Adversarial Risk Analysis: Analyses of Borel Games

David Banks and Francesca Petralia

Department of Statistical Science, Duke University

banks@stat.duke.edu, francesca.petralia@duke.edu

Shouqiang Wang, Fuqua School of Business

Duke University, shouqiang.wang@duke.edu

Abstract

Adversarial risk analysis offers a new solution concept in game theory. This paper explores its application to a range of simple gambling games, enabling comparison with minimax solutions for similar problems. We find that adversarial risk analysis has several attractive advantages: it is easier to compute, it takes account of asymmetric information, it corresponds better to human behavior, and it reduces to previous solutions in appropriate circumstances.

Acknowledgment: This work was partially supported by DARPA grant W911NF-09-1-0337 and through a subcontract with the Institute for Homeland Security Solutions (IHSS). IHSS is a research consortium established to conduct applied research in the social and behavioral sciences. The Human Factors Division is the DHS sponsor for IHSS. More information about IHSS is available at www.ihssnc.org. This paper is based upon work supported by the DHS under Contract No. HSHQDC-08-C-00100. Any opinions, findings, conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the official policy or position of the DHS or of the U.S. Government. The paper is approved for public release.

1 Introduction

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

ARA is an alternative to the solution concepts in traditional game theory. Camerer (2002) and Gintis (2009), summarizing on a large body of empirical work, criticize minimax and related solutions (including maximization of expected utility) as unreasonable descriptions of human decision-making. Worse, minimaxity can lead to sub-optimal solutions, in the sense that if one's opponent is not perfectly rational, then solutions based on that premise may be too pessimistic—by mitigating the worst conceivable scenario, one avoids better outcomes that correspond to choices a human opponent might realistically select. Additional problems are that minimax solutions can be difficult to compute, and often require strong and unreasonable assumptions about common knowledge.

Such concerns led Kadane and Larkey (1982) and Raiffa (1982, 2002) to propose decision analysis. It is controversial (cf. Harsanyi, 1982). One concise critique is due to Myerson (1999):

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

However, instead of following Myerson in defaulting back to Nash equilibria, we pursue an ARA solution concept.

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

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

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

We apply this ARA method to several famous gambling problems, based on various modifications of the Borel Game (1938), also known as *Le Relance*. This game may be viewed as a simplified form of poker. These applications enable comparison with results obtained from other solution strategies, and illustrate the key ideas. Section 2 treats the basic Borel Game with fixed bets. Section 3 extends this simple model to other versions of this that allow continuous bets, multiple players, and multiple rounds. Section 4 summarizes the results.

2 A Simple Game

Consider the following two-person game:

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

The Borel Game's assumption of a unit ante is not restrictive; one can rescale the solution to other cases by maintaining the ratio of the ante to the bet.

The Borel Game is a classical problem in game theory. Von Neumann and Morgenstern (1947) extended it to allow players to check; Bellman and Blackwell (1949) examined bluffing strategy; and Karlin and Restrepo (1957) generalized the problem to multiple players and/or multiple rounds and/or multiple increments of bet size. A recent review of the area, with modern proofs and language, is given in Ferguson and Ferguson (2003). Minimax solutions have been found for a wide range of variant games, which include multiple players, betting increments, and checking.

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

For all versions of ARA, Bart has observed that his draw is $X = x$; Lisa's draw $Y = y$ is unknown to him, although he knows its distribution. Let V_x be the amount won by Bart after drawing $X = x$. Table 1 shows the possible situations, depending on how each player decides to bet and the values of X and Y.

(Insert Table 1 about here.)

Using Table 1 and neglecting ties (an event with measure zero), the expected amount won by Bart, conditional on his draw, is

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

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

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

$$\begin{aligned} V_x &= -[1 - g(x)] + g(x)\mathbb{P}[\text{ Lisa folds} \mid \text{Bart bets }] \\ &\quad +(1+b)g(x)\mathbb{P}[\text{ Bart wins} \mid \text{Lisa calls }]\mathbb{P}[\text{ Lisa calls} \mid \text{Bart bets }] \\ &\quad -(1+b)g(x)\mathbb{P}[\text{ Bart loses} \mid \text{Lisa calls }]\mathbb{P}[\text{ Lisa calls} \mid \text{Bart bets }] \end{aligned} \tag{2.1}$$

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

2.1 Zero-Order ARA

In the simplest form of ARA, which is essentially the form proposed by Raiffa (1982) and Kadane and Larkey (1982), Bart does not explicitly try to model Lisa's thinking. Instead, he merely declares the subjective distribution $\pi(c)$ that he has over the value c at which Lisa will bet. [In this game, there is no advantage to Lisa from bluffing (Borel, 1938); she should have a fixed rule such that if $Y > c$, she bets, and otherwise she folds.] Bart's distribution on c takes no formal account of the fact that Lisa's rule might be informed by Bart's decision to bet, although this could be implicit in his personal elicitation.

