Two Applications of Adversarial Risk Analysis

by

Shouqiang Wang

Department of Statistical Science
Duke University

Date: ______________________

Approved:

____
David Banks, Supervisor

____
Fan Li

____
Robert Nau

Thesis submitted in partial fulfillment of the requirements for the degree of
Master of Science in the Department of Statistical Science
in the Graduate School of Duke University
2011

Abstract

Two Applications of Adversarial Risk Analysis

by

Shouqiang Wang

Department of Statistical Science
Duke University

Date: ________________________

Approved:

____
David Banks, Supervisor

____
Fan Li

____
Robert Nau

An abstract of a thesis submitted in partial fulfillment of the requirements for
the degree of Master of Science in the Department of Statistical Science
in the Graduate School of Duke University
2011

Abstract

Adversarial risk analysis (ARA) attempts to apply statistical methodology to game-theoretic problems and provides an alternative to the solution concepts in traditional game theory. Specifically, it uses a Bayesian model for the decision-making processes of one's opponents to develop a subjective distribution over their actions, enabling the application of traditional risk analysis to maximize the expected utility. This thesis applies ARA framework to network routing problems in an adversarial contexts and a range of simple Borel gambling games.

Contents

List of Tables

List of Figures

1

Introduction

Researchers at the interface of statistics and operations research have an opportunity to improve strategic decision making by combining tools from statistical risk analysis and game theory. Classical statistical risk analysis has largely overlooked adversarial situations instead, it has focused on the uncertainties associated with natural events and the concomitant costs. And classic game theory has largely overlooked realistic randomness instead, it is has usually assumed that all uncertainties are characterized by a common prior (Harsanyi, 1973). But this formulation makes untenable assumptions about how humans process information and cope with uncertainty.

Adversarial risk analysis (ARA) attempts to apply statistical methodology to game-theoretic problems (Rios Insua et al., 2009). Specifically, it uses a Bayesian model for the decision-making processes of one's opponents to develop a subjective distribution over their actions, enabling the application of traditional risk analysis to maximize the expected utility.

ARA is an alternative to the solution concepts in traditional game theory. Camerer (2003) and Gintis (2009), summarizing on a large body of empirical work, criticize minimax and related solutions (including maximization of expected utility) as unreasonable descriptions of human decision-making. Worse, minimaxity can lead to sub-optimal solutions, in the sense that if one's opponent is not perfectly rational, then solutions based on that premise may be too pessimistic—by mitigating the worst conceivable scenario, one avoids better outcomes that correspond to choices a human opponent might realistically select. Additional problems are that minimax solutions can be difficult to compute, and often require strong and unreasonable assumptions about common knowledge.

Such concerns led Kadane and Larkey (1982) and Raiffa (1982, 2002) to propose decision analysis as an alternative modeling approach and solution concept. It is controversial (Harsanyi, 1982). One concise critique is due to Myerson (1991):

> "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 Nash equilibria, we pursue an ARA solution concept.

The ARA strategy is a subset of a more general approach, sometimes called expected utility decision analysis, in which a player acts so as to maximize their expected utility under some kind of subjective belief about the probabilities of their opponents' choices. The ARA approach develops that subjective belief by using a family of models for the decision processes of an opponent. This family is indexed by how far ahead the player believes that his opponent thinks when making a strategic decision. If the opponent does not strategize (zero-order), then the opponent treats the player as purely random (nature). If the opponent attempts to model the player's thinking, then it is a first-order analysis; if the opponent models the player's model of the opponent's decision-making, then it is a second-order analysis, and so forth. The ARA approach is similar to a Bayesian version of level-k thinking (Stahl and Wilson, 1995). The level-k models are based in part on empirical evidence that players "look-ahead" one or two levels (depending upon the game) when deciding upon their actions.

Other authors have also developed Bayesian or quasi-Bayesian approaches, all of which aim at maximizing the expected utility. Velu and Iyer (2008a,b) perform a probabilistic analysis of the Traveler's Dilemma and the Prisoner's Dilemma. Banks and Anderson (2006) consider subjective distributions over payoff matrices, and examine the concomitant distribution of the minimax solutions. And Paté-Cornell and Guikema (2002) propose a risk analysis procedure in which an adversary strikes a target with probability proportional to the target's utility (this is an example of a zero-order ARA model).

But in all these efforts, the main obstacle to operationalizing decision analysis has been the lack of an explicit mechanism that allows a decision-maker to develop subjective probability distributions which adequately represent an opponent's behavior. ARA resolves that difficulty through a "mirroring" procedure, in which the decision-maker mimics the opponent's analysis, while taking account of the fact that the opponent may be simultaneously performing an analysis of the decision-maker's process. One obtains a probability distribution over the opponent's options, allowing traditional risk analysis to derive the action that maximizes expected utility (Rios Insua et al., 2009).

This thesis applies the ARA framework to address two different problems raised in real life. The first part of the thesis studies 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. The second part of the thesis explores a range of simple gambling games, based on various modifications of the Borel Game (Borel, 1938), also known as *Le Relance*. This game may be viewed as a simplified form of poker. These applications enable comparison with results obtained from other solution strategies, and illustrate the key ideas.

2

Network Routing in Counterterrorism

2.1 Motivation

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 (Nasar, 1998) 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 chapter, 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 the 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 Gutfraind et al. (2009), Israeli and Wood (2002), and arise generally in Woodruff (2002). More distantly, with quite different payoff functions but without the network structure, there is a relationship to Blotto games (Dresher, 1961; Roberson, 2006).

As pointed out in Chapter 1, 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 chapter addresses that difficulty through an *Adversarial Risk Analysis* (ARA) approach 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 argument 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.

The rest of this chapter is organized as follows. Section 2.2 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 2.3. Section 2.4 discusses computational issues. Section 5 provides conclusions.

2.2 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 (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.

2.2.1 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 $\mathscr{G} = (\mathscr{V}, \mathscr{R})$ where $\mathscr{V} = \{v_0 = S, v_1, \ldots, v_N, v_{N+1} = T\}$ is the set of vertices and $\mathscr{R} = \{r_1, \ldots, r_K\}$ is the set of all possible routes connecting S to T (that do not traverse the same edge twice). Obviously, we assume $\mathscr{R} \neq \emptyset$. Algebraically, $\mathscr{G}$ can be represented by an $N \times K$ incidence matrix (denoted by $\boldsymbol{G}$) whose rows are indexed by the elements in $\mathscr{V} \setminus \{S, T\}$ and whose columns are indexed by the elements in $\mathscr{R}$. The (i, j)th element of $\boldsymbol{G}$ is 1 if the ith vertex is visited on the jth route; otherwise, it is 0.

Figure 2.1 provides a toy example to illustrate these definitions.

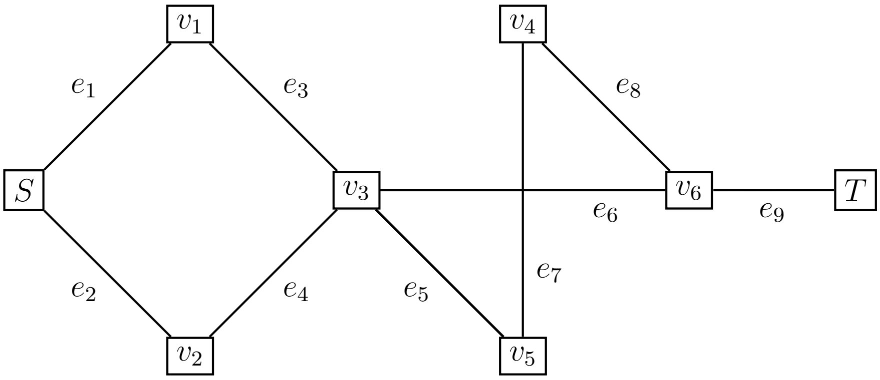

FIGURE 2.1: An example of a road network.

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

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

The corresponding incidence matrix is:

$$\boldsymbol{G} = \left(\begin{array}{c|cccc} & r_1 & r_2 & r_3 & r_4 \\ \hline v_1 & 1 & 1 & 0 & 0 \\ v_2 & 0 & 0 & 1 & 1 \\ v_3 & 1 & 1 & 1 & 1 \\ v_4 & 1 & 0 & 1 & 0 \\ v_5 & 1 & 0 & 1 & 0 \\ v_6 & 1 & 1 & 1 & 1 \end{array} \right). \tag{2.1}$$

The action space for the Defender is identical to $\mathscr{R}$. A *pure strategy* for the Defender is to choose a particular route in $\mathscr{R}$. These strategies can be represented by a (column) vector of K dimension, with $K-1$ zeroes and a one in the kth position to indicate that route k has been chosen. Let $\mathscr{D}$ 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 $\mathscr{D}$.

