# Adversarial Risk Analysis

Insua RIOS INSUA, Jesus RIOS, and David BANKS

Applications in counterterrorism and corporate competition have led to the development of new methods for the analysis of decision making when there are intelligent opponents and uncertain outcomes. This field represents a combination of statistical risk analysis and game theory, and is sometimes called *adversarial risk analysis*. In this article, we describe several formulations of adversarial risk problems, and provide a framework that extends traditional risk analysis tools, such as influence diagrams and probabilistic reasoning, to adversarial problems. We also discuss the research challenges that arise when dealing with these models, illustrate the ideas with examples from business, and point out relevance to national defense.

KEY WORDS: Auctions; Decision theory; Game theory; Influence diagrams.



## 1. INTRODUCTION

Game theory has long been considered impractical for risk management decision-making (Bier and Cox 2007). This viewpoint has recently become less dogmatic because:

- High-profile terrorist attacks have demanded significant national investment in protective responses, and there is public concern that not all of these investments are prudent or effective (Parnell et al. 2008).
- Key business sectors (especially finance, e-commerce, and software) have become much more mathematically sophisticated, and are now using this expertise to shape corporate strategy for auction bidding, timing of product release, lobbying efforts, and other decisions (cf., McAfee and McMillan 1996; Rothkopf 2007).
- Regulatory legislation must balance competing interests (for growth, environmental impact, safety) in a way that is credible and transparent (cf., Heyes 2000).
- The on-going arms race in cybersecurity means that the financial penalties for myopic protection are large and random (Killourhy, Maxion, and Tan 2004).

These challenges cross many fields (Statistics, Economics, Operations Research, Engineering, and so on) and are characterized by the fact that there are two or more intelligent opponents who make decisions for which the outcome is uncertain. Collectively, we call this problem area adversarial risk analysis (ARA).

Traditional statistical risk analysis grew in the context of nuclear reactor safety, insurance, and other applications in which loss was governed by chance (sometimes called Nature, and described as a neutral opponent) rather than the malicious (or self-interested) actions of intelligent actors. But in ARA one needs to have some model for the decision-making of all the participants. This model might be classically game-theoretical, with (noncooperative) Nash equilibria as the core concept (Myerson 1991) or it might be more psychological, reflecting either a Kadane-Larkey formulation (1982) or empirical studies of game behavior (Camerer 2003). Hausken (2002) provided additional insights on combining risk analysis and game theory.

Much of the new ARA literature involves counterterrorism. Banks and Anderson (2006) analyzed strategies for a smallpox attack by modeling the problem as a zero sum game with random payoffs and solving the game through both minimax and Bayesian approaches. Zhuang and Bier (2007) computed best responses and Nash equilibria as a basis for allocating resources against terrorism when the defender and attacker have different multiattribute utility functions, in situations of both simultaneous and sequential play. Brown, Carlyle, Salmeron, and Wood (2006) presented bilevel (max-min, min-max) and trilevel (min-max-min) optimization models for defender-attacker, attacker-defender, and defender-attacker-defender problems that may be framed as Stackelberg games. Kardes and Hall (2005) argued for the use of robust stochastic games to deal with counterterrorism. Paté-Cornell and Guikema (2002) provided an asymmetric prescriptive/descriptive approach in the light of modern negotiation analysis (cf., Raiffa 2002). A decision analysis viewpoint is adopted by, among others, von Winterfeldt and O'Sullivan (2006), who used decision trees to evaluate Man-Portable Air Defense Systems countermeasures; Pinker (2007), who applied influence diagrams to assess the deployment of various short-term countermeasures; and Parnell (2007), who used generic multiobjective decision trees and influence diagrams to evaluate bioterrorist threats. There is also a rich literature in political science regarding game theory and terrorism, though it places little emphasis on risk analysis aspects (e.g., Siqueira and Sandler 2006; Arce and Sandler 2007; Powell 2007).

This article focuses upon ARA in simple two-person conflicts, but we emphasize that the same basic ideas apply to corporate competition, government regulation, and cybersecurity. Our purpose is to provide a unifying perspective on ARA by discussing the current techniques, comparing their features, and identifying open challenges.

Section 2 describes a framework for conventional risk analysis that is generalized to ARA in Section 3. Section 4 assesses previous methodologies based on game theory, decision analysis, and negotiation analysis. We then present a novel approach that encompasses previous work and apply it to simple auction situations. Section 7 concludes with some open questions in this area of potential interest to the statistical community.

David Rios Insua is Professor of Statistics at Rey Juan Carlos University, Madrid, Spain (E-mail: *david.rios@urjc.es*). Jesus Rios is a Postdoctoral Research Fellow in Computer Science at the University of Luxembourg (E-mail: *jrios@samsi.info*). David Banks is Professor of the Practice of Statistics at Duke University, Durham, NC (E-mail: *banks@stat.duke.edu*). Most of this work was done as part of the SAMSI Program on Risk in 2006–2007. This material was based on work partially supported by the National Science Foundation under Grant DMS-0635449 to the Statistical and Applied Mathematical Sciences Institute. David Rios Insua acknowledges support from the Spanish Ministry of Science and Education, the E-DEMOCRACIA-CM, and the COST Algorithmic Decision Theory programs. The authors thank the referees for their helpful comments and discussion.


Journal of the American Statistical Association
June 2009, Vol. 104, No. 486, Review
DOI 10.1198/jasa.2009.0155

## 2. A FRAMEWORK FOR RISK ANALYSIS

This section reviews a schematic framework that formalizes standard risk analysis, assessment, and management methods as in Haimes (2004) or Bedford and Cooke (2001), adapted to the classic proposal of Kaplan and Garrick (1981). The framework uses influence diagrams to structure the problems—these are popular in the decision analysis and artificial intelligence communities, and encompass both simultaneous and sequential games (in which decision trees are also used). For simplicity, we shall assume that losses can be monetized as costs. All the participating agents are assumed to be expected utility maximizers (cf., French and Rios Insua 2000). In this context, risk analysis entails activities of risk assessment and risk management as described later.

Figure 1 shows an *influence diagram* (cf., Pearl 2005) that displays the simplest version of a nonadversarial risk management problem. An influence diagram is a directed acyclic graph with three kinds of nodes: decision nodes, shown as rectangles; uncertainty nodes, shown as ovals; and value nodes, shown as hexagons. Arrows into a value or uncertainty node indicate functional and probabilistic dependence, respectively. Thus, the utility function at the value node depends on its immediately preceding nodes and probabilities at a chance node are conditional on the values of its direct predecessors. Arrows into a decision node indicate that when the decision is made, the values of its preceding nodes are known. Influence diagrams display some of the same information that is shown in a decision tree (cf., Singpurwalla 2006, Sec. 2.9), but they are organized differently, show the utility function explicitly, and take a higher-level view of the problem.

In Figure 1 the rectangle represents the set $A$ of possible decisions or actions, the oval represents the random costs associated with the decisions, and the hexagon represents the net consequences, or values, in terms of the decision maker's utility function. It corresponds to a problem in which an organization has to make a decision $a$ from a set $A$ of choices. The cost $c$ that results from each decision is uncertain and is modeled through the density $\pi(c \mid a)$; this cost may reflect the fact that the outcome for a particular decision is uncertain, or that the cost associated with a particular outcome is uncertain, or both. The utility $u(c)$ of the cost is decreasing and typically nonlinear; costs are bad, and catastrophic costs are disproportionately bad. One seeks the decision that maximizes the expected utility

$$\psi = \max_{a \in \mathcal{A}} \left[ \psi(a) = \int u(c) \pi(c \mid a) dc \right]. \quad (1)$$

In practice, the cost for a particular action is complex and conditional on the outcome; it often includes fixed and random summands. The organization will typically perform a *risk assessment* to:

1. identify disruptive events $E_1, E_2, \ldots, E_k$ (these may be assumed to be mutually exclusive);
2. assess their probabilities of occurrence, $P(E_i \mid a) = q_i(a)$; and,
3. assess the cost $c_i$ conditional on the occurrence of $E_i$ and decision $a$ (these costs are typically random and the assessment may be a distribution).

It is convenient to let $E_0$ be the event that there are no disruptions, with probability $q_0(a)$. Figure 2 shows the influence diagram that extends the previous formulation to include a risk assessment that accounts for specific disruptive hazards and the additional random costs these may entail.

Let $\boldsymbol{q}(a)$ be the vector of probabilities corresponding to decision $a$ and let $\pi_i(c \mid a)$ be the cost density under decision $a$ if event $E_i$ occurs. Then, the density of the cost for decision $a$ is the mixture $\sum_{i=0}^{k} q_i(a)\, \pi_i(c \mid a)$. Once the risk assessment is performed, the organization wants the maximum expected utility decision, which is found by solving

$$\psi_r = \max_a \left[ \psi_r(a) = \sum_{i=0}^{k} q_i(a) \int u(c) \pi_i(c \mid a)\, dc \right]. \quad (2)$$

In some cases, the probabilities $q_i(a)$ are themselves uncertain (e.g., if one is combining elicited assessments from multiple experts). In that case, one can describe that uncertainty through a distribution $g(\boldsymbol{q}(a))$ on the unit simplex $S$ in $\mathrm{IR}^{k+1}$ and solve

$$\psi_r = \max_a \left[ \psi_r(a) = \int_S g(\boldsymbol{q}(a)) \times \left( \sum_{i=0}^{k} q_i(a) \int u(c) \pi_i(c \mid a)\, dc \right) d\boldsymbol{q}(a) \right]. \quad (3)$$

Note that this maximizes the utility with respect to uncertainty from two different sources—the randomness in the costs and the imprecise knowledge about the disruption probabilities.

Consider the difference $\psi - \psi_r$. This is nonnegative, as $\psi$ describes a decision problem that excludes costs associated with disruptive events, whereas $\psi_r$ relies upon risk assessment and is more realistic. To reduce this difference, organizations often undertake a risk management strategy. *Risk management* introduces an additional set of choices $M$, such as contingency plans or insurance policies; these tend to lower the costs associated with particular disruptions or lower the chance of

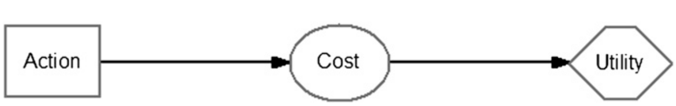


Figure 1. Basic influence diagram.

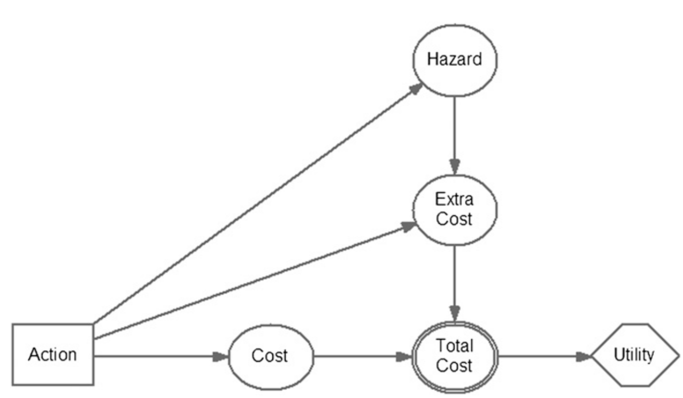


Figure 2. Influence diagram with risk assessment.

disruption or both. As an example of risk management in a nonadversarial situation, imagine that an engineering company is building a dam in a foreign country. The company considers two possible designs, these being the decisions in $A$. But the risk assessment indicates the possibilities of a national strike or logistic delays in critical supplies. The company therefore considers buying insurance (which would protect against the costs associated with both hazards) or employing extranational engineers (at more expense, but with less chance of striking) or doing neither. These choices are elements in $M$.

