

Adversarial Risk Analysis: An Overview

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 strategic 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 (ARA). It extends the decision analysis proposed by Kadane and Larkey (1982) and Raiffa (1982) [and rejected by Harsanyi (1982)].

Myerson (1991, p. 114) points up the issues 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 ARA. In some cases this may be viewed as a Bayesian version of Level- k thinking (Stahl and Wilson, 1995).

The ARA framework builds a model for the decision-making process of the opponents, and uses that to develop a subjective distribution on their actions. The model can be complex; e.g., it can be a mixture over several simpler models.

ARA conveniently partitions the uncertainty in the problem into

- aleatory uncertainty, which describes the randomness in the outcome conditional on the actions chosen;
- epistemic uncertainty, which describes Bayesian beliefs about the utilities, information and capabilities of an opponent; and
- concept uncertainty, which describes uncertainty about the solution concept that an opponent is using.

Parnell and Merrick (2011) compared Probabilistic Risk Analysis with various intelligent adversary methods, and preferred ARA, in large part because this division enables more transparent modeling.

We now explore ARA more formally.

2. Auctions

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

Daphne's utility function is linear in money and her personal valuation of the book is d_0 . If money is infinitely divisible, her choice set is $\mathcal{D} = \mathbb{R}^+$. so her expected utility from a bid of d is $(d_0 - d)\mathbb{P}[d > V^*]$. Thus Daphne should maximize her expected utility by bidding

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

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

2.1 Two-Person Auctions

Consider a two-person first-price independent private-value sealed-bid auction among risk-neutral opponents (hereafter, **auction**).

Specifically, suppose Daphne and Apollo are bidding for a first edition of the Theory of Games and Economic Behavior.

Aleatory uncertainty arises in this situation if the value of the book is a random variable. Perhaps it is damaged, or has marginalia by John Nash. So the profit or loss, conditional on the bids, is a random variable.

Epistemic uncertainty arises because neither opponent knows the value (or expected value) of the book to the other.

Concept uncertainty arises because Daphne does not know how Apollo will determine his bid. Will he be non-strategic? Will he seek a Bayes Nash equilibrium? Will he use level- k thinking?

To begin, we assume there is no aleatory uncertainty—both bidders know their personal value. We analyze the game from the perspective of Daphne, whose certain value for the book is x_0 .

She may think that Apollo is non-strategic, and that he bids some random fraction P of his true value V . As a Bayesian, Daphne has a subjective distribution f_1 over V and a subjective distribution f_2 over P . In that case her belief about the distribution of Apollo's bid is

$$G(y) = \mathbf{IP}[PV \leq y] = \int_0^\infty \int_0^{y/v} f_2(p) f_1(v) dp dv.$$

Then Daphne's optimal bid is

$$x^* = \mathbf{argmax}_{x \in \mathbb{R}^+} (x_0 - x)G(x)$$

since this maximizes her expected utility. If she wins, her profit is $(x_0 - x)$, and her subjective probability of winning is $G(x)$.

But Daphne may think Apollo is strategic. Perhaps he seeks a **Bayes Nash equilibrium** (BNE) solution.

The BNE formulation makes a strong common knowledge assumption: both Apollo and Daphne have distributions H_D and H_A for each other's valuation, and each knows both distributions and knows that the other knows them.

∞ This leads to solving a system of first-order ODEs. For an asymmetric auction, when $H_A \neq H_D$, no solution algorithm exists, although it is known that if H_A and H_D are differentiable then a unique solution exists and is also differentiable (LeBrun, 1999).

Previous attempts at solutions are based on the **backshooting algorithm**. But Fibich and Gavish (2011) have recently shown that all such algorithms are inherently unstable. Kirkegaard (2009) established results on crossing conditions in the solutions, and Hubbard et al. (2012) used these to provide a visual test, but their work fails in examples. Tim Au (2014) has an algorithm that succeeds, based on the limit of discretized bids and points of indifference.

From an ARA perspective, the common knowledge assumption can be replaced by something more reasonable. Daphne has a subjective opinion about the distribution H_D that she thinks Apollo has for her value, and she has a subjective opinion about H_A , the distribution she believes he thinks is her distribution for his value. The H_A and H_D represent her epistemic uncertainty.

In that framework, Apollo solves the BNE equations:

$$\mathbf{argmax}_{d \in \mathbb{R}^+} (D^* - d)G(d) \sim F$$

$$\mathbf{argmax}_{a \in \mathbb{R}^+} (A^* - a)F(a) \sim G.$$

where $D \sim H_D$ and $A \sim H_A$. The equilibrium solution gives G , her best guess, under the BNE solution concept, of the distribution for Apollo's bid.

Now Daphne should step outside the BNE framework and solve

$$x^* = \mathbf{argmax}_{x \in \mathbb{R}^+} (x_0 - x)G(x)$$

where x_0 is her true value. This is a **mirroring argument**.

Note: A nice feature of the ARA mirror equilibrium formulation is that it allows a new class of problems in n -person games. If Bob is also bidding for the book, then Daphne can have opinions about what Apollo thinks about Bob and what Apollo thinks Daphne thinks about Bob that are not expressible in the BNE common-knowledge framework.

As long as all of Daphne's opinions are coherent, then there is a solution that gives her best guess about the bidding distributions of each opponent, allowing her to find the solution that maximizes her expected utility.