In the road network example illustrated in Figure 2.1, the set of all pure strategies for the Defender is

$$\mathscr{D} = \{\boldsymbol{r}_1, \boldsymbol{r}_2, \boldsymbol{r}_3, \boldsymbol{r}_4\},$$

where the strategies correspond to vectors in the obvious way:

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

Notice that $\mathscr{D}$ (and hence the element $\boldsymbol{r}_i$) and $\mathscr{R}$ (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 $\mathscr{R}$ at random according to some probability distribution $Q = (q_1, \cdots, q_K)$; i.e., route $\boldsymbol{r}_k \in \mathscr{D}$ is chosen with probability q_k. Denote the set of all probability distributions over $\mathscr{D}$ as

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

Similarly, the set of all pure strategies for the Attacker, which we denote as $\mathscr{A}$, 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 $\mathscr{A}$ reflects this constraint. Each pure strategy $\boldsymbol{a} \in \mathscr{A}$ 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,

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

where $I = |\mathscr{A}|$ is the cardinality of $\mathscr{A}$. Obviously, $\mathscr{A}$ 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 2.1, if the Attacker can choose at most three IED locations, then

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

A pure strategy $\boldsymbol{a} = (0,1,1,0,0,1)'$ in $\mathscr{A}$ 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 $\mathscr{A}$ at random according to some probability distribution $P = (p_1, \cdots, p_I)$; i.e., $\boldsymbol{a}_i \in \mathscr{A}$ is chosen with probability p_i. Denote the set of all probability distributions over $\mathscr{A}$ as

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

2.2.2 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 2.1. 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* $\boldsymbol{G} = [g_{ij}]$ *is denoted by* $\boldsymbol{Y} = [Y_{ij}]$ *for* $i = 1, \ldots, N$ *and* $j = 1, 2, \ldots, K$*, 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* $\boldsymbol{X} = [X_{ij}]$ *for* $i = 1, \ldots, N$ *and* $j = 1, 2, \ldots, K$*, 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 $\boldsymbol{a}'\boldsymbol{Y}\boldsymbol{r}$ and $\boldsymbol{a}'\boldsymbol{X}\boldsymbol{r}$, respectively, when the Defender chooses route $\boldsymbol{r} \in \mathscr{D}$ and the Attacker chooses IED sites $\boldsymbol{a} \in \mathscr{A}$. If the Defender employs a mixed strategy $Q = (q_1, \cdots, q_K)$ and the Attacker uses $P = (p_1, \cdots, p_I)$, then the expected payoffs to the Defender and the Attacker are $\sum\limits_{k=1}^{K} \sum\limits_{i=1}^{I} p_i q_k \boldsymbol{a}_i'\boldsymbol{Y}\boldsymbol{r}_k$ and $\sum\limits_{k=1}^{K} \sum\limits_{i=1}^{I} p_i q_k \boldsymbol{a}_i'\boldsymbol{X}\boldsymbol{r}_k$, respectively.

In the road network example illustrated in Figure 2.1, 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:

$$\boldsymbol{Y} = \left(\begin{array}{c|cccc} & r_1 & r_2 & r_3 & r_4 \\ \hline v_1 & L_1 & L_1 & 0 & 0 \\ v_2 & 0 & 0 & L_2 & L_2 \\ v_3 & L_3 & L_3 & L_3 & L_3 \\ v_4 & L_4 & 0 & L_4 & 0 \\ v_5 & L_5 & 0 & L_5 & 0 \\ v_6 & L_6 & L_6 & L_6 & L_6 \end{array}\right), \tag{2.2a}$$

$$\boldsymbol{X} = \left(\begin{array}{c|cccc} & r_1 & r_2 & r_3 & r_4 \\ \hline v_1 & U_1 & U_1 & 0 & 0 \\ v_2 & 0 & 0 & U_2 & U_2 \\ v_3 & U_3 & U_3 & U_3 & U_3 \\ v_4 & U_4 & 0 & U_4 & 0 \\ v_5 & U_5 & 0 & U_5 & 0 \\ v_6 & U_6 & U_6 & U_6 & U_6 \end{array}\right). \tag{2.2b}$$

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

2.3 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 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 et al., 2006, 2008) is to elicit expert opinion (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 (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 (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 et al., 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; Gintis, 2009). A second issue is that it does not take account of all the information that is available (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.

2.3.1 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 2.2 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 $\mathscr{A}$, 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 $\mathscr{A}$. Then the Defender finds $\mathbb{E}_P[\tilde{\boldsymbol{a}}'] = \sum_{i=1}^{I} p_i \boldsymbol{a}_i$, where $\tilde{\boldsymbol{a}}$ 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 $\boldsymbol{r}^*$ (a pure strategy) such that

$$\boldsymbol{r}^* = \underset{\boldsymbol{r} \in \mathscr{D}}{\operatorname{argmax}} \, \mathbb{E}_P[\tilde{\boldsymbol{a}}'] \boldsymbol{Y} \boldsymbol{r},$$

where $\boldsymbol{Y}$ 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.)

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

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

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

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

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

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

In 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 $(\Omega, \mathscr{F}, \mathbb{P})$ 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 $\widetilde{\boldsymbol{X}}(\omega)$ (viewed as deterministic for the fixed ω) and that the Attacker forms a probability distribution

$Q[\cdot|\omega] \in \mathscr{Q}$ over the Defender's strategy space $\mathscr{D}$. Thus, the Defender believes that the Attacker will try to find

$$\operatorname*{argmax}_{P \in \mathscr{P}} \mathbb{E}_P[\tilde{\boldsymbol{a}}'] \widetilde{\boldsymbol{X}}(\omega) \mathbb{E}_{Q[\cdot|\omega]}[\tilde{\boldsymbol{r}}], \tag{2.3}$$

which elicits *one* mixed-strategy. Next, allowing ω to have non-unitary support , the probability distribution defined in (2.3) becomes a random vector on $(\Omega, \mathscr{F}, \mathbb{P})$ taking values in $\mathscr{P}$. Thus, given all the information $(\Omega, \mathscr{F}, \mathbb{P})$ that she has, the Defender will predict the Attacker's strategy to be

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

which is a probability distribution over $\mathscr{A}$.

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

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

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

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

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

$$P^* = \mathbb{E}_{\mathbb{P}}\left[\operatorname*{argmax}_{P \in \mathscr{P}} \mathbb{E}_P[\tilde{\boldsymbol{a}}'] \widetilde{\boldsymbol{X}}(\omega) \mathbb{E}_{Q^*[\cdot|\widetilde{\mathbf{Y}}(\omega)]}[\tilde{\boldsymbol{r}}]\right], \tag{2.4a}$$

$$Q^*[\cdot|\widetilde{\boldsymbol{Y}}(\omega)] = \operatorname*{argmax}_{Q \in \mathscr{Q}} \mathbb{E}_{P^*}[\tilde{\boldsymbol{a}}'] \widetilde{\boldsymbol{Y}}(\omega) \mathbb{E}_Q[\tilde{\boldsymbol{r}}], \quad \omega \in \Omega. \tag{2.4b}$$

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

$$\boldsymbol{r}^{\mathrm{ARA}} = \operatorname*{argmin}_{\boldsymbol{r} \in \mathscr{D}} \mathbb{E}_{P^*}[\tilde{\boldsymbol{a}}]' \boldsymbol{Y} \boldsymbol{r}, \tag{2.5}$$

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

Note that mixed-strategy ARA solution is given by $Q^{\mathrm{ARA}} = \operatorname*{argmin}_{Q \in \mathscr{Q}} \mathbb{E}_{P^*}[\tilde{\boldsymbol{a}}]' \boldsymbol{Y} \mathbb{E}_{Q}[\tilde{\boldsymbol{r}}]$. 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 $(\Omega, \mathscr{F}, \mathbb{P})$ 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 2.3.3, demonstrating that it is a well-defined solution concept.

2.3.2 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 $\widetilde{\boldsymbol{X}}(\omega)$ and $\widetilde{\boldsymbol{Y}}(\omega)$ are modeled as private information (*types*) for the corresponding players. Most importantly, the probability space $(\Omega, \mathscr{F}, \mathbb{P})$ 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) (Myerson, 1991), whose computation is achieved jointly by the Defender and the Attacker based on their common prior $(\Omega, \mathscr{F}, \mathbb{P})$. In contrast, within the ARA framework, the probability space $(\Omega, \mathscr{F}, \mathbb{P})$ is constructed based solely on the information available to the Defender. More importantly, $\widetilde{\boldsymbol{X}}(\omega)$ and $\widetilde{\boldsymbol{Y}}(\omega)$ are not interpreted as the private information of the players but rather are subjective assessments purely based on the Defender's information $(\Omega, \mathscr{F}, \mathbb{P})$ 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:

$$\begin{aligned}\underset{P\in\mathscr{P}}{\operatorname{argmax}}\,\mathbb{E}_P[\tilde{\boldsymbol{a}}']\overline{\boldsymbol{X}}\,\mathbb{E}_{\bar{Q}}[\tilde{\boldsymbol{r}}] &= \bar{P},\\ \underset{Q\in\mathscr{Q}}{\operatorname{argmax}}\,\mathbb{E}_{\bar{P}}[\tilde{\boldsymbol{a}}']\overline{\boldsymbol{Y}}\,\mathbb{E}_Q[\tilde{\boldsymbol{r}}] &= \bar{Q},\end{aligned}$$

where $\overline{\boldsymbol{X}} = \mathbb{E}_{\mathbb{P}}\left[\widetilde{\boldsymbol{X}}(\omega)\right]$ and $\overline{\boldsymbol{Y}} = \mathbb{E}_{\mathbb{P}}\left[\widetilde{\boldsymbol{Y}}(\omega)\right]$ are the *average* payoff matrices. Comparing these with equations (2.4a) and (2.4b), we notice that the order of taking expectation $\mathbb{E}_{\mathbb{P}}$ and optimization argmax are interchanged. Indeed, $\left\{\bar{P},\bar{Q}\right\}$ is, by definition, a (mixed strategy) Nash Equilibrium for a bimatrix game with payoff matrices $\overline{\boldsymbol{X}}$ for the row player and $\overline{\boldsymbol{Y}}$ for the column player. Thus the *averaged game* only uses the "mean" information of $\widetilde{\boldsymbol{X}}$ and $\widetilde{\boldsymbol{Y}}$.

