Adversarial Risk Analysis
Auctions, Poker, and Terrorism

David Banks
Duke University

1. Introduction

Classical game theory has focused upon situations in which outcomes are known. When uncertainty is addressed, it makes unreasonable assumptions about common knowledge (cf. Harsanyi, 1967/68a,b). Also, game theory makes unreasonable assumptions about human decision-making (Camerer, 2003).

Classical risk analysis has focused upon situations in which the hazards arise at random. This is appropriate for accident and life insurance, but it does not apply when hazards result from the actions of an intelligent adversary.

Corporate competition, federal regulation, and counterterrorism all entail game-theoretic problems with uncertain outcomes and partial information about the goals and actions of the opponents. This talk describes a Bayesian approach to adversarial risk analysis. It extends the method of Kadane and Larkey (1982) through the use of a **mirroring argument**.

2. Auctions

Suppose Apollo is bidding for a first edition of the Theory of Games and Economic Behavior. He is the only bidder, but the owner has set a secret reservation price v^* below which the book will not be sold. Apollo does not know v^*, and expresses his uncertainty as a subjective Bayesian distribution $F(v)$.

Apollo's utility function is linear in money and his personal valuation of the book is a^*. If money is infinitely divisible, his choice set is $\mathcal{A} = \mathbb{R}^+$. so his expected utility from a bid of a is $(a^* - a)\mathbb{P}[a > V^*]$. Thus Apollo should maximize his expected utility by bidding

$$a_0 = \mathbf{argmax}_{a \in \mathbb{R}^+}(a^* - a)F(a).$$

This is a standard approach in Bayesian auction theory (cf. Raiffa, 2002).

Now suppose that Apollo and Daphne are bidding against each other to own the first edition. Apollo needs to perform a game-theoretic calculation to find his subjective distribution F over Daphne's bid D_0. Then Apollo can maximize his expected utility by bidding $a_0 = \text{argmax}_{a \in \mathbb{R}^+}(a^* - a)F(a)$.

In order to find F, Apollo uses the fact that Daphne must make the symmetric calculation. This is the mirroring argument.

Specifically, suppose Daphne values the book at d^* and has distribution G on Apollo's bid a_0. Then Daphne would solve $d_0 = \text{argmax}_{d \in \mathbb{R}^+}(d^* - d)G(d)$; and symmetrically, to obtain $G(d)$, Daphne would need to mirror Apollo's calculation.

But Apollo cannot duplicate Daphne's calculation since he does not know her value for the book, nor the value she thinks Apollo puts on the book, nor the value she thinks Apollo believes is her value for the book. As a Bayesian, Apollo must express his uncertainty on all three quantities through distributions.

The notation becomes complicated; the following key is helpful:

- a^* is Apollo's value for the book
- D^* is Daphne's value for the book; since it is unknown to Apollo, he assigns it the distribution H_D
- A^* is the random variable that Apollo thinks Daphne uses to represent Apollo's value for the book; it has distribution H_A
- F is Apollo's belief about the distribution of Daphne's bid.
- D_0 is Daphne's bid
- G is Apollo's inference about Daphne's distribution on Apollo's bid.
- A_0 is Apollo's bid from Daphne's perspective.

These probabilities are all belong to Apollo; he imputes the beliefs that Daphne holds. If he is mistaken, he diminishes his chance of maximizing his gain.

To determine his bid a_0, Apollo needs F, the distribution of Daphne's bid. He knows that Daphne's bid D_0 should satisfy $D_0 = \text{argmax}_{d \in \mathbb{R}^+}(D^* - d)G(d)$ where D^* is Daphne's value (a random variable, to Apollo) for the book and $G(d)$ is Apollo's estimate of Daphne's probability that a bid of d exceeds Apollo's bid A_0.

And, to Daphne, $A_0 = \text{argmax}_{d \in \mathbb{R}^+}(A^* - a)F(a)$ where A^* is Daphne's belief about Apollo's value for the book and $F(a)$ is Apollo's estimate of Daphne's probability that a bid of a exceeds her bid D_0. Thus $D_0 \sim F$ and $A_0 \sim G$.