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

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

Also, routine calculation shows

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

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

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

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

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

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

2.1.1 Example 1: Beta priors.

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

$$\int_0^x (x-c)\pi(c)\, dc = xI_x(\alpha,\beta) - \frac{\alpha}{\alpha+\beta} I_x(\alpha+1,\beta) \tag{2.2}$$

where

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

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

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

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

For the agnostic case with $\alpha = \beta = 1$, equation (2.2) reduces to $\frac{1}{2}x^2$. So Bart will bet if and only if

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

When Bart thinks Lisa's step is close to 1 (i.e. $\alpha = 3$ and $\beta = 1$), equation (2.2) is equal to $\frac{1}{4}x^4$ and Bart will bet if and only if

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

And when Bart thinks Lisa's betting rule is close to 1/2, i.e. $\alpha = \beta = 2$, then equation (2.2) is equal to $x^3 - \frac{1}{2}x^4$ and Bart will bet if and only if

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

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

Figure 1 shows how the point at which Bart bets depends upon the value b of the bet, under the three different priors. The monotonicity in b implies that as the bet size increases, Bart should become more conservative.

(Insert Figure 1 about here.)

Note 1: Harsanyi (1967, 1968a, 1968b) describes a game-theoretic analysis that uses Bayesian information, but it requires assumptions about mutual information and mutual priors that are unreasonable in many contexts. For example, in this case, there is no reason to assume that Lisa would know the subjective distribution $\pi(c)$ that Bart places over her betting rule.

2.2 First-Order ARA

The first-order ARA is more interesting. Here Bart models Lisa's reasoning in order to develop his belief about her decision rule. Bart accomplishes this through a "mirroring" argument, in which he does the analysis he expects her to make, using subjective distributions to describe the quantities he does not know.

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

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

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

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

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

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

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

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

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

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

or

$$V_x = \begin{cases} -1 + g(x)(2\tilde{y} + b\tilde{y} - b) & \text{if } x \leq \tilde{y} \\ -1 + g(x)(2x + 2bx - b\tilde{y} - b) & \text{if } x > \tilde{y}. \end{cases} \tag{2.3}$$

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

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

Case I: Bart Believes that Lisa Plays Minimax.

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

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

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

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

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

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

and solve to find

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

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

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

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

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

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

and then calculate

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

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

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

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

This agrees with the traditional minimax value of the game.

Case II: Bart Believes that Lisa Is Rash.

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

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

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

The value of this ARA game to Bart is

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

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

Case III: Bart Believes that Lisa Is Conservative.

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

$$V_x = -(1-g(x)) + g(x)\tilde{y} + (1+b)g(x)(1-\tilde{y})\frac{x-\tilde{y}}{1-\tilde{y}} - (1+b)g(x)(1-\tilde{y})\left(1 - \frac{x-\tilde{y}}{1-\tilde{y}}\right). \quad (2.4)$$

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

$$V_x = -1 + g(x)\left[1 + \tilde{y} - (1+b)(1-\tilde{y})\right]. \quad (2.5)$$

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

The value V of this game is

$$\mathbb{E}[V_x] = \int_0^{\tilde{y}} [\tilde{y} - (1+b)(1-\tilde{y})]\, dx + \int_{\tilde{y}}^1 [\tilde{y} + (1+b)(x-\tilde{y}) - (1+b)(1-x)]\, dx. \tag{2.6}$$

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

Note 2: This analysis of the Borel Game extends immediately to situations in which the two players draw independently from a continuous distribution W with density w. In that case, the conditional distribution that Bart imputes to Lisa is

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

and Bart's bluffing function takes its step at

$$\tilde{x} = \frac{1}{2}\left[1 - \frac{1}{1+b}\frac{1+W(\tilde{y})}{1-W(\tilde{y}}\right].$$

The more useful extension, in which Bart and Lisa draw from a bivariate, possibly discrete distribution $W(x,y)$ (e.g., a deck of cards) is tedious but straightforward; Bart's distribution for Y is the conditional $W(y|X=x)$, and he knows that Lisa's analysis is symmetric.

Note 3: Some readers may be uncomfortable with our specificity in requiring Bart to assume that Lisa thinks his bluffing function is $\tilde{g}(x)$. They might argue that Bart could not guess it exactly—that it would be more reasonable to say that he has a subjective distribution over the set $\mathcal{G}$ of all possible bluffing functions. Admittedly, this level of introspective integration is not realistic, but in practice Bart can probably do well by imagining something simpler; for example, he could suppose that there is a 20% chance that Lisa has read Bellman and Blackwell (1949); a 30% chance that she thinks Bart bets whenever $X > 1/2$, and a 50% chance that she believes Bart bets with probability X. Of course, these judgments would be based entirely upon whatever personal knowledge Bart may have about Lisa, but averaging a few specific benchmark strategies of this kind should enable a practical approximation to Bart's true uncertainty about Lisa's belief.

2.2.1 Example: The $\tilde{g}$ is a power function.

Suppose that Bart believes that Lisa thinks his bluffing function has the form $g(x) = x^p$ for some fixed value $p > -1$. Then $\tilde{y} = \sqrt[p+1]{\frac{1}{2}\frac{b}{1+b}}$. When p goes to infinity $g(x)$ approaches a step function at 1; i.e., Bart bids with probability 1 if $x = 1$ and with probability zero otherwise. Large values of p imply that Lisa believes Bart tends to bet for large values of x, leading Lisa to fold more frequently and increasing Bart's expected payoff. For $b = 2$, the left panel in Figure 2 shows the minimum value of x for which

Bart should bet as a function of p. The right panel shows the value of the game, to Bart, as a function of p.

