Bayesian Borel Games

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 decision analysis of Kadane and Larkey (1982) and Raiffa (1982) through the use of a **mirroring argument**.

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

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

However, instead of following Myerson in defaulting back to game theory, we use the mirroring method. It may be viewed as a Bayesian version of Level-k thinking (Stahl and Wilson, 1995).

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

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 $d_0 = \operatorname{argmax}_{d \in \mathbb{R}^+} (d^* - d)G(d)$; 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.

The following figures illustrate the fixed point solution. (Note that the caption reverses the roles of Apollo and Daphne.) The starting points for H_D and H_A were distinct triangular distributions on $[0, 100]$.

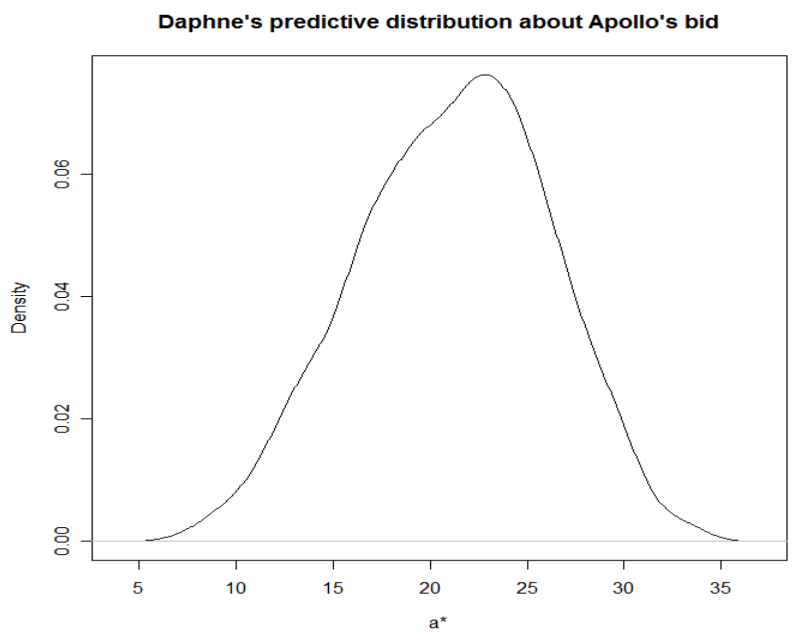

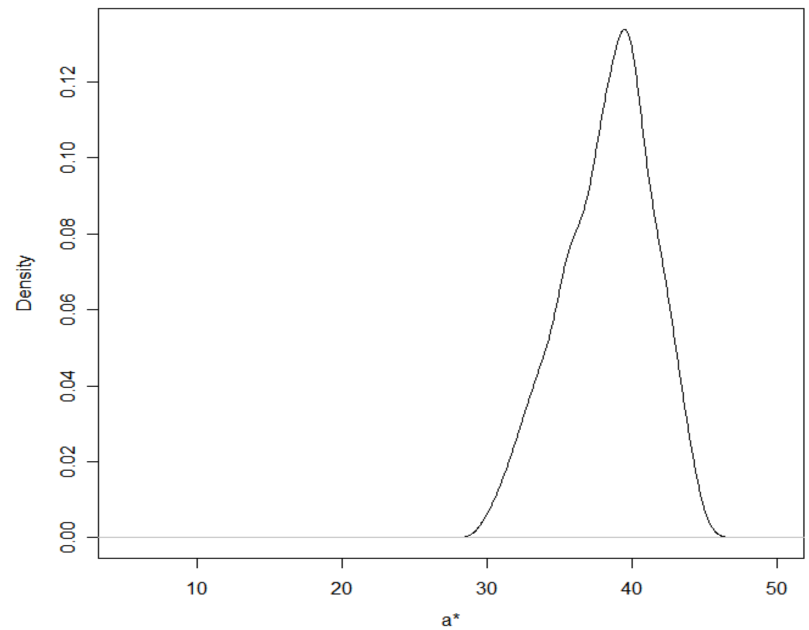

The left panel shows the third iterate; the right panel shows the tenth iterate.

