

REPORT DOCUMENTATION PAGE				Form Approved OMB NO. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 20-03-2011		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 1-Jul-2009 - 31-Dec-2010	
4. TITLE AND SUBTITLE Final Report: Adversarial Risk Analysis for Dynamic Network Routing				5a. CONTRACT NUMBER W911NF-09-1-0337	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER 9620AZ	
6. AUTHORS David Banks				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES Duke University Office of Research Support Duke University Durham, NC 27705 -				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211				10. SPONSOR/MONITOR'S ACRONYM(S) ARO	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) 56711-NS-DRP.2	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT Route selection in hostile terrain is a problem in adversarial risk analysis (ARA). This report summarizes work done to date that addresses the routing problem from the perspective of decision analysis, which takes account of both asymmetric uncertainties and the strategy of an intelligent opponent. The methodology is novel, and usefully different from traditional game theory or risk analysis. The main results of this funding are (1) a paper that is under second review by _Naval Research Logistics_ and (2) presentations at a number of conferences and colloquia.					
15. SUBJECT TERMS Bayesian, game theory, IEDs, risk analysis, routing games					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON David Banks
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 919-684-3743

Report Title

Final Report: Adversarial Risk Analysis for Dynamic Network Routing

ABSTRACT

Route selection in hostile terrain is a problem in adversarial risk analysis (ARA). This report summarizes work done to date that addresses the routing problem from the perspective of decision analysis, which takes account of both asymmetric uncertainties and the strategy of an intelligent opponent. The methodology is novel, and usefully different from traditional game theory or risk analysis. The main results of this funding are (1) a paper that is under second review by _Naval Research Logistics_ and (2) presentations at a number of conferences and colloquia.

List of papers submitted or published that acknowledge ARO support during this reporting period. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Number of Papers published in peer-reviewed journals: 0.00

(b) Papers published in non-peer-reviewed journals or in conference proceedings (N/A for none)

Number of Papers published in non peer-reviewed journals: 0.00

(c) Presentations

Invited presentations on ARA at Georgia Tech, the Naval Postgraduate School, the Probability and Statistics Conference, the King Juan Carlos University, the International Statistics in Business and Industry Conference, the Conference on Algorithmic Decision Making (sponsored by DIMACS at Rutgers), Pennsylvania State University, the Army Conference on Applied Statistics, Yale University, the IHSS Research Summit, the Cincinnati Chapter of the American Statistical Association, Duke University, the INFORMS conference at Monterey, the University of California at Davis, the University of California at Berkeley, and the Los Alamos National Laboratory. Other presentations of this research are scheduled.

Number of Presentations: 16.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts): 0

Peer-Reviewed Conference Proceeding publications (other than abstracts):

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts): 0

(d) Manuscripts

This paper was submitted to _Naval Research Logistics_ in November, 2009. We received a largely favorable referee report in July, 2010. The paper was revised and resubmitted in December, 2010 and we await the editor's decision.

The title of the paper is "Network Routing for Counterinsurgency: An Adversarial Risk Analysis Framework".

Number of Manuscripts: 1.00

Patents Submitted

Patents Awarded

Awards

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
Francesca Petralia	0.75
FTE Equivalent:	0.75
Total Number:	1

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	National Academy Member
David Banks	0.17	No
FTE Equivalent:	0.17	
Total Number:	1	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period:	0.00
The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:.....	0.00
The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:.....	0.00
Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):	0.00
Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:.....	0.00
The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense	0.00
The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:	0.00

Names of Personnel receiving masters degrees

<u>NAME</u> Shouqiang Wang
Total Number: 1

Names of personnel receiving PHDs

<u>NAME</u>
Total Number:

Names of other research staff

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

The main accomplishment has been the development of a new solution concept for the decision-analysis approach in a game theory, and the application of that concept to the problem of routing a convoy through a network in which intelligent adversaries may stage attacks that cause cumulative damage along the route. The work is described in a paper that is under review, but more importantly, the solution concept is gaining attention within the defense community (as reflected in work by other researchers, and success in obtaining further funding). The solution concept has been presented at a large number of conferences and seminars, and particularly at the Naval Postgraduate School and the Los Alamos National Laboratory, which are centers for this kind of work.

Technology Transfer

Final Report on DARPA Grant W911NF0910337:
Adversarial Risk Analysis for Dynamic Network Routing

David Banks, Department of Statistical Science, Duke University
banks@stat.duke.edu; 919-684-3743

Foreword: The U.S. military relies upon logistical networks. These networks carry the personnel and material that are critical for mission success. In many situations, strategic adversaries attempt to disrupt these networks, e.g., through use of IEDs. Previous work in this area used game theory to identify optimal routes, but that approach cannot incorporate probabilistic information from human intelligence or prior history. The Bayesian approach proposed in our research takes full advantage of such information.

List of Appendices, Illustrations and Tables

Appendix A: “Network Routing for Insurgency: An Adversarial Risk Analysis Framework”

Appendix B: Computer code.

Problem Statement

This project has combined tools from game theory, network analysis, Bayesian inference, and risk analysis. It represents an extension of the new field of adversarial risk analysis (cf. Rios Insua, Rios, and Banks, 2009) into the domain of secure routing across a network in which there is probabilistic information about threats and relevant historical information.

Suppose that one has a network of roads, and missions that require traversals of the network between various points at various times. There is an enemy that can deploy resources to set traps or to block segments of the network, and one must select a route and allocate security assets to improve the overall probability of mission success while controlling mission cost. To inform the choice of route and the allocation of security assets, one can use the observed record from previous traversals and beliefs or intelligence about enemy operations. In general, the selection of the route is determined in advance; in some cases that choice may be altered during the traversal, perhaps as a result of an obstacle or an attack.

This research has emphasized practical application, although some components entail academic research. To ensure relevance, we met with the leading researchers at the Naval Postgraduate School who work on modeling and mitigation of threats from improvised explosive devices (IEDs) and with scientists at Los Alamos National Laboratory who apply game theory to network interdiction.

Our primary goal was to develop new methodology for convoy routing that would allow decision-makers to make better use of the available information than traditional game theory permits. A secondary goal was to further the progress of adversarial risk analysis as a solution concept in the military and counterterrorism communities.

Important Results

The primary results of the proposed research are:

1. a theoretical framework for route selection that incorporates modeling of decision processes by an intelligent adversary;
2. a paper that describes the convoy-routing methodology in detail;
3. prototype software that handles non-trivial problems in routing choice from an adversarial risk analysis perspective;
4. a growing community of researchers who consider this approach a useful alternative to classical game theory.

These results are described in more detail below.

The theoretical framework for adversarial risk analysis (ARA) is laid out in Rios, Rios Insua, and Banks (2009), Banks Petralia and Wang (2011) and in the paper in the appendix of this report. Essentially, ARA provides an explicit method for implementing the decision analysis paradigm (cf. Kadane and Larkey, 1982; Raiffa, 1982). The key is a “mirroring” argument, in which the decision-maker constructs a model of the analysis that the opponent performs, and that model may include the fact that the opponent is building a corresponding model for the analysis of the decision-maker. This may be regarded as a Bayesian version of level-k thinking, as described in Stahl and Wilson (1995). The ARA approach is novel, and offers significant advantages compared to other methods for strategic planning.

The paper in the appendix applies the ARA approach to the convoy routing problem. Among its contributions are an existence proof of the mirroring fixed-point solution; an algorithm that robustly converges to that solution; a statement of sufficient conditions under which the ARA solution coincides with the Bayes Nash equilibrium solution; two worked examples; and a study of the computational requirements for the ARA solution.

The software used in this research is provided in the second appendix. It is written in MATLAB and runs on a personal computer. The software consists of several routines which collectively generate the solutions and tables for the manuscript in Appendix A.

ARA has been gaining attention in the military and counterterrorism communities. In the last three months I have been asked to referee two papers on ARA that were written by people who are not part of my research circle. The ARA perspective formed the basis for a workshop at the DIMACS center at Rutgers University, and has been featured in invited sessions at four statistics conferences (with two more scheduled for this summer). The reason for this success is probably a combination of some reasonably high-profile publications, a large number of seminars and colloquia at various institutions, and a nearly universal sense among military decision-makers that classical game theory is inadequate for the kinds of problems that arise in practice. The ARA approach may seem a little technical, but in fact it simply mathematizes the kinds of decision processes that normal people regularly employ.

Bibliography

- Arce, D. and T. Sandler (2007). Terrorist Signalling and the Value of Intelligence, *British Journal Political Science*, **37**, 573–586.
- Aubin, J. P. (1993). *Optima and Equilibria*, Springer-Verlag, New York, N.Y.
- Banks, D. and S. Anderson (2006). Game Theory and Risk Analysis in the Context of the Smallpox Threat, in *Statistical Methods in Counterterrorism*, A. Wilson, G. Wilson and D. Olwell, eds., Springer, New York, N.Y., pp. 9–22.
- Banks, D., Petralia, F., and Wang, S. (2011). “Adversarial Risk Analysis: Borel Games” (with discussion). To appear in *Applied Stochastic Models in Business and Industry*.
- Bayrak, H., and Bailey, M. D. (2008). ”Shortest Path Network Interdiction with Asymmetric Information,” *Networks*, **52**, 133–140.
- Berger, U. (2005). Fictitious Play in $2 \times n$ Games. *J. Econ. Theory*, **120**, 139-154.
- Brown, G. W. (1949). Some Notes on Computation of Games Solutions. Report P-78. The Rand Corporation.
- Brown, G., M. Carlyle, Salmeron, J. and K. Wood (2006). “Defending critical infrastructure,” *Interfaces*, **36**, 530–544.
- Brown, G., W.M. Carlyle, and R. Wood (2008). “Optimizing Department of Homeland Security Defense Investments: Applying Defender-Attacker(-Defender) Optimization to Terror Risk Assessment and Mitigation,” Appendix E, National Academies Press, Washington, D.C.
- Camerer, C. (2003). *Behavioral Game Theory*, Princeton University Press, Princeton, N.J.
- Dimitrov, N., Michalopoulos, D., Morton, D., Nehme, M., Pan, F., Popova, E., Schneider, E. and Thoreson, G. (2009). “Network Deployment of Radiation Detectors with Physics-Based Detection Probability Calculations,” *Annals of Operations Research*, to appear. DOI: 10.1007/s10479-009-0677-2.

- Dresher, M. (1961). *Games of Strategy: Theory and Applications*, RAND/Prentice Hall, Englewood Cliffs, N.J.
- Gintis, H. (2009). *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences*, Princeton University Press, Princeton, N.J.
- Goodie, A., Doshi, P., and Young, D. (2010). “Levels of Theory-of-Mind Reasoning in Competitive Games,” to appear in the *Journal of Behavioral Decision Making*.
- Gutfraind, A., Hagberg, A., and Pan, F. (2009). “Optimal Interdiction of Unreactive Markovian Evaders,” in *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, Willem-Jan van Hoeve and John N. Hooker (eds.), Springer, Heidelberg, pp. 102–116.
- Harsanyi, J. (1967a). “Games with Incomplete Information Played by ‘Bayesian’ Players, Part I: The Basic Model,” *Management Science*, **14**, 159–182.
- Harsanyi, J. (1967b). “Games with Incomplete Information Played by ‘Bayesian’ Players, Part II: Bayesian Equilibrium Points,” *Management Science*, **14**, 320–334.
- Harsanyi, J. (1967c). “Games with Incomplete Information Played by ‘Bayesian’ Players, Part III: The Basic Probability Distribution of the Game”, *Management Science*, **14**, 486–502.
- Harsanyi, J. (1982). “Subjective Probability and the Theory of Games: Comments on Kadane and Larkey’s Paper,” *Management Science*, **28**, 120–124.
- Hausken, K. (2002). “Probabilistic Risk Analysis and Game Theory”, *Risk Analysis*, **22**, 17–27.
- Israeli, E., and Wood, R. K. (2002). ”Shortest-Path Network Interdiction,” *Networks*, **40**, 97–111.
- Kadane, J. B., and Larkey, P. D. (1982). “Subjective Probability and the Theory of Games”, *Management Science*, **28**, 113–120; reply: 124.