(Insert Figure 2 about here.)

2.3 Second-Order ARA, and Higher-Order ARA

As the previous Myerson (1999) quote implied, the mirroring argument can be extended to as many levels as seems appropriate to the problem. In a two-person Borel game, the next level would have Bart assess the bluffing function that he believes Lisa thinks Bart is using to describe her guess about Bart's bluffing function. Then his calculation of Lisa's strategy would take account of of the fact that Lisa is mirroring Bart's thinking in order to develop her own playing rule.

The analysis is straightforward, but demands organized thinking and patience. In principle, one could address any (finite) level of analysis by means of symbolic calculation—one is just nesting a series of calculations. But it would require careful programming.

It is natural to ask whether there could be convergence to a solution (perhaps the minimax solution) as the order increases. This is an open question, but we suspect the answer is "no". The final rule depends only upon the assessment made at the highest level of mirroring. And it seems quite reasonable for the final rule to be any one of a wide range of strategies, depending upon his view of Lisa's psychology of play. If that is correct, then the influence of the highest order assessment does not diminish, or does not diminish enough to produce a unique solution.

For the Borel game, and its close relatives in the poker family, first-order ARA seems quite sufficient for nearly all practical circumstances. It is conceivable that, at the very highest levels of professional play, one might want to go further, but that is unclear. It certainly seems unlikely that poker professionals currently make detailed minimax calculations (Ferguson and Ferguson, 2007); if they do not do that, then perhaps they would not engage in high-order ARA reasoning either.

In the three-person game discussed in Section 4, the case for the practical advantage of a second-order ARA seems more compelling.

3 Continuous Bets

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

We use the following notation:

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

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

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

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

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

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

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

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

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

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

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

$$\begin{aligned}
\mathbb{E}_{g(x),h(\cdot|x)}\left[V_B|X=x\right] \quad = \quad & \underbrace{-(1-g(x))}_{\text{Bart folds}} + g(x)\Big\{\mathbb{E}_{h(\cdot|x)}\Big[\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa folds} \mid \text{Bart bets } B_x] \mid X = x\Big] \\
& + \mathbb{E}_{h(\cdot|x)}\Big[\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa loses} \mid \text{Bart bets } B_x] \cdot (1 + B_x) \mid X = x\Big] \\
& - \mathbb{E}_{h(\cdot|x)}\Big[\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa wins} \mid \text{Bart bets } B_x] \cdot (1 + B_x) \mid X = x\Big]\Big\}. \quad (3.1)
\end{aligned}$$

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

$$\{g^*(x), h^*(\cdot|x)\} \in \underset{g(x),h(\cdot|x)}{\operatorname{argmax}} \quad \mathbb{E}_{g(x),h(\cdot|x)}\left[V_B|X=x\right]. \quad (3.2)$$

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

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

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

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

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

Therefore, Bart believes Lisa will call if and only if

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

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

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

Then Lisa will call if and only if

$$Y \geq \tilde{y}^*(B_x) \stackrel{def}{=} \inf\left\{y \in [0, 1] : \int_0^y \tilde{f}(z|B_x)\, dz \geq \frac{B_x}{2(1 + B_x)}\right\}. \tag{3.3}$$

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

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

Consequently, Bart can compute the following quantities:

$$
\begin{aligned}
\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa folds} \mid \text{Bart bets } B_x] &= \tilde{y}^*(B_x); && (3.4)\\
\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa loses} \mid \text{Bart bets } B_x] &= \mathbb{P}[\tilde{y}^*(B_x) \le Y \le x|B_x] \\
&= [x - \tilde{y}^*(B_x)]^+; && (3.5)\\
\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa wins} \mid \text{Bart bets } B_x] &= \mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa calls} \mid \text{Bart bets } B_x] \\
&\quad -\mathbb{P}_{\tilde{f}(\cdot|B_x)}[\text{ Lisa loses} \mid \text{Bart bets } B_x] \\
&= 1 - \tilde{y}^*(B_x) - [x - \tilde{y}^*(B_x)]^+. && (3.6)
\end{aligned}
$$

Plugging (3.4), (3.5) and (3.6) into (3.1), we obtain

$$
\begin{aligned}
&\mathbb{E}_{f(x),g(\cdot|x)}[V_B|X = x] = \\
&\quad -(1 - g(x)) + g(x)\mathbb{E}_{h(\cdot|x)}\left[\tilde{y}^*(B_x) + 2[x - \tilde{y}^*(B_x)]^+(1 + B_x) - (1 - \tilde{y}^*(B_x))(1 + B_x)\right]. \quad (3.7)
\end{aligned}
$$

Lemma 3.1. *Suppose $\tilde{f}(\cdot|b)$ is positive and continuous in $b \in [\epsilon, K]$, then $\tilde{y}^*(b)$ is continuous in b.*