In principle, one could take the cross product of the sets $A$ and $M$ and then solve for $\psi_r$ over this extended set. But in practice, it is often helpful for managers to keep these distinct. The risk management solution remains the same,

$$\psi_m = \max_{(a,m)\in\mathcal{A}\times\mathcal{M}} \psi_r(a,\ m) \tag{4}$$

where, in an obvious extension of previous notation,

$$\psi_r(a,m) = \sum i\ q_i(a,m) \int u(c)\pi_i(c \mid a,m)\, dc.$$

As before, a slightly more complicated formula applies when there is uncertainty in the $q_i(a, m)$ probabilities. Figure 3 shows the influence diagram for a risk management problem.

Because risk management extends the set of choices, then $\psi_m \geq \psi_r$, but both are still less than $\psi$. The problem described in (4) can be viewed as an example of a sequential decision problem and could be represented as a decision tree; first one picks the design, and then one picks a choice in $M$, with the corresponding uncertainty nodes. In general, problems of this kind require dynamic programming, because the decision at each time period depends on the likely future outcomes, deeper in the tree. But trees can be a problematic representation, because the choice sets need not be discrete (as in our example) but could be continuous (e.g., if the construction company can choose among insurance policies with infinitely divisible prices). Clearly, additional complexity arises if there are many levels of sequential investment, if one allocates risk management resources according to a portfolio analysis that constrains total expenses, and if there are multiattribute utility functions.

## 3. ADVERSARIAL RISKS: MODELING

We now consider the situation in which there are adversaries whose actions affect each other's risks. Assume that there are just two opponents (Apollo and Daphne). Their decision problems are structurally similar: both can take actions that affect the costs of the other, and both seek to maximize their expected utilities. The sets of actions for Apollo and Daphne are, respectively, denoted by $\mathcal{A}$ and $\mathcal{D}$; their utility functions are $u_A(\cdot)$ and $u_D(\cdot)$; and their collection of probabilities about outcomes are $\mathcal{P}_{\mathrm{A}}$ and $\mathcal{P}_{\mathrm{D}}$.

In this kind of situation, Apollo and Daphne may have different utility functions and different probability assessments of the costs. Each player knows their own utilities but the other's utilities and probability assessments may be unknown. A simple example is a two-party auction; neither party knows what utility the other puts upon the object of the bidding. Furthermore, players may also be uncertain about their own object valuation, as in procurement auctions in which bidders are uncertain about their costs and, therefore, their benefits associated with the execution of the auctioned project or service. Similarly, a situation with different probability assessments might arise in counterterrorism, where both parties could have intelligence that leads to very different estimates for the probability of successful attack under different Attacker/Defender choices.

To illustrate these ideas in the context of a two-person procurement auction for a construction contract, we extend the traditional influence diagram in Figure 4 to show the interaction between the decisions of Apollo and Daphne in an adversarial situation. Essentially, the presence of an adversary changes the risk analysis probabilities from being conditioned on one event (Nature's actions) to being conditioned on two events (Nature and the choices made by the opponent).

In this procurement example the roles are symmetric, as indicated in the diagram, although in, say, a counterterrorism situation, the diagram would probably be asymmetric. One of the nodes is common (i.e., the node labeled ''Hazard''); it

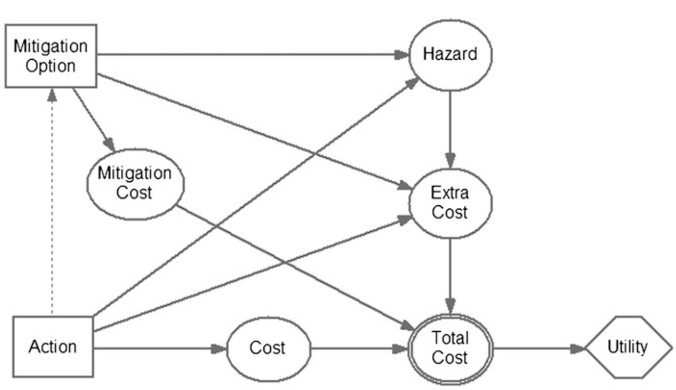


Figure 3. Influence diagram with risk management.

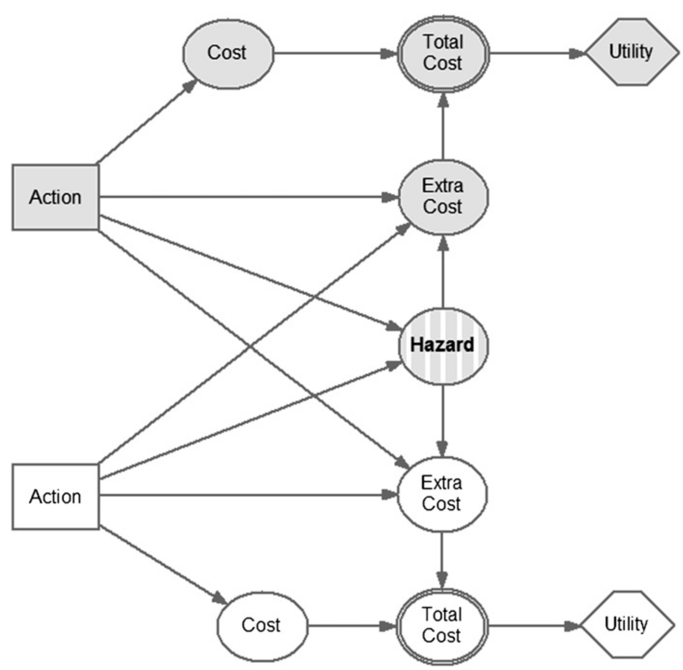


Figure 4. Symmetric adversarial risk influence diagram.

reflects (e.g., the mutual chance of weather delays in a construction project leading to additional costs). Other cost nodes (i.e., ‘‘Cost’’ and ‘‘Extra Cost’’) are not common; these represent the random costs associated with Apollo’s and Daphne’s respective firms, and could be very different. Similarly, the utility functions are not common; a small firm has much more at stake in a given bid than a large firm.

For the ARA case, the expected utilities of Apollo and Daphne depend upon the actions of both. Specifically, extending (1), the utility that Apollo expects from choosing action $a \in \mathcal{A}$ when Daphne makes decision $d \in \mathcal{D}$ is

$$\psi_A(a, d) = \int u_A(c)\pi_A(c \mid a, d)dc,$$

where $\pi_A(c \mid a, d) \in \mathcal{P}_A$ represents Apollo’s beliefs about the distribution of his costs for the decision pair $(a, d)$. As in (3), Apollo can identify events $E_1, \ldots, E_k$ that affect his costs and include their impact explicitly:

$$\psi_A(a,\ d) = \int_S g_A(\boldsymbol{q}(a,\ d)) \times \left(\sum_i q_i(a,\ d) \int u_A(c)\pi_A(c \mid a, d, E_i)\, dc\right) d\boldsymbol{q}(a, d), \tag{5}$$

where $g_A(\cdot)$, $\boldsymbol{q}_A(a, d)$, $\pi_A(c \mid a, d, E_i) \in \mathcal{P}_A$. The expected utility $\psi_D(a, d)$ for Daphne is analogous. In this framework, the key remaining problem is to determine how Apollo and Daphne make their decisions.

As written, this description of ARA applies to normal form games, in which players make simultaneous decisions. But it also applies to sequential (or ‘‘extensive form’’) games, such as Stackelberg games in which the Leader and Follower alternate their moves, or games in which agents act asynchronously. The influence diagram representation encompasses all of these, but alternating games are usually represented through decision trees. The risk analysis must condition on the choices that have already been made when calculating expected utilities. The tree form is conceptually transparent, but often the tree of unmade decisions that drives the risk calculation is elaborate, and simplifying assumptions are needed. A special case of this arose in Equation (4), where $\mathcal{M}$ was used to denote a set of decisions that were made separately. Additionally, in realistic applications there are usually more than two adversaries, requiring an even more complex analysis.

## 4. GAME THEORY FOR ARA

The previous section emphasized the risk analysis aspect of ARA. This section emphasizes the adversarial aspect, focusing on specific solution concepts.

### 4.1 Nash Equilibria Analysis

The typical game theory approach to predict Apollo and Daphne’s choice of actions is to find the Nash equilibria for their expected utilities. This is a combination of choices such that no unilateral deviation from the choice can improve a player’s situation. Essentially, each player makes the choice that helps them the most, while taking account of the fact that the other will do the same. Thus, the best choice for Apollo depends upon the action taken by Daphne, and conversely. The typical result is that neither party gets the outcome they would like, but they mutually avoid the worst results.

Under reasonable conditions at least one equilibrium exists, but often there are many. Sometimes equilibrium is a pure strategy, in which the choices of both players are determined; but often an equilibrium corresponds to a ‘‘mixed strategy’’ in which Apollo and Daphne choose among their available actions with fixed, mutually known, probabilities. When the sets $\mathcal{A}$ and $\mathcal{D}$ are finite, then the payoff table with choice $a$ in Apollo’s row and choice $d$ in Daphne’s column has entry $(\Psi_A(a, d), \Psi_D(a, d))$, the expected utilities for both. A sequence of linear programming problems can be solved to eventually find Nash equilibrium (Lemke and Howson 1964).

The simplest example is a two-person zero sum game. Suppose Apollo and Daphne are playing a game in which Apollo has two possible actions ($A_1$ and $A_2$), as does Daphne ($D_1$ and $D_2$). For each pair of actions, there is a payoff; the payoff is the amount that, say, Daphne wins, which equals the amount that Apollo loses. This is shown in the following payoff table.

| | $A_1$ | $A_2$ |
|---|---|---|
| $D_1$ | $a$ | $b$ |
| $D_2$ | $c$ | $d$ |

Suppose $b > a > c$. In that case, Daphne will choose $D_1$ and Apollo will choose $A_1$. The $(D_1, A_1)$ choice is a *saddlepoint solution*, because Daphne cannot improve her outcome by switching, nor can Apollo. The $(D_1, A_2)$ choice is not a solution because Apollo would be better off with $(D_1, A_1)$; some thought shows that each of the other possible choices is similarly improvable for one or both players, and that the value of $d$ is irrelevant.

Now suppose that $a \geq d > b \geq c$. The payoff table has no saddlepoint (i.e., there is no dominating strategy). In this case the minimax theorem shows that the solution is a mixed strategy, in which Daphne chooses $D_1$ with probability $(d - c)/(a + b - d - c)$ and Apollo chooses $A_1$ with probability $(d - b)/(a + b - d - c)$. This maximizes each of their expected gains (or minimizes their expected losses) given that their opponent is trying to do the same, and neither player can deviate from this solution without lowering their expected return.