We summarize this discussion in the following proposition:

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

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

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 $\mathbb{P}\left[\cdot|\boldsymbol{X}\right]$ to denote the conditional probability $\mathbb{P}\left[\cdot|\widetilde{\boldsymbol{X}}(\omega) = \boldsymbol{X}\right]$ for any given deterministic payoff matrix $\boldsymbol{X}$ and a similar convention applies to $\mathbb{P}\left[\cdot|\boldsymbol{Y}\right]$. By definition, the family of distributions

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

and

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

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

$$\operatorname*{argmax}_{P\in\mathscr{P}} \mathbb{E}_P[\boldsymbol{a}']\widetilde{\boldsymbol{X}}\mathbb{E}_{\mathbb{P}}\left[\mathbb{E}_{Q^{**}[\cdot|\widetilde{\mathbf{Y}}]}[\tilde{\boldsymbol{r}}]\middle|\widetilde{\boldsymbol{X}}\right] = P^{**}[\cdot|\widetilde{\boldsymbol{X}}], \tag{2.6a}$$

$$\operatorname*{argmax}_{Q\in\mathscr{Q}} \mathbb{E}_{\mathbb{P}}\left[\mathbb{E}_{P^{**}[\cdot|\widetilde{\mathbf{X}}]}[\tilde{\boldsymbol{a}}']\middle|\widetilde{\boldsymbol{Y}}\right]\widetilde{\boldsymbol{Y}}\mathbb{E}_Q[\boldsymbol{r}] = Q^{**}[\cdot|\widetilde{\boldsymbol{Y}}]. \tag{2.6b}$$

Comparing (2.4a)-(2.4b) in Definition 2 with the equations (2.6a)-(2.6b), 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. ARA approach involves 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 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 $(i = 1, 2)$. Graphically, it can be represented by the following graph:

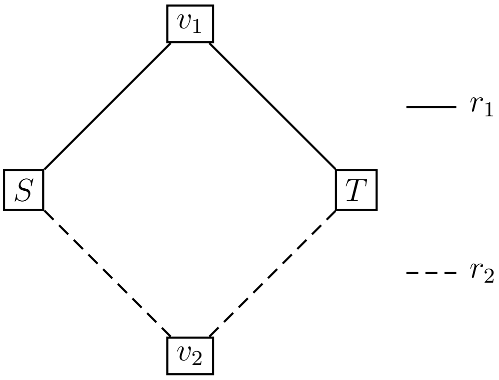

FIGURE 2.2: The road network represented by the toy example.

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

$$\widetilde{\boldsymbol{X}}(\omega_1) = \boldsymbol{X}_1, \quad \widetilde{\boldsymbol{X}}(\omega_2) = \boldsymbol{X}_2, \quad \widetilde{\boldsymbol{X}}(\omega_3) = \boldsymbol{X}_2,$$

$$\widetilde{\boldsymbol{Y}}(\omega_1) = \boldsymbol{Y}_1, \quad \widetilde{\boldsymbol{Y}}(\omega_2) = \boldsymbol{Y}_1, \quad \widetilde{\boldsymbol{Y}}(\omega_3) = \boldsymbol{Y}_2,$$

where we take

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

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

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

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

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

$$P^{**}[\cdot|\boldsymbol{X}_1] = \left(\frac{23}{78}, \frac{55}{78}\right), \quad P^{**}[\cdot|\boldsymbol{X}_2] = \left(\frac{4}{13}, \frac{9}{13}\right),$$

$$Q^{**}[\cdot|\boldsymbol{Y}_1] = (0.4, 0.6), \quad Q^{**}[\cdot|\boldsymbol{Y}_2] = \left(\frac{38}{65}, \frac{27}{65}\right).$$

- 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 $\mathbb{E}_{\mathbb{P}}\left[P^{**}[\cdot|\widetilde{\boldsymbol{X}}]\right] \neq P^*$.

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.

2.3.3 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 (2.4a) and (2.4b). To that end, we assume $\Omega = \{\omega_1, \ldots, \omega_L\}$ is finite (this could easily be relaxed) and that

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

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

$$V_{a,r}^{(l)} := \boldsymbol{a}'\widetilde{\boldsymbol{X}}(\omega_l)\boldsymbol{r}, \quad l = 1, 2, \ldots, L,$$

$$W_{a,r}^{(l)} := \boldsymbol{a}'\widetilde{\boldsymbol{Y}}(\omega_l)\boldsymbol{r}, \quad l = 1, 2, \ldots, L,$$

$$\overline{W}_{a,r}^{(l,m)} := \pi_m W_{a,r}^{(l)}, \quad l, m = 1, 2, \ldots, L.$$

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

With this notion, the fixed point equations (2.4a) and (2.4b) defining the probability distributions $P^* = (P_a^*)_{a\in\mathscr{A}} \in \mathscr{P}$ and $\left\{Q^{*(l)} = (Q_r^{*(l)})_{r\in\mathscr{D}} \in \mathscr{Q} : l = 1, \ldots, L\right\}$ can now be written in the following form:

$$P^* = \sum_{l=1}^{L} \pi_l \left(\underset{(P_a)\in\mathscr{P}}{\operatorname{argmax}} \sum_{a\in\mathscr{A}, r\in\mathscr{D}} P_a V_{a,r}^{(l)} Q_r^{*(l)} \right) \tag{2.7}$$

$$Q^{*(l)} = \underset{(Q_r)\in\mathscr{Q}}{\operatorname{argmax}} \sum_{a\in\mathscr{A}, r\in\mathscr{D}} P_a^* W_{a,r}^{(l)} Q_r, \quad l = 1, \ldots, L. \tag{2.8}$$

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

$$P^{*(l)} = \underset{(P_a)\in\mathscr{P}}{\operatorname{argmax}} \sum_{a\in\mathscr{A}, r\in\mathscr{D}} P_a V_{a,r}^{(l)} Q_r^{*(l)}, \quad l = 1, \ldots, L, \tag{2.9}$$

$$Q^{*(l)} = \underset{(Q_r)\in\mathscr{Q}}{\operatorname{argmax}} \sum_{\substack{a\in\mathscr{A}, r\in\mathscr{D} \\ 1\le m\le L}} P_a^{*(m)} \overline{W}_{a,r}^{(l,m)} Q_r, \quad l = 1, \ldots, L. \tag{2.10}$$

Then $P^* = \sum_{l=1}^{L} \pi_l P^{*(l)}$ and $\left\{Q^{*(l)} : l = 1, \ldots, L\right\}$ satisfy (2.7) *and* (2.8) *and hence constitute a mirroring fixed point.*

Proof: Multiplying (2.9) with π_l and summing over $l = 1, \ldots, L$, we immediately recover (2.7). In (2.10), we notice, by definition, that for any $Q_r \in \mathscr{Q}$,

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

which yields (2.8). $\square$

Lemma 3. *The fixed point of the system* (2.9) *and* (2.10) *exists.*

Proof: For any l and any fixed $Q = (Q_r)_{r\in\mathscr{D}}$, the term $\sum\limits_{a\in\mathscr{A},r\in\mathscr{D}} P_a V^{(l)}_{a,r} Q_r$ in (2.9) is linear (and hence concave) in the decision variables $P = (P_a)_{a\in\mathscr{A}}$. Similarly, for any l and any fixed $\left\{P^{(m)} = (P^{(m)}_a)_{a\in\mathscr{A}} : m = 1,\ldots,L\right\}$, the term $\sum\limits_{\substack{a\in\mathscr{A},r\in\mathscr{D}\\1\le m\le L}} P^{(m)}_a \overline{W}^{(l,m)}_{a,r} Q_r$ is also linear (and hence concave) in the decision variables $Q = (Q_r)_{r\in\mathscr{D}}$. Also, all the feasible sets of decision variables are convex compact sets in a finite Euclidean space. Indeed, for any $(\hat{P}^{(l)}_a)_{a\in\mathscr{A}} \in \mathscr{P}$ and $(\check{P}^{(l)}_a)_{a\in\mathscr{A}} \in \mathscr{P}$, we have

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

for any α_1, α_2 such that $\alpha_1 + \alpha_2 = 1$ and hence $\alpha_1 \hat{P}^{(l)}_a + \alpha_2 \check{P}^{(l)}_a \in \mathscr{P}$. A similar argument applies to $\mathscr{Q}$. Therefore, by the Nash Fixed-Point Theorem (c.f. Aubin, 1993), the fixed point exists for the system (2.9) and (2.10). $\square$

As a remark, we point out that the fixed point to the system (2.9) and (2.10) can mathematically be regarded as the (mixed-strategy) Nash Equilibrium of a game where there are L P-players (choosing $\{P^{(l)} \in \mathscr{P} : l = 1,\ldots,L\}$) with payoff functions $\sum\limits_{a\in\mathscr{A},r\in\mathscr{D}} P^{(l)}_a V^{(l)}_{a,r} Q^{(l)}_r$ and there are L Q-players (choosing $\{Q^{(l)} \in \mathscr{Q} : l = 1,\ldots,L\}$) with the payoff functions $\sum\limits_{\substack{a\in\mathscr{A},r\in\mathscr{D}\\1\le m\le L}} P^{(m)}_a \overline{W}^{(l,m)}_{a,r} Q^{(l)}_r$.