Apollo must find his personal belief about F by solving:

$$\begin{aligned} \mathbf{argmax}_{d \in \mathbb{R}^+}(D^* - d)G(d) &\sim F \\ \mathbf{argmax}_{a \in \mathbb{R}^+}(A^* - a)F(a) &\sim G. \end{aligned}$$

The distributions for D^* and A^* are H_D and H_A, respectively.

Once Apollo has F, he solves $a_0 = \text{argmax}_{a \in \mathbb{R}^+}(a^* - a)F(a)$ to determine his bid.

To solve this system of equations, one iteratively alternates between the two equations until convergence:

1. Select F_0 and G_0 arbitrarily.
2. Simulate a large number of samples from H_A, and solve the argmax problem under G_i. The distribution of those solutions gives F_{i+1}.
3. Simulate a large number of samples from H_D, and solve the argmax problem under F_{i+1}. The distribution of those solutions gives G_{i+1}.
4. If some convergence threshold δ is satisfied (e.g., $\|F_i - F_{i+1}\| < \delta$ and $\|G_i - G_{i+1}\| < \delta$), then stop. Otherwise, return to step 2.

In simulation, this iterative solution has always converged. But one wants a fixed-point theorem, and the key issue is to show this iteration is a contraction operator. For a finite dimensional space (roughly corresponding to bids in pennies, rather than infinitely divisible money), I think this can be done in terms of Gauss-Siedel systems of equations.

This framework allows Apollo to incorporate secret information.

For example, suppose Apollo alone knows that the book was owned by Sir Ronald Fisher, with annotations in his hand. In that case, his personal value a^* is high, but his distribution for Daphne's value, H_D, will concentrate on much smaller values.

Similarly, he might know that Daphne knows the provenance of the book but thinks that Daphne believes, falsely, that Apollo does not. In that case H_D will give concentrate on large values, but Apollo's belief about what Daphne thinks is his value for the book, H_A, will concentrate on small values.

In principle, one could go into an infinite regress:

Apollo thinks that Daphne thinks that

Apollo thinks that Daphne thinks that . . .

But for human reasoning, it is probably quite reasonable to stop at the third step, with the distribution H_A for A^*, as described in the mirroring analysis.

3. Gambling: A Prmitive Poker

Pokeresque games have received considerable attention in the game theory literature. Early work by von Neumann and Morgenstern (1947) and Borel (1938) developed solutions under various simplifying assumptions. More recently, Ferguson and Ferguson (2008) provide approximate analyses pertinent to more complex games, such as Texas hold'em.

In the following, we assume that Bart and Lisa play a game in which each secretly and independently draws a $\mathcal{U}[0,1]$ random number. Each must ante an amount $a = 1$. First, Bart examines his number X and decides whether to bet b or fold. Then Lisa examines her Y and decides whether to bet b or fold. If both players bet, they compare their draws to determine who wins the pot. Otherwise, the first person to fold forfeits his or her ante.

Let V_x be the amount Bart wins. The table shows the possible outcomes:

V_x	**Bart's Decision**	**Lisa's Decision**	**Outcome**
-1	**fold**		
1	**bet**	**fold**	
1+b	**bet**	**bet**	$X > Y$
-(1+b)	**bet**	**bet**	$X < Y$

From the table, the expected amount won by Bart, given his draw $X = x$, is:

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

Bart must use mirroring to find a subjective distribution for the probabilities, based on the adversarial analysis he expects Lisa to perform.

Assume that Bart uses a "bluffing function" $g(x)$; given x, he bets with probability $g(x)$. Then

$$\begin{aligned}\mathbb{E}[V_x] &= -[1-g(x)] + g(x)\mathbb{P}[\text{ Lisa folds } | \text{ Bart bets }] \\ &\quad +(1+b)g(x)x\mathbb{P}[\text{ Lisa bets } | \text{ Bart bets }] \\ &\quad -(1+b)g(x)(1-x)\mathbb{P}[\text{ Lisa bets } | \text{ Bart bets }].\end{aligned}$$

For optimal play, Bart needs to find $\mathbb{P}[$ Lisa bets $|$ Bart bets $]$.