For this simple example, one can solve the game using algebra. In more complicated tables, where players have more than two choices, or in which there are more than two players, linear programming is required. A particularly important case is the nonzero sum game, in which each pair of actions has two payoffs, one for Daphne and one for Apollo, and the gain/loss for Daphne is not equal to the loss/gain for Apollo. In these games it may be possible for both players to gain or both to lose. Such problems are usually more realistic, but can be very difficult to solve.

Zhuang and Bier (2007) modeled a nonzero sum game in which a government must spend resources on countermeasures against both terrorism and natural disasters. Generic utility functions are assumed with appropriate risk aversion and marginal return features. The terrorists aim at maximizing their expected utility by choosing an optimal level of effort to devote to each target, whereas the government aims at maximizing its expected utility by choosing an optimal level of defensive investments for each target. No explicit resource constraints are assumed for either party; these are implicit in the options considered. For this set-up, Zhuang and Bier (2007) considered two approaches, both under the assumption of common knowledge. In the first model, the government and the terrorists decide simultaneously and solve for a Nash equilibrium; that is, they seek ($a^*$, $d^*$) such that

$$\psi_A(a^*, d^*) = \max_{a \in \mathcal{A}} \psi_A(a, d^*) \quad \text{and} \quad \psi_D(a^*, d^*) = \max_{d \in \mathcal{D}} \psi_D(a^*, d).$$

The notation is as before; just consider the terrorists as the Attacker with the same notation as for Apollo and the the government as the Defender with the same notation as for Daphne. For the second model, Zhuang and Bier (2007) considered a sequential game with first-mover advantage. Here the Defender chooses a defensive investment and then the Attacker selects an action after observing the defensive investments. This is a special case of a Stackelberg game (Gibbons 1992) and is sometimes called a "Defend-Attack" model. In this set-up the concept of best response is especially relevant. The Attacker solves

$$a^*(d) = \text{argmax}_{a \in \mathcal{A}} \psi_A(a, d),$$

where $d$ is known. The Defender knows the kind of calculation the Attacker will make and must solve

$$d^* = \text{argmax}_{d \in \mathcal{D}} \psi_D(d, a^*(d)),$$

with $a^*(d)$ as previously. Bier, Oliveros, and Samuelson (2007) applied this approach to a range of problems in public policy decision-making.

ARA with alternating moves is also treated in Brown et al. (2006) and Brown, Carlyle, and Wood (2008), who consider not only Defend-Attack scenarios, but also Attack-Defend and Defend-Attack-Defend cases. In the Defend-Attack-Defend model, the Defender makes initial investments, some of which may entail deploying resources for future use. The Attacker observes these investments and chooses a response, and then the Defender uses the deployed resources to recover after the attack. These authors solve the corresponding nested optimization problems, allowing:

incorporation of explicit resource constraints on the Attacker and the Defender;
(implicit) calculation of the Nash equilibrium when the payoff table contains performance measures such as the mean, the median, or the 99th percentile of the Defender's distribution on the random cost.

The Defend-Attack-Defend model replaces the random outcomes with the performance measure, and then uses uses mathematical programming to solve the game theory problem:

$$\begin{aligned} d_2^*(d_1, a) &= \text{argmax}_{d_2 \in \mathcal{D}_2} \psi_D(d_1, a, d_2) \\ a^*(d_1) &= \text{argmax}_{a \in \mathcal{A}} \psi_A(d_1, a, d_2^*(d_1, a)) \\ d_1^* &= \text{argmax}_{d_1 \in \mathcal{D}_1} \psi_D(d_1, a^*(d_1), d_2^*(d_1, a^*(d_1))), \end{aligned}$$

where $\mathcal{D}_1$ and $\mathcal{D}_2$ are the sets of decisions available to the Defender at the initial and recovery phases, respectively. These decision sets can incorporate cost constraints. There are scaling issues as the complexity of the constraints grows and the sizes of the decision sets increase.

A different approach is taken in Kardes (2005), who considered robust stochastic games. These involve dynamic models in discrete time, using robustness to address the uncertainty in expert probability assessments. This uncertainty is related to robust Bayesian methods (cf., Rios Insua and Ruggeri, 2000), in which rather than specifying distributions $g_A$ and $\pi_A$ in (5) to find expected utilities, one puts constraints on the set of distributions and calculates expected utilities for worst-case distributions. These are then used in computing Nash equilibria.

## 4.2 Expected Best Choice Analysis

The preceding methods handle the uncertainty in the risk analysis by replacing the random cost with its expectation (or sometimes the median cost, or the 99th percentile of a cost distribution). This approach to uncertainty permits more scalable calculation and allows portfolio analysis to account for resource constraints, but it is at best an approximation. Taking expectations and computing Nash equilibria do not commute, and this section treats methods that reverse the order in the previous section. In general, one wants to make the decision that has the best average outcome, rather than the best decision for the average outcomes.

To address this issue, Banks and Anderson (2006) described two approaches in the context of a counterbioterrorism scenario in which there are three kinds of smallpox attacks and four possible defenses, all pairs of which involve uncertain costs/damages. The first step is a risk analysis that elicits probabilities from experts to determine a joint distribution $\pi(T)$ for entries in the payoff table whose rows and columns represent the actions available to the adversaries. Then, in the first approach, they:

1. generate many random payoff tables according to the elicited distribution $\pi(T)$; these payoff tables represent random realizations of the unknown table that nature has chosen for the game the opponents must play;
2. solve the Nash equilibrium problem for each table in the sample, finding the best action for each opponent and the corresponding payoffs; and
3. identify the action that has the best average payoff, thus estimating the play that is the best expected value for the unknown true payoff matrix.

Obviously, the entries in the payoff table $T$ are unlikely to be independent. The dependence structure is obtained from subject-matter experts during the risk analysis stage.

The second approach in Banks and Anderson (2006) is similar to the first, except that after the risk analysis produces a distribution on the payoff tables, the next step is to put a distribution on the actions of the opponent. This enables the decision-maker to select the action that maximizes the expected payoff with respect to the prior on the adversary's action. In their counterterrorism example, this approach reflects the belief that terrorists do not use classical game theory, and that there is relevant intelligence information that makes some of their choices more likely than others. This prior could incorporate expert judgment about the Attacker's utility function, rationality, and resources. This analysis is a Bayesian version of game theory, as proposed in Kadane and Larkey (1982). It is similar to the approach in Raiffa (1982); see Kadane (1993) for historical background.

Broadly speaking, Bayesian approaches to game theory are less conservative than classical game theory. The Nash equilibrium (in a zero-sum game) assumes that the opponent invariably chooses the action that most advantages them, given their calculation of their opponent's choice. This forces each player to defend against the worst-case scenario, leading to solutions that are typically expensive and often unrealistic. The Bayesian approach allows the analyst to use subjective beliefs about the constraints and goals of the opponents to better forecast their decisions. Bayesian game theory can be approached in several ways, but the common element is that each party must mirror the decision-making process of the other, placing distributions over the actions of their adversary that reflect symmetric reasoning on the part of their opponent.

### 4.3 Decision Analysis

Some ARA methods avoid explicit Nash solution; instead this gets built into probabilities about the actions of opponents. The second approach in Banks and Anderson (2006) is one example, but there are several others. These typically employ decision trees or influence diagrams.

As an example, the Department of Homeland Security has used a tree-based analysis for risk assessment in counter-bioterrorism. That effort was strongly criticized by a National Academies panel (cf., Parnell et al. 2008). That investigation motivated Parnell (2007) to propose a more rigorous decision tree approach that expresses uncertainty through probabilities of terrorist choices. Similarly, Pinker (2007) and von Winterfeldt and O'Sullivan (2006) used, respectively, influence diagrams and decision trees from the point of view of the Defender, who somehow assesses probabilities for the Attacker's actions.

Paté-Cornell and Guikema (2002) addressed the problem of assigning probabilities to terrorist decisions by using separate influence diagrams, one for the Attacker and one for the Defender. They use the influence diagram of the Attacker, with probabilities and utilities of the Attacker from the point of view of the Defender, to assess the expected utilities of the Attacker's actions and to estimate the probabilities of these actions. Specifically, the Defender estimates the Attacker's perceived expected utility of different attacks. These are then renormalized as probabilities and multiplied by a "base rate" that is the probability that each of the possible attacks is undertaken in a unit time period. These then feed the Defender's influence diagram, enabling computation of the estimated optimal countermeasures. However, from a game-theoretical perspective, this approach seems flawed. The renormalization of expected utilities is problematic (e.g., if it is equally easy to steal \$5 or \$10, it seems unreasonable to say there is 1/3 chance that a person will choose to steal the \$5). Also, it is unclear how to combine this approach with intelligence information or resource constraints. And there is no explicit "mirroring" of the thought processes of each adversary about the other, only a one-sided assessment by the Defenders. This asymmetric prescriptive/descriptive approach is in the spirit of Raiffa's (2002) approach to games.

The influence diagram in Figure 5 represents an example of this decision theoretical approach for a Defender-Attacker problem in which the Attacker chooses after observing the Defender's decision. We have converted the Attacker's decision node to a chance node and (implicitly) removed the Attacker's other chance nodes, those not shared with the Defender. Once we have assessed $p_D(a \mid d) \in \mathcal{P}_D$, which are the Defender's beliefs about the Attacker's response given the Defender's actions, we determine the Defender's decision $d^*$ that has the largest expected utility with respect to the uncertainty in both $a$ and the cost:

$$d^* = \operatorname{argmax}_{d \in \mathcal{D}} \sum_{a \in \mathcal{A}} p_D(a \mid d) \int u_D(c) \pi_D(c \mid a, d) dc.$$

A sampling-based approach, as in Banks and Anderson (2006), could be used to solve this. For complex applications one could use augmented probability simulation, as in Bielza, Muller, and Rios Insua (1999), or numerical integration. However, when model uncertainty dominates, it seems faster and more robust to use sampling procedures that automatically focus on the high-probability events.

The main obstacle, as pointed out already in Kadane and Larkey (1982), is the assessment of $p_D(a \mid d)$. Paté-Cornell and Guikema (2002) suggested taking $p_D(a \mid d) \propto \psi_A(a, d)$, so that actions with larger expected utility to the Attacker have larger probability of being selected. But this proposal does not take proper account of the fact that the (idealized) Attacker is an expected utility maximizer and thus would certainly choose the optimal action $a^*(d)$ (a choice that could be divined by the Defender, if the Defender knows the Attacker's utilities and risk analysis). This is why Harsanyi (1982) objected to the decision theory approach as contrary to the spirit of game theory, because the assessment of the adversaries' actions should be based on an analysis of their rational behavior. The standard rebuttal is that Harsanyi assumes full and common knowledge of the game by both players; the use of $p_D(a \mid d)$ is the Defender's way of expressing uncertainty about the utilities and risk analysis of the Attacker. One could think of assessing $p_D(a \mid d)$ by standard probability assessment methods (cf.,

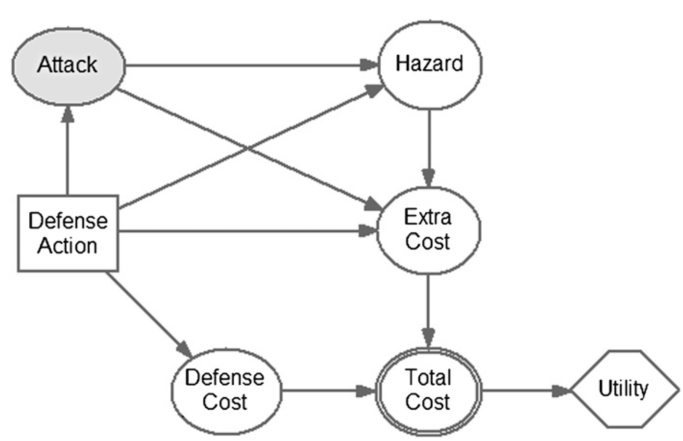


Figure 5. Transformed adversarial risk influence diagram.