Combining Lemma 2 and Lemma 3, we obtain the main result.

Theorem 4. *The mirroring fixed point defined by* (2.4a) *and* (2.4b) *exists.*

2.4 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 and 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 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 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 computationally-friendly solution concept compared with the classical BNE, where Bayesian updating of the players' belief systems is computationally a daunting task.

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

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

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

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

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

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

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

2.B Compute the empirical mean:

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

where $\boldsymbol{e}_k$ represents the K-dimensional vector with 1 in the kth 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

$$\boldsymbol{r}^* = \underset{\boldsymbol{r} \in \mathscr{D}}{\text{argmin}}\, \mathbb{E}_{P^*}[\tilde{\boldsymbol{a}}]' \boldsymbol{Y} \boldsymbol{r},$$

where $\boldsymbol{Y}$ 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.

More specifically, we sequentially concatenate 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.) In the Toy Example there are $2J$ potential locations to attack and 2^J possible routes.

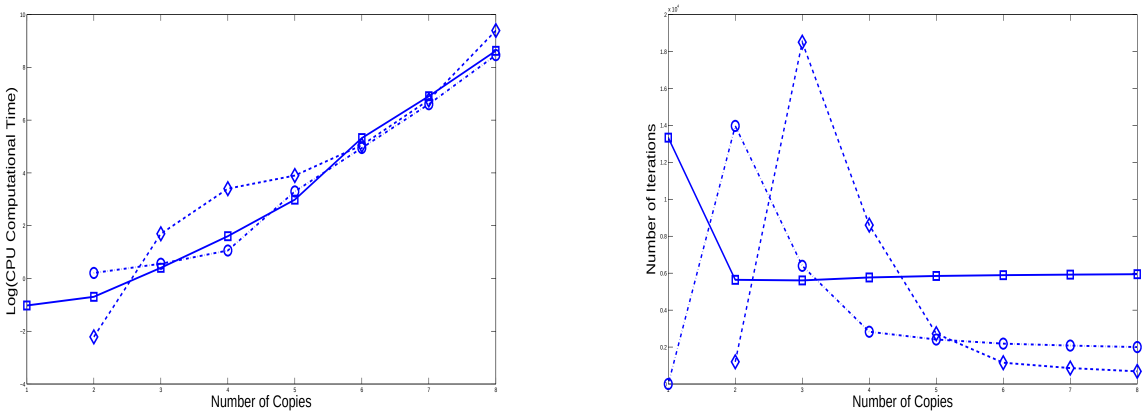

(a) Log(CPU Computational Time).

(b) Number of Iterations.

FIGURE 2.3: Computational complexity with respect to the network size $J = 1, \ldots, 8$. $-\square-$, $-\circ-$ and $-\diamond-$ represent the cases when the maximum number of locations can be attacked is 1, 2 and 3, respectively.

We record the CPU compute time for $J = 1, 2, \ldots, 8$ and plot the logarithm of that time against the number of copies J in Figure 2.3(a). As can be seen, the computational time scales exponentially with respect to the network size. Figure 2.3(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 range, it takes longer to converge while 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 algorithm does not scale with the network size.

We now turn to the road network given by Figure 2.1. Suppose the payoff matrix for the Defender is given by (2.2a) and that of the Attacker is given by (2.2b). Furthermore, we assume that their payoff matrices are of zero-sum; i.e., $L_i = -U_i$ for $i = 1, \ldots, 6$. Also, suppose

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

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 $\mathscr{D} = \{\boldsymbol{r}_1, \boldsymbol{r}_2, \boldsymbol{r}_3, \boldsymbol{r}_4\}$.

- Suppose the Attacker's budget can afford at most only *one* IED attack, so that the Attacker's action space is $\mathscr{A} = \{\emptyset, v_1, \cdots, v_6\}$. 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 2.1.

Table 2.1: 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

- Now suppose the Attacker's budget can afford at most *two* IED attacks, so that the Attacker's action space is $\mathscr{A} = \{\emptyset, v_1, \cdots, v_6, \{v_1, v_2\}, \cdots, \{v_5, v_6\}\}$. After 3907 iterations of the main oracle in the ARA Algorithm, we obtain P^* in Table 2.2.

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

Table 2.2: 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

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.

2.5 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.

3

Analysis of Borel Games

3.1 Motivation

Consider the following two-person game:

> **Borel Game (Borel, 1938):** Both Bart and Lisa must ante a single unit to play. Each picks a private number, independently, from the uniform distribution on [0, 1]. Then Bart either folds or bets a fixed amount b. If Bart bets, then Lisa either folds or calls an amount b. If Lisa calls, then the player with the larger number wins the pot.

The Borel Game's (also known as *Le Relance*) assumption of a unit ante is not restrictive; one can rescale the solution to other cases by maintaining the ratio of the ante to the bet.

The Borel Game is a classical problem in game theory. von Neumann and Morgenstern (1944) extended it to allow players to check; Bellman and Blackwell (1949) examined bluffing strategy; and Karlin and Restrepo (1957) generalized the problem to multiple players and/or multiple rounds and/or multiple increments of bet size. A recent review of the area, with modern proofs and language, is given in Ferguson and Ferguson (2003). Minimax solutions have been found for a wide range of variant games, which include multiple players, betting increments, and checking.

We apply the ARA method to several famous gambling problems, based on various modifications of the Borel Game. This game may be viewed as a simplified form of poker. These applications enable comparison with results obtained from other solution strategies, and illustrate the key ideas. In this chapter, we will derive the ARA solution from Bart's perspective. Section 3.2 treats the basic Borel Game with fixed bets. Section 3.3 extends this simple model to the version that allows continuous bets. Section 3.4 summarizes the results.

3.2 Fixed Bets

We briefly treat a naive zero-level version, in that Bart assumes that Lisa does not attempt to analyze the problem from Bart's perspective. We then move to more sophisticated versions, in which Bart examines Lisa's decision-making in greater depth. We note that Bart reasons as a Bayesian, and in general he assumes that Lisa is Bayesian too (but that is not essential—all he requires is a model for her decision-making). Also, it is convenient (but not necessary) to assume that Bart has a utility function that is linear in money, and thinks that Lisa does too, i.e. both players are risk neutral.

For all versions of ARA, Bart has observed that his draw is $X = x$; Lisa's draw $Y = y$ is unknown to him, although he knows its distribution. Let V_x be the amount won by Bart after drawing $X = x$. Table 3.1 shows the possible situations, depending on how each player decides to bet and the values of X and Y.

Table 3.1: Amount won by Bart, with two players.

V_x	Bart's Decision	Lisa's Decision	Bart's Win Condition
-1	fold		
1	bet	fold	
1+b	bet	call	$x > Y$
-(1+b)	bet	call	$x < Y$

Using Table 3.1 and neglecting ties (an event with measure zero), the expected amount won by Bart, conditional on his draw, is

$$\begin{aligned} V_x &= -\mathbb{P}[\text{ Bart folds }] + \mathbb{P}[\text{ Bart bets and Lisa folds }] \\ &\quad +(1+b)\mathbb{P}[\text{ Lisa calls and loses }] - (1+b)\mathbb{P}[\text{ Lisa calls and wins }]. \end{aligned}$$

Bart wants to maximize his expected utility, and so he seeks the play that maximizes V_x.

Bart needs to find a "bluffing function" $g(x)$; given $X = x$, he bets with probability $g(x)$. Then

$$\begin{aligned} V_x &= -[1-g(x)] + g(x)\mathbb{P}[\text{ Lisa folds } \mid \text{ Bart bets }] \\ &\quad +(1+b)g(x)\mathbb{P}[\text{ Bart wins } \mid \text{ Lisa calls }]\mathbb{P}[\text{ Lisa calls } \mid \text{ Bart bets }] \\ &\quad -(1+b)g(x)\mathbb{P}[\text{ Bart loses } \mid \text{ Lisa calls }]\mathbb{P}[\text{ Lisa calls } \mid \text{ Bart bets }] \qquad (3.1) \end{aligned}$$

In order to derive his optimal play, Bart needs to understand both $\mathbb{P}[\text{ Lisa calls } \mid \text{ Bart bets }]$ and $\mathbb{P}[\text{ Bart wins } \mid \text{ Lisa calls }]$.

3.2.1 Zero-Order ARA

In the simplest form of ARA, which is essentially the form proposed by Raiffa (1982) and Kadane and Larkey (1982), Bart does not explicitly try to model Lisa's thinking. Instead, he merely declares the subjective distribution $\pi(c)$ that he has over the value c at which Lisa will bet.[1] Bart's subjective probability assessment on c takes no formal account of the fact that Lisa's rule might be informed by Bart's decision to bet, although this could be implicit in his personal elicitation.

Given $\pi(c)$, with mean μ, Bart's guess about the probability that Lisa will bet is

$$\mathbb{P}[\text{ Lisa bets }] = \int_0^1 \int_c^1 \pi(c)dx\,dc = \int (1-c)\pi(c)\,dc = 1-\mu.$$

---

[1] In this game, there is no advantage to Lisa from bluffing (Borel, 1938); she should have a fixed rule such that if $Y > c$, she bets, and otherwise she folds.