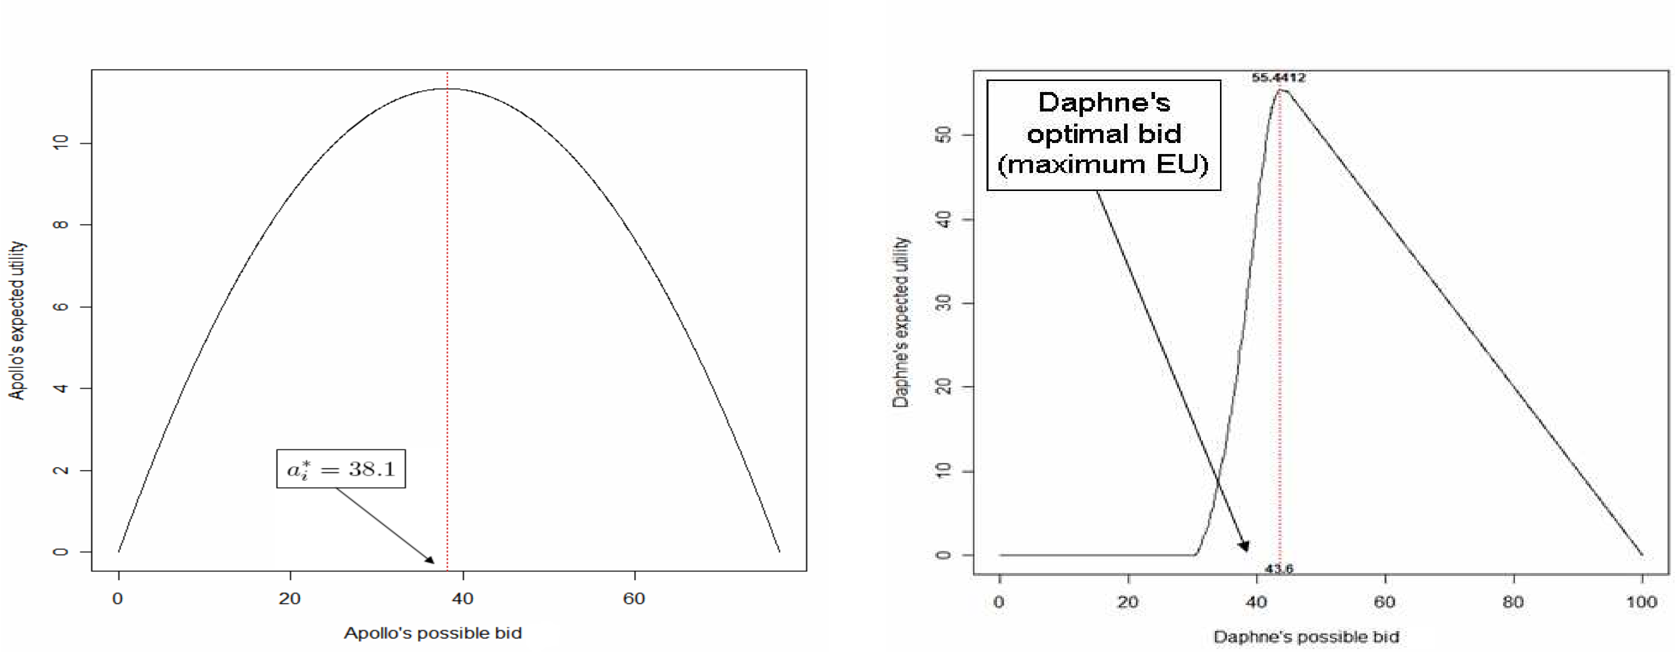

These panels show the result of a algorithm. The left is the expected utility Daphne believes Apollo thinks he will get from a given bid. The right shows the expected utility that Daphne will receive from a given bid.

Caveat: I am not certain that these figures are correct.

Note: 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. La Relance: A Primitive Version of 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, assume that Bart and Lisa play a game in which each privately 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 four possible situations:

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:

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

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

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

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 $\mathbb{P}[X \leq y| \text{ Bart bet }] = \tilde{F}(y)$, 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. So the expected value of the game for Bart, given $X = x$, is:

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

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

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

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

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.