- Krishna, V. (1992). Learning in Games with Strategic Complementarities. Mimeo. Harvard University.
- Lemke, C.E., and Howson, J.T. (1964). "Equilibrium Points of Bimatrix Games," *Journal of the Society for Industrial and Applied Mathematics*, **12**, 413-423.
- Morris, S. (1995). "The Common Prior Assumption in Economic Theory," *Economics and Philosophy*, **11**, 227-253.
- Morris, S., and Shin, H. S. (1998). "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks," *Economic Review*, **88**, 587-597.
- Myerson, R. (1991). *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge MA.
- Nasar, S. (1998). *A Beautiful Mind*, Simon & Schuster, New York, N.Y.
- O'Hagan, A., Buck, C., Daneshkhah, A., Eiser, J., Garthwaite, P., Jenkinson, D., Oakley, J., Rakow, T. (2006). *Uncertain Judgements: Eliciting Experts' Probabilities*, Wiley, New York, N.Y.
- Papadimitriou, C. (2001). "Algorithms, Games, and the Internet." In: *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing*. 749-753.
- Paté-Cornell, E., and Guikema, S. (2002). "Probabilistic Modeling of Terrorist Threats: A Systematic Analysis Approach to Setting Priorities Among Countermeasures," *Military Operations Research*, **7**, 5-23.
- Raiffa, H. (1982). *The Art and Science of Negotiation*, Harvard University Press, Cambridge, MA.
- Raiffa, H. (2002). *Negotiation Analysis*, Harvard University Press, Cambridge, MA.
- Raiffa, H., and Schlaifer, R. (1961). *Applied Statistical Decision Theory*, MIT Press, Cambridge, MA.

- Rios Insua, D., Rios, J., and Banks, D. (2009). “Adversarial Risk Analysis,” *Journal of the American Statistical Association*, **104**, 841–854.
- Roberson, B. (2006). “The Colonel Blotto Game,” *Economic Theory*, **29**, 1–24.
- Robinson, J. (1951). “An Iterative Method of Solving a Game,” *Annals of Mathematics*, **52**(2), 296–301.
- Stahl, D., and Wilson, P. (1995). “On Players’ Model of Other Players,” *Games and Economic Behavior*, **10**, 218–254.
- Topkis, D. (1998). *Supermodularity and Complementarity*, Princeton University Press, Princeton, NJ.
- Washburn, A., and Wood, R. K. (1995). “Two-Person Zero-Sum Games for Network Interdiction,” *Operations Research*, **43**, 243–251.
- Woodruff, D. (2002). *Network Interdiction and Stochastic Integer Programming*, D. Woodruff, editor. Kluwer Academic Publishers, Boston, MA.

Appendix A: Network Routing for Insurgency: An Adversarial Risk Analysis Framework

Network Routing for Insurgency: An Adversarial Risk Analysis Framework

Shouqiang Wang, Fuqua School of Business
Duke University, shouqiang.wang@duke.edu

David Banks, Department of Statistical Science
Duke University, banks@stat.duke.edu

Abstract

Problems in counterterrorism and corporate competition have prompted research that attempts to combine statistical risk analysis with game theory in ways that support practical decision-making. This paper applies these methods of adversarial risk analysis to the problem of selecting a route through a network in which an opponent chooses vertices for ambush. The motivating application is convoy routing across a road network when there may be improvised explosive devices and imperfect intelligence about their locations.

1 Introduction

An important class of problems in game theory pertains to routing choice through a network when an intelligent adversary is attempting to block passage. A famous example is *Nash*, a game invented by John Nash (cf. Nasar, 1998, p. 77) in which two opponents compete to create an unbroken path from North to South or from East to West, respectively, on a board tiled by hexagons. In real life, analogous problems arise when corporations attempt

to impede each other’s access to critical resources or distribution links. In this paper, the motivating application is a game involving convoy routing on a road network with improvised explosive devices (IEDs).

In our application, the use of IEDs does not necessarily block passage. Rather, the IEDs cause random amounts of damage. The mission fails if the cumulative damage exceeds the value of the convoy, but it is more realistic to suppose that the convoy commander wants to select a route that minimizes total damage, whereas the insurgents want to locate IEDs so as to maximize the damage. In this framework, one has a game between the Defender (the convoy commander) and the Attacker (the insurgents). We suppose this is a normal form game, in that both the selection of the entire route and the decision about the siting of the IEDs are made in advance; i.e., the Defender does not alter the route based upon outcomes that occur along the route, and the Attacker does not plant IEDs in real-time as the route choice is revealed.

Traditional game theorists have treated somewhat similar games. Normal form network interdiction games have been studied by Dimitrov et al. (2009), Bayrak and Bailey (2008), and Washburn and Wood (1995), although these papers make somewhat different assumptions and have different payoff criteria. Extensive form games, in which the route or the IED placement or both are decided adaptively, have been studied by Guttfraind, Hagberg and Pan (2009), Israeli and Wood (2002), and arise generally in Woodruff (2002). More distantly, without the network structure, there is a relationship to Blotto games (cf. Dresher, 1961; Roberson, 2006).

Our concern is that traditional game theory is well known to be an unreliable guide to human behavior (cf. Camerer, 2003). Kadane and Larkey (1982) and Raiffa (1982) proposed decision analysis as an alternative; this is controversial departure from the customary equilibrium solution concepts (e.g., Harsanyi, 1982), but it has appealing advantages in terms of plausibility; Stahl and Wilson (1995) report that people appear to make decisions based on probabilistic models for their opponent’s reasoning. The main difficulty in operationalizing the decision analysis approach is that there has been little exploration of the mechanism whereby a decision-maker formulates the subjective probability distributions that represent their opponent’s behavior.

This paper addresses that difficulty by describing a “mirroring” procedure, in which the decision-maker mimics the opponent’s analysis, taking account of the fact that the opponent is simultaneously performing a symmetric study of the decision-maker’s analysis. The result is a probability distribution over the opponent’s options, and the decision-maker then selects the action (a route) that maximizes expected utility.

We use the term *Adversarial Risk Analysis* (ARA) to describe methods that combine decision analysis with an explicit model for the strategic thinking of one’s opponent, usually through a mirroring argument. The rest of this chapter is organized as follows. Section 2 sets up the model and notation in the convoy routing context. A detailed examination of the ARA solution concept, with connections to related ideas in classical game theory, is given in Section 3. Section 4 discusses computational issues. Section 5 provides conclusions.

2 Model and Notation

Suppose there is a Defender and an Attacker. For narrative simplicity, we will use the female gender when referring to the Defender and the male gender for the Attacker. The Defender wants to start at a location S in a fixed network, and chooses a route that leads to a terminus T . At certain locations in the network, it is possible for the Attacker to cause damage, for example, by placing improvised explosive devices (IEDs). The Defender may have historical information and military intelligence on the placement of the IEDs, which can be encoded in the Attacker’s utility function assessed by her; these can guide her selection of a route. Similarly, the Attacker may possess historical information and intelligence regarding previous routing choices, which can be encoded in the Defender’s loss function as estimated by him; these can guide his choice of where to place a fixed number of available IEDs.

In this game, once the path is selected, the convoy traverses that path to the end, possibly absorbing cumulative damage along the way. The Defender wants to choose the route that will minimize the expected cumulative damage (and, more generally, to discover whether the minimum expected damage exceeds either the value of the convoy or the value of the mission, as appropriate, in which case the trip should be canceled). Similarly, the

Attacker wants to place a fixed number of IEDs, or other kinds of ambush, so as to maximize the expected cumulative damage (and, more generally, to discover whether the maximum expected utility of the inflicted damage is less than the cost of placing the IEDs, in which case no IEDs should be sited). This formulation is different from the resource allocation problems encountered in some interdiction and Colonel Blotto games (e.g., Dresher, 1961; Roberson, 2006); the Defender is picking a single path rather than dividing resources among battlefields, and the Attacker gets no increase in damage from placing multiple IEDs at the same location.

Because the road network hardly changes during a short time window, we assume the underlying structure is static. Of course, in the real world, the Defender might find that a traffic jam precludes passage on a particular street. We do not attempt to model such situations, which would require a more complex dynamic analysis.

Similarly, the convoy routing problem is modeled as a one-shot game. The Defender chooses the entire route at the outset and will continue the chosen route as planned even if there are one or more IED ambushes. And the Attacker places all the IEDs before the convoy sets out, rather than planting new ones adaptively as the path of the convoy is realized. Our model complies with most standard convoy operations in Iraq and Afghanistan, and could serve as a starting point for dynamic extensions.

Our main task in this section is to formulate the network routing problem as a *two-player simultaneous game with private information* in which the payoff depends upon the structure of the roadway graph.

2.1 Routes Through an Undirected Graph

For a road network, we use a vertex to denote a candidate location for the Attacker to place an IED and road segment between vertices is denoted as an undirected edge (i.e., traffic can move in both directions). This set-up has no substantive effect on the analysis except to simplify the notation. Thus, the network is an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{R})$ where $\mathcal{V} = \{v_0 = S, v_1, \dots, v_N, v_{N+1} = T\}$ is the set of vertices and $\mathcal{R} = \{r_1, \dots, r_K\}$ is the set of all possible routes connecting S to T (that do not traverse the same edge twice). Obviously, we may assume $\mathcal{R} \neq \emptyset$. Algebraically, \mathcal{G} can be represented by an $N \times K$ incidence matrix

(denoted by \mathbf{G}) whose rows are indexed by the elements in $\mathcal{V} \setminus \{S, T\}$ and whose columns are indexed by the elements in \mathcal{R} . The (i, j) th element of \mathbf{G} is 1 if the i th vertex is visited on the j th route; otherwise, it is 0.

Figure 1 provides a toy example to illustrate these definitions.

(Insert Figure 1 here.)

There are four possible routes that can connect S to T :

$$\begin{aligned} r_1 &: S \rightarrow v_1 \rightarrow v_3 \rightarrow v_5 \rightarrow v_4 \rightarrow v_6 \rightarrow T \\ r_2 &: S \rightarrow v_1 \rightarrow v_3 \rightarrow v_6 \rightarrow T \\ r_3 &: S \rightarrow v_2 \rightarrow v_3 \rightarrow v_5 \rightarrow v_4 \rightarrow v_6 \rightarrow T \\ r_4 &: S \rightarrow v_2 \rightarrow v_3 \rightarrow v_6 \rightarrow T. \end{aligned}$$

The corresponding incidence matrix is:

$$\mathbf{G} = \begin{pmatrix} & r_1 & r_2 & r_3 & r_4 \\ v_1 & 1 & 1 & 0 & 0 \\ v_2 & 0 & 0 & 1 & 1 \\ v_3 & 1 & 1 & 1 & 1 \\ v_4 & 1 & 0 & 1 & 0 \\ v_5 & 1 & 0 & 1 & 0 \\ v_6 & 1 & 1 & 1 & 1 \end{pmatrix}. \quad (1)$$

The action space for the Defender is identical to \mathcal{R} . A *pure strategy* for the Defender is to choose a particular route in \mathcal{R} . These strategies can be represented by a (column) vector of K dimension, with $K - 1$ zeroes and a one in the k th position to indicate that route k has been chosen. Let \mathcal{D} consist of the K standard basis elements of \mathbb{R}^K . Then the set of all pure strategies for the Defender can be denoted as \mathcal{D} .