So Bart must "mirror" the thinking that Lisa will perform in deciding whether to bet. He knows that Lisa's opinion about X is updated by the knowledge that Bart decided to bet. Further, suppose Bart has a subjective belief that Lisa thinks that his bluffing function is $\tilde{g}(x)$. In that case, Lisa should calculate the conditional density of X, given that Bart decided to bet, as

$$\tilde{f}(x) = \frac{\tilde{g}(x)}{\int \tilde{g}(z)\,dz}.$$

Note: 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 always bets if it is greater), then the posterior distribution on X is truncated below the X value corresponding to x_0 and the weight is reallocated proportionally to values above x_0.

From this analysis, Bart believes that Lisa calculates her probability of winning as $\mathbb{P}[X \leq y|$ Bart bet $] = \tilde{F}(y)$, where $Y = y$ is unknown to Bart. And thus Bart believes that Lisa will bet if the expected value of her return V_y from betting b is greater than the loss of a that results from folding; i.e., Lisa would bet if

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

So Bart believes Lisa will bet if and only if $\tilde{F}(y) \geq b/2(1+b)$.

Set $\tilde{y} = \inf\{y : \tilde{F}(y) \geq b/2(1+b)\}$. The probability that Lisa has drawn $Y > \tilde{y}$ is $1-\tilde{y}$ and this is the probability that she bets. Letting $q = 1-\tilde{y}$, the expected value of the game for Bart, given $X = x$, is:

$$\mathbb{E}[V_x] = -[1-g(x)] + g(x)(1-q) + (1+b)g(x)xq - (1+b)g(x)(1-x)q.$$

Thus Bart should choose $g(x)$ to maximize $\mathbb{E}[V_x]$.

Bart's expected value has the form $-1 + cg(x)$, where

$$c = 2 - 2q + 2qx + 2qxb - qb.$$

To maximize the expectation, Bart should make $g(x)$ as small as possible when c is negative (i.e., $g(x) = 0$), but as large as possible when c is positive (i.e., $g(x) = 1$). Thus the optimal $g(x)$ is a step function. It implies that Bart should never bluff, no matter what he believes about the playing strategy used by Lisa.

To find $\tilde{x}$, the point at which $g(x)$ jumps from 0 to 1, Bart must solve $c = 0$ to get

$$\tilde{x} = \frac{qb + 2q - 2}{2qb + 2q}$$

As a sanity check, if $b = 0$ then Lisa should always bet. So $q = 1$, thus $\tilde{x} = 0$, properly implying that Bart also always bets.

The expected value of the game, to Bart, is

$$\int_0^1 E[V_x]\, dx = -\tilde{x} + \int_{\tilde{x}}^1 (1 - q) + (1 + b)qx\, dx.$$

The framework discussed is common to analyses by Borel (1938) and Bellman and Blackwell (1949). The von Neumann and Morgenstern analysis allows players to check, and so it is not directly comparable.

	optimal strategy	**game value (Bart)**
Borel	**bet if** $x > \left(\frac{b}{b+2}\right)^2$	$-\left(\frac{b}{b+2}\right)^2$
B & B	**bet if** $x > 1 - \frac{2}{2+b}$	$-\left(\frac{b}{b+2}\right)^2$
	else bluff with probability $\frac{2}{(2+b)^2}$	
Bayesian	**bet if** $x > 1 - \frac{qb+2}{2(qb+q)}$	**complex**

Bart to exploit weaknesses he thinks exist in Lisa's strategy. If Bart is (approximately) correct about Lisa's belief regarding his bluffing function, then Bart can improve his value of the game—unless Lisa thinks he is playing optimally, as per Bellman and Blackwell, in which case he cannot improve their solution.

Such exploitation is standard when humans play poker, but it is unavailable from traditional game theory.

4. Counterterrorism

Counterterrorism clearly entails adversarial risk analysis. But it has not been properly addressed by DHS. Parnell et al. (2008) criticize the decision theory used by DHS in a NRC review—specifically, the DHS analyses assumed that terrorists behave randomly, like weather, rather than purposively.