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

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

must be continuous in b. □

Summing up the previous results, we obtain

Theorem 3.1. *For any $x \in [0, 1]$ and given $\tilde{f}(\cdot|b)$ positive and continuous in $b \in [\epsilon, K]$, let $\tilde{y}^*(b)$ be defined as in (3.3). Also let*

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

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

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

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

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

3.1 Example: Lisa has a step-function posterior.

We now provide a simple illustration of how to apply Theorem 3.1 to obtain Bart's first-order ARA solution of the Borel game with continuous bets. In this subsection, we assume $\tilde{f}(\cdot|b)$ is of the following form:

$$\tilde{f}(x|b) = \begin{cases} \frac{1+K}{1+b} & \text{if } 0 \leq x \leq \frac{1+b}{1+K} \\ 0 & \text{otherwise.} \end{cases} \tag{3.8}$$

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

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

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

1. For $x < \frac{\epsilon^2+2(1+K)\epsilon}{4(1+K)(1+\epsilon)}$, then $b^*(x) = \epsilon$ and $\Delta^*(x) = -\frac{\epsilon^2}{2(1+K)} + (2x-1)(\epsilon+1) < -1$. By Theorem 3.1, $g^*(x) = 1$; i.e., Bart will fold w.p. 1. There is no need to specify $h^*(\cdot|x)$.

2. For $\frac{\epsilon^2+2(1+K)\epsilon}{4(1+K)(1+\epsilon)} \leq x < \frac{1}{2} + \frac{\epsilon}{2(1+K)}$, then $b^*(x) = \epsilon$ and $\Delta^*(x) = -\frac{\epsilon^2}{2(1+K)} + (2x-1)(\epsilon+1) \geq -1$. By Theorem 3.1, $g^*(x) = 1$ and $h^*(b|x) = \delta(b - \epsilon)$, i.e. Bart will bet ϵ w.p. 1.

3. For $\frac{1}{2} + \frac{\epsilon}{2(1+K)} \leq x < \frac{1}{2} + \frac{K}{2(1+K)}$, then $b^*(x) = 2(1+K)x - (1+K)$ and $\Delta^*(x) = \frac{1+K}{2}(2x-1)^2 + (2x-1) \geq -1$. By Theorem 3.1, $g^*(x) = 1$ and $h^*(b|x) = \delta(b - (2(1+K)x - (1+K)))$; i.e., Bart will bet $2(1+K)x - (1+K)$ w.p. 1.

4. For $x \geq \frac{1}{2} + \frac{K}{2(1+K)}$, then $b^*(x) = K$ and $\Delta^*(x) = -\frac{K^2}{2(1+K)} + (2x-1)(K+1) \geq -1$. Then, by Theorem 3.1, $g^*(x) = 1$ and $h^*(b|x) = \delta(b - K)$; i.e., Bart will bet K w.p. 1.

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

(Insert Figure 3 about here.)

3.2 Three Players

Suppose that Bart, Lisa, and Milhouse are playing *Le Relance*. Each antes one unit and then independently draws a uniform random number on $[0, 1]$. Bart sees $X = x$ and may bet b units or fold. Lisa knows $Y = y$ and Bart's decision; she may then bet b or fold. Milhouse knows $Z = z$ and both Lisa's and Bart's decision, and he may fold or bet b. At this point, those players who did not fold compare their numbers, and the highest number wins the pot.

This analysis derives the optimal play for Bart, but similar calculation could be performed from the perspective of any of the other players. Table 2 shows the possible outcomes for Bart, conditional on $X = x$.

(Insert Table 2 about here.)

The zero-order ARA is straightforward. Bart has a subjective distribution for Lisa's bluffing function $g_L(y)$; he also has subjective distributions $\pi(c \mid \text{Bart's action, Lisa's action})$ on the value of Z for which Milhouse will bet. He performs a Bayesian risk analysis to determine the strategy that maximizes his expected winnings, and that analysis does not explicitly account for the strategic thinking of either Milhouse or Lisa.

Conceivably, Bart might do a "half-order" ARA. Here he would mirror the thinking of one of Lisa and Milhouse, while using a non-adversarial subjective distribution for the the decision rule of the other. This might be appropriate if, for example, he believes that Lisa is smart and strategic, while Milhouse plays according to some simple rule-of-thumb. Kadane (2009, p. 241) points out that Poe's Detective Dupin tells of school-boy who drew such distinctions in a game of 'even-and-odd'.

In the interest of brevity, this section focuses upon first-order ARA, for the case in which Bart mirrors the thinking of both Lisa and Milhouse (but one could perform similar analyses from the perspectives of each of the other players). The complexity of the problem requires a bit of notation. The function $u^I_J(x)$ is what Bart believes player I thinks player J assumes is Bart's bluffing function. This may depend on preceding play, in which case the conditioning is denoted by the usual vertical stroke. The distribution that Bart believes player I thinks player J holds for Bart's draw X is written as $U^I_J(x)$, and when indicated, this is conditional on preceding play. Similarly, $v^I_J(y)$ and $w^I_J(z)$ are what Bart believes player I thinks player J guesses to be the bluffing functions for Lisa and Milhouse, respectively; and $V^I_J(y)$ and $W^I_J(z)$ are the corresponding distributions for their draws. The I and J may range over B, L, and M, for Bart, Lisa, and Milhouse. In some combinations the notation is redundant: $u^B_B(x)$ is what Bart believes Bart thinks Bart thinks is Bart's bluffing function. However, as shown below, there are other situations that necessarily require the distinctions this notation can draw.