O'Hagan et al. 2006), but we prefer methods that take account of the game theoretical structure. Bier et al. (2007) provided extended consideration of these issues.

### 4.4 A Critique of ARA

The ARA framework described in Section 4 accommodates many different views of game theory. But game theory is an unreliable guide to human behavior (Camerer 2003) and, until recently, an unpopular basis for policy decisions (Bier and Cox 2007). This section reviews some main concerns. See also Rothkopf (2007) for a related discussion concerning the role of game theory and decision analysis in auctions.

First, most versions of game theory assume that the opponents not only know their own payoffs, preferences, beliefs, and possible actions, but also those of their opponent. Moreover, when there is uncertainty in the game, it is assumed that players have common probabilities over the uncertain variables. This strong common knowledge assumption allows a symmetric joint normative analysis in which players try to maximize their expected utilities and expect the other players to do the same, and therefore their decisions can be anticipated. It leads to Nash equilibrium concepts in static games (in which both opponents choose their actions simultaneously) and backward induction solutions in dynamic ones (in which the decisions of the opponents alternate in time, or occur on more complicated schedules). However, in real life, players do not typically have full knowledge of their opponent's objectives, beliefs, and possible moves, and this problem is aggravated when participants try to conceal information, as would occur in terrorism or business (cf., Bier et al. 2007).

From a policy standpoint, game theory methods can lead to social dilemmas. Heal and Kunreuther (2006) studied the implementation of security measures in interconnected networks (e.g., baggage screening for airline passengers). In their example, security increases with investments in risk reduction by the network members. However, each member is better off if he defects but the others contribute. A defector contributes nothing but enjoys the benefits of investments by the other network members, so defection is the selfish optimal strategy. However, if everyone defects, the result is worse for each player than if they were all to cooperate. For this reason, third party regulators are needed to impose mechanisms to ensure security investments. But creating such mechanisms is difficult; in complex games there may be multiple Nash equilibria, and the regulators should expect unintended outcomes.

Classical noncooperative game theory assumes that opponents do not communicate. But in adversarial situations, communication allows players to acquire common knowledge and influence the views of the opponent. Communication in policy applications can expand the set of actions and reduce the chance of socially inferior equilibria. On the other hand, communication among corporations allows collusion with competitors for mutual benefit, even when the market is designed to promote competition. There are games in which a player is better off without communication, see Raiffa (2002) (e.g., if an Attacker can choose an action that hurts himself but hurts the opponent more, then the Defender is vulnerable to threats). Communicating threats and bluffs can leverage disproportionate investment in protective strategies by the Defender (an issue that arises in connection with the U.S. response to Al-Qaeda).

Despite these problems, leaders must have good tools to manage risks. ARA needs to advise a player on the best strategic decision given knowledge and beliefs about how their opponents will behave. This information about the opponent's objectives is based on expert knowledge and past behavior. Probability is used to represent uncertainty about adversarial thinking, but the elicitation should rest on a game-theory view that models the opponent's analysis.

## 5. A UNIFIED BAYESIAN FRAMEWORK FOR ARA

This section proposes an analysis that reflects both the elements of ARA and the issues raised in the preceding critique. Our objective here is to provide a unified framework for analysis aimed at prescribing advice to one of the participants. The key issue in our framework is the assessment of the probabilities of adversaries' actions under the assumption that adversaries are expected utility maximizers. Thus, the probabilities on the adversary's possible actions stem from our uncertainty about the adversary's decision problem.

### 5.1 The Game-Theoretical Analysis

We describe first the game-theoretical approach to ARA, through two simple cases. The first one is a simultaneous-move game in which Apollo and Daphne each have just two options, respectively $\mathcal{A} = \{a_1, a_2\}$ and $\mathrm{D} = \{d_1, d_2\}$ and the only uncertainty is a binary outcome $S$ (say, success or failure for Apollo). The utility functions over the costs are $u_A(c_A)$ and $u_D(c_D)$, with costs dependent on actions of both parties. This problem can be represented by a coupled influence diagram for Apollo and Daphne jointly, and also as a decision tree, both shown in Figure 6.

Solving this problem requires probability assessments over the costs, conditional on $(a, d, S)$; and about $S$, conditional on $(a, d)$. Apollo and Daphne may have different assessments: for example, for success, these are $p_A(S = 1 \mid a, d)$ and $p_D(S = 1 \mid a, d)$, respectively. For Apollo, the expected utility obtained with $(a, d)$ is

$$\psi_A(a,d) = p_A(S=0 \mid a,d) \int u_A(c_A)\pi_A(c_A \mid a,d,S=0)dc_A + p_A(S=1 \mid a,d) \int u_A(c_A)\pi_A(c_A \mid a,d,S=1)dc_A \tag{6}$$

and symmetrically for Daphne. The Nash equilibrium solution $(a^*, d^*)$ satisfies

$$\psi_A(a^*,d^*) \geq \psi_A(a,d^*) \ \forall a \in \mathcal{A} \ \text{and} \ \mathbf{\Psi}_D(a^*,d^*) \geq \mathbf{\Psi}_D(a^*,d), \forall d \in \mathcal{D}$$

and this may require randomized strategies.

The second case is a sequential decision game, say a Defend-Attack situation, in which Daphne chooses a defense in $\mathcal{D}$ and then Apollo, having observed the defense, chooses an attack in $\mathcal{A}$. The consequences for both players depend on the success of his attack. The new influence diagram adds an arc to show that Daphne's choice is observed by Apollo, and the decision tree

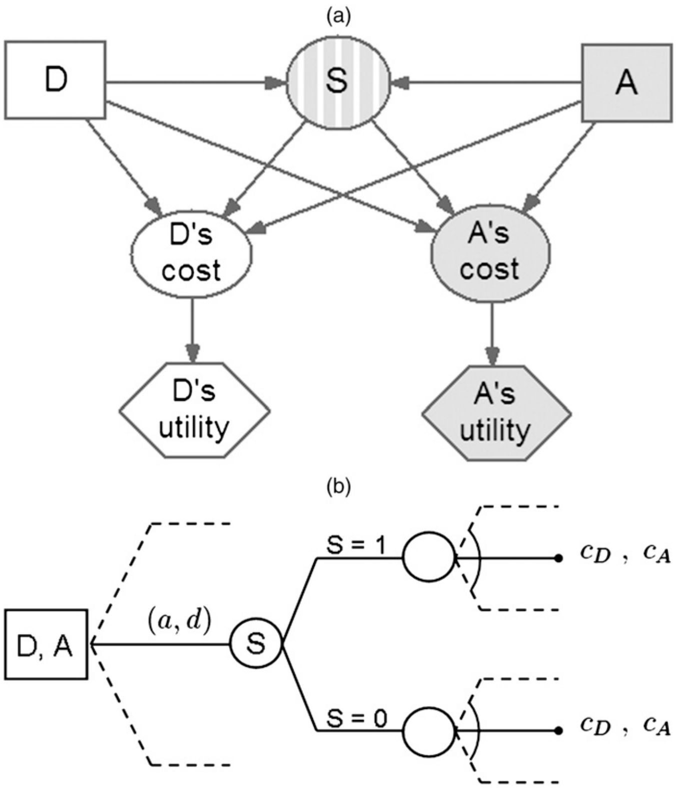


Figure 6. The two player simultaneous decision game.

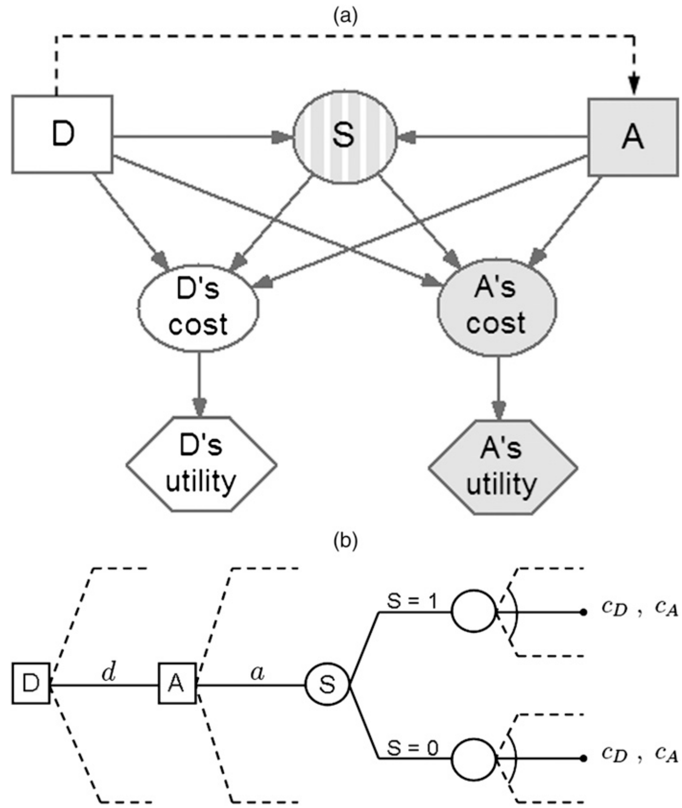


Figure 7. The two-player sequential decision game.

also reflects this sequencing, as shown in Figure 7. The solution does not require Apollo to know Daphne's probabilities and utilities because he observes Daphne's actions, but Daphne needs to model Apollo's. The expected utilities of the players at node $S$ in Figure 7 are computed as in (6) and Apollo's best attack for each of Daphne's defenses is

$$a^*(d) = \text{argmax}_{a \in \mathcal{A}} \psi_A(a, d). \tag{7}$$

Knowing this, Daphne's best defense is

$$d^* = \text{argmax}_{d \in \mathcal{D}} \psi_D(a^*(d), d).$$

The solution ($a^*(d^*)$,$d^*$) is a Nash equilibrium. Daphne has an advantage; she is the first to move and, ceteris paribus, has larger expected utility than in a simultaneous game—if she can accurately model Apollo's $p_A$, $u_A$, and $\pi_A$. Bier et al. (2007) discussed this in detail, pointing out that sometimes disclosing information about one's defenses against terrorism can have deterrent effects.