In the road network example illustrated in Figure 1, the set of all pure strategies for the Defender is

$$\mathcal{D} = \{\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3, \mathbf{r}_4\},$$

where the strategies correspond to vectors in the obvious way:

$$\begin{aligned}\mathbf{r}_1 &= (1, 0, 0, 0)', \\ \mathbf{r}_2 &= (0, 1, 0, 0)', \\ \mathbf{r}_3 &= (0, 0, 1, 0)', \\ \mathbf{r}_4 &= (0, 0, 0, 1)'. \end{aligned}$$

Notice that \mathcal{D} (and hence the element \mathbf{r}_i) and \mathcal{R} (and hence the element r_i) actually refer to the same object with the former as an algebraic representation and the latter as a graphical representation. To facilitate later exposition, we call it a *mixed strategy* for the Defender if she chooses among the routes in \mathcal{R} at random according to some probability distribution $Q = (q_1, \dots, q_K)$; i.e., route $\mathbf{r}_k \in \mathcal{D}$ is chosen with probability q_k . Denote the set of all probability distributions over \mathcal{D} as

$$\mathcal{Q} = \left\{ Q = (q_1, \dots, q_K) : \sum_{k=1}^K q_k = 1 \text{ and } q_k \geq 0 \text{ for } k = 1, \dots, K \right\}.$$

Similarly, the set of all pure strategies for the Attacker, which we denote as \mathcal{A} , consists of all possible combinations of locations at which the Attacker may site IEDs. Usually, because of financial, time or human resource constraints, the Attacker cannot attack all of the vertices; the subset \mathcal{A} reflects this constraint. Each pure strategy $\mathbf{a} \in \mathcal{A}$ can now be algebraically represented by a binary (column) vector of length N , with ones corresponding to the vertices at which IEDs are placed and zeroes corresponding to vertices without IEDs. Formally,

$$\mathcal{A} = \{ \mathbf{a}_i : \mathbf{a}_i = (\alpha_{1i}, \dots, \alpha_{Ni})' \text{ where } \alpha_{ni} = 0 \text{ or } 1 \text{ for } n = 1, \dots, N \text{ and } i = 1, \dots, I \},$$

where $I = |\mathcal{A}|$ is the cardinality of \mathcal{A} . Obviously, \mathcal{A} is a subset in $\subseteq \{0, 1\}^N$. This formulation assumes, innocuously, that IEDs cannot be placed at S or T , the start or terminus of the route.

In the road network example illustrated in Figure 1, if the Attacker can choose at most three IED locations, then

$$\mathcal{A} = \{ \mathbf{a} : \mathbf{a} = (\alpha_1, \dots, \alpha_N)' : \text{for } \alpha_i = 0 \text{ or } 1 \text{ and } \sum_{i=1}^N \alpha_i \leq 3 \}.$$

A pure strategy $\mathbf{a} = (0, 1, 1, 0, 0, 1)'$ in \mathcal{A} occurs if the Attacker makes a non-randomized choice to place IEDs at (and only at) vertices v_2, v_3 and v_6 .

The mixed strategies for the Attacker are obtained by choosing the elements in \mathcal{A} at random according to some probability distribution $P = (p_1, \dots, p_I)$; i.e., $\mathbf{a}_i \in \mathcal{A}$ is chosen with probability p_i . Denote the set of all probability distributions over \mathcal{A} as

$$\mathcal{P} = \left\{ P = (p_1, \dots, p_I) : \sum_{i=1}^I p_i = 1 \text{ and } p_i \geq 0 \text{ for } i = 1, \dots, I \right\}.$$

2.2 Payoff Structures

The next step is to describe the payoff structures. In applications, these will be complex, and the Attacker and Defender will have only imperfect knowledge of their opponent's valuations. For adversarial advantage, both sides need to model the payoff functions of their opponents.

The Defender realizes gains associated with the successful arrival of the convoy at T , but costs associated with damage to the convoy. Additionally, there are costs specific to the route—long routes cost more in time and fuel, and some routes may block traffic and antagonize the city population. For the Attacker, the gains are the (probably unknown and hence random) value of inflicting damage, and perhaps political capital. The Attacker's costs are the resources needed for the attack. In practice, neither the Defender nor the Attacker will have precise knowledge of their opponent's payoff function, but previous experience and intelligence information will enable subjective probability assessments.

To formalize this, we make two assumptions:

Assumption 1: It is possible to express the utility function of both the Attacker and Defender in commensurate scalar units; i.e., the values can be monetized.

Assumption 2: The total payoff to the Defender and the Attacker is the sum of their corresponding incremental payoffs across the entire route.

Our sense is that the assumptions are generally realistic. Assumption 1 is quite standard in the literature. Assumption 2 could be inadequate when, say, the Defender's utility function is nonlinear in monetized loss, but it provides a good approximation for most situations.

To illustrate these assumptions, suppose that the Attacker places IEDs at v_1 and v_3 in the road network described in Figure 1. If the Defender follows route r_1 (respectively, route r_2), then Assumption 1 asserts that there exist two real numbers ℓ_{11} and ℓ_{31} (respectively, ℓ_{12} and ℓ_{32}) measuring the payoffs incurred to the Defender at these two locations, and there exist two real numbers u_{11} and u_{31} (respectively, u_{12} and u_{32}) that measure the payoffs for the Attacker at these two locations. In general, we do not need to require any particular relationship between the ℓ_{ij} 's or between the u_{ij} 's. As we indicated at the beginning of this section, the values of the ℓ_{ij} 's and u_{ij} 's incorporate not only payoff-relevant information such as military intelligence beliefs about values of targets but also network-dependent information such as the cost associated with a particular vertex or route. Assumption 2 asserts that the total payoff for the Defender is $\ell_{11} + \ell_{31}$ and the total payoff for the Attacker is $u_{11} + u_{31}$ if the Defender chooses route r_1 and the attacks occur at v_1 and v_3 .

Under these assumptions, one can define the payoff matrices for the Defender and the Attacker.

Definition 1 *The payoff matrix for the Defender corresponding to the graph incidence matrix $\mathbf{G} = [g_{ij}]$ is denoted by $\mathbf{Y} = [Y_{ij}]$ for $i = 1, \dots, N$ and $j = 1, 2, \dots, K$, where Y_{ij} is a numerical value (possibly a random variable) representing the payoff the Defender receives if she chooses route j and there is an attack at vertex i . Similarly, the payoff matrix for the Attacker is $\mathbf{X} = [X_{ij}]$ for $i = 1, \dots, N$ and $j = 1, 2, \dots, K$, where X_{ij} is a numerical value (possibly a random variable) representing the payoff the Attacker receives if he attacks at vertex i when the Defender chooses route j .*

Using this notation, the payoffs to the Defender and the Attacker are $\mathbf{a}'\mathbf{Y}\mathbf{r}$ and $\mathbf{a}'\mathbf{X}\mathbf{r}$, respectively, when the Defender chooses route $\mathbf{r} \in \mathcal{D}$ and the Attacker chooses IED sites $\mathbf{a} \in \mathcal{A}$. If the Defender employs a mixed strategy $Q = (q_1, \dots, q_K)$ and the Attacker uses $P = (p_1, \dots, p_I)$, then the expected payoffs to the Defender and the Attacker are $\sum_{k=1}^K \sum_{i=1}^I p_i q_k \mathbf{a}'_i \mathbf{Y} \mathbf{r}_k$ and $\sum_{k=1}^K \sum_{i=1}^I p_i q_k \mathbf{a}'_i \mathbf{X} \mathbf{r}_k$, respectively.

In the road network example illustrated in Figure 1, if we assume that the payoffs are both vertex-dependent but are route-independent, then the payoff matrices for the Defender

and the Attacker are:

$$\mathbf{Y} = \left(\begin{array}{c|cccc} & r_1 & r_2 & r_3 & r_4 \\ \hline v_1 & L_1 & L_1 & 0 & 0 \\ v_2 & 0 & 0 & L_2 & L_2 \\ v_3 & L_3 & L_3 & L_3 & L_3 \\ v_4 & L_4 & 0 & L_4 & 0 \\ v_5 & L_5 & 0 & L_5 & 0 \\ v_6 & L_6 & L_6 & L_6 & L_6 \end{array} \right), \quad \mathbf{X} = \left(\begin{array}{c|cccc} & r_1 & r_2 & r_3 & r_4 \\ \hline v_1 & U_1 & U_1 & 0 & 0 \\ v_2 & 0 & 0 & U_2 & U_2 \\ v_3 & U_3 & U_3 & U_3 & U_3 \\ v_4 & U_4 & 0 & U_4 & 0 \\ v_5 & U_5 & 0 & U_5 & 0 \\ v_6 & U_6 & U_6 & U_6 & U_6 \end{array} \right). \quad (2)$$

In this particular example, if the Defender chooses route $\mathbf{r}_1 = (1, 0, 0, 0)'$ and the Attacker selects vertices v_2 , v_3 and v_6 as IED sites, so $\mathbf{a} = (0, 1, 1, 0, 0, 1)'$, then the total payoff for the Defender will be $\mathbf{a}'\mathbf{Y}\mathbf{r}_1 = L_3 + L_6$ and the total payoff for the Attacker will be $\mathbf{a}'\mathbf{X}\mathbf{r}_1 = U_3 + U_4$.

3 Game Theory and Adversarial Risk Analysis

Given the previous formulation of the problem, there are two major solution strategies: game-theoretic and decision-analytic. The former seeks an equilibrium in which no agent, acting alone, can improve their outcome; the latter attempts to maximize an agent's expected utility. Both of these can be implemented in multiple ways. We shall use an adversarial risk analysis (ARA) approach, which is a particular kind of decision analysis. But first we give a brief review of some alternatives.

Classical game theory would treat the routing problem as a normal form game. There is a small difficulty in that the entries in the payoff matrices are mutually unknown, and so must be represented by random variables. Harsanyi (1967a,b,c) devised a Bayesian Nash equilibrium solution, in which Nature makes a prior move to randomly determine the "type" of the player, each type having its own distinct belief about the payoff matrix, and where the distribution with which Nature assigns types is *common knowledge* to both players. Alternatively, the *global game* solution concept (cf. Morris and Shin, 1998) imagines that players observe noisy correlated signals about the random payoff matrices. Global game solutions are attractive in part because they are unique, unlike the multiple Nash equilibria

that provide little prescriptive guidance. But these methods, and others, require strong assumptions such as common knowledge that are not tenable in the context of convoy routing.

To avoid this problem, a recent approach (Brown, Carlyle, Salmeron, and Wood, 2006; Brown, Carlyle and Wood, 2008) is to elicit expert opinion (cf. O’Hagan et al., 2006) and use this to develop a joint distribution for the random variables. Then one replaces the entries in the tables by their expected values and computes the Nash equilibrium solution.

