for intrusion detection in mobile ad hoc networks. *International Journal of Network Security* **2**(2) 131–137.

[129] Paté-Cornell, M. E., S. Guikema. 2002. Probabilistic modeling of terrorist threats: A systems analysis approach to setting priorities among countermeasures. *Military Operations Research* **7**(4) 5–20.

[130] Pita, J., H. Bellamane, M. Jain, C. Kiekintveldt, J. Tsai, F. Ordonez, M. Tambe. 2009. Security applications: Lessons of real-world deployment. *ACM SIGecom Exchanges* **8**(2).

[131] Powell, R. 2007. Defending against terrorist attacks with limited resources. *American Political Science Review* **101**(3) 527–541.

[132] Powell, M. J. D. The BOBYQA algorithm for bound constrained optimization without derivatives. *DAMTP Technical Report 2009/NA06*. Cambridge, England.

[133] Qian, P. Z., C. J. Wu. 2009. Sliced space-filling designs. *Biometrika* **96**(4) 945–956.

[134] Qian, P. Z. 2009. Nested Latin hypercube designs. *Biometrika* **96**(4) 957–970.

[135] Rai, B. K., R. Sarin. 2007. Generalized contest success functions. *Economic Theory* **40**(1) 139–149.

[136] Rao, J. S., M. Sobel. 1980. Incomplete Dirichlet integrals with applications to ordered uniform spacings. *Journal of Multivariate Analysis* **10** 603–610.

[137] Revelt, D., K. Train. 1998. Mixed logit with repeated choices: Households' choice of appliance efficiency level. *Review of Economics and Statistics* **LXXX**(4) 647–657.

[138] Richardson L. 2007. *What Terrorists want: Understanding the Enemy, Containing the Threat.* Random House Trade Paperbacks, Reprint Edition.

[139] Rinnooy Kan, A. H. G., G. T. Timmer. 1987. Stochastic global optimization methods. *Mathematical Programming* **39** 27–78.

[140] Rios, L. M., N. V. Sahinidis. 2012. Derivative-free optimization: A review of algorithms and comparison of software implementations. *Journal of Global Optimization* 1–47.

[141] Robinson, S. M. 1993a. Shadow prices for measures of effectiveness, I: Linear model. *Operations Research* **41**(3) 518–535.

[142] Robinson, S. M. 1993b. Shadow prices for measures of effectiveness, II: General model. *Operations Research* **41**(3) 536–548.

[143] Rosoff, H., R. John. 2009. Decision analysis by proxy for the rational terrorist. *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI-09), Quantitative Risk Analysis for Security Applications.* Retrieved September 16, 2011, http://teamcore.usc.edu/QRASA-09/.

[144] Rosoff, H., D. von Winterfeldt. 2007. A Risk and Economic Analysis of Dirty Bomb Attacks on the Ports of Los Angeles and Long Beach. *Risk Analysis* **27**(3) 533–546.

[145] Rowan, T. 1990. *Functional Stability Analysis of Numerical Algorithms*. Ph.D. thesis, Department of Computer Sciences, University of Texas at Austin, Austin, TX.

[146] Rubin B., Rubin J., eds. 2002. *Anti-American Terrorism and the Middle East*. Oxford University Press, New York, NY.

[147] Rust, R., K. Lemon, V. Zeithaml. 2004. Return on marketing: Using customer equity to focus marketing strategy. *Journal of Marketing* **68**(1) 109–127.

[148] Sahinidis, N. V., M. Tawarmalani. 2011. GAMS solver manual on BARON. http://www.gams.com/dd/docs/solvers/baron.pdf

[149] Sandler, T., H. E. Lapan. 1988. The calculus of dissent: An analysis of terrorists' choice of targets. *Synthese* **76**(2) 245–261.

[150] Sandler, T., K. Siqueira. 2009. Games and terrorism: Recent developments. *Simulation and Gaming* **40**(2) 164–192.

[151] Sándor, Z., M. Wedel. 2002. Profile construction in experimental choice designs for mixed logit models. *Marketing Science* **21**(4) 455–475.

[152] Sarabando, P., L. C. Dias. 2009. Multiattribute choice with ordinal information: A comparison of different decision rules. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans* **39** (3) 545–554.

[153] Schoemaker, P. J. H., C. D. Waid. 1982. An experimental comparison of different approaches to determining weights in additive utility models. *Management Science* **28** 182–196.

[154] Seppäläinen, T. 2010. *Lecture Notes on Basics of Stochastic Analysis.* University of Wisconsin-Madison Department of Mathematics, WI. Retrieved on January 1, 2012, http://www.math.wisc.edu/ seppalai/sa-book/notes.pdf.

[155] Shan, X., J. Zhuang. 2012. Hybrid defensive resource allocations in the face of partially strategic attackers, and comparison with fully endogenous and exogenous model. *European Journal of Operational Research* **228**(1) 262–272.

[156] Shapiro, A. 1993. Asymptotic behavior of optimal solutions in stochastic programming. *Mathematics of Operations Research* **18**(4) 829–845.

[157] Shapiro, A. 2003. Monte Carlo sampling methods. A. Ruszczynski, A. Shapiro, eds. *Handbooks in Operations Research and Management Science*, 10, Elsevier, 353–425.

[158] Shapiro, A., D. Dentcheva, A. Ruszczyński. 2009. *Lectures on Stochastic Programming: Modeling and Theory.* SIAM-Society for Industrial and Applied Mathematics, Philadelphia, PA.

[159] Shapiro, A., T. Homem-de-Mello. 2000. On the rate of convergence of optimal solutions of Monte Carlo approximations of stochastic programs. *SIAM Journal of Optimization* **11**(1) 70–86.