Bart's mirroring analysis starts with Milhouse (M), the last decision-maker. Milhouse can observe four things: Both Bart (B) and Lisa (L) fold, both bet, Bart bets and Lisa folds, or Bart folds and Lisa bets. From Bart's perspective, he need not analyze the situations in which he folds—he knows his payoff will be -1.

So first suppose that both Bart and Lisa bet. Then, if Milhouse bets with probability $g_M(z)$, the value of game to Milhouse is:

$$V_z = -1 + g_M(z)\,\{1 + 2(1+b)\mathbb{P}[M \text{ wins }] - (1+b)\mathbb{P}[M \text{ loses }]\}$$

where

$$\mathbb{P}[M \text{ wins }] = \mathbb{P}[X < z \,|\, B \text{ bet }]\mathbb{P}[Y < z \,|\, B, L \text{ bet }].$$

If Bart believes that

- Milhouse thinks Bart's bluffing function is $u_B^M(x)$, and
- Milhouse thinks Lisa's bluffing function, conditional on Bart betting, is $v_L^M(y|B \text{ bet })$

then

$$\mathbb{P}[M \text{ wins }] = \mathbb{P}[X < z \,|\, B \text{ bet }]\mathbb{P}[Y < z|B, L \text{ bet }] = U_B^M(z)V_L^M(z \,|\, B \text{ bet })$$

where

$$\begin{aligned} U_B^M(x) &= \frac{\int_0^z u_B^M(w)\,dw}{\int_0^1 u_B^M(w)\,dw} \\ V_L^M(y \,|\, B \text{ bet} &= \frac{\int_0^z v_L^M(w \,|\, B \text{ bet })\,dw}{\int_0^1 v_L^M(w \,|\, B \text{ bet })\,dw}. \end{aligned}$$

Bart believes Milhouse will bet in order to maximize the value of his game, so Bart believes that Milhouse bets if and only if $Z > \tilde{z}_1$, where $\tilde{z}_1$ solves

$$1 + 2(1+b)\mathbb{P}[M \text{ wins }] - (1+b)\mathbb{P}[M \text{ loses }] = 0.$$

When $g_M(z)$ jumps at $\tilde{z}_1$, then Milhouse bets only when his expected winnings are greater than the loss of 1 unit from folding (conditional on both Bart and Lisa betting).

If Bart bets but Lisa folds, then similar arguments show that Bart believes Milhouse will bet only when $Z > \tilde{z}_2$, where $\tilde{z}_2$ solves

$$1 + (2+b)\mathbb{P}[M \text{ wins }] - (1+b)\mathbb{P}[M \text{ loses }]$$

with

$$\mathbb{P}[M \text{ wins }] = U_B^M(z).$$

So, for the cases relevant to Bart's decision, Bart has derived the rules he believes Milhouse will use.

Now Bart considers Lisa's situation. He need not calculate Lisa's expected payoff when he folds, since he knows he would lose 1 unit. So assume Lisa observes that Bart bets. If she then bets with probability $g_L(z)$, the expected value of her game is

$$\begin{aligned} V_y = & -1 + g_L(y)\left\{1 + 2(1+b)\mathbb{P}[M \text{ bets and } L \text{ wins }]\right. \\ & +(2+b)\mathbb{P}[M \text{ folds and } L \text{ wins }] - (1+b)\mathbb{P}[M \text{ bets and } L \text{ loses }] \\ & \left. -(1+b)\mathbb{P}[M \text{ folds and } L \text{ loses }]\right\}. \end{aligned}$$

To discover her betting rule, Bart must analyze what Lisa thinks about $\mathbb{P}[M \text{ bets and } L \text{ wins }]$ and $\mathbb{P}[M \text{ folds and } L \text{ wins }]$. (Understanding these is sufficient since for any two events A and B, $\mathbb{P}[A \text{ and } B^c] = \mathbb{P}[A] - \mathbb{P}[A \text{ and } B]$, where the superscript denotes complementation.)

Bart believes that Lisa thinks that Milhouse's bluffing function, when both she and Bart bet, is $w_M^L(z \,|\, B, L \text{ bet })$. Also, Bart believes that Lisa thinks that Bart's bluffing function is $u_B^L(x)$. So Lisa should calculate

$$\begin{aligned} \mathbb{P}[M \text{ bets and } L \text{ wins }] &= \mathbb{P}[L \text{ wins } \,|\, B, M \text{ bet }]\mathbb{P}[M \text{ bets } \,|\, B, L \text{ bet }] \\ &= U_B^L(y) W_M^L(y \,|\, B, L \text{ bet }) \end{aligned}$$

and

$$\begin{aligned} \mathbb{P}[M \text{ folds and } L \text{ wins }] &= \mathbb{P}[L \text{ wins } \,|\, B \text{ bets }]\mathbb{P}[M \text{ folds } \,|\, B, L \text{ bet }] \\ &= U_B^L(y) \int_0^1 (1 - \int_0^1 w_M^L(z \,|\, B, L \text{ bet })\, dz). \end{aligned}$$

Using these quantities, Bart believes Lisa will solve to find the betting rule, a step function at $\tilde{y}$, that maximizes the value of her game.

