

Risk Analysis in Support of Intelligence Analysis

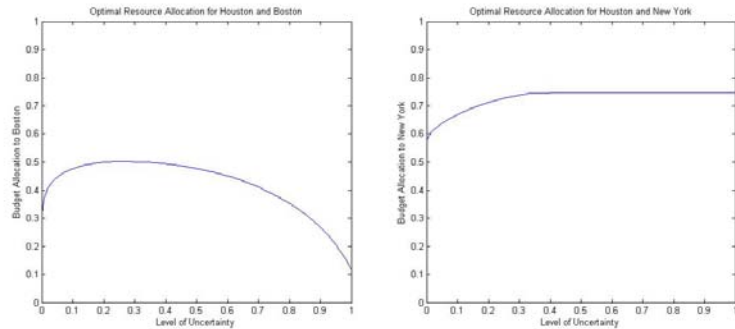
Modeling Area: Risk Assessment

Principal Investigator: Vicki Bier

Institution: University of Wisconsin-Madison

Other Investigators: Chen Wang (Graduate student)

Brief Description: The main purpose of this research project is to develop methods for identifying the benefits of good intelligence about the terrorists' goals and motivations. The proposed model will use Bayesian analysis, optimization, and game theory to explore how optimal resource allocations might vary in the face of greater or lesser uncertainty, in a dynamic environment in which the defender can update her knowledge about the goals and motivations of multiple heterogeneous attackers based on intelligence gained from observing previous attacks (or attempted attacks). Thus, it combines intelligence analysis with risk analysis.



As the defender's uncertainty about the attacker's preference increases, the optimal resource allocation can involve either greater or lesser hedging.

Objectives: The objectives of this research are: (a) to develop a methodology for quantifying the benefits of intelligence about terrorists' goals and motivations by allowing the defender to update her knowledge based on previous (successful or attempted) attacks; (b) to investigate how such intelligence affects the defender's optimal resource allocations; and (c) to determine the conditions under which intelligence is likely to be especially valuable.

Interfaces to other Center Projects: This project builds on and extends the results of previous work on innovations in risk and economic modeling of counter-terrorism. In particular, it is an extension of Bier et al. (2007), reflecting previous CREATE-funded work on defensive resource allocation in the face of uncertainty about attacker preferences. It also relates to ongoing foundational work on the interface between risk analysis and intelligence analysis in homeland security.

Interfaces to non-Center Projects: None at present

Major Products and Customers: The major deliverable of this work will be a methodology for updating the attacker weights on various attributes (e.g., fatalities, property losses, infrastructure damage, attack difficulty, etc.), by observing the choice of targets in previous attacks (or attempted attacks). The major customers for this work will be the intelligence community—including for example the Office of Intelligence and Analysis, and the Command, Control, and Interoperability Division of the Directorate for Science and Technology, within the Department of Homeland Security.

We also hope to develop guidelines on the circumstances under which intelligence is likely to be especially valuable (e.g., when the value of additional information is likely to be high). This could help contribute to cost-effective allocation of defensive resources, through more targeted intelligence collection and dissemination.

Technical Approach: This project will use Bayesian analysis, optimization, and game theory to identify the optimal allocation of defensive resources in the face of intelligence about the attacker's preferences. We model it first as a sequential game (defender allocates her resources, attacker observes the resource allocations and chooses an attack strategy), and then as a dynamic game (with repeated interactions over time). To specify the defender's uncertainty about the attacker's target valuations, we assume that targets are valued according to a multi-attribute utility function with m attributes, and that the list of attributes is known to both the attacker and the defender.

We allow the defender to be uncertain about the weights placed by the attacker on the various attributes. We model this uncertainty by a Dirichlet prior distribution. Choice of this functional form allows us to vary the extent of the defender's uncertainty by changing a single parameter, while holding the expected values of the attribute weights constant (and ensuring that they sum to one). Interestingly, preliminary results to date indicate that as the defender's uncertainty about the attacker's preferences increases, the optimal resource allocation can involve either greater or lesser hedging. Moreover, the optimal total expected loss to the defender is not always increasing in the extent of the defender's uncertainty, suggesting that intelligence is sometimes of little or no value.

In the dynamic (i.e., repeated) game, the defender updates her knowledge of the attacker's attribute weights by Bayesian updating after each observed attack (or attempted attack). This results in a truncated Dirichlet posterior distribution, in which possible attribute weights that are inconsistent with the observed choice of target are excluded. In subsequent extensions, we plan to allow for multiple different types of attackers—e.g., expert computer hackers versus “script kiddies,” state actors (traditional military opponents) versus non-state actors (terrorists), Islamic terrorists versus animal-rights or ecological terrorists, etc. In such cases, the defender must defend against multiple types of attackers (described by different Dirichlet distributions for their attribute weights) simultaneously.

Major Milestones and Dates: The first four milestones below should be achievable within one year. Additional milestones (shown in italics) are conditional on receipt of an additional year of funding.

1. Establish the basic theoretical framework, and complete the analysis for a two-target, two-attribute game -- by September 2008.
2. Develop computational techniques for the N-target, N-attribute case -- by February 2009.
3. Apply the resulting model to a realistic attacker objective function, based on attributes and weights generated by Beitel et al. (2004), for a large number of attributes -- by June 2009.
4. Complete a paper using risk analysis to quantify the benefits of intelligence in counterterrorism resource allocation -- by September 2009.
5. *Extend the model to the case of multiple different attackers, starting with the analysis for a two-attacker model -- by January 2010.*
6. *Develop computational techniques for the N-attacker case and apply it to realistic data -- by July 2010.*
7. *Complete another paper documenting the results for multiple heterogeneous attackers -- by September 2010.*