## 5.2 The Bayesian Analysis

We now weaken the common knowledge assumptions for the sequential decision game. (In Section 6 we shall also weaken it for simultaneous games, using the example of an auction.) Suppose Daphne does not know ($p_A$, $u_A$, $\pi_A$). Daphne's influence diagram, in Figure 8, no longer has the hexagonal value node with Apollo's information. And her decision tree denotes uncertainty about Apollo's attack by replacing $\boxed{A}$ with Ⓐ

To solve this game, Daphne needs more than $p_D(S \mid a, d)$, $\pi_D(c_D \mid a, d, S)$ and $u_D(c_D)$. She also needs $p_D(a \mid d)$, which is her assessment of the probability that Apollo will choose attack $a$ after observing that Daphne has chosen defense $d$. To find $p_D(a \mid d)$, she must estimate Apollo's utility function and his probabilities about both success $S$, conditional on $(a, d)$, and the associated costs $c_A$, conditional on $(a, d, S)$, and consequently compute the required probability. Eliciting these assessments requires Daphne to analyze the problem from Apollo's perspective, and this is one of the most difficult aspects in the Bayesian approach of ARA, as we describe.

First, Daphne must put herself in Apollo's shoes, and mirror his decision problem. Figure 9 represents Apollo's problem (as seen by Daphne). Daphne must analyze Apollo's problem taking into account that he is an expected utility maximizer. Thus, she will use all the information and judgment that she can about Apollo's utilities and probabilities. Instead of using point estimates for $\pi_A$, $p_A$, and $u_A$ to find Apollo's optimal decision $a^*(d)$ as in (7), Daphne's uncertainty about Apollo's decision should derive from her uncertainty about Apollo's ($\pi_A$, $p_A$, $u_A$). Specifically, a Bayesian strategy for expressing this uncertainty puts a distribution $F$ over ($\pi_A$, $p_A$, $u_A$), inducing a distribution on Apollo's expected utility $\psi_A(a, d)$ defined in (6). Thus, assuming Apollo is rational, Daphne finds

$$p_D(a \mid d) = \mathrm{IP}_F[a = \text{argmax}_{x \in \mathcal{A}} \Psi_A(x, d)]$$

where

$$\begin{aligned} \Psi_A(a, d) = {} & P_A(S = 0 \mid a, d) \int U_A(c_A) \Pi_A(c_A \mid a, d, S = 0) dc_A \\ & + P_A(S = 1 \mid a, d) \int U_A(c_A) \Pi_A(c_A \mid a, d, S = 1) dc_A \end{aligned} \tag{8}$$

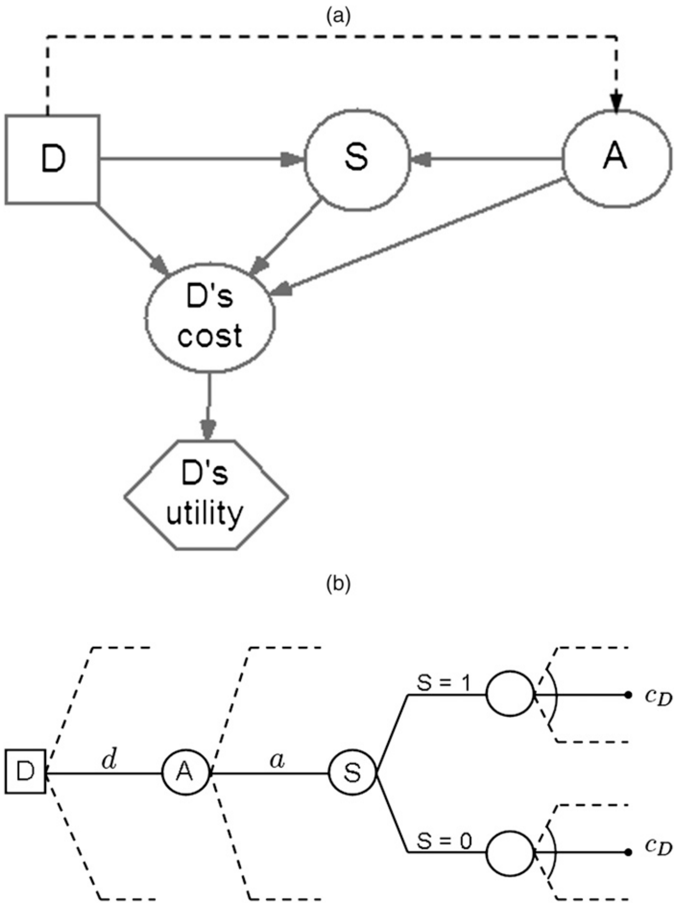


Figure 8. The decision problem as seen by Daphne.

for $(\Pi_A, P_A, U_A) \sim F$. She can use Monte Carlo simulation to approximate $p_D(a \,|\, d)$ by drawing $n$ samples $\left\{\left(\pi_A^i, p_A^i, u_A^i\right)\right\}_{i=1}^n$ from $F$ and setting

$$p_D(a \,|\, d) \approx \frac{\#\{a = \text{argmax}_{x \in \mathcal{A}}\, \psi_A^i(x, d)\}}{n}. \quad (9)$$

Once Daphne has completed these assessments, she can solve the problem. Her expected utilities at node Ⓢ in Figure 8 are

$$\psi_D(a, d) = p_D(S = 0 \,|\, a, d) \int u_D(c_D)\pi_D(c_D \,|\, a, d, S = 0)dc_D + p_D(S = 1 \,|\, a, d) \int u_D(c_D)\pi_D(c_D \,|\, a, d, S = 1)dc_D.$$

Working up the tree, her estimated expected utilities at node Ⓐ are

$$\psi_D(d) = \psi_D(a_1, d)p_D(a_1 \,|\, d) + \psi_D(a_2, d)p_D(a_2 \,|\, d).$$

Finally, her optimal decision is $d^* = \text{argmax}_{d \in \mathcal{D}}\, \psi_D(d)$.

We summarize the previous discussion with the following procedure to find a recommendation for Daphne in the Defend-Attack model.

1. Assess $u_D$, $p_D$, $\pi_D$, from Daphne.
2. Assess $F = (U_A, P_A, \Pi_A)$, describing Daphne's uncertainty about $u_A$, $p_A$, $\pi_A$.
3. For each $d$, simulate to assess $p_D(a \,|\, d)$ as follows:
   (a) generate $(u^i{}_A, p^i{}_A, \pi^i{}_A) \sim F$, $i = 1, \ldots, n$

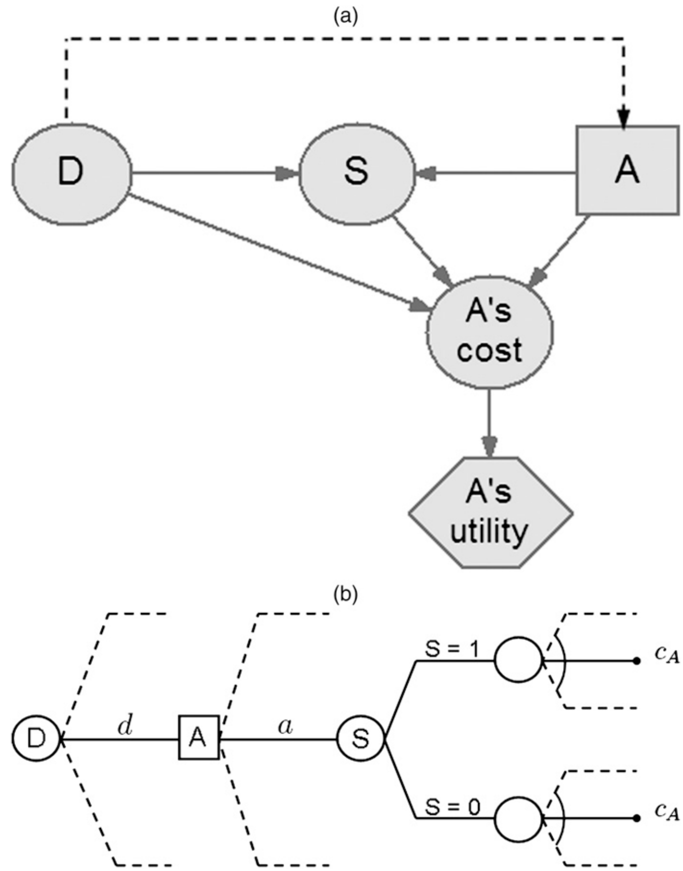


Figure 9. Daphne's analysis of Apollo's problem.

   (b) solve $a_i^*(d) = \text{argmax}_{a \in \mathcal{A}}\, \psi_A^i(a, d)$
   (c) approximate $\hat{p}_D(a \,|\, d) = \#\{a = a_i^*(d)\}/n$
1. solve Daphne's problem

$$d^* = \text{argmax}_{d \in \mathcal{D}}\, \psi_D(a_1, d)\, \hat{p}_D(a_1 \,|\, d) \;+\; \psi_D(a_2, d)\hat{p}_D(a_2 \,|\, d).$$

Note that, in terms of classic game theory, the solution $d^*$ for the sequential game need not correspond to a Nash equilibrium. For example, assume there is a third party who knows Daphne's true $(\pi_D, p_D, u_D)$ and her beliefs $F$ about Apollo's utilities and probabilities, as well as knowing Apollo's true $(\pi_A, p_A, u_A)$ and his beliefs $G$ about Daphne. That party would then be able to predict the game, identifying the decisions chosen by each player. However, this omniscient prediction would not be the Nash equilibrium computed based on the true $(\pi_D, p_D, u_D)$ and $(\pi_A, p_A, u_A)$. Because the players lack full and common knowledge, their choices are unlikely to coincide with those made in the traditional game theory formulation.

## 6. BIDDING IN AUCTIONS

As a simple, realistic, and specific case of ARA, we consider two applications in auctions. The first is nonadversarial but introduces the basic ideas. The second describes the adversarial case. Moreover, it illustrates the application of the ARA framework to a simultaneous decision making problem in which the assessment of probabilities on the adversary's actions needs to be more elaborate than in the sequential decision game from Section 5.

### 6.1 One Sealed Bid, Unknown Reservation Price

Suppose Daphne is bidding for a certain object. She is the only bidder, but the owner has set a secret reservation price $v$ below which the object will not be sold. Daphne does not know $v$, and expresses her uncertainty through a distribution $F(v)$. Daphne's utility function in money is $u_D(\cdot)$ and her personal valuation of the auctioned object is $v_D$. Her choice set is $\mathcal{D} = \mathrm{IR}^+$ and her expected utility for a bid of $d \in \mathcal{D}$ is $u_D(v_D - d)$ $\mathrm{IP}[d > V]$. Thus, by standard decision theory (cf., Raiffa 2002), Dahpne should maximize her expected utility by bidding $d^* = \mathrm{argmax}_{d \in \mathcal{D}} u_D(v_D – d)F(d)$.

### 6.2 Two Sealed Bids, the Highest Bid Wins

Suppose now that Daphne and Apollo are bidding against each other. Each knows their own valuation of the auctioned object but does not know the valuation of the other. Each submits their bid in a sealed envelope without knowing the other's bid, and the winner is the highest bidder. This simultaneous decision-making situation is shown in the influence diagram in Figure 10.

Harsanyi's (1967) approached leads to the solution concept of Bayes'-Nash equilibrium for games with incomplete information, based on the assumption that players share a common prior, which in this case requires that players disclose, inter alia, their true beliefs about the other player's valuation. Thus, Daphne's probabilistic assessment of Apollo's valuation and Apollo's probabilistic assessment of Daphne's valuation would be common knowledge. Only under this assumption it is possible to compute the solution.