When $x \le \tilde{y}$, 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 if $x > \tilde{x} = [b(1+\tilde{y})]/[2(1+b)]$, folds if $x < \tilde{x}$, and may do as he pleases when $x = \tilde{x}$.

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

The expected value of the game, to Bart, is $V = \int_0^1 V_x \, dx$. Its value depends on his belief about Lisa's play.

Case I: Bart Believes that Lisa Plays Minimax.

The traditional minimax solution has $\tilde{y} = b/(b+2)$. In that case it is known 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 $V = -b^2/(b+2)^2$; he is disadvantaged by the sequence of play.

In contrast, the ARA analysis finds that when Lisa uses the minimax threshold $\tilde{y} = b/(b+2)$, then Bart may bet or not, as he pleases, when $x \leq \tilde{x}$. This is slightly different from the minimax solution.

The difference arises because, if Lisa knows that Bart's bluffing function does not bet with probability $2/(b+2)$ when $x \leq b/(b+2)$, then she can improve her expected value for the game by changing the threshold at which she calls.

In the minimax game, Bart's bluff pins Lisa down, preventing her from using a more profitable rule. But for either game, the value for Bart is unchanged: $-\left(\frac{b}{b+2}\right)^2$.

Case II: Bart Believes that Lisa Is Rash.

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

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

where $\tilde{x} = [b(1+\tilde{y})]/[2(1+b)]$.

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 value of this ARA game is strictly larger than the minimax value.

Case III: Bart Believes that Lisa Is Conservative.

Suppose Bart believes that Lisa is risk averse, calling with $\tilde{y} > b/(b+2)$. Then

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

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

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

For $\tilde{y} > b/(b+2)$, the quantity in the square brackets is strictly positive. Thus, when $x < \tilde{y}$, Bart should bet.

The value V of this game is

$$V = \int_0^{\tilde{y}} \tilde{y} - (1+b)(1-\tilde{y}) + \int_{\tilde{y}}^1 \tilde{y} + (1+b)(x-\tilde{y}) - (1+b)(1-x).$$

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: 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

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

If Bart and Lisa draw from a bivariate, possibly discrete distribution $W(x, y)$ (e.g., a deck of cards) then the analysis is trivial (in G. H. Hardy's sense): Bart's distribution for Y is the conditional $W(y|X = x)$, and he knows that Lisa's analysis is symmetric.

Note: Some may be uncomfortable with the specificity in requiring Bart to assume that Lisa thinks his bluffing function is $g\tilde{(x)}$. They might argue that Bart could not guess that 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. But when Bart integrates over that space with respect to his subjective distribution, he then obtains the $\tilde{g}$ that he needs for this analysis.

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 $\tilde{y} = \sqrt[p+1]{\frac{1}{2}\frac{b}{1+b}}$. 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.