But the operations of computing expectations and Nash equilibria do not commute. The best solution to the average game need not be the solution that is, on average, the best. In particular, information encoded in the player’s payoff functions may get lost in the averaging procedure; this is related, in a different context than game theory, to the value of information (cf. Raiffa and Schlaifer, 1961, Part II, chap. 1). To resolve this, one could simulate many payoff tables at random, according to the joint distribution over costs obtained from the experts, find the minimax solution for each, and then choose the action that has the largest average payoff (cf. Banks and Anderson, 2006). A slightly different approach is to simulate outcomes according to the joint distribution on costs, use probabilities on the actions of opponents obtained from some method (e.g., our mirroring argument, or the relative utility of Paté-Cornell and Guikema, 2002), and then choose the action that maximizes the expected payoff (Rios Insua, Rios, and Banks, 2009).

Classical game theory, based on some sort of equilibrium solution concept, has many critics. One issue is that it does a poor job of describing human behavior (Camerer, 2003; also the game-theorist Gintis, 2009). A second issue is that it does not take account of all the information that is available (cf. Arce and Sandler, 2007; Hausken, 2002). In our example, it is quite plausible that the Defender would have highly reliable (but not perfect) intelligence about where the Attacker has sited IEDs; however, that information would not affect the computation of the classical Nash equilibrium, nor is that information employable by any of the alternative equilibrium concepts without substantial additional machinery and assumptions.

Decision analysis arose in response to these concerns (Raiffa, 1982; Kadane and Larkey, 1982). It was and remains controversial; Harsanyi (1982) laid out the main arguments

against it. Raiffa (2002) provides a more recent account of the issues.

In general, decision analysis uses expert judgment to place a subjective distribution over the actions of one's opponent, and then makes the choice that maximizes one's expected value. The advantage is that it provides a natural unification of several types of uncertainty (the random payoffs from a given pair of Attacker-Defender actions, the uncertainty about the utility functions, and the behavioral uncertainty about the opponent's choice). But it suffers the usual criticism of Bayesian subjectivity.

A second disadvantage is that decision analysis is difficult to operationalize. Previous work is largely opaque about the process through which the distribution on the opponent's choices is developed. In the context of infrastructure protection, when the Defender must choose which assets to protect, Paté-Cornell and Guikema (2002) propose taking the subjective distribution to be proportional to the Attacker's utility function evaluated at the possible targets, but this is a "zero-order" analysis in that it precludes consideration of strategy by the Attacker. Banks and Anderson (2006) apply decision analysis to the prospect of a bioterrorist attack with smallpox, but their distributions have only a first-order grounding in strategy. It might apply when the opponent has some meager cunning, but is not adequate in general.

Myerson (1991, p. 114) points up this problem clearly:

"A fundamental difficulty may make the decision-analytic approach impossible to implement, however. To assess his subjective probability distribution over the other players' strategies, player i may feel that he should try to imagine himself in their situations. When he does so, he may realize that the other players cannot determine their optimal strategies until they have assessed their subjective probability distributions over i 's possible strategies. Thus, player i may realize that he cannot predict his opponents' behavior until he understands what an intelligent person would rationally expect him to do, which is, of course, the problem that he started with. This difficulty would force i to abandon the decision analytic approach and instead undertake a game-theoretic approach, in which he tries to solve all players' decision problems simultaneously."

However, instead of following Myerson in defaulting back to game theory, the remainder of this section describes a new way to address the problem. Our approach is similar to a Bayesian implementation of Level- k thinking (Stahl and Wilson, 1995), for a single value of k . It also has similarities to a partially-observable Markov decision process approach developed by Goodie, Doshi and Young (2010) to model different levels of reasoning about an opponent’s thinking.

3.1 The Mirroring Argument and ARA Solution

Decision analysis emphasizes formulating the decision making process from a *single* decision maker’s perspective and models adversarial situations through subjective assessments about the opponents’ behaviors. We describe the decision analysis solution of the routing problem described in Section 2 from the standpoint of the Defender.

The Defender must use her belief about the payoff functions and other available information to develop a subjective prediction for the Attacker’s behavior, i.e., his (mixed) strategy; this is her probability distribution over \mathcal{A} , the set of all pure strategies for the Attacker. Her distribution should reflect the fact that the Attacker is performing a similar analysis regarding the Defender’s strategy, although (since she is not telepathic) she must employ a subjective Bayesian model that describes the Attacker’s thinking. We use the term *mirroring* to refer to the process of modeling an opponent’s decision-making.

To formalize the idea, suppose that the Defender has somehow constructed a probability distribution P over the Attacker’s strategy space \mathcal{A} . Then the Defender finds $\mathbb{E}_P[\tilde{\mathbf{a}}'] = \sum_{i=1}^I p_i \mathbf{a}_i$, where $\tilde{\mathbf{a}}$ is Attacker’s choice of IED sites, which is unknown and thus random to the Defender. And, as an expected utility maximizer, the Defender’s problem is to select the route \mathbf{r}^* (a pure strategy) such that

$$\mathbf{r}^* = \operatorname{argmax}_{\mathbf{r} \in \mathcal{D}} \mathbb{E}_P[\tilde{\mathbf{a}}'] \mathbf{Y} \mathbf{r},$$

where \mathbf{Y} is the (actual) payoff matrix for the Defender.

In order to construct P , we describe the mirroring argument. The following key to the notation is helpful. (As a notational convention, we place a tilde on the top of a deterministic quantity to denote the corresponding random quantity.)

$(\Omega, \mathcal{F}, \mathbb{P})$: the probability space which models *the Defender's* information set (a generic outcome is denoted as ω);

\mathbf{Y} : the Defender's privately known loss matrix (which is unknown to the Attacker);

$\widetilde{\mathbf{X}}(\omega)$: the random matrix on $(\Omega, \mathcal{F}, \mathbb{P})$ which the Defender uses to model the Attacker's payoff matrix;

$\widetilde{\mathbf{Y}}(\omega)$: the random matrix on $(\Omega, \mathcal{F}, \mathbb{P})$ that the Defender uses to describe the Attacker's beliefs about the Defender's payoff matrix;

$\tilde{\mathbf{a}}$: the random vector that *the Defender uses* to model the Attacker's decision (a mixed strategy);

$\tilde{\mathbf{r}}$: the random vector that *the Defender uses* to model the Attacker's belief about the Defender's decision (a mixed strategy).

In the this setting, all uncertainties are described from the perspective of the Defender. We notice that no common knowledge assumption is imposed. In particular, the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is assessed purely based on the Defender's historical data, military intelligence, expert opinions, and so forth, about which the Attacker may have no knowledge. A generic outcome $\omega \in \Omega$ can be interpreted as the *state of the situation* from the Defender's perspective.

First suppose that the Defender has a point-mass prior that the state of the situation is ω . Then the Defender believes that the Attacker's payoff matrix is given by $\widetilde{\mathbf{X}}(\omega)$ (viewed as deterministic for the fixed ω) and that the Attacker forms a probability distribution $Q[\cdot|\omega] \in \mathcal{Q}$ over the Defender's strategy space \mathcal{D} . Thus, the Defender believes that the Attacker will try to find

$$\operatorname{argmax}_{P \in \mathcal{P}} \mathbb{E}_P[\tilde{\mathbf{a}}'] \widetilde{\mathbf{X}}(\omega) \mathbb{E}_{Q[\cdot|\omega]}[\tilde{\mathbf{r}}], \quad (3)$$

which elicits *one* mixed-strategy. Next, allowing ω to have non-unitary support, the probability distribution defined in (3) becomes a random vector on $(\Omega, \mathcal{F}, \mathbb{P})$ taking values in \mathcal{P} . Thus, given all the information $(\Omega, \mathcal{F}, \mathbb{P})$ that she has, the Defender will predict

the Attacker's strategy to be

$$\mathbb{E}_{\mathbb{P}} \left[\operatorname{argmax}_{P \in \mathcal{P}} \mathbb{E}_P[\tilde{\mathbf{a}}'] \widetilde{\mathbf{X}}(\omega) \mathbb{E}_{Q[\cdot|\omega]}[\tilde{\mathbf{r}}] \right] \in \mathcal{P},$$

which is a probability distribution over \mathcal{A} .

On the other hand, knowing the state of the situation $\omega \in \Omega$ and a prediction P of the Attacker's strategy, *the Defender believes that the Attacker believes* that the Defender will try to solve

$$\operatorname{argmax}_{Q \in \mathcal{Q}} \mathbb{E}_P[\tilde{\mathbf{a}}'] \widetilde{\mathbf{Y}}(\omega) \mathbb{E}_Q[\tilde{\mathbf{r}}], \quad \forall \omega \in \Omega,$$

which yields a random vector on $(\Omega, \mathcal{F}, \mathbb{P})$ taking values in \mathcal{Q} .

These predicted probability distributions will be consistent if they satisfy the following definition:

Definition 2 *A probability distribution $P^* \in \mathcal{P}$ and a family of probability distributions $\{Q^*[\cdot|\tilde{\mathbf{Y}}(\omega)] \in \mathcal{Q} : \omega \in \Omega\}$ constitute a mirroring fixed point based on the Defender's information set $(\Omega, \mathcal{F}, \mathbb{P})$ if they simultaneously satisfy*

$$P^* = \mathbb{E}_{\mathbb{P}} \left[\operatorname{argmax}_{P \in \mathcal{P}} \mathbb{E}_P[\tilde{\mathbf{a}}'] \widetilde{\mathbf{X}}(\omega) \mathbb{E}_{Q^*[\cdot|\tilde{\mathbf{Y}}(\omega)]}[\tilde{\mathbf{r}}] \right], \quad (4a)$$

$$Q^*[\cdot|\tilde{\mathbf{Y}}(\omega)] = \operatorname{argmax}_{Q \in \mathcal{Q}} \mathbb{E}_{P^*}[\tilde{\mathbf{a}}'] \widetilde{\mathbf{Y}}(\omega) \mathbb{E}_Q[\tilde{\mathbf{r}}], \quad \omega \in \Omega. \quad (4b)$$

A route \mathbf{r}^{ARA} is said to be a pure-strategy adversarial risk analysis (ARA) solution for the Defender whose actual payoff matrix is \mathbf{Y} if

$$\mathbf{r}^{\text{ARA}} = \operatorname{argmin}_{\mathbf{r} \in \mathcal{Q}} \mathbb{E}_{P^*}[\tilde{\mathbf{a}}'] \mathbf{Y} \mathbf{r}, \quad (5)$$

where P^* is obtained as the fixed point in the mirroring analysis.

Note that mixed-strategy ARA solution is given by $Q^{\text{ARA}} = \operatorname{argmin}_{Q \in \mathcal{Q}} \mathbb{E}_{P^*}[\tilde{\mathbf{a}}'] \mathbf{Y} \mathbb{E}_Q[\tilde{\mathbf{r}}]$. In our case, where the payoff functions are linear in the decision variables, the mixed-strategy solution reduces to the pure-strategy case.

Also note that the mirror fixed point defined above is asymmetric in terms of how the Defender's information $(\Omega, \mathcal{F}, \mathbb{P})$ is used. Knowing the state of the situation, the Defender's strategy should be conditioned on that information. In contrast, the prediction

for the Attacker should not depend on the information that is *only* available to the Defender. This asymmetry reflects the basic starting point of ARA framework: information is not equally available to all decision makers and should be used to favor of the party who owns the information and who is conducting the analysis as well. An unsurprising consequence is that, in general, the party with the best information obtains better outcomes from its own perspective.

The differences and connections between the ARA approach and the traditional game-theoretical approach are discussed in the next subsection. We also prove the existence of the mirroring fixed point in Section 3.3, demonstrating that it is a well-defined solution concept.