Also, routine calculation shows

$$\begin{aligned} \mathbb{P}[\text{ Bart wins with } x \mid \text{Lisa calls }] &= \int_0^x \int_c^x \pi(c)\, dx dc \\ &= \int_0^x (x-c)\pi(c)\, dc. \end{aligned}$$

Suppose that, given $X = x$, Bart bets with probability $g(x)$. Set $\gamma = (1+b)(1-\mu)$. Then the expected value of the game (for Bart) is

$$\begin{aligned} V_x &= -(1-g(x)) + g(x)\mu + g(x)\gamma \int_0^x (x-c)\pi(c)\, dc - g(x)\gamma \left[1 - \int_0^x (x-c)\pi(c)\, dc\right] \\ &= -1 + g(x)\left[1 + \mu - (1+b) + 2(1+b)\int_0^x (x-c)\pi(c)\, dc\right]. \end{aligned}$$

This expression has the form $V_x = -1 + g(x)m(x)$, so Bart's optimal strategy is to bet when

$$m(x) = \left[1 + \mu - \gamma + 2\gamma \int_0^x (x-c)\pi(c)\, dc\right] > 0$$

and fold when $m(x) < 0$; he may do as he pleases when $m(x) = 0$. (If $\pi(c)$ is continuous, then equality occurs with probability zero.)

Example 1: Beta priors.

As an illustration, suppose that $\pi(c)$ has a beta distribution with parameters α and β. Then

$$\int_0^x (x-c)\pi(c)\, dc = xI_x(\alpha,\beta) - \frac{\alpha}{\alpha+\beta} I_x(\alpha+1,\beta) \tag{3.2}$$

where

$$I_x(\alpha,\beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \int_0^x t^{\alpha-1}(1-t)^{\beta-1}\, dt$$

is the regularized incomplete beta function. For α and β positive integers,

$$I_x(\alpha,\beta) = \sum_{j=\alpha}^{\alpha+\beta-1} \binom{\alpha+\beta-1}{j} x^j (1-x)^{\alpha+\beta-1-j}.$$

A few reasonable cases are $\alpha = \beta = 1$, in which Bart assumes nothing about the location of the step in Lisa's betting function; $\alpha = \beta = 2$, so Bart assumes that Lisa's step is near $y = 1/2$; and $\alpha = 3, \beta = 1$, so Bart assumes that Lisa is conservative, preferring to call when Y is relatively large.

For the agnostic case with $\alpha = \beta = 1$, equation (3.2) reduces to $\frac{1}{2}x^2$. So Bart will bet if and only if

$$x > \left(1 - \frac{3}{1+b}\right)^{1/2}.$$

When Bart thinks Lisa's step is close to 1 (i.e. $\alpha = 3$ and $\beta = 1$), equation (3.2) is equal to $\frac{1}{4}x^4$ and Bart will bet if and only if

$$x > \left(2 - \frac{14}{1+b}\right)^{1/4}.$$

And when Bart thinks Lisa's betting rule is close to 1/2, i.e. $\alpha = \beta = 2$, then equation (3.2) is equal to $x^3 - \frac{1}{2}x^4$ and Bart will bet if and only if

$$x^3 - \frac{1}{2}x^4 > \frac{1}{2}\left(1 - \frac{3}{1+b}\right),$$

which has no simple solution. However, the left-hand side is increasing in $(0, 1)$, and so Bart's bluffing function is still a step function with its step in $(0, 1)$.

Figure 3.1 shows how the point at which Bart bets depends upon the value b of the bet, under the three different priors. The monotonicity in b implies that as the bet size increases, Bart should become more conservative.

3.2.2 First-Order ARA

The first-order ARA is more interesting. Here Bart models Lisa's reasoning in order to develop his belief about her decision rule. Bart accomplishes this through a "mirroring"

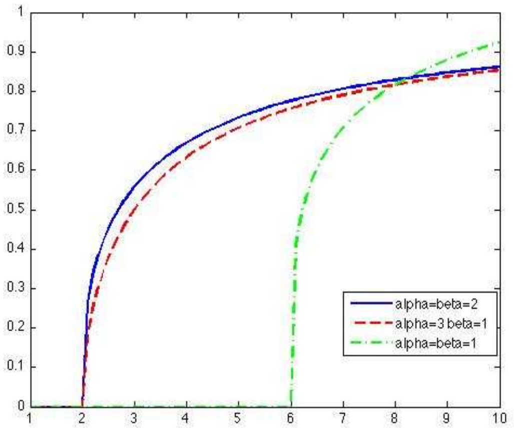

FIGURE 3.1: The threshold value beyond which Bart should bet as a function of b.

argument, in which he does the analysis he expects her to make, using subjective distributions to describe the quantities he does not know.

Mirroring can have different levels of sophistication; there is an analogy with the Level-k approach in game theory (Stahl and Wilson, 1995). In contrast, the zero-order ARA had no explicit adversarial component; it was simply Bayesian risk analysis, and strategy was not relevant.

With first-order ARA, Bart knows that Lisa's opinion about his value of X is updated by the knowledge that Bart decided to bet. Further, suppose Bart has a subjective opinion that Lisa thinks that his bluffing function is $\tilde{g}(x)$. In that case, Bart believes that a Bayesian Lisa would calculate her conditional density of X, given that Bart decided to bet, as $\tilde{f}(x) = \tilde{g}(x)/\int \tilde{g}(w)\,dw$. Note that if $\tilde{g}$ is a step function (i.e., Lisa believes that Bart does not bet if x is less than some value x_0, but he always bets if it is greater), then her posterior distribution on X is truncated below x_0 and the weight is reallocated proportionally to values above x_0.

From this perspective, Bart believes that Lisa will calculate her probability of winning, conditional on Bart's bet, as

$$\mathbb{P}[X \leq y \,|\, \text{Bart bets}\,] = \tilde{F}(y) = \int_0^y \tilde{f}(z)\,dz,$$

where y is unknown to Bart. And thus Bart believes that Lisa will call if the expected value of her return V_y from betting b is greater than the loss of one unit that results from folding; i.e., Lisa would call if

$$(1+b)\tilde{F}(y) - (1+b)[1 - \tilde{F}(y)] \geq -1.$$

Solving this shows that Bart believes Lisa will call if her $\tilde{F}(y) > b/[2(1+b)]$, fold if $\tilde{F}(y) < b/[2(1+b)]$, and do as she pleases when $\tilde{F}(y) = b/[2(1+b)]$.

Let $\tilde{y} = \inf\{y : \tilde{F}(y) > b/[2(1+b])\}$. The probability that Lisa has drawn $Y > \tilde{y}$ is $1 - \tilde{y}$ and this is Bart's best guess about the probability that she calls. And if Lisa has drawn a number larger than $\tilde{y}$, then $\mathbb{P}[\text{ Bart wins } | \text{ Lisa calls }] = [(x-\tilde{y})/(1-\tilde{y})]^+$, where $[\,\cdot\,]^+$ takes the value zero if the argument is negative.

So Bart believes the expected value of his game, given $X = x$, is:

$$V_x = -[1 - g(x)] + g(x)\tilde{y} + (1+b)g(x)[x-\tilde{y}]^+ - (1+b)g(x)(1 - \tilde{y} - [x - \tilde{y}]^+)$$

or

$$V_x = \begin{cases} -1 + g(x)(2\tilde{y} + b\tilde{y} - b) & \text{if } x \leq \tilde{y} \\ -1 + g(x)(2x + 2bx - b\tilde{y} - b) & \text{if } x > \tilde{y}. \end{cases} \tag{3.3}$$

Bart should choose $g(x)$ to maximize V_x.

When $x \leq \tilde{y}$, then Bart bets if $\tilde{y} > b/(b+2)$, he folds if $\tilde{y} < b/(b+2)$ and he may do as he pleases when $\tilde{y} = b/(b+2)$. When $x > \tilde{y}$, then Bart bets when $x > \tilde{x} = [b(1+\tilde{y})]/[2(1+b)]$, he folds when $x < \tilde{x}$, and he may do as he pleases when $x = \tilde{x}$. So there are three cases, depending on the value of $\tilde{y}$.

Case I: Bart Believes that Lisa Plays Minimax

The traditional minimax solution has $\tilde{y} = b/(b+2)$: Bart should bet if $x > \tilde{y}$, and he should bet with probability $2/(b+2)$ when $x \leq \tilde{y}$ (von Neumann and Morgenstern, 1944). The value of the game (to Bart) is $V = -b^2/(b+2)^2$; so he is disadvantaged by the sequence of play.

In contrast, our ARA analysis shows that when Lisa uses the minimax threshold $\tilde{y} = b/(b+2)$, then Bart should bet if $x > \tilde{x}$, where simple algebra shows that $\tilde{x} = b/(b+2)$, as usual. But *he may bet or not, as he pleases, when* $x \leq \tilde{x}$. This is slightly different from the minimax solution.

The reason for the discrepancy is that if Lisa knows that Bart's bluffing function does not bet with probability $2/(b+2)$ when $x \leq b/(b+2)$, then she can improve her expected value for the game by changing the threshold at which she calls. To see this, suppose Bart uses the ARA solution strategy and chooses to bet if and only if $x > b/(b+2)$. Then Bart's bluffing function is

$$g(x) = \begin{cases} 0 & \text{if } 0 \leq x \leq b/(b+2) \\ 1 & \text{if } b/(b+2) < x \leq 1. \end{cases}$$

If Lisa knew Bart's rule, or could derive it, then Lisa would calculate

$$\tilde{F}(y) = \begin{cases} 0 & \text{if } 0 \leq y \leq b/(b+2) \\ \frac{y - \frac{b}{b+2}}{1 - \frac{b}{b+2}} & \text{if } b/(b+2) < x \leq 1 \end{cases}$$

and solve to find

$$\tilde{y} = \tilde{F}^{-1}\left(\frac{b}{2(b+1)}\right) = \frac{b}{b+1}.$$

Thus Lisa would not call when $b/(b+2) < y < b/(b+1)$. She could improve the expected value of her game by shifting the threshold at which she calls to a slightly larger number.