[160] Shocker, A. D., V. Srinivasan. 1973. Linear programming techniques for multidimensional analysis of preferences. *Psychometrika* **38** (6) 337–369.

[161] Shocker, A. D., V. Srinivasan. 1979. Multiattribute approaches for product concept evaluation and generation: A critical review. *Journal of Marketing Research* **16** 159–180.

[162] Shlyakhter, A.I., D. M. Kammen, C. L. Broido, R. Wilson. 1994. Quantifying the credibility of energy projections from trends in past data: the U.S. Energy Sector. *Energy Policy*, February 1994, 119-130.

[163] Siqueira, K., T. Sandler. 2006. Terrorists versus the government: Strategic interaction, support, and sponsorship. *Journal of Conflict Resolution* **50**(6) 1–21.

[164] Skaperdas, S. 1996. Contest success functions. *Economic Theory* **7**(2) 283–290.

[165] Stewart, T. J. 1996. Robustness of additive value function methods in MCDM. *Journal of Multi-Criteria Decision Analysis* **5** 301–309.

[166] Sticha, P. J., D. Buede, R. L. Rees. 2005. Apollo: An analytical tool for predicting a subject's decision making. *International Conference on Intelligence Analysis Proceedings*.

[167] Stillwell, W. G., D. A. Seaver, W. Edwards. 1981. A comparison of weight approximation techniques in multiattribute utility decision making. *Organizational Behavior and Human Performance* **28** 62–77.

[168] Streetman, S. 2010. A scalable approach to adversary modeling for terrorism risk analysis. *Society for Risk Analysis Annual Meeting*. Salt Lake City, Utah.

[169] Tierney, L. 1994. Markov chains for exploring posterior distributions. *The Annals of Statistics* **22** 1701–1762.

[170] Tierney, L. 1996. Introduction to general state-space Markov chain theory. W. R. Gilks, S. Richardson, D. J. Spiegelhalter, eds. *Markov Chain Monte Carlo in Practice.* Chapman and Hall, London, UK, 59–74.

[171] Tversky, A., D. Kahneman. 1974. Judgment under uncertainty: Heuristics and biases. *Science* **185**(4157) 1124–1131.

[172] Van Slyke, R. M., R. J. Wets. 1969. L-shaped linear programs with applications to optimal control and stochastic linear programming. *SIAM Journal of Applied Mathematics* **17** 638–663.

[173] Viswanathan, J., I. E. Grossmann. 1990. A combined penalty function and outer-approximation method for MINLP optimization. *Computers & Chemical Engineering* **14**(7) 769–782.

[174] von Winterfeldt, D., W. Edwards. 1986. *Decision Analysis and Behavioral Research.* Cambridge University Press, New York.

[175] Wang, C., V. M. Bier. 2011. Target-hardening decisions based on uncertain multiattribute terrorist utility. *Decision Analysis* **8**(4) 286–302.

[176] Wang, C., V.M. Bier. 2012. Optimal defensive allocations in the face of uncertain terrorist preferences, with an emphasis on transportation. *Homeland Security Affairs,* DHS Center of Excellence Science and Technology Student Papers, No.4.

[177] Wang, C., V. M. Bier. 2013a. Expert elicitation of adversary preferences using ordinal judgments. *Operations Research* **61**(2) 372–385.

[178] Wang, C., V. M. Bier. 2013b. Quantifying adversary capabilities to inform target-hardening decisions. Submitted to *Risk Analysis*.

[179] Wang, C., V. M. Bier. 2013c. Construction of multiattribute utility functions from ordinal preference rankings. Under revision for L.A. Cox, Jr., J.J. Cochran, eds. *Breakthroughs in Decision and Risk Analysis*.

[180] Washburn, A., Wood, K., 1995. Two-person zero sum games for network interdiction. *Operations Research* **43** 243–251.

[181] Watson, S. R., D. M. Buede. 1987. *Decision Synthesis.* Cambridge University Press, England.

[182] Wein, L. M., A. H. Wilkins, M. Baveja, S. E. Flynn. 2006. Preventing the importation of illicit nuclear materials in shipping containers. *Risk Analysis* **26**(5) 1377–1393.

[183] Westerlund, T., R. Pörn. 2002. Solving pseudo-convex mixed integer optimization problems by cutting plane techniques. *Optimization and Engineering* **3**(3) 253–280.

[184] Williams, H. P. 1999. *Model Building in Mathematical Programming* 4th ed. John Wiley & Sons, Chichester, UK.

[185] Willis, H. H., A. R. Morral, T. K. Kelly, J. J. Medby. 2005. *Estimating Terrorism Risk.* RAND Corporation, Santa Monica, CA.

[186] Willis, H. H. 2007. Guiding Resource Allocations Based on Terrorism Risk. *Risk Analysis* **27**(3) 597–606.

[187] Willis, H. H. 2008. *Challenges of Applying Risk Management to Terrorism Security Policy*. RAND Corporation, Santa Monica, CA.

[188] Woo, G. 2002. Quantifying insurance terrorism risk. Prepared for the National Bureau of Economic Research Meeting, Cambridge, MA.

[189] Wood, R.K. 1993. Deterministic network interdiction. *Mathematical and Computer Modelling* **17** 1–18.

[190] Yuille, A., A. Rangarajan. 2003. The concave-convex procedure (CCCP). *Neural Computation* **15**(4) 915–936.

[191] Zhuang, J., V. M. Bier. 2007. Balancing terrorism and natural disasters âŁ“ Defensive strategy with endogenous attacker effort. *Operations Research* **55**(5) 976–991.

[192] Zhuang, J., V. M. Bier, O. Alagoz. 2010. Modeling secrecy and deception in a multiple-period attacker-defender signaling game. *European Journal of Operational Research* **203**(2) 409–418.