3.2 ARA versus Classical Game-Theoretical Solution Concepts

We compare the ARA solution concept with those of traditional game theory. When finding Nash equilibria, the random payoff matrices $\widetilde{\mathbf{X}}(\omega)$ and $\widetilde{\mathbf{Y}}(\omega)$ are modeled as private information (*types*) for the corresponding players. Most importantly, the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is assumed to be *common knowledge* (Harsanyi's common prior assumption; cf. Morris, 1995), which fundamentally limits its applicability (e.g., as in our counter-insurgency example).

Under the common prior assumption, one can define a commonly used solution concept, the *Bayesian Nash Equilibrium* (BNE) (c.f. Myerson (1991)), whose computation is achieved jointly by the Defender and the Attacker based on their common prior $(\Omega, \mathcal{F}, \mathbb{P})$. In contrast, within the ARA framework, the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is constructed based solely on the information available to the Defender. More importantly, $\widetilde{\mathbf{X}}(\omega)$ and $\widetilde{\mathbf{Y}}(\omega)$ are not interpreted as the private information of the players but rather are subjective assessments purely based on the Defender's information $(\Omega, \mathcal{F}, \mathbb{P})$ and, hence, the computation of the mirroring fixed point is conducted *only* by the Defender. The subjective probability assessment by a single party versus the preassumption of common prior among all the parties is the first philosophical difference between the ARA solution and the classical game theory approach.

The second conceptual difference between the ARA solution and the classical game

theory approach lies in the distinction between prediction versus implementation. Game theory solutions are regarded as a prediction *as well as* a decision rule for all the players; while in the ARA framework, the mirroring fixed point serves only as a (subjective) prediction of the opponents' behavior. The decision maker's decision rule is determined subsequent to and separately from the mirroring process.

The ARA approach also has computational advantages. Calculating the BNE usually requires conditional probability elicitation, which is avoided within the ARA framework. Of course, typically the Defender must still perform a difficult elicitation—this is an ubiquitous challenge in BNE and decision analysis.

Mathematically, Definition 2 asserts that the mirroring fixed point is the simultaneous best responses *on average* (from the Defender's perspective) between the Defender and the Attacker as opposed to the simultaneous best responses *of an "averaged game"*, which would correspond to the following system of fixed point equations:

$$\begin{aligned}\arg\max_{P \in \mathcal{P}} \mathbb{E}_P[\tilde{\mathbf{a}}'] \overline{\mathbf{X}} \mathbb{E}_{\bar{Q}}[\tilde{\mathbf{r}}] &= \bar{P}, \\ \arg\max_{Q \in \mathcal{Q}} \mathbb{E}_{\bar{P}}[\tilde{\mathbf{a}}'] \overline{\mathbf{Y}} \mathbb{E}_Q[\tilde{\mathbf{r}}] &= \bar{Q},\end{aligned}$$

where $\overline{\mathbf{X}} = \mathbb{E}_{\mathbb{P}}[\tilde{\mathbf{X}}(\omega)]$ and $\overline{\mathbf{Y}} = \mathbb{E}_{\mathbb{P}}[\tilde{\mathbf{Y}}(\omega)]$ are the *average* payoff matrices. Comparing these with equations (4a) and (4b), we notice that the order of taking expectation $\mathbb{E}_{\mathbb{P}}$ and optimization $\arg\max$ are interchanged. Indeed, $\{\bar{P}, \bar{Q}\}$ is, by definition, a (mixed strategy) Nash Equilibrium for a bimatrix game with payoff matrices $\overline{\mathbf{X}}$ for the row player and $\overline{\mathbf{Y}}$ for the column player. Thus the *averaged game* only uses the “mean” information of $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{Y}}$.

We summarize this discussion in the following proposition:

Proposition 1 *If $\tilde{\mathbf{X}}(\omega) \equiv \mathbf{X}$ and $\tilde{\mathbf{Y}}(\omega) \equiv \mathbf{Y}$ with probability one, then the mirroring fixed point $(P^*, Q^*[\cdot|\tilde{\mathbf{Y}}(\omega)])$ coincides with a (mixed strategy) Nash Equilibrium (\bar{P}, \bar{Q}) for a bimatrix game with payoff matrices \mathbf{X} for the row player and \mathbf{Y} for the column player, i.e. with probability one,*

$$P^* = \bar{P}, \quad Q^*[\cdot|\tilde{\mathbf{Y}}(\omega)] = \bar{Q}.$$

Proposition 1 establishes the connection between the ARA and the traditional game-theoretical approaches in the absence of uncertainty. When there is uncertainty, we want

to identify the relationship between the mirroring fixed point and BNE solutions. Because classical game theory models all the players jointly, as explained above, each of the players will use Bayes rule to update their belief simultaneously based on their private information when they are making decisions. In terms of notation, we use $\mathbb{P}[\cdot|\mathbf{X}]$ to denote the conditional probability $\mathbb{P}[\cdot|\widetilde{\mathbf{X}}(\omega) = \mathbf{X}]$ for any given deterministic payoff matrix \mathbf{X} and a similar convention applies to $\mathbb{P}[\cdot|\mathbf{Y}]$. By definition, the family of distributions

$$\mathbb{P}[\cdot|\widetilde{\mathbf{X}}] := \left\{ P^{**}[\cdot|\mathbf{X}] \in \mathcal{P} : \mathbf{X} \text{ is in the support of } \widetilde{\mathbf{X}}(\omega) \right\}$$

and

$$\mathbb{P}[\cdot|\widetilde{\mathbf{Y}}] := \left\{ Q^{**}[\cdot|\mathbf{Y}] \in \mathcal{Q} : \mathbf{Y} \text{ is in the support of } \widetilde{\mathbf{Y}}(\omega) \right\}$$

constitute a BNE for the game where $(\Omega, \mathcal{F}, \mathbb{P})$ is common knowledge, provided that they satisfy

$$\operatorname{argmax}_{P \in \mathcal{P}} \mathbb{E}_P[\mathbf{a}'\widetilde{\mathbf{X}}] \mathbb{E}_{\mathbb{P}} \left[\mathbb{E}_{Q^{**}[\cdot|\widetilde{\mathbf{Y}}]}[\tilde{\mathbf{r}}] \middle| \widetilde{\mathbf{X}} \right] = P^{**}[\cdot|\widetilde{\mathbf{X}}], \quad (6a)$$

$$\operatorname{argmax}_{Q \in \mathcal{Q}} \mathbb{E}_{\mathbb{P}} \left[\mathbb{E}_{P^{**}[\cdot|\widetilde{\mathbf{X}}]}[\tilde{\mathbf{a}}'] \middle| \widetilde{\mathbf{Y}} \right] \widetilde{\mathbf{Y}} \mathbb{E}_Q[\mathbf{r}] = Q^{**}[\cdot|\widetilde{\mathbf{Y}}]. \quad (6b)$$

Comparing (4a)-(4b) in Definition 2 with the equations (6a)-(6b), we find that the mirroring fixed point and BNE solutions are fundamentally different. In ARA, the Defender solely possesses knowledge of the underlying information set, so, for each realized state of the situation, she can evaluate and optimize the payoff functions with respect to her decision. There is no Bayesian updating based on privately observed information as in BNE. Secondly, in BNE, players have to first average out the payoff functions according to their updated beliefs and hence pick the strategy; whereas, in the mirroring fixed point solution, the Defender will first pick the strategy and hence average out to form the prediction for the Attacker whose strategy cannot depend on the information that is only available to the Defender. Therefore, in general, there is no correspondence between ARA mirroring fixed point and BNE.

To illustrate the differences between the predictions made by the mirroring fixed point and the traditional game theoretical approach, we consider the following toy example where exact calculation can be performed.

Toy Example: The simplest nontrivial road network is one where the Defender (the column player) has two routes r_1 and r_2 to choose from and the Attacker (the row player) has two attacking strategies a_1 and a_2 , where a_i represents an attack at a location along the route r_i ($i = 1, 2$). This example is represented by the following graph:

(Insert Figure 2 here)

Suppose, from the Defender's perspective, that the information set is $\Omega = \{\omega_1, \omega_2, \omega_3\}$ and that

$$\mathbb{P}(\omega_1) = \frac{1}{2}, \quad \mathbb{P}(\omega_2) = \frac{1}{3}, \quad \mathbb{P}(\omega_3) = \frac{1}{6}.$$

The payoff matrices are given by

$$\begin{aligned} \widetilde{\mathbf{X}}(\omega_1) &= \mathbf{X}_1, & \widetilde{\mathbf{X}}(\omega_2) &= \mathbf{X}_2, & \widetilde{\mathbf{X}}(\omega_3) &= \mathbf{X}_2, \\ \widetilde{\mathbf{Y}}(\omega_1) &= \mathbf{Y}_1, & \widetilde{\mathbf{Y}}(\omega_2) &= \mathbf{Y}_1, & \widetilde{\mathbf{Y}}(\omega_3) &= \mathbf{Y}_2, \end{aligned}$$

where we take

$$\mathbf{X}_1 = \begin{bmatrix} 9 & 2 \\ 3 & 6 \end{bmatrix}, \quad \mathbf{X}_2 = \begin{bmatrix} 10 & 1 \\ 3 & 7 \end{bmatrix}, \quad \mathbf{Y}_1 = \begin{bmatrix} 1 & 8 \\ 7 & 4 \end{bmatrix}, \quad \mathbf{Y}_2 = \begin{bmatrix} 0 & 9 \\ 7 & 3 \end{bmatrix}.$$

In this example we can now compute the mixed-strategy solutions under three different solution concepts:

- For the *average game* with averaged payoff matrices,

$$\overline{\mathbf{X}} = \begin{bmatrix} 19/2 & 3/2 \\ 3 & 13/2 \end{bmatrix}, \quad \overline{\mathbf{Y}} = \begin{bmatrix} 5/6 & 49/6 \\ 7 & 23/6 \end{bmatrix},$$

the Nash Equilibrium $\{\bar{P}, \bar{Q}\}$ is given by

$$\bar{P} = \left(\frac{19}{63}, \frac{44}{63} \right), \quad \bar{Q} = \left(\frac{11}{21}, \frac{10}{21} \right).$$

- For the Bayesian game where $(\Omega, \mathcal{F}, \mathbb{P})$ is assumed to be common knowledge for both players whose privately known payoff matrices are $\widetilde{\mathbf{X}}$ and $\widetilde{\mathbf{Y}}$, respectively, the Bayesian Nash Equilibrium $\{P^{**}[\cdot|\mathbf{X}_1], P^{**}[\cdot|\mathbf{X}_2], Q^{**}[\cdot|\mathbf{X}_1], Q^{**}[\cdot|\mathbf{X}_2]\}$ is given by

$$\begin{aligned} P^{**}[\cdot|\mathbf{X}_1] &= \left(\frac{23}{78}, \frac{55}{78} \right), & P^{**}[\cdot|\mathbf{X}_2] &= \left(\frac{4}{13}, \frac{9}{13} \right), \\ Q^{**}[\cdot|\mathbf{Y}_1] &= (0.4, 0.6), & Q^{**}[\cdot|\mathbf{Y}_2] &= \left(\frac{38}{65}, \frac{27}{65} \right). \end{aligned}$$

- For the ARA framework, the mirroring fixed point $\{P^*, Q^*\}$ is given by

$$P^* = (0.3, 0.7), \quad Q^*[\cdot|\omega_1] = Q^*[\cdot|\omega_2] = (0.4, 0.6), Q^*[\cdot|\omega_3] = (1, 0).$$

We note that $\mathbb{E}_{\mathbb{P}} \left[P^{**}[\cdot|\widetilde{\mathbf{X}}] \right] \neq P^*$.