However, suppose Bart used the an ARA solution that coincides with the minimax rule. For example, he could play the admissible rule of Ferguson and Ferguson (2003):

$$g(x) = \begin{cases} 0 & \text{if } 0 \leq x \leq [b/(b+2)]^2 \\ 1 & \text{if } [b/(b+2)]^2 < x \leq 1. \end{cases}$$

If Lisa knew that this was Bart's bluffing function, then she would find

$$\tilde{F}(y) = \begin{cases} 0 & \text{if } 0 \leq y \leq [b/(b+2)]^2 \\ \frac{y - \left(\frac{b}{b+2}\right)^2}{1 - \left(\frac{b}{b+2}\right)^2} & \text{if } [b/(b+2)]^2 < x \leq 1 \end{cases}$$

and then calculate

$$\tilde{y} = \tilde{F}^{-1}\left(\frac{b}{2(b+1)}\right) = \frac{b}{b+2}.$$

Thus, when Bart plays an ARA rule that is also minimax rule, Lisa's optimal play is forced to take its step at $b/(b+2)$.

Of course, when Bart believes that Lisa plays her minimax strategy, calling if and only if $y > b/(b+2)$, then the value of the game under any ARA solution is

$$\begin{aligned} \int_0^1 V_x \, dx &= \int_0^{\tilde{y}} -1 + g(x)(2\tilde{y} + b\tilde{y} - b) \, dx + \int_{\tilde{y}}^1 -1 + 2x + 2bx - b\tilde{y} - b \, dx \\ &= -\left(\frac{b}{b+2}\right)^2. \end{aligned}$$

This agrees with the traditional minimax value of the game.

Case II: Bart Believes that Lisa Is Rash

Suppose that Bart's analysis leads him to think that Lisa is a little reckless, calling with $\tilde{y} < b/(b+2)$. Then the previous ARA shows that his bluffing function should be

$$g(x) = \begin{cases} 0 & \text{if } 0 \le x \le \max\{\tilde{y}, \tilde{x}\} \\ 1 & \text{if } \max\{\tilde{y}, \tilde{x}\} < x \le 1 \end{cases}$$

where $\tilde{x} = [b(1+\tilde{y})]/[2(1+b)]$. Simple algebra shows that if $\tilde{y} < b/(b+2)$, then $\tilde{x} > \tilde{y}$.

The value of this ARA game to Bart is

$$\begin{aligned} V &= -\int_0^{\tilde{x}} dx + \int_{\tilde{x}}^1 (-1 + 2x + 2bx - b\tilde{y} - b) \, dx \\ &= b\tilde{x} - b\tilde{y}(1 - \tilde{x}) - (1+b)\tilde{x}^2. \end{aligned}$$

The minimax value of the game is $-b^2/(b+2)^2$. Extensive but straightforward manipulation shows that *the value of this ARA game is strictly larger than the minimax value.* In particular, when Lisa always calls (i.e., $\tilde{y} = 0$), then the value of the game is $b^2/(4+4b)$; so, if Bart is confident that Lisa is reckless, then it is possible for his game to have a positive value, despite the disadvantage of going first.

Case III: Bart Believes that Lisa Is Conservative

Now suppose that Bart believes that Lisa is too pessimistic, calling with $\tilde{y} > b/(b+2)$. It is simple to show that this implies that $\tilde{x} < \tilde{y}$. When $x > \tilde{y}$, then

$$V_x = -(1-g(x)) + g(x)\tilde{y} + (1+b)g(x)(1-\tilde{y})\frac{x-\tilde{y}}{1-\tilde{y}} - (1+b)g(x)(1-\tilde{y})\left(1 - \frac{x-\tilde{y}}{1-\tilde{y}}\right). \quad (3.4)$$

When $x > \tilde{y}$, Bart's optimal play is to bet. On the other hand, when $x < \tilde{y}$, Bart's payoff is

$$V_x = -1 + g(x)\left[1 + \tilde{y} - (1+b)(1-\tilde{y})\right]. \quad (3.5)$$

And for $\tilde{y} > b/(b+2)$, the quantity in the square brackets is strictly positive. Thus, when $x < \tilde{y}$, the optimal $g(x)$ is a constant equal to 1.

The value V of this game is

$$\mathbb{E}[V_x] = \int_0^{\tilde{y}} \left[\tilde{y} - (1+b)(1-\tilde{y})\right] dx + \int_{\tilde{y}}^1 \left[\tilde{y} + (1+b)(x-\tilde{y}) - (1+b)(1-x)\right] dx. \quad (3.6)$$

Solving the integral shows $V = -b\tilde{y} + \tilde{y}^2(1+b)$. This value is increasing in $\tilde{y}$ for $\tilde{y} > b/(2+b)$ and it is equal to the minimax value at $\tilde{y} = b/(b+2)$. Thus *the value of the ARA game when Lisa is conservative is strictly larger than the minimax value.*

3.3 Continuous Bets

This section considers a modification of the Borel Game, in which Bart is not constrained to bet a fixed amount b, but may bet a random amount on some interval $(\epsilon, K]$ (perhaps as part of a bluff). This is a difficult problem for the minimax strategy; Karlin and Restrepo (1957) obtain a solution when the minimum bet is one unit and there are a finite number of possible larger bids. Ferguson and Ferguson (2007) report unpublished work by W. H. Cutler in 1976 that addresses the case of continuous bets in the context of the poker endgame. This section examines the case of continuous bet sizes from a first order ARA perspective.

We use the following notation:

ϵ, K: the lower and upper bounds of the bets Bart can choose, if he decides to bet; i.e. $[\epsilon, K]$ is Bart's betting strategy space, where $0 < \epsilon \ll K$ (usually ϵ is a very small positive number).

$g(x)$: the probability that Bart decides to bet after learning $X = x$.

$h(b|x)$: a probability density on $[\epsilon, K]$ that Bart will use to select his bet *conditional on* his decision to bet.

B_x: a random variable with value in $[\epsilon, K]$ representing Bart's betting strategy after he learns $X = x$.

We use $\mathbb{P}_{h(\cdot|x)}[\cdot]$ and $\mathbb{E}_{h(\cdot|x)}[\cdot]$ to stand for the probability or expectation computed using the probability measure induced by the density $h(\cdot|x)$.

In first-order ARA, Bart will "mirror" Lisa's opinion about his value of X *given that* she observes Bart's bet $B_x = b$. Formally, we have

$\tilde{g}(x)$: Bart's belief about Lisa's assessment on the probability that he decides to bet after he learns $X = x$.

$\tilde{h}(b|x)$: Bart's belief about Lisa's assessment of the probability density on $[\epsilon, K]$ that Bart will use to apply his bet *conditional on* his decision to bet.

$\tilde{f}(x|b)$: This is Bart's belief about Lisa's *posterior* probability assessment of the density for X after she observes that Bart bets and bets b. If Lisa is rational, it must be

$$\tilde{f}(x|b) = \frac{\tilde{h}(b|x)\tilde{g}(x)}{\int_0^1 \tilde{h}(b|z)\tilde{g}(z)\, dz}.$$

Give the above notation, we can write Bart's expected payoff given $X = x$ and his strategy $g(x), h(\cdot|x)$:

$$
\begin{aligned}
& \mathbb{E}_{g(x),h(\cdot|x)}\left[V_B|X=x\right] \\
= \; & \underbrace{-(1-g(x))}_{\text{Bart folds}} + g(x)\left\{\mathbb{E}_{h(\cdot|x)}\left[\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa folds } | \text{ Bart bets } B_x] \,|\, X=x\right]\right. \\
& +\mathbb{E}_{h(\cdot|x)}\left[\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa loses } | \text{ Bart bets } B_x]\cdot(1+B_x)\,|\,X=x\right] \\
& \left.-\mathbb{E}_{h(\cdot|x)}\left[\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa wins } | \text{ Bart bets } B_x]\cdot(1+B_x)\,|\,X=x\right]\right\}. \qquad (3.7)
\end{aligned}
$$

Hence, the first-order ARA solution from Bart's point of view, denoted by $\{g^*(x), h^*(\cdot|x)\}$, is

$$
\{g^*(x), h^*(\cdot|x)\} \in \underset{g(x),h(\cdot|x)}{\operatorname{argmax}} \quad \mathbb{E}_{g(x),h(\cdot|x)}\left[V_B|X=x\right]. \qquad (3.8)
$$

In order to solve for $\{g^*(x), h^*(\cdot|x)\}$, we first study Lisa's strategy and then roll back.