Counterterrorism applications are hard. There are serious issues of scalability, decentralized and tiered decision-making, and resource constraints that require portfolio analysis.

It appears that traditional game theory is inadequate and/or unrealistic. A Bayesian approach may give better results. It also is better able to incorporate judgment, exploit soft counterintelligence, and use historical data.

4.1 The Smallpox Decision

In 2002, U.S. policy-makers were concerned that terrorists might launch a bioterrorist attack with smallpox. They considered three scenarios:

- A major attack, involving multiple cities and weaponized virus.
- A minor attack, similar to the anthrax letters.
- No smallpox attack.

And the kinds of defenses under consideration were:

- Stockpiling vaccine
- Stockpiling vaccine and increasing biosurveillance
- Stockpile, increase biosurveillance, inoculate all first responders
- Inoculate essentially everyone.

To explore these possible defenses, the U.S. held a series of public meetings and sought scientific advice. D. A. Henderson went around the country giving lectures, analysts weighed intelligence, political support was calculated, op-ed pieces got written, and so forth.

The CDC recommended against widespread inoculation. The vaccine kills about 1-3 people per million, and about 3 people per 100,000 undergo extensive hospitalization. About 30% of recipients miss the next day of work.

But President Bush and other political leaders wanted to buy and use massive amounts of vaccine.

In this setting we describe a tabletop exploration of the game theory and risk analysis approach. The analysis was performed at the FDA, but not disclosed at the time. A description was given later by Banks and Anderson (2006).

In game-theoretic terms, the payoff matrix for this problem is shown below, where C_{ij} represents the cost (or benefit) to the U.S., and D_{ij} represents the corresponding cost to the terrorists.

	No Attack	**Minor Attack**	**Major Attack**
Stockpile	(C_{11}, D_{11})	(C_{12}, D_{12})	(C_{13}, D_{13})
Biosurveillance	(C_{21}, D_{12})	(C_{22}, D_{22})	(C_{23}, D_{23})
First Responders	(C_{31}, D_{31})	(C_{32}, D_{32})	(C_{33}, D_{33})
Mass Inoculation	(C_{41}, D_{41})	(C_{42}, D_{42})	(C_{43}, D_{43})

Note: Ideally, the option of not even stockpiling vaccine could have been part of this table. However, FDA management ruled against that exploration.

A classical game theory person would use the minimax theorem to find the optimal play for U.S. policy-makers. But this overlooks many problems.

Some of the problematic assumptions of classic game theory include:

- **Perfect knowledge of the payoff matrix entries.** But this might be improved by risk analysis.
- **A zero-sum game.** For this application, the zero-sum formulation is probably reasonably accurate.
- **Rational minimaxing opponents.** Kahnemann and Tversky (1972) show that human decision-making is not rational. Bayesian versions of game theory may apply better.
- **One-time games with simultaneous moves.** The matrix represents a normal-form version of game theory; the extensive form applies to sequential alternating moves. Dynamic programming is probably a better tool.
- **Sufficient resources for all attacks.** Some of the options are much more expensive than others. Additionally, terrorists must have regular visible successes to ensure a continual supply of committed volunteers. These circumstances suggest that portfolio analysis should be used.

The costs in the payoff table are unknown. Risk analysis can find reasonable distributions for these random costs. Consider C_{11} in the top-left cell. It represents the cost to the U.S. of stockpiling vaccine when, in fact, no smallpox attack is made.

$$C_{11} = \textbf{cost to test diluted Dryvax} + \textbf{cost to test Aventis vaccine} + \textbf{cost to make 209 x 106 doses} + \textbf{cost to produce VIG} + \textbf{logistic/storage/device costs.}$$

The Dryvax and Aventis vaccines had been kept in storage for years. Their residual potency was unclear; testing was needed. Two experts pooled their opinions and decided that the testing cost might be uniformly distributed between \$2 million and \$5 million.