Now Bart has what he needs to determine the strategy that maximizes his expected value. He should bet with $X = x$ when the expected value is greater than -1, the amount that he loses by folding; i.e., he bets when

$$\begin{aligned} -1 \;<\; & 2(1+b)\mathbb{P}[B \text{ wins and } L, M \text{ bet }] + (2+b)\mathbb{P}[B \text{ wins and only } L \text{ bets }] \\ & +(2+b)\mathbb{P}[B \text{ wins and only } M \text{ bets }] - (1+b)\mathbb{P}[B \text{ loses and } L, M \text{ bet }] \\ & -(1+b)\mathbb{P}[B \text{ loses and only } L \text{ bets }] - (1+b)\mathbb{P}[B \text{ loses and only } M \text{ bets }]. \end{aligned}$$

He knows that Lisa bets when $Y > \tilde{y}$, and that Milhouse bets when $Z > \tilde{z}_1$ and Lisa bets, or when $Z > \tilde{z}_2$ and Lisa folds.

From this information, given $X = x$,

$$\begin{aligned} \mathbb{P}[B \text{ wins and } L, M \text{ bet }] &= \mathbb{P}[B \text{ wins } \,|\, L, M \text{ bet }]\mathbb{P}[L, M \text{ bet }] \\ &= \mathbb{P}[B \text{ wins } \,|\, L, M \text{ bet }](1-\tilde{y})(1-\tilde{x}_1) \\ &= \frac{[x - \max\{\tilde{y}, \tilde{z}_1\}]^+}{(1 - \max\{\tilde{y}, \tilde{z}_1\})}(1-\tilde{y})(1-\tilde{z}_1). \end{aligned}$$

Similarly,

$$\begin{aligned}
\mathbb{P}[B \text{ wins and only } L \text{ bets }] &= \frac{[x-\tilde{y}]^+}{1-\tilde{y}}(1-\tilde{y})\tilde{z}_1 \\
\mathbb{P}[B \text{ wins and only } M \text{ bets }] &= \frac{[x-\tilde{z}_2]^+}{1-\tilde{z}_2}\tilde{y}(1-\tilde{z}_2).
\end{aligned}$$

Each of these terms can be computed, and thus Bart can find the value of X such that he should bid for larger values, and fold for smaller values.

The previous ARA analysis may seem more complex than it is. Most of the effort is simply a matter of bookkeeping. G. H. Hardy would have said that it is trivial, in the sense that the calculation may be carried out in a straightforward manner, despite the fact that it would become tedious with larger number of players. Similarly, the calculation for higher-order ARA is trivial.

Note 4: The three-player game explicitly opens the door to incoherent thinking. It is entirely possible to specify a set of bluffing functions held by, say, Milhouse, which are inconsistent. Of course, it is well known that humans often maintain incoherent beliefs, and if Bart is insightful enough to correctly capture the incoherencies in Milhouse's opinions, then he should be able to take better advantage of Milhouse. From the ARA perspective, it would be wrong to require Bart to assume that his opponents hold mutually consistent beliefs when, in fact, it is Bart's honest and best judgment that they do not.

Note 5: In a second-order ARA, Bart would impute more calculation to each of the players. For example, this first-order analysis declared that $\tilde{v}_L^M(y|B \text{ bet })$ was Bart's belief about what Milhouse thought was Lisa's bluffing function, conditional on her observing that Bart bets. But if Bart felt that Milhouse was more thoughtful, then the second-order ARA would require that Milhouse derive $\tilde{v}_L^M(y|B \text{ bet })$ from what Milhouse thinks are Lisa's beliefs about Bart's and Milhouse's bluffing functions, by selecting the function that maximizes her expected value. But if one feels that one's opponents are so smart and invested that they analyze strategy this deeply, then as a matter of practicality it is better to seek out more stupid people with whom to play next.

4 Conclusions

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

> . . . it is in this theory [the theory of games] that they [minimax solutions] are more justifiable. *But even here in practice you would prefer to maximize your expected gain.* You would probably use minimax solutions when you had a fair degree of belief that your opponent was a good player. Even when you use the minimax solution you may be maximizing your expected gain since you may already have worked out the details of the minimax solution,

> and you would probably not have time to work out anything better once a game had started. [Note the implicit use here of "type II rationality" which is especially pertinent to avoid an "infinite regress".] To attempt to use a method other than the minimax method would then lead to too large a probability of a large loss, especially in a game like poker.

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

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

This paper has laid out the ARA approach in the context of one of the classical problems in game theory, the Borel game. It is a rich example, both mathematically and historically; the research pedigree includes key contributions by Borel, von Neumann and Morgenstern, Bellman and Blackwell, Karlin and Restrepo, and, more recently, Ferguson and Ferguson. (Note: Chris Ferguson won five World Series of Poker events, including the 2008 NBC National Championship).