The following approach seems more realistic than Harsanyi's. We analyze the game from the perspective of Daphne, who knows her value $v_D$ for the object and has a (secret) probabilistic assessment of Apollo's valuation $v_A \sim V_A$. Daphne has to choose her bid $d$. If this is bigger than Apollo's bid $a$, she wins, obtaining a utility $u_D(v_D – d)$; if it is smaller ($d < a$), she gets 0, as reflected in Figure 11(a). Therefore, as in Section 6.1, the problem she must solve is

$$\max_d u_D(v_D - d)\mathrm{IP}_D(d > \mathrm{A} \mid d).$$

The key issue is assessing Daphne's probability $\mathrm{IP}_D(d > A \mid d)$ of winning for each of her possible bids. This assessment would be based on her prediction about Apollo's bid, represented by the probability distribution $\pi_D(a)$, as $\mathrm{IP}_D(d > A \mid d) = \int_{-\infty}^{d} \pi_D(a)\, da$. She might assess $\pi_D(a)$ from historical data on the behavior of bidders, as in Keefer, Beckley, and Back (1991), or to use a noninformative distribution.

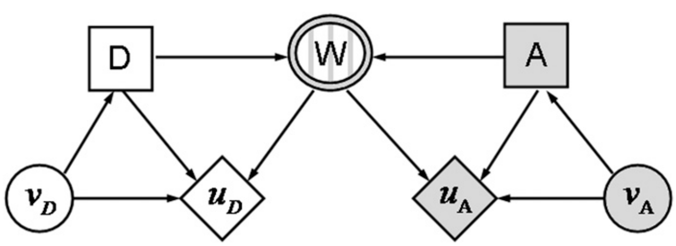


Figure 10. ID of the sealed bid auction problem.

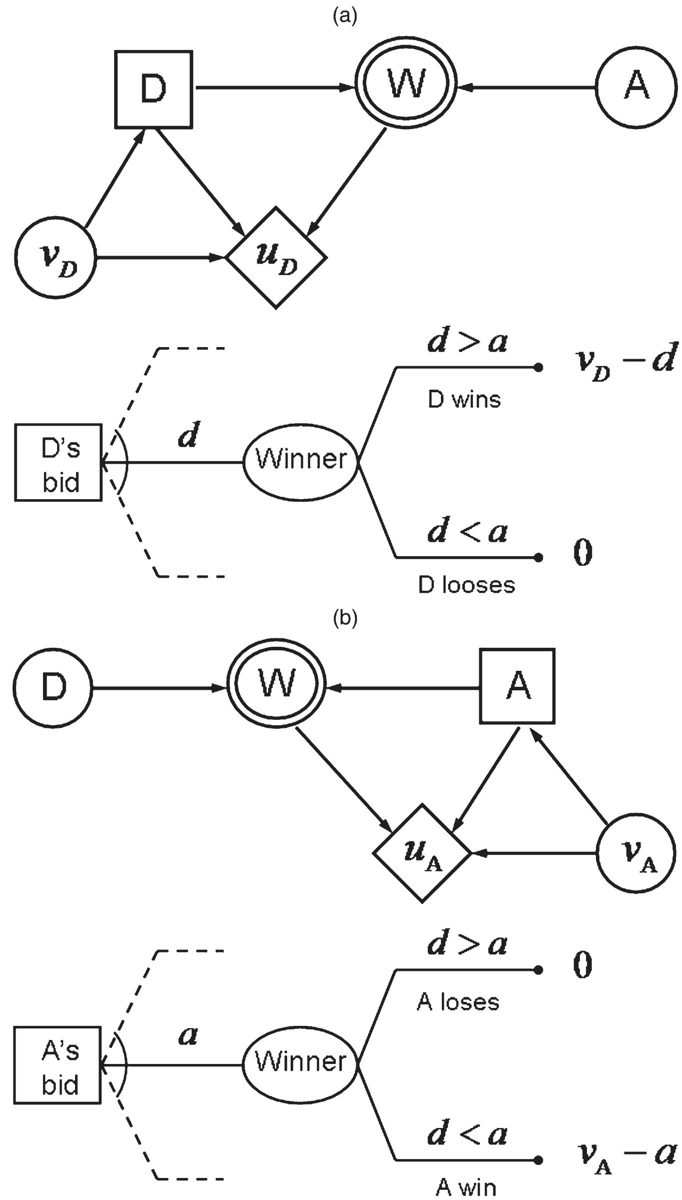


Figure 11. Auction analysis from Daphne's perspective. Influence diagram (top) and decision tree (bottom) representations.

Another possibility is to assess $\pi_D(a)$ judgmentally, as in Section 5.2, through an analysis of Apollo's bidding decision problem from Daphne's perspective, shown in Figure 11(b). To simplify the discussion assume that both Apollo and Daphne are risk neutral in the range of interest, so that their utility functions $u_D$ and $u_A$ are linear. Then, Daphne knows that Apollo would solve $\max_a (v_A – a)\ \mathrm{IP}_A(a > D \mid a)$.

Daphne needs to know Apollo's solution, but she cannot solve this since she does not know $v_A$ and $\mathrm{IP}_A(a > D \mid a)$. But her beliefs about $v_A$ are modeled through $V_A$, and she can elicit her subjective assessment of $\mathrm{IP}_A(a > D \mid a)$ for Apollo's distribution $\pi_A$ on her bid $d$, written as $\mathrm{IP}_A(a{>}D \mid a) = \int_{-\infty}^{a} \pi_A(d)\, \mathrm{d}d$. This assessment might incorporate information from data on Daphne's bids that, she thinks, are available to him, or expert opinion, or both; or she could use a noninformative distribution for $\pi_A(d)$. The detailed ARA analysis can help Daphne to elicit

$\pi_A(d)$ through the identification of variables that may affect Apollo's guess about her bid.

Daphne knows that, symmetrically, Apollo will analyze her problem as in Figure 11(a) and solve

$$\max_d\ (\hat{v}_D - d)\ \mathrm{IP}_D(d > \hat{A} \mid d),$$

where $\hat{v}_D$ represents Apollo's assessment of Daphne's valuation $v_D$. Because Apollo's estimate is unknown to Daphne in her analysis of Apollo's analysis of her problem, the $\hat{v}_D$ is (to Daphne) a random variable $\hat{V}_D$. Similarly, Daphne's assessment of Apollo's bid, as elicited from Apollo's perspective, is the random variable $\hat{A}$. Thus, implicitly, the elicitation of $D \sim \pi_A(d)$ depends on $\hat{V}_D$ and $\hat{A}$, both assessed by Daphne.

If Daphne goes one more step further in the ARA analysis, she finds that, in turn, the assessment of $\hat{A}$ depends on (1) $\hat{V}_A$, the random variable that Daphne thinks Apollo uses to represent Daphne's estimate of his valuation $v_A$, and (2) the distribution of $\hat{D}$, representing what Daphne thinks Apollo thinks about what Daphne thinks Apollo thinks about her bid. To avoid an infinite regress, she stops here and uses a heuristic approach, assessing the distribution $D \sim \pi_A(d)$ based on the relevant $\hat{V}_D$ and $\hat{V}_A$ distributions that have been identified, but disregarding the more complex model that describes a thinking-about-what-the-other-is-thinking-about kind of analysis. That level of detail is difficult to handle and is not a realistic reflection of how people think about auctions.

In this framework, we assume that Apollo expects Daphne to make a bid $d$ that is a function $f(\hat{v}_D, \hat{v}_A)$, but Daphne has some uncertainty about it. We explore the case in which Apollo expects Daphne to bid

$$\min(\alpha\ \hat{v}_D, \beta\ \hat{v}_A), \qquad (10)$$

where $\alpha, \beta \in (0, 1)$, $\hat{v}_D$ is Apollo's estimate of her valuation $v_D$, and $\hat{v}_A$ is Daphne's estimate of Apollo's valuation $v_A$, as assessed by Apollo. The intuitive interpretation of Equation (10) is that Daphne's bidding behavior consists of submitting a bid that takes into account (1) Daphne's profit in terms of a proportion $(1 - \alpha)$ of her valuation of the object $\hat{v}_D$; (2) Apollo's profit in terms of a proportion $(1 - \beta)$ of his object valuation $\hat{v}_A$; and (3) that she will not overbid Apollo's bid. Note that (1) and (2) are consistent with standard linear bidding functions based on available information (cf., McAfee and McMillan 1987 or Aliprantis and Chakrabarti 2000).

Daphne's uncertainty about her assessment of the distributions over $(\alpha, \hat{v}_D, \beta, \hat{v}_A)$, as well as her uncertainty about the accuracy of the proposed heuristic for her analysis of Apollo's problem, could be described through a hierarchical model with a new parameter $\sigma$ that models her certitude; for example, by arguing that $\pi_A(d \mid \alpha, \hat{v}_D, \beta, \hat{v}_A, \sigma)$ has a normal distribution with mean $\min(\alpha\ \hat{v}_D, \beta\ \hat{v}_A)$ and standard deviation $\sigma$, truncated on $[0, \hat{v}_D]$, with $(\alpha, \hat{v}_D, \beta, \hat{v}_A, \sigma) \sim \Pi_A$. As $\sigma$ gets larger, Daphne approaches a noninformative distribution for $\pi_A(d)$.

We have reduced Daphne's ARA about $D$ with probability density function $\pi_A(d)$ to the elicitation of the distribution $\Pi_A$ representing her uncertainty over the unknown parameters. All these quantities can be obtained directly from Daphne through her analysis of Apollo's problem, yielding a specific distribution for $D$. Note that Daphne's distribution $\Pi_A$ might reflect her opinions about (e.g., how Apollo uses his valuation $v_A$ of the object to assess his estimate of $\hat{v}_D$), and, therefore, $\Pi_A$ would depend upon $V_A$. We illustrate this in the following numerical example.

The Monte Carlo solution estimates $\mathrm{IP}_D(d > A \mid d)$ by looping, for $i = 1, \ldots, n$ iterations,

1. draw $v_A^i \sim V_A$;
2. draw $\omega^i = (\alpha^i, \hat{v}_D^i, \beta^i, \hat{v}_A^i, \sigma^i) \sim \Pi_A \mid v_A^i$;
3. set $D_i \mid \omega^i \sim N\big(\min\big(\alpha^i\ \hat{v}_D^i, \beta^i\ \hat{v}_A^i\big), \sigma^i\big)$ truncated on $[0, \hat{v}_D^i]$, with $\pi_A^i(d_i \mid \omega^i)$ its probability density function;
4. solve $a_i^* = \mathrm{argmax}_a (v_A^i - a)\ \mathrm{IP}_{\pi_A^i}(D_i < a \mid a, \omega^i)$, where

$$\mathrm{IP}_{\pi_A^i}(D_i < a \mid a, \omega^i) = \int_{-\infty}^{a} \pi_A^i(d_i \mid \omega^i)\ \mathrm{d}d_i.$$

Then use $\hat{\mathrm{IP}}_D(d > A^* \mid d) = \#\{d > a_i^*\} / n$ as an approximation to Daphne's probability of winning conditional on her bid $d$.