In summary, a fundamental difference between ARA and traditional game theory is whether the assumption of common knowledge is imposed or not. For the motivating problem of convoy routing, it is unrealistic to assume that the Defender and the Attacker have the same common information. But it is entirely reasonable to imagine that the players in such games have relevant probabilistic knowledge about their opponent's actions, derived from military intelligence and previous IED placements. Therefore we recommend the ARA solution concept in problems of this kind.

3.3 Existence of the Mirroring Fixed Point

In order to show that the ARA approach yields a well-defined solution concept, we need to prove the existence of the mirroring fixed point defined by (4a) and (4b). To that end, we assume $\Omega = \{\omega_1, \dots, \omega_L\}$ is finite (this could easily be relaxed) and that

$$\mathbb{P}[\omega_l] = \pi_l, \quad l = 1, 2, \dots, L.$$

For any fixed pair \mathbf{a} and \mathbf{r} , define the quantities

$$\begin{aligned} V_{a,r}^{(l)} &:= \mathbf{a}' \widetilde{\mathbf{X}}(\omega_l) \mathbf{r}, \quad l = 1, 2, \dots, L, \\ W_{a,r}^{(l)} &:= \mathbf{a}' \widetilde{\mathbf{Y}}(\omega_l) \mathbf{r}, \quad l = 1, 2, \dots, L, \\ \overline{W}_{a,r}^{(l,m)} &:= \pi_m W_{a,r}^{(l)}, \quad l, m = 1, 2, \dots, L. \end{aligned}$$

Here, for cosmetic reasons, we drop the boldface on \mathbf{a} and \mathbf{r} when they are used as subscripts.

With this notion, the fixed point equations (4a) and (4b) defining the probability distributions $P^* = (P_a^*)_{a \in \mathcal{A}} \in \mathcal{P}$ and $\left\{ Q^{*(l)} = (Q_r^{*(l)})_{r \in \mathcal{D}} \in \mathcal{Q} : l = 1, \dots, L \right\}$ can now be

written in the following form:

$$P^* = \sum_{l=1}^L \pi_l \left(\operatorname{argmax}_{(P_a) \in \mathcal{P}} \sum_{a \in \mathcal{A}, r \in \mathcal{D}} P_a V_{a,r}^{(l)} Q_r^{*(l)} \right) \quad (7)$$

$$Q^{*(l)} = \operatorname{argmax}_{(Q_r) \in \mathcal{Q}} \sum_{a \in \mathcal{A}, r \in \mathcal{D}} P_a^* W_{a,r}^{(l)} Q_r, \quad l = 1, \dots, L. \quad (8)$$

Lemma 1 Suppose L probability distributions $\{P^{*(l)} \in \mathcal{P} : l = 1, 2, \dots, L\}$ over \mathcal{A} and L probability distributions $\{Q^{*(m)} \in \mathcal{Q} : m = 1, 2, \dots, L\}$ over \mathcal{D} satisfy the following system of fixed-point equations:

$$P^{*(l)} = \operatorname{argmax}_{(P_a) \in \mathcal{P}} \sum_{a \in \mathcal{A}, r \in \mathcal{D}} P_a V_{a,r}^{(l)} Q_r^{*(l)}, \quad l = 1, \dots, L, \quad (9)$$

$$Q^{*(l)} = \operatorname{argmax}_{(Q_r) \in \mathcal{Q}} \sum_{a \in \mathcal{A}, r \in \mathcal{D}} P_a^{*(m)} \overline{W}_{a,r}^{(l,m)} Q_r, \quad l = 1, \dots, L. \quad (10)$$

Then $P^* = \sum_{l=1}^L \pi_l P^{*(l)}$ and $\{Q^{*(l)} : l = 1, \dots, L\}$ satisfy (7) and (8) and hence constitute a mirroring fixed point.

Proof: Multiplying (9) with π_l and summing over $l = 1, \dots, L$, we immediately recover (7).

In (10), we notice, by definition, that for any $Q_r \in \mathcal{Q}$,

$$\sum_{m=1}^L P_a^{*(m)} \overline{W}_{a,r}^{(l,m)} Q_r = \sum_{m=1}^L P_a^{*(m)} \pi_m W_{a,r}^{(l)} Q_r = P_a^* W_{a,r}^{(l)} Q_r,$$

which yields (8). \square

Lemma 2 The fixed point of the system (9) and (10) exists.

Proof: For any l and any fixed $Q = (Q_r)_{r \in \mathcal{D}}$, the term $\sum_{a \in \mathcal{A}, r \in \mathcal{D}} P_a V_{a,r}^{(l)} Q_r$ in (9) is linear (and hence concave) in the decision variables $P = (P_a)_{a \in \mathcal{A}}$. Similarly, for any l and any fixed $\{P^{(m)} = (P_a^{(m)})_{a \in \mathcal{A}} : m = 1, \dots, L\}$, the term $\sum_{a \in \mathcal{A}, r \in \mathcal{D}} P_a^{(m)} \overline{W}_{a,r}^{(l,m)} Q_r$ is also linear (and hence concave) in the decision variables $Q = (Q_r)_{r \in \mathcal{D}}$. Also, all the feasible sets of decision variables are convex compact sets in a finite Euclidean space. Indeed, for any $(\hat{P}_a^{(l)})_{a \in \mathcal{A}} \in \mathcal{P}$ and $(\check{P}_a^{(l)})_{a \in \mathcal{A}} \in \mathcal{P}$, we have

$$\sum_{a \in \mathcal{A}} \left(\alpha_1 \hat{P}_a^{(l)} + \alpha_2 \check{P}_a^{(l)} \right) = \underbrace{\alpha_1 \sum_{a \in \mathcal{A}} \hat{P}_a^{(l)}}_{=1} + \underbrace{\alpha_2 \sum_{a \in \mathcal{A}} \check{P}_a^{(l)}}_{=1} = 1, \quad l = 1, \dots, L$$

for any α_1, α_2 such that $\alpha_1 + \alpha_2 = 1$ and hence $\alpha_1 \hat{P}_a^{(l)} + \alpha_2 \check{P}_a^{(l)} \in \mathcal{P}$. A similar argument applies to \mathcal{Q} . Therefore, by the Nash Fixed-Point Theorem (c.f. Aubin, 1993), the fixed point exists for the system (9) and (10). \square

As a remark, we point out that the fixed point to the system (9) and (10) can mathematically be regarded as the (mixed-strategy) Nash Equilibrium of a game where there are L P -players (choosing $\{P^{(l)} \in \mathcal{P} : l = 1, \dots, L\}$) with payoff functions $\sum_{a \in \mathcal{A}, r \in \mathcal{D}} P_a^{(l)} V_{a,r}^{(l)} Q_r^{(l)}$ and there are L Q -players (choosing $\{Q^{(l)} \in \mathcal{Q} : l = 1, \dots, L\}$) with the payoff functions $\sum_{\substack{a \in \mathcal{A}, r \in \mathcal{D} \\ 1 \leq m \leq L}} P_a^{(m)} \overline{W}_{a,r}^{(l,m)} Q_r^{(l)}$.

Combining Lemma 1 and Lemma 2, we obtain the main result.

Theorem 1 *The mirroring fixed point defined by (4a) and (4b) exists.*

4 Algorithm Structure

To solve an ARA problem, the key step is to compute the probability distribution P^* in the mirroring fixed point (the $Q^*[\cdot|\omega]$ is auxillary, and not of primary interest to the analyst). Given P^* , the rest of the procedure is simply an optimization problem with the mirroring fixed point as an input. In the traditional game-theoretic literature on equilibria for *games with complete information*, there are many algorithms for computing (mixed-strategy) Nash equilibria, but each algorithm usually applies to a particular class of games with some specific structure. For example, three standard algorithms and their domains of success are:

1. The Fictitious Play (FP) scheme proposed by Brown (1949) has been proven to work for zero-sum games (Robinson, 1951), $2 \times n$ games (Berger, 2005) and supermodular games with diminishing returns (Krishna, 1992), among others.
2. The Round Robin (or tatônnement) scheme—a stronger modification of FP based on *best response dynamics*—has been proven to work for supermodular games (Topkis, 1998).

3. The Lemke-Howson scheme based on simplex methods has been shown to solve the class of bimatrix games (Lemke and Howson, 1964).

In general, the search for algorithms that solve for Nash equilibria (i.e., the Nash-type fixed points) is still an open and active research area (Papadimitriou, 2000). In particular, there are many counterexamples and open questions regarding the convergence of the above algorithms beyond the stated families of games. To the best of our knowledge, no iterative algorithms have been proposed to solve for equilibria for *games with incomplete information*.

In our situation, due to the randomness in the payoff matrices in order to model the asymmetric information between the players, we propose an iterative algorithm to compute the probability distribution P^* in the ARA mirroring fixed point, following the spirit of the FP scheme. Our algorithm incorporates statistical simulation in each iteration step, which is an effective and natural approach in the ARA setting. This is due to the nature of such fixed points: first one optimizes, and then one takes expectations. The main difference between our algorithm and the FP scheme is that we just need to keep track of and update one of the empirical distributions, namely that for P^* , and update $Q[\cdot|\omega]$ with a (pure-strategy) best response dynamics (i.e., we don't need to keep track of or update the empirical distribution of the Defender's play). Therefore, our algorithm will only estimate P^* but not $Q^*[\cdot|\omega]$, because only P^* is needed in the third step, when the Defender selects the route that maximizes her expected utility.

This algorithm also demonstrates a computational advantage of the ARA mirroring fixed point as a solution concept. For traditional games of complete information, the FP algorithm must update all of the players' mixed strategies while we just need to keep track of the Attacker's mixed strategy in our ARA algorithm. Compared with the classical BNE solution concept for games of incomplete information, the ARA mirroring fixed point is a computation-friendly solution concept compared with the classical BNE, where Bayesian updating of the players' belief systems is computationally a more challenging task.

The specifics of the iterative algorithm for the ARA mirroring fixed point solution are as follows:

ARA Algorithm: Let the primitives be $(\Omega, \mathcal{F}, \mathbb{P})$, $\mathbf{a} \in \mathcal{A}$ and $\mathbf{r} \in \mathcal{D}$.

1. *Initialize.* The Defender starts with a probability distributions P_0 over \mathcal{A} .

2. *Iterate.* Given distributions P_t , simulate M samples from $(\Omega, \mathcal{F}, \mathbb{P})$.

2.A For each sample $\omega \in \Omega$, compute

$$k_t^*(\omega) = \operatorname{argmax}_{1 \leq k \leq K} \left\{ \left[\mathbb{E}_{P_t}[\tilde{\mathbf{a}}]' \tilde{Y}(\omega) \right]_k \right\},$$

where $[\cdot]_k$ represents the k th element of a vector.

2.B Compute the empirical mean:

$$R_t \leftarrow \frac{1}{M} \sum_{\omega \in \Omega} \operatorname{argmax}_{P \in \mathcal{P}} \mathbb{E}_P[\tilde{\mathbf{a}}]' \tilde{X}(\omega) \mathbf{e}_{k_t^*(\omega)},$$

where \mathbf{e}_k represents the K -dimensional vector with 1 in the k th component and 0 in all other components.

2.C Update:

$$P_{t+1} \leftarrow \frac{t}{t+1} P_t + \frac{1}{t+1} R_t.$$

2.D If $P_t - P_{t+1}$ is sufficiently small with respect to a suitable metric, terminate the iteration and set P^* to be the terminating estimate P_t . Otherwise, repeat all of Step 2.