The FDA said that new vaccine production was fixed by contract at \$512 million. The VIG cost might be N(\$100 million, \$20 million^2). Logistics costs might be N(\$940 million, \$100 million^2).

Other terms that experts assessed include:

- The value of a human life was treated as fixed, at $750,000. (This follows the DOT human capital model; non-market methods tend to give higher values.)
- The number of key personnel to be inoculated: this was guessed to be uniform between .5 million and .6 million.
- The number of smallpox cases per attack: this was guessed to be gamma with mean 10 and sd 100.
- The cost to treat one smallpox case: normal with mean $200,000 and sd $50,000.
- The economic costs of an attack: gamma with mean $5 billion, sd $10 billion.

Expert elicitation is always problematic, but for practical reasons it has been embraced by DHS, CREATE, and other counterterrorism analysts.

Note: The different costs in the matrix are correlated. If the stockpiling costs turn out to be higher than expected, those higher costs should also appear as a summand in every cell in that same row. The payoff table is a random matrix with a complex correlation structure.

Given the joint distribution for bivariate entries in the table, some people replace the entries by their expected values and find the Nash equilibrium (Brown, Carlyle, and Wood, 2008). But computing Nash equilibria and taking expectations do not commute—they find the best solution to the average problem, as opposed to the solution that is, on average, the best.

Banks and Anderson (2006) avoid this by generating many tables at random, computing the Nash equilibrium for each, and tallying the number of times a particular defense investment is chosen.

They also consider a Bayesian analysis, using a simple mirroring argument, to find a subjective probability of a smallpox attack given that it is public knowledge that the U.S. has chosen one of the four smallpox defense options.

The comparison of the minimax solution with the Bayesian solution finds important differences.

The following table gives estimates of the probabilities of a smallpox attack given terrorist knowledge of which protective investment the U.S. made. No input was provided by an intelligence analyst working in this area, but the overall magnitude seems reasonable.

	No Attack	Minor Attack	Major Attack
Stockpile	.95	.040	.010
Biosurveillance	.96	.035	.005
First Responders	.96	.039	.001
Mass Inoculation	.99	.009	.001

If the U.S. were to only stockpile vaccine, then the probability of no smallpox attack is .95, the probability of a single attack is .04, and the probability of multiple attacks is .01. And so forth.

The Bayesian loss criterion multiplies the probabilities in each row above by the corresponding costs in the corresponding row of the payoff table, and then sums across the columns. The best defense has the smallest expected loss.

We ran 100 simulations of the payoff table and found the best decisions under the minimax and Bayesian criteria.

	Mimimax	Bayes
Stockpile	13	99
Biosurveillance	21	0
First Responders	27	0
Mass Inoculation	39	1

If we had been allowed to consider the option of not even stockpiling vaccine, that might have been preferred by the Bayesian analysis.

One can do sensitivity analysis on the influence of different assumptions about the probability of attack or different models for generating the random payoff table.

Exploration found a probability table that made the Bayesian criterion give results similar to the minimax criterion:

	No Attack	Minor Attack	Major Attack
Stockpile	.70	.20	.10
Biosurveillance	.80	.15	.05
First Responders	.85	.10	.05
Mass Inoculation	.90	.05	.05

	Mimimax	Bayes
Stockpile	13	15
Biosurveillance	21	29
First Responders	27	40
Mass Inoculation	39	16

5. Conclusions

Adversarial risk analysis is an important combination of tools from game theory and statistical risk analysis. In particular, Bayesian versions of ARA provide attactive alternatives to the known deficiencies of current methods.

In auctions, gambling, and counterterrorism, agents often have mental models for the decision-processes of their opponents. If those models are correct, then there is the opportunity to improve on the innate pessimism of minimax solutions.

A key component is the need to properly handle many different kinds of uncertainty, which arise in different parts of the analysis. The toy problems considered in this talk point to some of the issues, in the context of normal form and extensive form games.

For more complex applications, say in counterterrorism, work is needed to account for the effects of decentralized and tiered decision-making, constraints on the resources of the opponents, and sensitivity analysis to the Bayesian beliefs and elicitations.