If Lisa folds, her payoff is -1. And given that Bart bets B_x and that Bart believes that Lisa will form the posterior assessment $\tilde{f}(\cdot|b)$ on his X, then after Lisa learns $Y = y$, Bart believes that Lisa's assessment of her probability of winning is

$$
\mathbb{P}_{\tilde{f}(\cdot|B_x)}[X \le Y \,|\, B_x, Y=y] = \int_0^y \tilde{f}(z|B_x)\,dz.
$$

So Bart believes that Lisa is, by calling, expecting a payoff of

$$
\begin{aligned}
V_y \;=\; & \mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa wins } | B_x, Y=y, \text{ Lisa calls }]\cdot(1+B_x) \\
& -\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa loses } | B_x, Y=y, \text{ Lisa calls }]\cdot(1+B_x) \\
=\; & \mathbb{P}_{\tilde{f}(\cdot|B_x)}[X\le Y\,|\,B_x,Y=y]\cdot(1+B_x) - \left\{1-\mathbb{P}_{\tilde{f}(\cdot|B_x)}[X\le Y\,|\,B_x,Y=y]\right\}\cdot(1+B_x) \\
=\; & 2\mathbb{P}_{\tilde{f}(\cdot|B_x)}[X\le Y\,|\,B_x,Y=y]\cdot(1+B_x)\cdot(1+B_x) - (1+B_x) \\
=\; & 2(1+B_x)\int_0^y \tilde{f}(z|B_x)\,dz - (1+B_x).
\end{aligned}
$$

Therefore, Bart believes Lisa will call if and only if

$$-1 \leq 2(1+B_x)\int_0^y \tilde{f}(z|B_x)\,dz - (1+B_x).$$

Since $\tilde{f}(z|B_x) \geq 0$, then for all $y \geq \tilde{y}^*(B_x)$ we must have

$$\int_0^y \tilde{f}(z|B_x)\,dz \geq \int_0^{\tilde{y}^*}(B_x)\tilde{f}(z|B_x)\,dz \geq \frac{B_x}{2(1+B_x)}.$$

Then Lisa will call if and only if

$$Y \geq \tilde{y}^*(B_x) \overset{def}{=} \inf\left\{ y \in [0,1] : \int_0^y \tilde{f}(z|B_x)\,dz \geq \frac{B_x}{2(1+B_x)} \right\}. \tag{3.9}$$

Hence, Bart believes that the probability Lisa will call after he bets B_x should be

$$\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa calls } \mid \text{ Bart bets } B_x] = \mathbb{P}[Y \geq \tilde{y}^*(B_x) \mid B_x] = 1 - \tilde{y}^*(B_x).$$

Consequently, Bart can compute the following quantities:

$$\begin{aligned}
\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa folds} \mid \text{Bart bets } B_x] &= \tilde{y}^*(B_x); && (3.10)\\
\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa loses} \mid \text{Bart bets } B_x] &= \mathbb{P}[\tilde{y}^*(B_x) \leq Y \leq x|B_x] \\
&= [x - \tilde{y}^*(B_x)]^+; && (3.11)\\
\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa wins} \mid \text{Bart bets } B_x] &= \mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa calls} \mid \text{Bart bets } B_x] \\
&\quad -\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa loses} \mid \text{Bart bets } B_x] \\
&= 1 - \tilde{y}^*(B_x) - [x - \tilde{y}^*(B_x)]^+. && (3.12)
\end{aligned}$$

Plugging (3.10), (3.11) and (3.12) into (3.7), we obtain

$$\begin{aligned}
&\mathbb{E}_{f(x),g(\cdot|x)}[V_B|X=x] = -(1-g(x)) \\
&\quad +g(x)\mathbb{E}_{h(\cdot|x)}\left[\tilde{y}^*(B_x) + 2[x-\tilde{y}^*(B_x)]^+(1+B_x) - (1-\tilde{y}^*(B_x))(1+B_x)\right]. \qquad (3.13)
\end{aligned}$$

Lemma 5. *Suppose* $\tilde{f}(\cdot|b)$ *is positive and continuous in* $b \in [\epsilon, K]$*, then* $\tilde{y}^*(b)$ *is continuous in* b*.*

Proof. The continuity and positivity of $\tilde{f}(\cdot|b)$ in b implies the continuity of $\int_0^y \tilde{f}(z|b)\,dz$ in (y, b). The positivity of $\tilde{f}(\cdot|b)$ implies the (global) one-to-one condition specified in Jittorntrum (1978). Hence, $\tilde{y}^*(b)$, as the (unique) solution of the following equation:

$$\int_0^y \tilde{f}(z|b)\,dz - \frac{b}{2(1+b)} = 0, \quad b \in [\epsilon, K],$$

must be continuous in b. ☐

Summing up the previous results, we obtain

Theorem 6. *For any* $x \in [0,1]$ *and given* $\tilde{f}(\cdot|b)$ *positive and continuous in* $b \in [\epsilon, K]$*, let* $\tilde{y}^*(b)$ *be defined as in* (3.9)*. Also let*

$$\begin{aligned} b^*(x) &\in \operatorname*{argmax}_{b\in[\epsilon,K]} \quad \tilde{y}^*(b) + 2(x - \tilde{y}^*(b))^+(1+b) - (1 - \tilde{y}^*(b))(1+b), \\ \Delta^*(x) &\overset{def}{=} \max_{b\in[\epsilon,K]} \quad \tilde{y}^*(b) + 2(x - \tilde{y}^*(b))^+(1+b) - (1 - \tilde{y}^*(b))(1+b). \end{aligned}$$

Then, the first-order ARA solution from Bart's perspective is given by

$$\begin{aligned} g^*(x) &= \begin{cases} 0 & \text{if } \Delta^*(x) < -1 \\ 1 & \text{if } \Delta^*(x) \geq -1; \end{cases} \\ h^*(b|x) &= \delta(b - b^*(x)), \end{aligned}$$

where $\delta(\cdot)$ *is the Dirac delta function.*

In other words, when he observes $X = x$, Bart will fold with probability 1 if $\Delta^*(x) < -1$ and bet $b^*(x)$ with probability 1 if $\Delta^*(x) \geq -1$. Of course, the regularity condition requiring that $\tilde{f}(\cdot|b)$ be positive and continuous in $b \in [\epsilon, K]$ is sufficient but not necessary.

3.3.1 Example: Lisa has a step-function posterior.

We now provide a simple illustration of how to apply Theorem 6 to obtain Bart's first-order ARA solution of the Borel game with continuous bets. In this subsection, we assume $\tilde{f}(\cdot|b)$ is of the following form:

$$\tilde{f}(x|b) = \begin{cases} \frac{1+K}{1+b} & \text{if } 0 \leq x \leq \frac{1+b}{1+K} \\ 0 & \text{otherwise.} \end{cases} \tag{3.14}$$

It is easy to see that $\tilde{y}^*(b) = \frac{b}{2(1+K)}$, and

$$\begin{aligned} &\tilde{y}^*(b) + 2(x - \tilde{y}^*(b))^+(1+b) - (1 - \tilde{y}^*(b))(1+b) \\ &= \begin{cases} -\frac{b^2}{2(1+K)} + (2x-1)(b+1) & \text{if } b \leq 2(1+K)x \\ \frac{b^2}{2(1+K)} - \frac{K}{1+K}b - 1 & \text{if } b > 2(1+K)x. \end{cases} \end{aligned}$$

Assume that ϵ is small enough that $\frac{\epsilon^2+2(1+K)\epsilon}{4(1+K)(1+\epsilon)} < \frac{1}{2} + \frac{\epsilon}{2(1+K)}$. Consider the following cases:

1. For $x < \frac{\epsilon^2+2(1+K)\epsilon}{4(1+K)(1+\epsilon)}$, then $b^*(x) = \epsilon$ and $\Delta^*(x) = -\frac{\epsilon^2}{2(1+K)} + (2x-1)(\epsilon+1) < -1$. By Theorem 6, $g^*(x) = 1$; i.e., Bart will fold w.p. 1. There is no need to specify $h^*(\cdot|x)$.

2. For $\frac{\epsilon^2+2(1+K)\epsilon}{4(1+K)(1+\epsilon)} \leq x < \frac{1}{2} + \frac{\epsilon}{2(1+K)}$, then $b^*(x) = \epsilon$ and $\Delta^*(x) = -\frac{\epsilon^2}{2(1+K)} + (2x-1)(\epsilon+1) \geq -1$. By Theorem 6, $g^*(x) = 1$ and $h^*(b|x) = \delta(b-\epsilon)$, i.e. Bart will bet ϵ w.p. 1.

3. For $\frac{1}{2} + \frac{\epsilon}{2(1+K)} \leq x < \frac{1}{2} + \frac{K}{2(1+K)}$, then $b^*(x) = 2(1+K)x - (1+K)$ and $\Delta^*(x) = \frac{1+K}{2}(2x-1)^2 + (2x-1) \geq -1$. By Theorem 6, $g^*(x) = 1$ and $h^*(b|x) = \delta(b-(2(1+K)x-(1+K)))$; i.e., Bart will bet $2(1+K)x - (1+K)$ w.p. 1.

4. For $x \geq \frac{1}{2} + \frac{K}{2(1+K)}$, then $b^*(x) = K$ and $\Delta^*(x) = -\frac{K^2}{2(1+K)} + (2x-1)(K+1) \geq -1$. Then, by Theorem 6, $g^*(x) = 1$ and $h^*(b|x) = \delta(b-K)$; i.e., Bart will bet K w.p. 1.

To summarize, we plot Bart's first-order ARA strategy as a function of his draw $X = x$ in Figure 3.2.