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

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

The ARA approach is, of course, more general in application than the Borel game considered here. Rios Insua, Rios, and Banks (2009) apply it to an auction problem. And Wang and Banks (2010) apply ARA to the problem of routing convoys through a city in which there may be improvised explosive devices. In both cases, the essential step is to build a model for the the decision processes of one's opponent, and then use subjective probability distributions to describe all unknown quantities. We are hopeful that ARA may provide a more useful tool than the minimax solution concept in addressing the kinds of practical problems, with asymmetric and uncertain information, that often arise in practice.

References

Banks, D. and S. Anderson (2006). Game theory and risk analysis in the context of the smallpox threat, in *Statistical Methods in Counterterrorism*, A. Wilson, G. Wilson and D. Olwell, 9-22.

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

Borel, E. (1938). "*Traité du Calcul des Probabilités et ses Applications*," Vol. IV, Fascicule 2, *Applications aux Jeus des Hazard*, Gautier-Villars, Paris.

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

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

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

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

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

Harsanyi, J. (1967) Games with incomplete information played by "Bayesian" players, I-III. Part I. The basic model. *Management Science*, 14, 3, 159-182.

Harsanyi, J. (1968a) Games with incomplete information played by "Bayesian" players, I-III. Part II. Bayesian equilibrium points. *Management Science*, 14, 5, 320-334.

Harsanyi, J. (1968b) Games with incomplete information played by "Bayesian" players, I-III. Part III. The basic probability distribution of the game. *Management Science*, 14, ?, 486-502.

Harsanyi, J. (1982) Subjective probability and the theory of games: Comments on Kadane and Larkey's paper, *Management Science*, 28, 2, 120-124.

Jittorntrum, K. (1978). An Implicit Function Theorem. *Journal of Optimization Theory and Application*, **25**, 575–577.

Kadane, J.B. (2009). Bayesian thought in early modern detective stories: Monsier Lecoq, C. Auguste Dupin, and Sherlock Holmes. *Statistical Science*, **24**, 238–243.

Kadane, J.B. and P.D. Larkey (1982) Subjective probability and the theory of games, *Management Science*, 28, 2, 113-120; reply: 124.

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

Larkey, P.D., Kadane, J.B., Austin, R., and Zamir, S. (1997). Skill in games. *Managemement Science*, **43**, 596–609.

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

Pate-Cornell, E. and S. Guikema (2002) Probabilistic modeling or terrorist threats: a systematic analysis approach to setting priorities among countermeasures, *Military Operations Research*, 7, 5-23.

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

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

Rios Insua, D., Rios, J., and Banks, D. (2009). "Adversarial Risk Analysis," *Journal of the American Statistical Association*, **104**, 841–854.

Stahl, D. O., and Wilson, P. W. (1995). "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, **10**, 218–254.

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

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

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

Wang, S., and Banks, D. (2011). "Network Routing in Counterterrorism: An Adversarial Risk Analysis Framework," submitted to *Naval Research Logistics*.

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

Table 1: Amount won by Bart, with two players.

V_x	Bart	Lisa	Milhouse	Bart's Win Condition
-1	folds			
2	bets	folds	folds	
2 +b	bets	bets	folds	$x > Y$
2 +b	bets	folds	bets	$x > Z$
-(1+b)	bids	bids	bids	$x < \max\{Y, Z\}$
-(1+b)	bids	bids	folds	$x < Y$
-(1+b)	bids	folds	bids	$x < Z$
2(1+b)	bids	bids	bids	$x > \max\{Y, Z\}$

Table 2: Amount won by Bart, with three players.

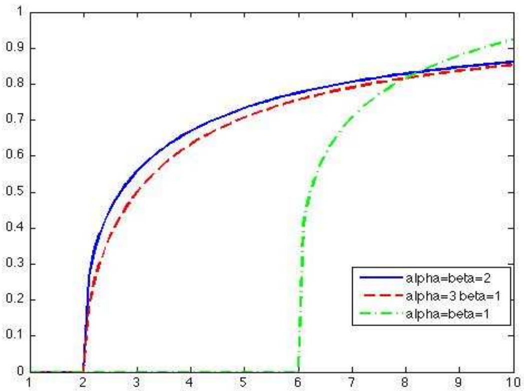

Figure 1: This shows, for three models with beta priors, how the smallest value at which Bart should bet depends upon the size of the bet.