3. At termination, the Defender chooses the action

$$\mathbf{r}^* = \operatorname{argmin}_{\mathbf{r} \in \mathcal{D}} \mathbb{E}_{P^*}[\tilde{\mathbf{a}}]' \mathbf{Y} \mathbf{r},$$

where \mathbf{Y} is the Defender's true loss matrix.

The purpose of the second step of the algorithm is to calculate the Defender's subjective probability assessment about the Attacker's choices, so that the Defender can optimize the route selection in the third step.

For the Toy Example, the above algorithm converges and correctly finds the theoretical result $P^* = (0.3, 0.7)$. However, a general proof of convergence is unresolved.

In order to investigate the computational complexity of this algorithm, we conducted an experiment based on the Toy Example proposed in Section 3.2 on a Lenovo ThinkPad

X201 personal laptop with a 2.67GHz Intel Core 2 i7 CPU processor and 2.92GB of RAM, running 32-bit Windows 7. The calculations were done using Matlab with a single processor; the run times were estimated using Matlab's Profiler utility.

Specifically, we extend the Toy Example by concatenating J copies of the the simple road network together with the associated payoff matrices. (These J payoff matrices are independently and identically distributed according to the probability distributions specified in the Toy Example. Because of the finiteness of the distribution, we can obtain the exact probability distribution to use in the computation instead of simulating the distribution.) This extension has $2J$ potential locations to attack and 2^J possible routes.

(Insert Figure 3 here.)

We record the CPU compute time for $J = 1, 2, \dots, 8$ and plot the logarithm of that time against the number of copies J in Figure 3(a). As can be seen, the computational time scales exponentially with respect to the network size. Figure 3(b) records the number iterations of the main oracle in the ARA algorithm needed in order to meet the convergence criterion: $\|P_t - P_{t+1}\|_\infty < 0.00001$. As may be seen, when the network size and the number of attacks are of the same order, the algorithm takes longer to converge; also, the number of iterations decreases and then levels off as the size of the network increases. This shows that the computational time increases because of algebraic operations involving matrices and vectors of larger dimension, not because more iterations of the main oracle are needed. In this sense, the computational complexity of the algorithm does not scale with the network size.

We now turn to the road network given by Figure 1. Suppose the payoff matrices for the Defender and the Attacker are given by (2). Furthermore, we assume that their payoff matrices are of zero-sum; i.e., $L_i = -U_i$ for $i = 1, \dots, 6$. Also, suppose

$$U_i \stackrel{i.i.d.}{\sim} \text{Binomial}(10, 0.5), \quad i = 1, \dots, 6.$$

The positive quantity U_i can be interpreted as the gain obtained by an IED attack at vertex i while $L_i < 0$ is the loss to the Defender from an IED attack at vertex i . So the Attacker wants to maximize his cumulative gains and the Defender wants to minimize her cumulative loss. The Defender's action space is simply $\mathcal{D} = \{\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3, \mathbf{r}_4\}$.

- Suppose the Attacker’s budget can afford at most only *one* IED attack, so that the Attacker’s action space is $\mathcal{A} = \{\emptyset, v_1, \dots, v_6\}$. After 3598 iterations of the main oracle in the ARA Algorithm (with convergence criterion again set to $\|P_t - P_{t+1}\|_\infty < 0.00001$), we obtain P^* in Table 1.

(Insert Table 1 here.)

- Now suppose the Attacker’s budget can afford at most *two* IED attacks, so that the Attacker’s action space is $\mathcal{A} = \{\emptyset, v_1, \dots, v_6, \{v_1, v_2\}, \dots, \{v_5, v_6\}\}$. After 3907 iterations of the main oracle in the ARA Algorithm, we obtain P^* in Table 2.

(Insert Table 2 here.)

Our simulation-based modified FP algorithm appears to work well in solving the ARA mirroring fixed-point problem. Although it is well-known that the FP algorithm in traditional game theory only converges to the (mixed-strategy) Nash Equilibrium for some classes of games, we suspect it converges more generally in ARA problems. One reason is its correct performance across a range of numerical experiments. A second reason is that the ARA solution only requires update of the Attacker’s mixed strategy. But the general conditions for convergence of FP remain an open question in the literature.

5 Conclusion

Adversarial routing problems are of longstanding interest in game theory, and the emergence of IEDs as a threat to military convoys has underscored their importance. Classical approaches to such problems make unrealistic assumptions about shared information; in practice, many decision-makers have imperfect knowledge of the utility function and resources of their opponents (and perhaps even of their own). Additionally, there may be relevant history and military intelligence, which, although unreliable and probabilistic, should inform the analysis.

Bayesian methods express such uncertainty through subjective probability elicitation. This leads to the creation of an ARA solution concept, in which one encodes personal

uncertainty through a mirroring argument in order to calculate a subjective distribution over the actions of the opponent. The mirroring argument hinges upon the construction of a model for the analysis that one’s opponent is performing, in which personal probability is used to model all quantities unknown to the decision-maker. This framework allows the decision-maker to find the distribution that expresses her belief about the choices of her opponent. She can then make the decision that maximizes her expected value.

The validity of the ARA approach depends upon the existence of a fixed-point solution, which we show to exist. We also provide a modified Fictitious Play algorithm which, in our experience, successfully converges to the fixed point.

More broadly, the ARA approach is a viable solution concept with several attractive features:

- ARA aims a maximizing expected utility; it is hard to imagine a circumstance in which a decision-maker would not want to do this.
- ARA naturally incorporates soft information, of the kind that is nearly always available in real-world problems.
- ARA integrates different kinds of uncertainty, including uncertainty about random outcomes conditional on the choices that are made, uncertainty about the decision process that produces the opponent’s choices, and uncertainty about how the opponent values different outcomes.
- ARA specifically addresses asymmetric information, which is certainly more realistic than the traditional assumptions, especially the common knowledge formulation.
- ARA explicitly models the decision processes of one’s opponent, which focuses analytic attention on a key psychological aspect of strategic games that is too often overlooked.
- As a minor virtue, ARA is straightforward to compute. The algorithm is fast and perhaps faster than BNE solutions for comparably complex problems.

These are substantial advantages over the other solution concepts reviewed in this paper.

In most practical situations, one's opponent is not a supercomputer armed with perfect knowledge that has been programmed to find Nash equilibria. Instead, one's opponent suffers all the cognitive frailties of human beings, and, if these are properly accounted for, it opens the door to superior play. ARA offers a solution concept that directly models the opponent's decision process. If the decision-maker has an accurate model, then she should be able to achieve better outcomes.

6 Acknowledgment

This work by David Banks was partially supported by DARPA grant W911NF-09-1-0337.

References

- Arce, D. and T. Sandler (2007). Terrorist Signalling and the Value of Intelligence, *British Journal Political Science*, **37**, 573–586.
- Aubin, J. P. (1993). *Optima and Equilibria*, Springer-Verlag, New York, N.Y.
- Banks, D. and S. Anderson (2006). Game Theory and Risk Analysis in the Context of the Smallpox Threat, in *Statistical Methods in Counterterrorism*, A. Wilson, G. Wilson and D. Olwell, eds., Springer, New York, N.Y., pp. 9–22.
- Bayrak, H., and Bailey, M. D. (2008). "Shortest Path Network Interdiction with Asymmetric Information," *Networks*, **52**, 133–140.
- Berger, U. (2005). Fictitious Play in $2 \times n$ Games. *J. Econ. Theory*, **120**, 139–154.
- Brown, G. W. (1949). Some Notes on Computation of Games Solutions. Report P-78. The Rand Corporation.
- Brown, G., M. Carlyle, Salmeron, J. and K. Wood (2006). "Defending critical infrastructure," *Interfaces*, **36**, 530–544.

- Brown, G., W.M. Carlyle, and R. Wood (2008). “Optimizing Department of Homeland Security Defense Investments: Applying Defender-Attacker(-Defender) Optimization to Terror Risk Assessment and Mitigation,” Appendix E, National Academies Press, Washington, D.C.
- Camerer, C. (2003). *Behavioral Game Theory*, Princeton University Press, Princeton, N.J.
- Dimitrov, N., Michalopoulos, D., Morton, D., Nehme, M., Pan, F., Popova, E., Schneider, E. and Thoreson, G. (2009). “Network Deployment of Radiation Detectors with Physics-Based Detection Probability Calculations,” *Annals of Operations Research*, to appear. DOI: 10.1007/s10479-009-0677-2.
- Dresher, M. (1961). *Games of Strategy: Theory and Applications*, RAND/Prentice Hall, Englewood Cliffs, N.J.
- Gintis, H. (2009). *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences*, Princeton University Press, Princeton, N.J.
- Goodie, A., Doshi, P., and Young, D. (2010). “Levels of Theory-of-Mind Reasoning in Competitive Games,” to appear in the *Journal of Behavioral Decision Making*.
- Gutfraind, A., Hagberg, A., and Pan, F. (2009). “Optimal Interdiction of Unreactive Markovian Evaders,” in *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, Willem-Jan van Hoeve and John N. Hooker (eds.), Springer, Heidelberg, pp. 102–116.
- Harsanyi, J. (1967a). “Games with Incomplete Information Played by ‘Bayesian’ Players, Part I: The Basic Model,” *Management Science*, **14**, 159–182.
- Harsanyi, J. (1967b). “Games with Incomplete Information Played by ‘Bayesian’ Players, Part II: Bayesian Equilibrium Points,” *Management Science*, **14**, 320–334.
- Harsanyi, J. (1967c). “Games with Incomplete Information Played by ‘Bayesian’ Players, Part III: The Basic Probability Distribution of the Game”, *Management Science*, **14**, 486–502.

- Harsanyi, J. (1982). "Subjective Probability and the Theory of Games: Comments on Kadane and Larkey's Paper," *Management Science*, **28**, 120–124.
- Hausken, K. (2002). "Probabilistic Risk Analysis and Game Theory", *Risk Analysis*, **22**, 17–27.
- Israeli, E., and Wood, R. K. (2002). "Shortest-Path Network Interdiction," *Networks*, **40**, 97–111.
- Kadane, J. B., and Larkey, P. D. (1982). "Subjective Probability and the Theory of Games", *Management Science*, **28**, 113–120; reply: 124.
- Krishna, V. (1992). Learning in Games with Strategic Complementarities. Mimeo. Harvard University.
- Lemke, C.E., and Howson, J.T. (1964). "Equilibrium Points of Bimatrix Games," *Journal of the Society for Industrial and Applied Mathematics*, **12**, 413–423.
- Morris, S. (1995). "The Common Prior Assumption in Economic Theory," *Economics and Philosophy*, **11**, 227–253.
- Morris, S., and Shin, H. S. (1998). "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks," *Economic Review*, **88**, 587–597.
- Myerson, R. (1991). *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge MA.
- Nasar, S. (1998). *A Beautiful Mind*, Simon & Schuster, New York, N.Y.
- O'Hagan, A., Buck, C., Daneshkhah, A., Eiser, J., Garthwaite, P., Jenkinson, D., Oakley, J., Rakow, T. (2006). *Uncertain Judgements: Eliciting Experts' Probabilities*, Wiley, New York, N.Y.
- Papadimitriou, C. (2001). "Algorithms, Games, and the Internet." In: *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing*. 749–753.