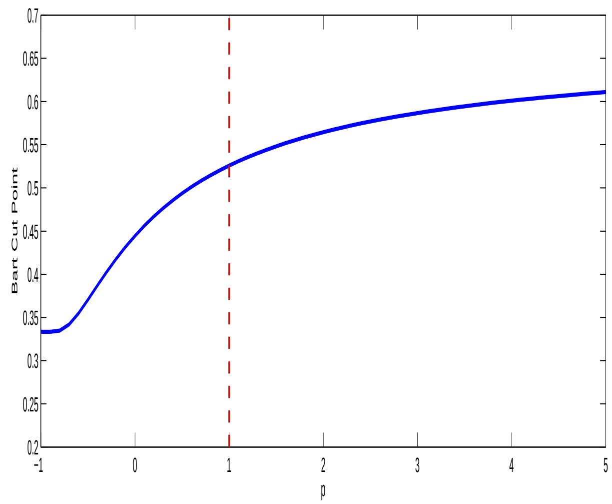

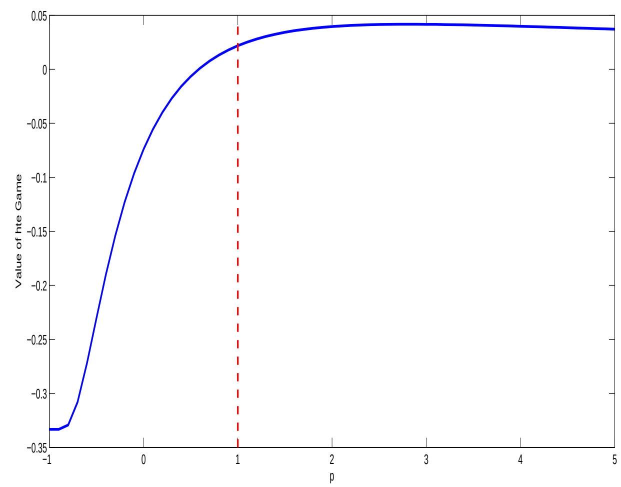

The left panel shows, for $b = 2$, the minimum value of x at which Bart should bet as a function of p. The right panel shows the game value, to Bart, as a function of p.

Continuous Bets

Consider a modification of the Borel Game, in which Bart is not constrained to bet any amount on some interval $(\epsilon, K]$.

Define 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 <u>conditional on</u> his decision to bet.

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

Let $\mathbb{P}_{h(\cdot|x)}[\cdot]$ and $\mathbb{E}_{h(\cdot|x)}[\cdot]$ denote the probability and expectation computed using the probability measure induced by the density $h(\cdot|x)$.

Bart must "mirror" Lisa's analysis given that she observes Bart's bet $B_x = b$. Define

$\tilde{g}(x)$**:** Bart's belief about Lisa's belief of the probability that he decides to bet with $X = x$.

$\tilde{h}(b|x)$**:** Bart's belief about Lisa's belief of the density on $[\epsilon, K]$ that Bart uses to bet.

$\tilde{f}(x|b)$**:** Bart's belief about Lisa's posterior density for X after she observes that he bets b:

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

Given $g(x)$ and $h(\cdot|x)$, then $V_x = \mathbb{E}_{g(x),h(\cdot|x)}\left[V_B | X = x\right]$:

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

Bart's first-order ARA solution is

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

To solve for $\{g^*(x), h^*(\cdot|x)\}$, he studies Lisa's strategy and rolls back.

Bart believes Lisa will form the posterior assessment $\tilde{f}(\cdot|b)$ on his X, so for $Y = y$, Bart believes Lisa thinks her probability of winning is

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

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

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

So 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

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

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

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

Now Bart is able to compute the following quantities:

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

Combining these expressions shows:

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

Theorem: For $x \in [0,1]$ and given $\tilde{f}(\cdot|b)$ positive and continuous in $b \in [\epsilon, K]$, let

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

Then, Bart's first-order ARA solution is

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

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

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

Example: Lisa has a step-function posterior.

To illustrate the use of the theorem to find the ARA solution in a Borel game with continuous bets, suppose $\tilde{f}(\cdot|b)$ is of the following form:

$$\tilde{h}(x|b) = \begin{cases} \frac{1+K}{1+b} & \textbf{if } 0 \le x \le \frac{1+b}{1+K} \\ 0 & \textbf{otherwise}. \end{cases}$$

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

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

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 the theorem, $g^*(x) = 1$; 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 the theorem, $g^*(x) = 1$ 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 the theorem, $g^*(x) = 1$ 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 the Theorem, $g^*(x) = 1$ and $h^*(b|x) = \delta(b-K)$; i.e., Bart will bet K w.p. 1.