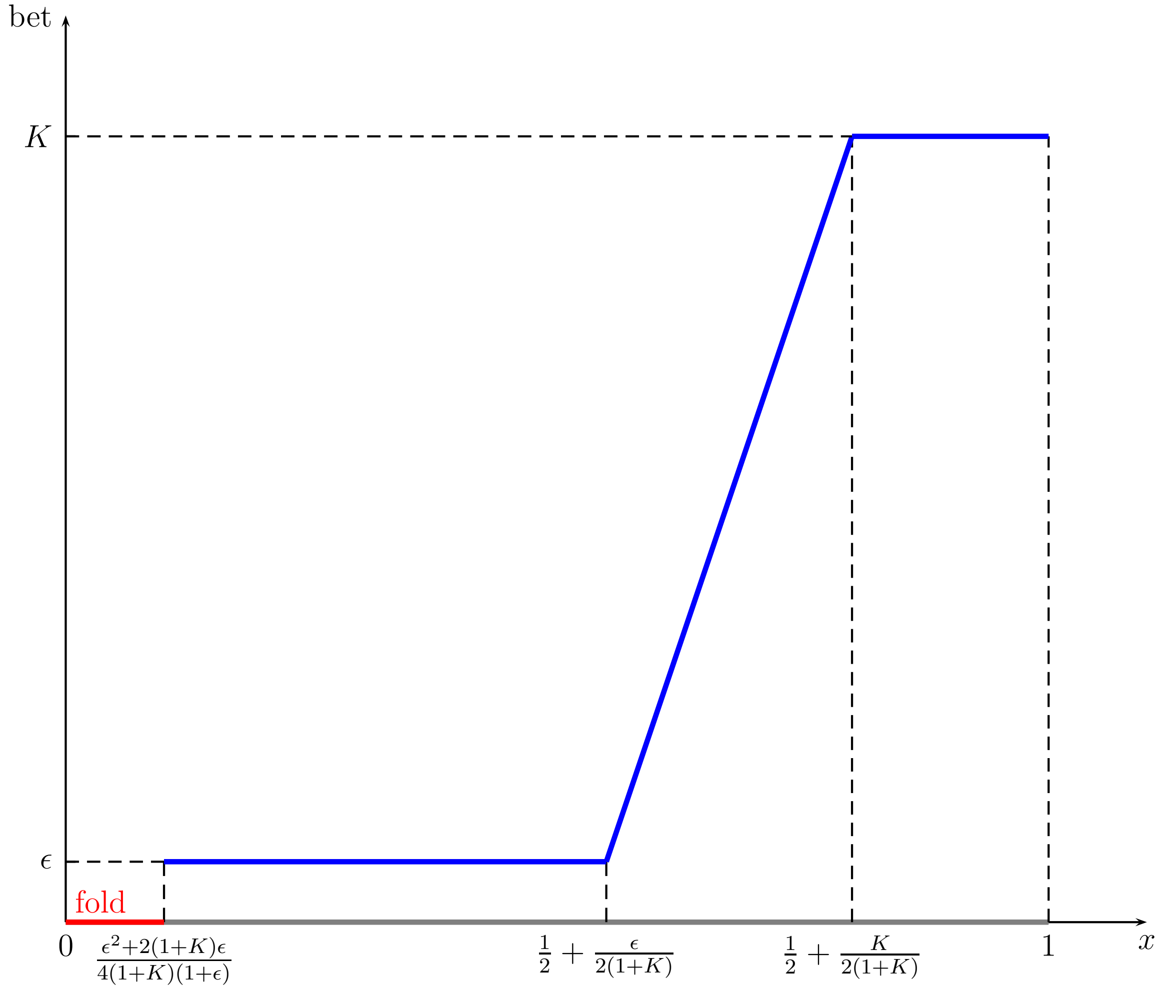

FIGURE 3.2: Bart's first-order ARA strategy as a function of x for $\tilde{h}(\cdot|b)$ given in (3.14).

3.4 Conclusion

It is not novel to suggest that the analysis of games should use a solution concept based on maximizing expected utility rather than minimizing the maximum loss. This was suggested by Raiffa (1982) and Kadane and Larkey (1982). Even earlier, (Good, 1952, p.114) raised the issue:

> . . . it is in this theory [the theory of games] that they [minimax solutions] are more justifiable. *But even here in practice you would prefer to maximize your expected gain.* You would probably use minimax solutions when you had a fair degree of

> belief that your opponent was a good player. Even when you use the minimax solution you may be maximizing your expected gain since you may already have worked out the details of the minimax solution, and you would probably not have time to work out anything better once a game had started. [Note the implicit use here of "type II rationality" which is especially pertinent to avoid an "infinite regress".] To attempt to use a method other than the minimax method would then lead to too large a probability of a large loss, especially in a game like poker.

Such arguments, combined with the fact that people's behavior does not conform with game theoretic prescriptions, warrants the exploration of the alternative in this paper.

If one wants to maximize one's expected utility, then one needs a probability distribution over the actions of an opponent. Our contribution has been to develop an explicit mechanism for developing this distribution. The mirroring argument derives from a model for how humans think about the world; for some problems, they are casual and impressionistic, but in other cases they focus more deeply and imagine the strategic reasoning of their adversary. Traditional risk analysis provides a computational framework for calculation, once one has a model of the decision processes of one's opponent, and we call this combination of mirroring and risk analysis *adversarial risk analysis* (ARA).

Our examination of the Borel game has found the two-person solution of lower-order ARA and the solution with continuous bet sizes. These were challenging problems in the early days of game theory. The ARA approach seems somewhat simpler, at least conceptually, although high-order ARA can be tedious.

A second advantage of ARA is that it allows the decision-maker to use information that is not permitted in the minimax context. For example, Bart may have played many previous rounds of a Borel game with Lisa, and formed an accurate opinion of her style; it seems foolish for him to ignore this for future games in favor of a minimax solution.

Bibliography

Arce, D. and Sandler, T. (2007), "Terrorist Signalling and the Value of Intelligence," *British Journal Political Science*, 37, 573–586.

Banks, D. and Anderson, S. (2006), "Game Theory and Risk Analysis in the Context of the Smallpox Threat," in *Statistical Methods in Counterterrorism*, eds. A. Wilson, G. Wilson, and D. Olwell, pp. 9–22, New York, N.Y., Springer.

Bayrak, H. and Bailey, M. D. (2008), "Shortest Path Network Interdiction with Asymmetric Information," *Networks*, 52, 133–140.

Bellman, R. and Blackwell, D. (1949), "Some Two-Person Games Involving Bluffing," *Proceedings of the National Academy of Sciences*, 35, 600–605.

Berger, U. (2005), "Fictitious Play in $2 \times n$ Games," *J. Econ. Theory*, 120, 139–154.

Borel, E. (1938), "*Traité du Calcul des Probabilités et ses Applications*," Applications aux Jeus des Hazard, Vol. IV.

Brown, G., Carlyle, W., and Wood, R. (2008), *Optimizing Department of Homeland Security Defense Investments: Applying Defender-Attacker(-Defender) Optimization to Terror Risk Assessment and Mitigation*, chap. Appendix E, National Academies Press, Washington, D.C.

Brown, G. M., Carlyle, M., Salmeron, J., and Wood, K. (2006), "Defending critical infrastructure," *Interfaces*, 36, 530–544.

Brown, G. W. (1949), "Some Notes on Computation of Games Solutions," Tech. Rep. Report P-78, The Rand Corporation.

Camerer, C. (2003), *Behavioral Game Theory*, Princeton University Press, Princeton, N.J.

Dresher, M. (1961), *Games of Strategy: Theory and Applications*, RAND/Prentice Hall, Englewood Cliffs, N.J.

Ferguson, C. and Ferguson, T. S. (2003), "On the Borel and von Neumann Poker Models," *Game Theory and Applications*, 9, 17–32.

Ferguson, C. and Ferguson, T. S. (2007), "The Endgame in Poker," in *Optimal Play—Mathematical Studies of Games and Gambling*, eds. S. Ethier and W. Eadington, pp. 79–106, Institute for the Study of Gambling and Commercial Gaming.

Gintis, H. (2009), *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences*, Princeton University Press, Princeton, N.J.

Good, I. J. (1952), "Rational Decisions," *Journal of the Royal Statistical Society, Series B*, 14, 107–114.

Gutfraind, A., Hagberg, A., and Pan, F. (2009), *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, chap. Optimal Interdiction of Unreactive Markovian Evaders, pp. 102–116, Springer, Heidelberg.

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. (1973), "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points," *International Journal of Game Theory*, 2, 1–23.

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.

Karlin, S. and Restrepo, R. (1957), "Multistage Poker Models," in *Contribution to the Theory of Games, III*, pp. 337–363, Princeton University Press.

Krishna, V. (1992), "Learning in Games with Strategic Complementarities," Mimeo.

Lemke, C. and Howson, J. (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., P., G., Jenkinson, D., Oakley, J., and Rakow, T. (2006), *Uncertain Judgements: Eliciting Experts' Probabilities*, Wiley, New York, N.Y.

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, 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.

Velu, C. and Iyer, S. (2008a), "The Rationality of Irrationality for Managers: Returns-Based Beliefs and the Traveler's Dilemma," Cambridge Working Papers in Economics 0826, Faculty of Economics,University of Cambridge.

Velu, C. and Iyer, S. (2008b), "The Returns Based Beliefs and the Prisoner's Dilemma," Cambridge Working Papers in Economics 0854, Faculty of Economics,University of Cambridge.

von Neumann, J. and Morgenstern, O. (1944), *Theory of Games and Economic Behavior*, Princeton University Press.

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*, Kluwer Academic Publishers, Boston, MA.