- Paté-Cornell, E., and Guikema, S. (2002). “Probabilistic Modeling of Terrorist Threats: A Systematic Analysis Approach to Setting Priorities Among Countermeasures,” *Military Operations Research*, **7**, 5–23.
- Raiffa, H. (1982). *The Art and Science of Negotiation*, Harvard University Press, Cambridge, MA.
- Raiffa, H. (2002). *Negotiation Analysis*, Harvard University Press, Cambridge, MA.
- Raiffa, H., and Schlaifer, R. (1961). *Applied Statistical Decision Theory*, MIT Press, Cambridge, MA.
- Rios Insua, D., Rios, J., and Banks, D. (2009). “Adversarial Risk Analysis,” *Journal of the American Statistical Association*, **104**, 841–854.
- Roberson, B. (2006). “The Colonel Blotto Game,” *Economic Theory*, **29**, 1–24.
- Robinson, J. (1951). “An Iterative Method of Solving a Game,” *Annals of Mathematics*, **52**(2), 296–301.
- Stahl, D., and Wilson, P. (1995). “On Players’ Model of Other Players,” *Games and Economic Behavior*, **10**, 218–254.
- Topkis, D. (1998). *Supermodularity and Complementarity*, Princeton University Press, Princeton, NJ.
- Washburn, A., and Wood, R. K. (1995). “Two-Person Zero-Sum Games for Network Interdiction,” *Operations Research*, **43**, 243–251.
- Woodruff, D. (2002). *Network Interdiction and Stochastic Integer Programming*, D. Woodruff, editor. Kluwer Academic Publishers, Boston, MA.




Figure 1: An example of a road network.


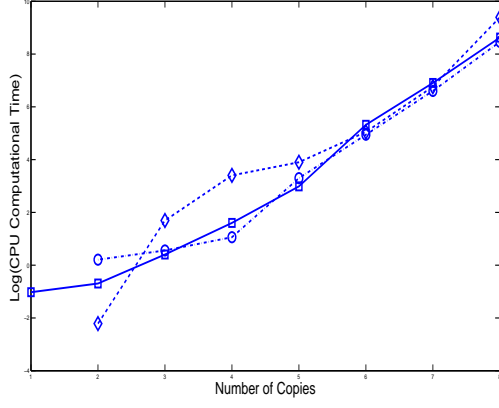



Figure 2: The road network represented by the toy example.



(a) Log(CPU Computational Time).



(b) Number of Iterations.

Figure 3: Computational complexity with respect to the network size $J = 1, \dots, 8$. $-\square-$, $-\circ-$ and $-\diamond-$ represent the cases when the maximum number of locations can be attacked is 1, 2 and 3, respectively.

Table 1: The mirroring fixed point P^* when the Attacker can afford at most one IED.

Attack	\emptyset	v_1	v_2	v_3	v_4	v_5	v_6
P^*	.0000	.1241	.1241	.4275	.0001	.0000	.3243

Table 2: The mirroring fixed point P^* when the Attacker can afford at most two IEDs.

Attack	\emptyset	v_1	v_2	v_3	v_4	v_5	v_6	v_1, v_2	v_1, v_3	v_1, v_4	v_1, v_5
P^*	.0000	.0000	.0000	.0000	.0000	.0000	.0000	.0000	.1699	.0000	.0000
Attack	v_1, v_6	v_2, v_3	v_2, v_4	v_2, v_5	v_2, v_6	v_3, v_4	v_3, v_5	v_3, v_6	v_4, v_5	v_4, v_6	v_5, v_6
P^*	.1334	.1699	.0000	.0000	.1334	.0000	.0000	.3933	.0000	.0000	.0000

Appendix B: Computer Code

Code that defines the example problem.

```
G=[1 1 0 0;  
    0 0 1 1;  
    1 1 1 1;  
    1 0 1 0;  
    1 0 1 0;  
    1 1 1 1];  
  
% number of vertices and routes  
[numV, numR]=size(G);  
  
numSample=1000000;  
epsilon=0.00001;  
  
numAttack=1;  
  
%%% binominal parameter  
baseN=10*ones(1,6);  
baseP=0.5*ones(1,6);  
  
% baseM=5*ones(1,6);  
% baseQ=0.6*ones(1,6);  
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%  
  
A=[];  
A=[A; zeros(1,numV)];  
numA=1;  
for j=1:numAttack  
    attack=nchoosek(1:numV,j);  
    numAtt=nchoosek(numV,j);
```

```

numA=numA+numAtt;
for i=1:numAtt
    Atemp=zeros(1,numV);
    Atemp(1,attack(i,:))=1;
    A=[A; Atemp];
end
end

Pa=(1/numA)*ones(1,numA); %%% we want to find the fixed point
Patemp=-ones(1,numA);%% row vector
t=0;
while (max(abs(Pa-Patemp))>=epsilon)
    Patemp=Pa;
    pay=binornd(repmat(baseN,numSample,1),repmat(baseP,numSample,1));
    c=G'*diag(Patemp*A)*pay';
    [C,K]=min(c);
    c=A*(pay'.*G(:,K));
    [C,I]=max(c);
    Pa=t/(t+1)*Patemp+(1/(1+t))*histc(I,1:numA)/numSample;
    t=t+1;
end

```

Code that defines the Toy Problem.

```

function [X,Y]=randXY(i)

X1=[9 2; 3 6];
X2=[10 1; 3 7];
Y1=[1 8; 7 4];
Y2=[0 9; 7 3];

```

```

switch i
    case 1
        X=X1;
        Y=Y1;
    case 2
        X=X2;
        Y=Y1;
    case 3
        X=X2;
        Y=Y2;
end

```

Code that defines the attack space.

```

%% given the number of vertices in the network and the number of attacks,
%% formulate the Attacker's strategy space and return the number of
%% strategies

```

```

function [A,numA]=AttackSpace(numV,numAttack)

```

```

A=[];
A=[A; zeros(1,numV)];
numA=1;
for j=1:numAttack
    attack=nchoosek(1:numV,j);
    numAtt=nchoosek(numV,j);
    numA=numA+numAtt;
    for i=1:numAtt
        Atemp=zeros(1,numV);

```

```

        Atemp(1,attack(i,:))=1;
        A=[A; Atemp];
end
end

```

Code that implements the fictitious play algorithm to find the mirroring fixed-point solution.

```

numSample=1000000;
epsilon=0.00001;
%%% number of copies that can be concatenated
%numCopy=2;
%numAttack=2;
%%%%%%%%%%%%%%
maxAttack=3;

maxCopy=8;

for numAttack=1:maxAttack
    for numCopy=1:maxCopy
        if numAttack==3 && numCopy==1
            continue
        end
        fprintf('numAttack=%d, numCopy=%d\n',numAttack,numCopy);
        [run_time,P,t]=Routing_Toy(numSample,epsilon,numCopy,numAttack);
        savefile=['numAttack=' int2str(numAttack) 'numCopy=' int2str(numCopy)];
        save(savefile,'run_time', 'P', 't');
    end
end
end

```

```

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
maxAttack=3;

maxCopy=8;

RunTime=zeros(maxCopy,maxAttack);
iterations=zeros(maxCopy,maxAttack);

for numAttack=1:maxAttack
    for numCopy=1:maxCopy
        if numAttack==3 && numCopy==1
            continue
        end
        savefile=['numAttack=' int2str(numAttack) 'numCopy=' int2str(numCopy)];
        x=load(savefile,'run_time','t');
        RunTime(numCopy,numAttack)=x.run_time;
        iterations(numCopy,numAttack)=x.t;
    end
end

figure;
plot(1:maxCopy,RunTime(:,1), '-','marker','s');
hold on;
plot(1:maxCopy,RunTime(:,2), '-.','marker','o');
hold on;
plot(2:maxCopy,RunTime(2:end,3), '--','marker','x');
xlabel('Number of Copies','FontSize',20)
ylabel('CPU Computational Time ','FontSize',20)

figure;

```



```

plot(1:maxCopy,iterations(:,1), '-','marker','s');
hold on;
plot(1:maxCopy,iterations(:,2), '-.','marker','o');
hold on;
plot(2:maxCopy,iterations(2:end,3), '--','marker','x');
xlabel('Number of Copies','FontSize',20)
ylabel('Number of Iterations','FontSize',20)

```

Code that produced the concatenated networks used as input for the study on computation time.

```

%% given the base matrix and numCopy, create the concatenated matrix and
%% return the dimension parameters too

```

```

function [GX,GY,indexSample,V,R]=concatCopy(n,numCopy,numV,numR)

```

```

%[numV,numR]=size(baseG);

```

```

%n= number of samples

```

```

V=numCopy*numV;

```

```

R=numR^numCopy;

```

```

%SampleIndex=zeros(numCopy,n^numCopy);

```

```

GX=zeros(V,R,n^numCopy); % this will be the network node-route incidence matrix

```

```

GY=zeros(V,R,n^numCopy);

```

```

if numCopy>1

```

```

    X = cell(1, numCopy);

```

```

    [X{:}] = ndgrid(1:numR);

```

```

X = X(end : -1 : 1);
y = cat(numCopy+1, X{:});
indexG = reshape(y, numR^numCopy, numCopy);
indexG=indexG';
%n= number of outcomes
X = cell(1, numCopy);
[X{:}] = ndgrid(1:n);
X = X(end : -1 : 1);
y = cat(numCopy+1, X{:});
indexSample = reshape(y, n^numCopy, numCopy);
indexSample=indexSample';

for sample=1:n^numCopy
    for column=1:numR^numCopy
        vx=[];
        vy=[];
        for row=1:numCopy
            [X,Y]=randXY(indexSample(row,sample));
            vx=[vx; X(:,indexG(row,column))];
            vy=[vy; Y(:,indexG(row,column))];
        end
        GX(:,column,sample)=vx;
        GY(:,column,sample)=vy;
    end
end
else
    indexSample=1:n;
    for sample=1:n
        [GX(:, :, sample),GY(:, :, sample)]=randXY(sample);
    end
end

```

```

    end
end

```

Code that generated the table on computation time.

```

function [run_time,P,t]=Routing_Toy(epsilon,numCopy,numAttack)

%% the basic parameters
% numSample=1000000;
% epsilon=0.00001;
% %%% number of copies that can be concatenated
% numCopy=2;
%%% the probability distribution of the states of situations
phi=[1/2 1/3 1/6];
n=length(phi);
%mnrnd(numSample,phi)
%%% max number of attacks that are allowed
% numAttack=2;
%%%%
numV=2*numCopy; %% number of vertices
%%%%
[A,numA]=AttackSpace(numV,numAttack);
%%% A will be the strategy space of attacks
P=(1/numA)*ones(1,numA); % the initial probability
% distribution over the attacker's strategy space,
% we will use this to store the result
Ptemp=-ones(1,numA);
% q=0.5;
% qtemp=-1;

[GX,GY,indexSample,V,R]=concateCopy(n,numCopy,2,2);

```

```

W=zeros(1,n^numCopy);
for i=1:n^numCopy
    Z=1;
    for copy=1:numCopy
        Z=Z*phi(indexSample(copy,i));
    end
    W(i)=Z;
end
t=0;
run_time=cputime;
while (max(abs(P-Ptemp))>=epsilon)
    %sample=mnrnd(numSample,phi,numCopy);
    Ptemp=P;
    P=zeros(1,numA);
    for i=1:n^numCopy
        X=GX(:, :, i);
        Y=GY(:, :, i);
        [C,K]=max(Ptemp*A*Y);
        Z=A*X;
        [C,I]=max(Z(:,K));
        P(I)=P(I)+W(i);
    end
    P=t/(t+1)*Ptemp+(1/(t+1))*P;
    %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
    t=t+1;
end
run_time=cputime-run_time;

```