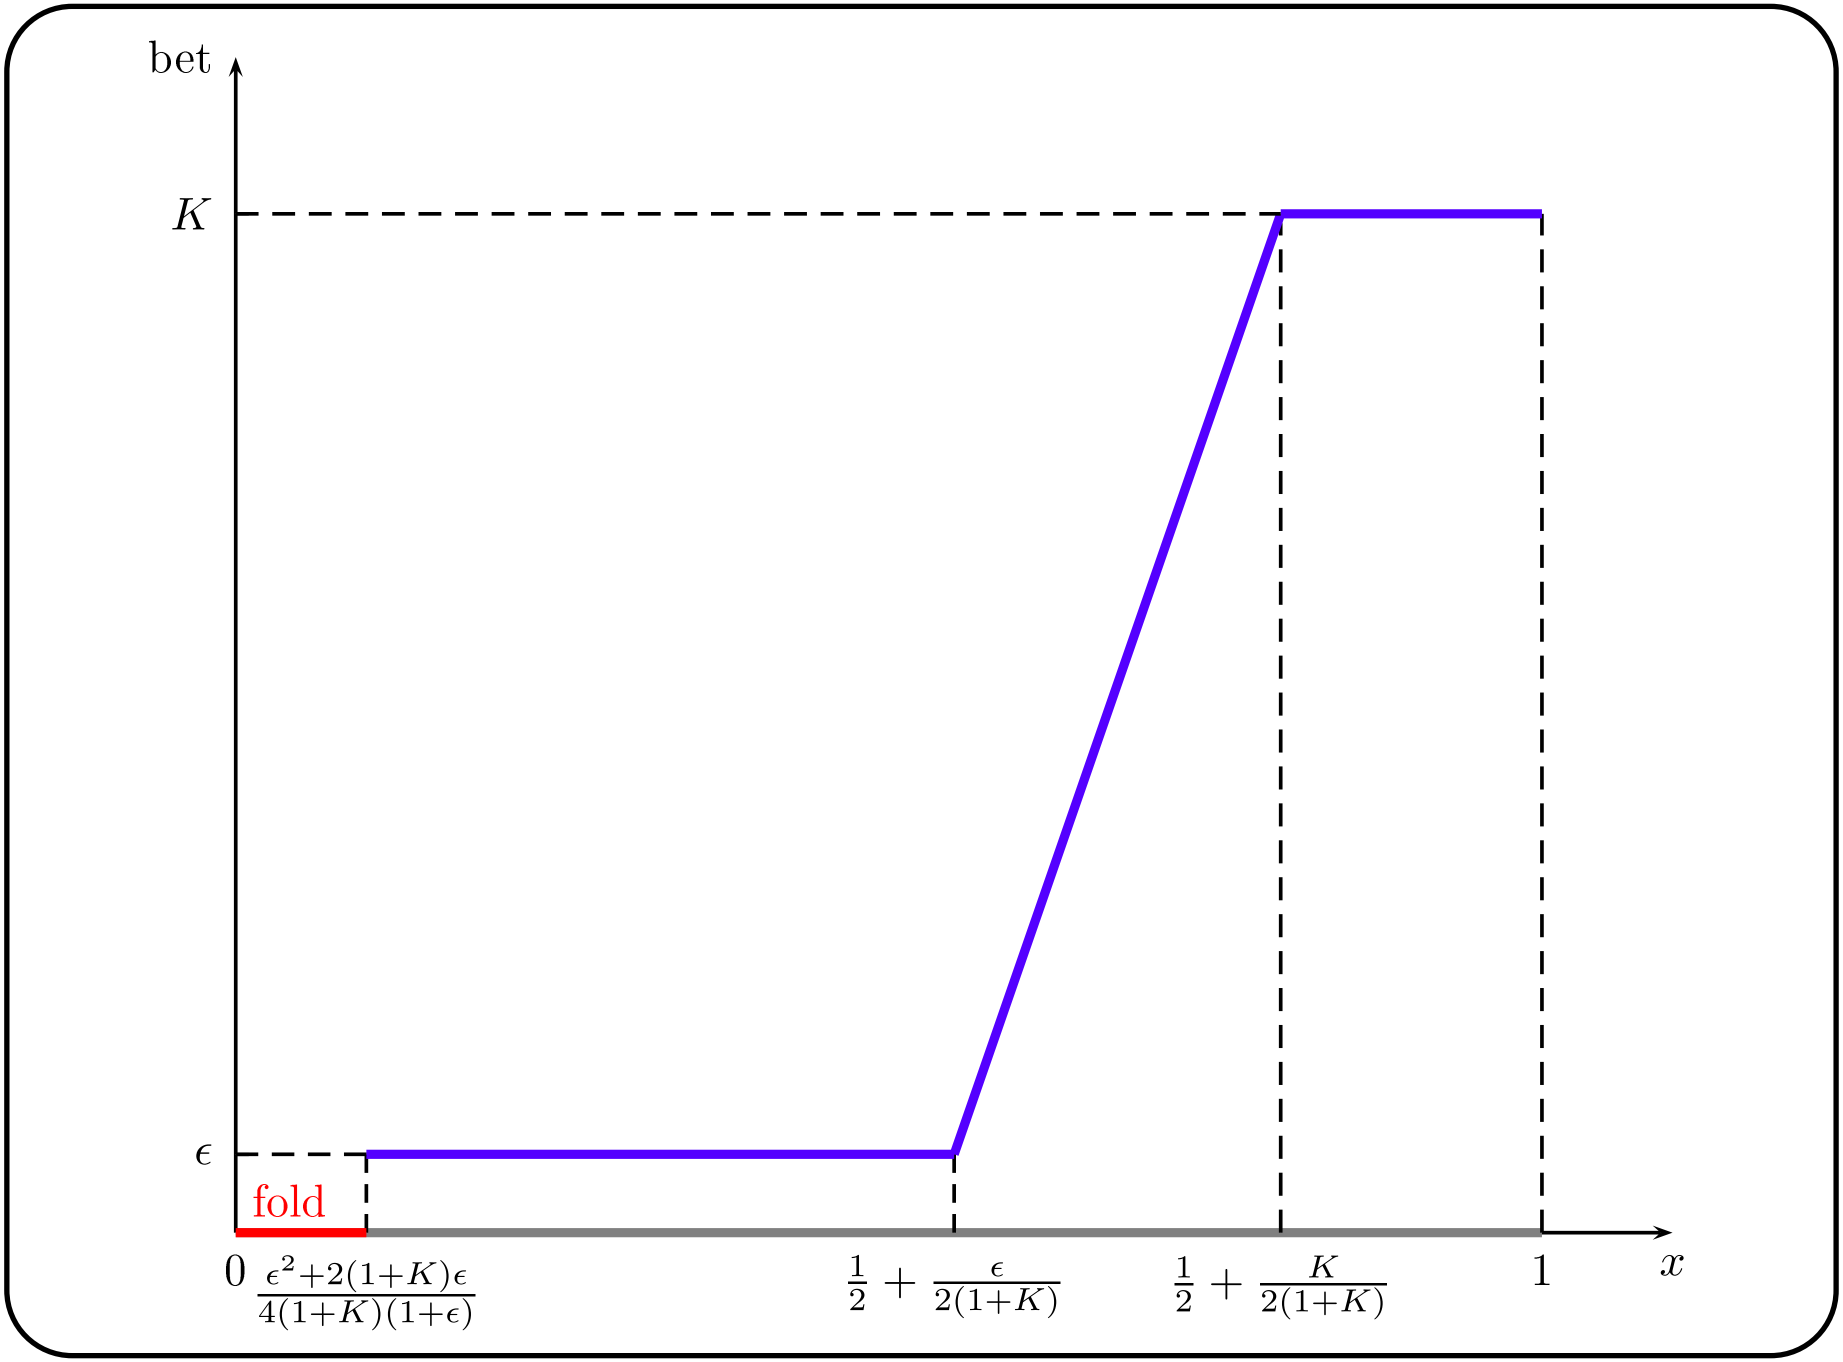
bet
K
ϵ
fold
0
$\frac{\epsilon^2+2(1+K)\epsilon}{4(1+K)(1+\epsilon)}$
$\frac{1}{2}+\frac{\epsilon}{2(1+K)}$
$\frac{1}{2}+\frac{K}{2(1+K)}$
1
x

4. Conclusions

The ARA approach has a number of attractive features:

- It is simpler to calculate than Nash equilibria. Sort of.
- It can take advantage of soft information.
- Its decisions that seem closer to the kind of strategizing that humans use.

In particular, for the Borel game, it is notable that many things that are difficult or still unresolved are straightforward (if tedious). The minimax solution was found by Borel; Bellman & Blackwell extended it to a game with two levels of bet, as did von Neumann and Morgenstern. Karlin and Restrepo (1957) obtained 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. And there are no good minimax solutions for games with dependent non-uniform distributions.