To sum up, Daphne's analysis should be:

1. Find $v_D$, her known valuation of the object.
2. Assess $F = (V_A, \Pi_A)$ describing Daphne's uncertainties in her analysis of Apollo's problem.
3. Estimate (by Monte Carlo as before) $\mathrm{IP}_D(d > A \mid d)$ through $\hat{\mathrm{IP}}_D(d > A^* \mid d)$.
4. Solve $d^* = \arg\max_d\ (v_D - d)\ \hat{\mathrm{IP}}_D(d > A^* \mid d)$.

It is a bit more complicated if the utility functions must also be assessed, but the strategy is straightforward.

Note that the elicitation for Daphne's bid $D$ could be simplified by just asking Daphne for point estimates for the quantities $(\alpha, \hat{v}_D, \beta, \hat{v}_A, \sigma)$, in which case we would not model her certitude in her numerical assessments or in the accuracy of (10) as a heuristic representation of bidding behavior. We illustrate this approach with an example.

*A Numerical Example.* Assume that we are able to obtain from Daphne the following subjective assessments:

$v_D = 100$: The value of the object for her.

$v_A \sim V_A$: Daphne believes that Apollo's object valuation must be in a range between 60 (min) and 90 (max), and most likely is 80 (mode). Based on this information, she fits a triangular distribution.

$\hat{v}_D \sim \hat{V}_D$: This is Apollo's estimate of Daphne's valuation $v_D$, as assessed by Daphne. Daphne believes that Apollo thinks her valuation of the auctioned object is around 100 units higher than his $(v_A)$, with an error range between $-5$ and 5 units. She also believes that all values within this range are equally likely, so $\hat{V}_D \mid v_A = v_A + 100 + U$ where $U$ is uniform on $[-5, 5]$.

$\hat{v}_A \sim \hat{V}_A$: This is Daphne's estimate of $v_A$ as derived by Apollo. Daphne believes that Apollo thinks that she believes that his valuation of the object is 50 units lower than his $(v_A)$, with an error range between $-5$ and 5 units. She also believes that all values within this range are equally likely. Thus $\hat{V}_A \mid v_A = v_A - 50 + V$ where $V$ is uniform on $[-5, 5]$.

$1-\alpha$ and $1-\beta$: These are profit value proportions used by Apollo in his analysis of her bidding problem when he thinks about how she analyzes bidding behavior. Daphne's subjective belief is that these parameters are $1–\alpha = 1–\beta = 0.3 + W$, for $W$ uniform on $[-0.05, 0.05]$.

$\sigma = 1$: This measures Daphne's confidence in the heuristic model and second order parameter assessment.

To solve this problem, for each $i = 1, \ldots, n$, we simulate $v_A^i \sim V_A$ and $\omega^i \sim \Pi_A \mid v_A^i$, set $D_i \mid \omega^i = N\big(\min\big(\alpha^i\,\hat{v}_D^i, \beta^i\,\hat{v}_A^i\big), \sigma^i\big)$, truncated on $[0, \hat{v}_D^i]$, and solve Apollo's optimization problem

$$a_i^* = \operatorname{argmax}_{a\in\mathcal{A}}(v_A^i - a)\ \mathrm{IP}_A(a > D_i \mid a, \omega^i).$$

Figure 12(a) illustrates the optimization problem solved at one of the iterations, including Apollo's expected utility for each of his possible bids along with his optimal bid $a_i^* = 23.4$ at that iteration. After $n = 1000$ iterations we obtain a sample $\{a_i^*,\ i = 1, \ldots, 1000\}$ from $A^* = \operatorname{argmax}_{a\in \mathrm{A}}(V_A - a)\mathrm{IP}_A(a > D \mid a)$, which represents Daphne's predictive distribution on Apollo's bid. Its probability density function $\pi_D(a)$ has been estimated from the sample through a kernel density estimator shown in Figure 12(b). Finally, we solve Daphne's decision problem

$$d^* = \operatorname{argmax}_{d\in\mathcal{D}}(v_D - d)\ \hat{\mathrm{IP}}_D(d > A^* \mid d),$$

where $\hat{\mathrm{IP}}_D(d > A^* \mid d)$ is the proportion of simulated $a_i^*$ below a given $d : \#\{d > a_i^*\}/1000$.

Figure 12(c) plots the Monte Carlo estimate of Daphne's expected utility for each of her possible bids. Her maximum expected utility bid is $d^* = 30.6$ with (estimated) expected utility of 68.5.

Should Daphne be unable to provide the information needed to assess a density $\pi_A(d)$ for $D$, or if she is not confident with the heuristic avoidance of the infinite regression, then she could use a noninformative prior to describe the distribution of $D$. We now run the same numerical example as before but with $\sigma = 100$, an approximation of the noninformative case. Figure 13 summarizes the results, parallel to the previous description. We note that when $\pi_A^i(d_i \mid \omega^i) \approx U(0, \hat{v}_D^i)$ at each iteration $i$. Therefore, Apollo's expected utility for each of his possible bids $a$ is $(v_A^i - a)\,a\,/\,\hat{v}_D^i$, corresponding to the parabola shown in Figure 13(a). This result and, consequently, Daphne's maximum expected utility bid, are independent of $\hat{V}_A$, $\alpha$ and $\beta$. This is appropriate when a noninformative distribution is used.

## 7. DISCUSSION

We have reviewed key approaches in adversarial risk analysis (ARA), a new research area that is becoming prominent in the context of terrorism, but which also arises in bidding and corporate competition. ARA entails both statistical problems in risk analysis and a game-theoretical perspective in anticipating the actions of one's opponents. Additionally, we have described a Bayesian approach to the ARA problem, drawing upon previous work by Kadane and Larkey (1982) and Raiffa (1982, 2002), and modifying influence diagrams to express decision processes for multiple agents.

Real problems are extremely complex. For example, counterterrorism involves thousands of possible decisions, and there are large uncertainties associated with the goals and resources

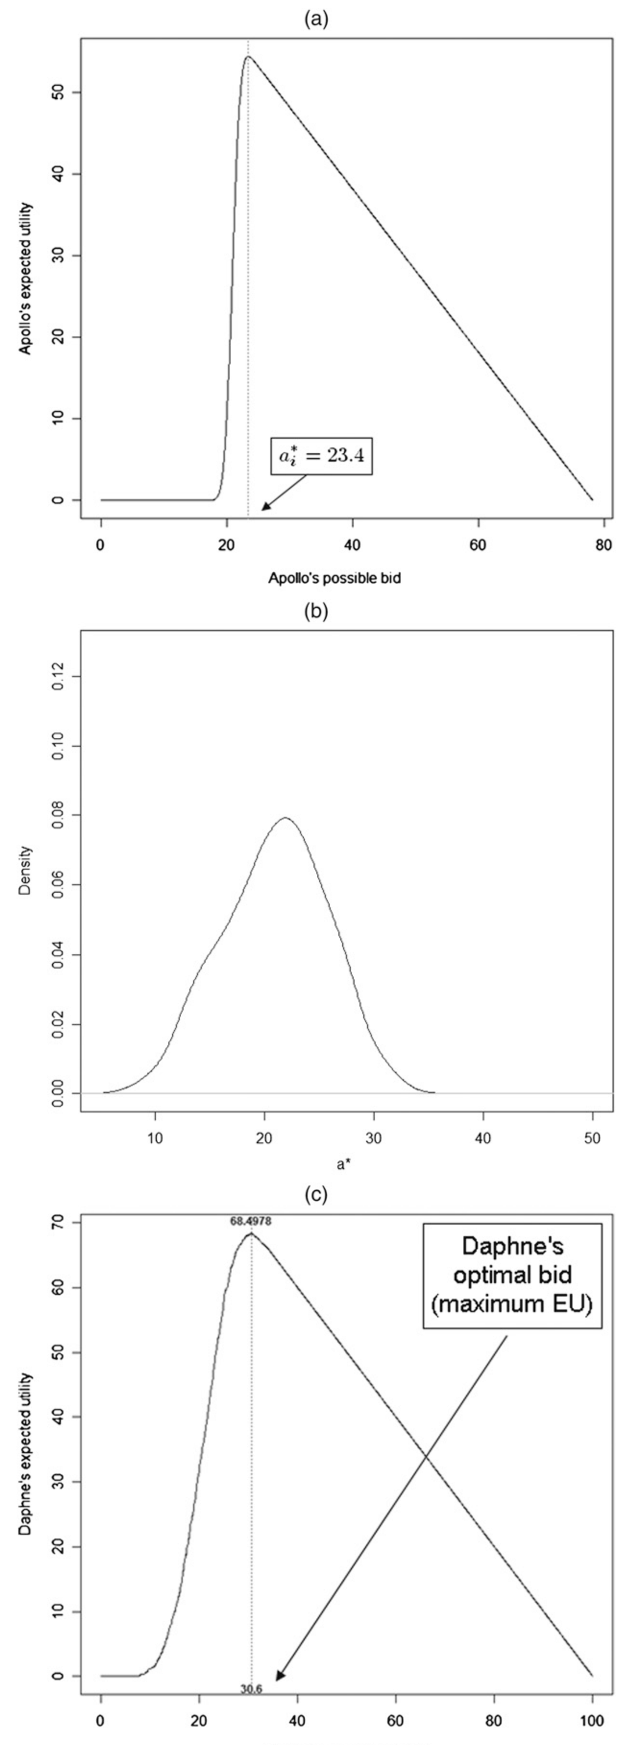


Figure 12. Numerical example. (a) Solving Apollo's optimization problem at the $i$th iteration, (b) Kernel density estimation of $\pi_D(a)$, and (c) Daphne's expected utility function.