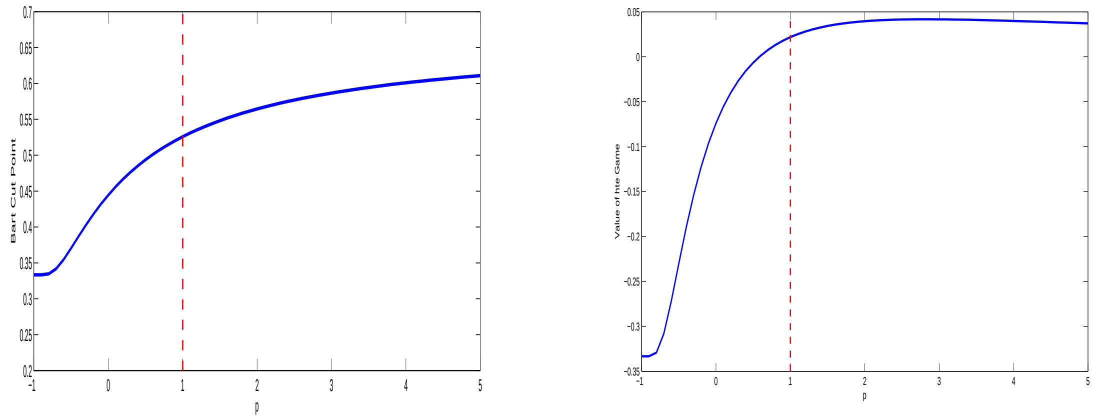

Figure 2: Bart's minimum value to bet (left) and the value of the game (right), as functions of p.

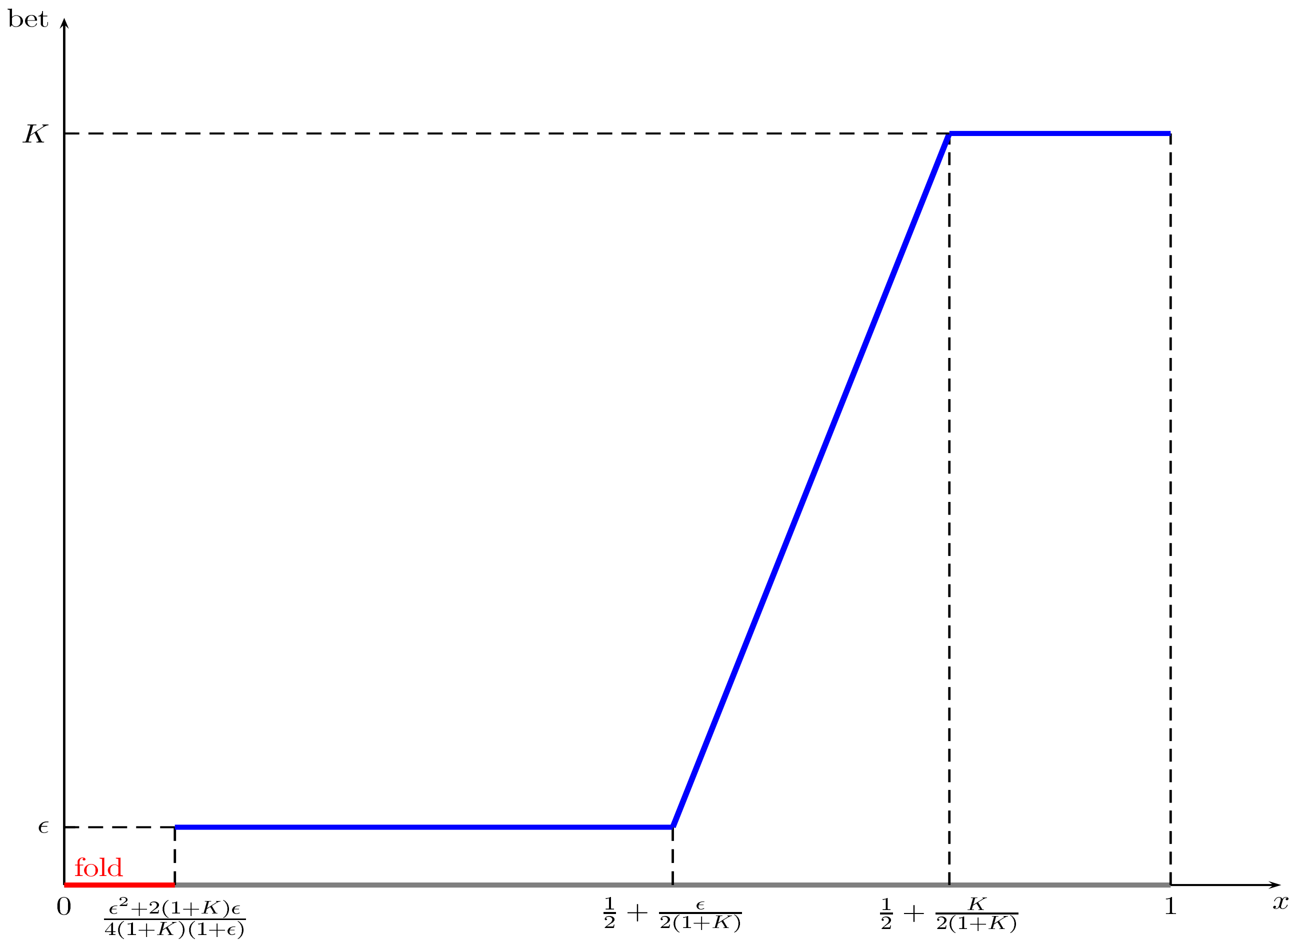

Figure 3: Bart's first-order ARA strategy as a function of x for $\tilde{f}(\cdot|b)$ given in (3.8)
.