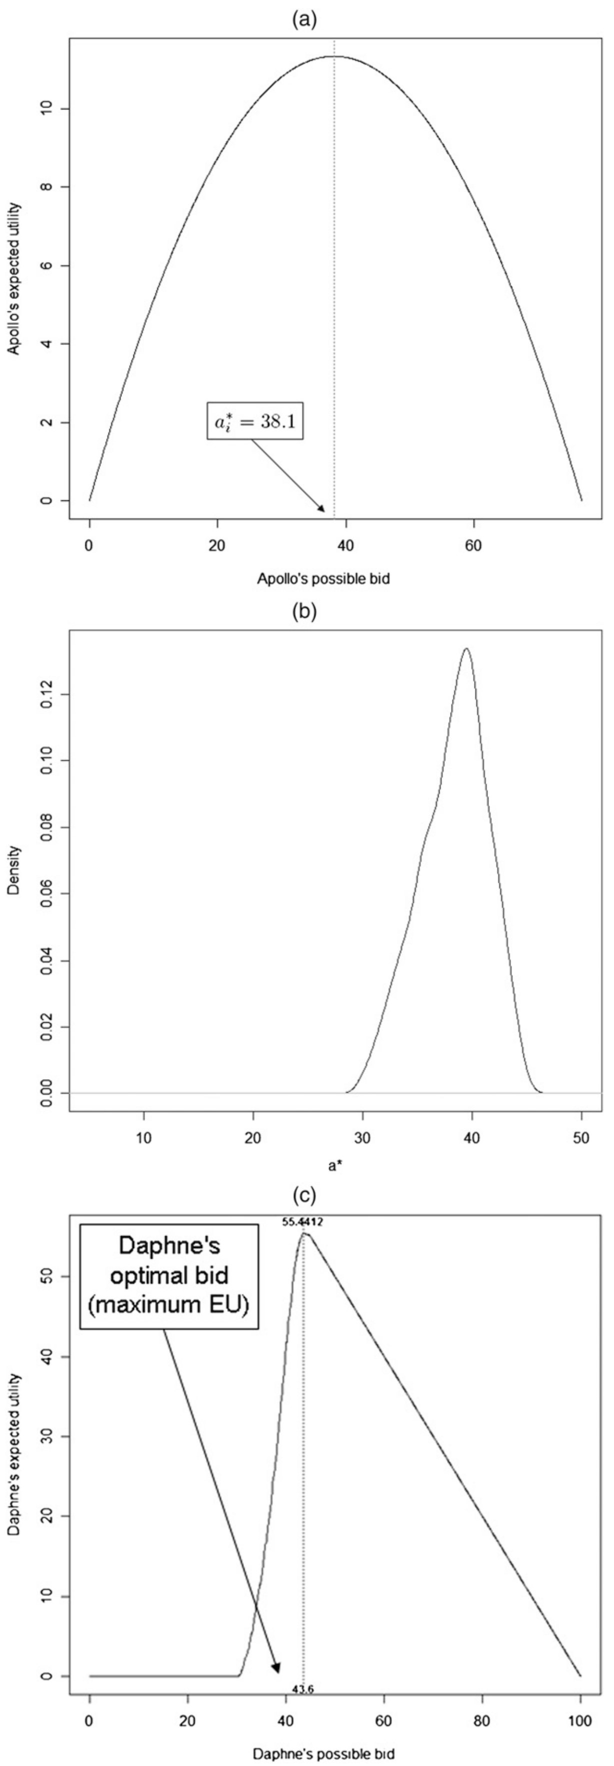


Figure 13. Noninformative case ($\sigma = 100$). (a) Solving Apollo's optimization problem at the $i$th iteration, (b) Kernel density estimation of $\pi_D(a)$, and (c) Daphne's expected utility function.

of the terrorists. For this reason, we have focused on developing a general formulation of the ARA problem and making comparisons among the proposed solution strategies, while keeping attention on the scalability of the analyses.

This article focused on two-person games, but the discussion directly extends to $n$-person games. We did not explicitly discuss cooperative games, but the formulation is general enough so that all previous results apply. However, when there are $n$ players and cooperation is allowed, the number of possible actions increases combinatorially, quickly posing computational challenges.

We briefly mentioned problems with resource constraints of various kinds. This leads directly to the analysis of portfolios of decisions (e.g., wise investment in counterterrorism requires tuning investments in many different, partially interacting and overlapping defensive measures). The models and solution concepts would remain the same, but the implementation would have to be very different. One cannot address the problem by considering action sets that consist of every possible funding allocation.

A second kind of problem is that game theory solutions are generally not decomposable. For example, if different federal agencies each use game-theoretical thinking to make counterterrorism investments, they are extremely unlikely to select the solutions that would result from a higher-level analysis that took the interests of all of the agencies into account together. Concretely, the Transportation Security Administration can invest in training and personnel and protocols, but they cannot invest in counterintelligence, but if the Central Intelligence Agency (CIA) were to invest more effectively in counterintelligence, it might obviate the need for some TSA expenditures. This article does not address decentralized decision making, but it is often an important issue in practice and deserves more attention.

A third issue is the timing of the moves of the opponents. The article carefully treats several kinds of staging, but is not exhaustive. We focus on simultaneous play, and sketch the formulation of the Defend-Attack model. The strategy for strictly alternating games is a simple extension, but the harder and more realistic problem concerns asynchronous play, coupled with the unexpected emergence of new information. This would require dynamic programming methods, and is likely to be intractable. Alternatively, the case of continual evolution over time might be formulated in terms of systems of stochastic differential equations (cf., Yeung and Petrosyan 2005). This could include the use of negotiation analysis to enlarge the set of actions and enable better outcomes for all; it also opens the door to the study of misinformation. The issue of multiobjective decision-making was not addressed, except insofar as these goals can be tied together in a multiattribute utility function.

ARA is a new branch of collaborative statistics. The problems are numerous and the applications are important. We believe that the Bayesian perspective has important contributions to make in this arena, and that our formulation is more realistic than the traditional Nash equilibrium analysis in Operations Research or the ad hoc decisions that are commonly made in practice by federal agencies and corporate executives.

## DISCLAIMER

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

*[Received May 2008. Revised November 2008.]*

## REFERENCES


Aliprantis, C., and Chakrabarti, S. (2000), *Games and Decision Making*, New York: Oxford University Press.

Arce, D., and Sandler, T. (2007), ''Terrorist Signalling and the Value of Intelligence,'' *British Journal of 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, Springer, New York, pp. 9–22.

Bedford, T., and Cooke, R. (2001), *Probabilistic Risk Analysis: Foundations and Methods*, Cambridge, U.K: Cambridge University Press.

Bielza, C., Muller, P., and Rios Insua, D. (1999), ''Decision Analysis by Augmented Probability Simulation,'' *Management Science,* 45, 995–1008.

Bier, V., and Cox, L. A., Jr. (2007), ''Probabilistic Risk Analysis for Engineered Systems. in *Advances in Decision Analysis*, eds. W. Edwards, D. von Winterfeldt, and R. Mles, New York: Cambridge University Press, pp. 279–301.

Bier, V., Oliveros, S., and Samuelson, L. (2007), ''Choosing What to Protect,'' *Journal of Economic Literature,* 9, 563–587.

Brown, G., Carlyle, M., Salmeron, J., and Wood, K. (2006), ''Defending Critical Infrastructure,'' *Interfaces,* 36, 530–544.

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

Camerer, C. (2003), *Behavioral Game Theory*, Princeton, NJ: Princeton University Press.

French, S., and Rios Insua, D. (2000), *Statistical Decision Theory*, London: Hodder Arnold Publishing.

Gibbons, R. (1992), *A Primer in Game Theory*, Harlow: Pearson Education Ltd.

Haimes, Y. (1999), *Risk Modeling, Assessment, and Management*, New York: Wiley.

Harsanyi, J. (1967), ''Games With Incomplete Information Played by ''Bayesian'' Players, I-III. Part I. The Basic Model,'' *Management Science,* 14, 159–182.

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.

Heal, G., and Kunreuther, H. (2006), ''You Can Only Die Once: Interdependent Security in an Uncertain World. in *The Economic Impacts of Terrorist Attacks*, eds. H. W. Richardson, P. Gordon, and J. E. Moore III, Northampton, MA: Edward Elgar Publishers.

Heyes, A. (2000), ''Implementing Environmental Regulation: Enforcement and Compliance,'' *Journal of Regulatory Economics,* 17, 107–129.

Kadane, J. B. (1993), ''Several Bayesians: A Review,'' *Test,* 2, 1–32.

Kadane, J. B., and Larkey, P. D. (1982), ''Subjective Probability and the Theory of Games,'' *Management Science,* 28, 113–120, Rreply: p. 124.

Kaplan, S., and Garrick, J. (1981), ''On a Qualitative Definition of Risk,'' *Risk Analysis,* 1, 11–27.

Kardes, E. (2005), ''Robust Stochastic Games and Applications to Counter-Terrorism Strategies, CREATE report.

Kardes, E. and R. Hall (2005), ''Survey of Literature on Strategic Decision Making in Presence of Adversaries, CREATE report.

Keefer, D., Beckley, F., and Back, H. (1991), ''Development and Use of a Modeling System to Aid a Major Oil Company in Allocating Bidding Capital,'' *Operations Research,* 39, 28–41.

Killourhy, K., Maxion, R., and Tan, K. (2004), "A Defense-Centric Taxonomy Based on Attack Manifestations," in *International Conference on Dependable Systems & Networks* (DSN-04), Los Alamitos, CA: IEEE Computer Society Press, pp. 102–111.

Lemke, C. E., and Howson, J. T., Jr. (1964), "Equilibrium Points of Bimatrix Games," *Journal of the Society for Industrial and Applied Mathematics,* 12, 413–423.

McAfee, R., and McMillan, J. (1996), ''Auctions and Bidding,'' *Journal of Economic Literature,* 25, 699–738.

McAfee, R. P., and McMillan, J. (1996), ''Competition and Game Theory,'' *Journal of Marketing Research,* 33, 263–267.

Myerson, R. (1991), *Game Theory: Analysis of Conflict*, Cambridge, MA: Harvard University Press.

O'Hagan, A., Buck, C., Daneshkhah, A., Eiser, J., Garthwaite, P., Jenkinson, D., Oakley, J., and Rakow, T. (2006), *Uncertain Judgements: Eliciting Experts' Probabilities*, New York: Wiley.

Parnell, G. (2007), ''Multi-Objective Decision Analysis. in *Handbook of Science and Technology for Homelan Security*, ed. Voeller, New York: Wiley.

Parnell, G., Banks, D., Borio, L., Brown, G., Cox, L. A., Gannon, J., Harvill, E., Kunreuther, H., Morse, S., Pappaioanou, M., Pollack, S., Singpurwalla, N., and Wilson, A. (2008), *Report on Methodological Improvements to the Department of Homeland Security's Biological Agent Risk Analysis*, Washington DC: National Academies Press.

Paté-Cornell, E., and Guikema, S. (2002), ''Probabilistic Modeling or Terrorist Threats: A Systematic Analysis Approach to Setting Priorities among Countermeasures,'' *Military Operations Research,* 7, 5–23.

Pearl, J. (2005), ''Influence Diagrams - Historical and Personal Perspectives,'' *Decision Analysis,* 2, 232–234.

Pinker, E. J. (2007), ''An Analysis of Short-Term Responses to Threats of Terrorism,'' *Management Science,* 53, 865–880.

Powell, R. (2007), "Allocating Defensive Resources With Private Information About Vulnerability," *The American Political Science Review,* 101, 799–809.

Raiffa, H. (1982), *The Art and Science of Negotiation*, Cambridge, MA: Harvard University Press.

Raiffa, H. (2002), *Negotiation Analysis*, Cambridge, MA: Harvard University Press.

Rios Insua, D., and Ruggeri, F. (2000), *Robust Bayesian Analysis*, New York: Springer-Verlag.

Rothkopf, M. (2007), ''Decision Analysis: The Right Tool for Auctions,'' *Decision Analysis,* 4, 167–172.

Singpurwalla, N. (2006), *Reliability and Risk*, Chichester, UK: Wiley.

Siqueira, K., and Sandler, T. (2006), ''Terrorists vs the Government: Strategic Intervention, Support and Sponsorship,'' *The Journal of Conflict Resolution,* 50, 878–898.

von Winterfeldt, D., and O'Sullivan, T. M. (2006), ''Should We Protect Commercial Airplanes Against Surface-to-Air Missile Attacks by Terrorists,'' *Decision Analysis,* 3, 63–75.

Yeung, D., and Petrosyan, L. (2005), *Cooperative Stochastic Differential Games*, New York: Springer-Verlag.

Zhuang, J., and Bier, V. (2007), ''Balancing Terrorism and Natural Disasters. Defensive Strategy with Endogenous Attack Effort,'' *Operations Research,* 55, 976–991.