Note: In terms of concept uncertainty, we first took Apollo to be non-strategic, and then assumed he used the BNE concept. In practice, Daphne might have probability p_1 that he is non-strategic, probability p_2 that he uses BNE, probability p_3 that he is a level-1 reasoner, and so forth. (There are many more possible solution concepts.)

She would then solve her decision theory problem under each scenario, and form the mixture distribution $G(x)$ with each solution component weighted by the p_i and then solve

$$x^* = \operatorname{argmax}_{x \in \mathbf{R}_+} (x_0 - x)G(x)$$

A fourth solution concept is **level- k thinking**. If Daphne is a level-0 thinker, she bids non-strategically. If she is a level-1 thinker, she believes Apollo is a level-0 thinker, and makes her best response given her subjective assessment of the probabilities. If she is a level-2 thinker, she believes Apollo is a level-1 thinker, and so forth.

The “I think that you think that I think ...” reasoning becomes intricate. (Recall Vizzini’s analysis of the iocaine powder in The Princess Bride). An example will be more clear: Suppose Daphne is a level-2 thinker. She believes Apollo is a level-1 thinker who thus believes that she is non-strategic.

Specifically, assume her subjective belief is that Apollo thinks her value for the book has distribution $F_1(v)$ supported on $[\$100, \$200]$ and that she bids a proportion of her value $F_2(p) = p^9$, $0 \leq p \leq 1$. Then

$$g(y) = \int_0^\infty f_1(v) f_2(y/v) \frac{1}{v} dv = \int_0^\infty g_1(v) 9(y/v)^8 v^{-1} dv \propto y^8$$

so $G(y) = (y/200)^9$.

Apollo's best response is to bid x^* such that

$$x^* = \operatorname{argmax}_{x \in \mathbf{R}_+} (X_0 - x)G(x)$$

where X_0 is his true value (a random variable to Daphne). He should take the derivative, set it to 0, and solve:

$$0 = \frac{d}{dx} [(X_0 - x)G(x)] = 9 \frac{x^8}{200^9} (X_0 - x) - \left(\frac{x}{200}\right)^9.$$

So Apollo's bid should be 90% of his true value X_0 .

Daphne does not know Apollo's true value, but suppose she thinks it has the triangular distribution on [\$140, \$200] with peak at \$170. Since Apollo should bid 90% of his true value, Daphne believes that his bid will be a random variable with triangular distribution $F(x)$ that is supported on [\$126, \$180] with peak at \$153.

Finally, for y_0 her (known) true value for the book, Daphne solves

$$y^* = \operatorname{argmax}_{y \in \mathbf{R}_+} (y_0 - y)F(y).$$

2.2 More Than Two Bidders

An important advantage of ARA is that it enables a more nuanced treatment of many-player games. Specifically, the ARA formulation allows one to frame fresh problems in auction theory when there are more than two bidders, by permitting asymmetric models for how each opponent views the others.

If Bonnie is a level-1 thinker, then she assumes that Alvin and Clyde are non-strategic, and there is no novelty in the analysis. She has distributions over the non-strategic bids of each, and chooses her bid according to the maximum of those. Specifically, she has a subjective distribution F_A over Alvin's bid A and a subjective distribution F_C over Clyde's bid C , and she calculates the distribution F of $\max\{A, C\}$. Then she makes the bid

$$b^* = \operatorname{argmax}_{b \in \mathbb{R}^+} (b_0 - b)F(b),$$

where b_0 is her true value for the book.

Now suppose Bonnie is a level-2 thinker. She thinks that Alvin has a belief about the distribution of her bid and also Clyde's bid; similarly, she thinks Clyde has a distribution for her bid and for Alvin's. Let $F_{IJ}(x)$ be what Bonnie thinks player I thinks is the distribution for player J 's bid, and $G_{IJ}(x)$ be her belief about what player I thinks is the distribution for player J 's value.

Her level-2 analysis assumes both Alvin and Clyde are level-1 thinkers who believe their opponents are level-0 thinkers, then knowing F_{IJ} directly determines G_{IJ} .

The level-2 ARA formulation means that Bonnie thinks Alvin will make the bid $a^* = \max\{a_B^*, a_C^*\}$ for

$$\begin{aligned} a_B^* &= \operatorname{argmax}_{a \in \mathbb{R}^+} (a_0 - a) \mathbf{IP}[B^* < a] \\ a_C^* &= \operatorname{argmax}_{a \in \mathbb{R}^+} (a_0 - a) \mathbf{IP}[C^* < a], \end{aligned}$$

where a_0 is Alvin's true value, B^* is a random variable whose distribution is Alvin's opinion about Bonnie's bid, and C^* is a random variable whose distribution is Alvin's opinion about Clyde's bid.

Bonnie does not know a_0 , and she does not know Alvin's distributions for the bids, but as a Bayesian, she has a subjective opinion about these. She regards a_0 as a random variable with distribution G_{BA} , and her best guess is that B^* and C^* have distributions F_{AB} and F_{AC} , respectively.

In order to find F_{AB} , Bonnie uses the fact that Alvin thinks she is a level-0 thinker. He views her as non-strategic, and thus thinks her bid follows some probability distribution, perhaps an unknown proportion of her unknown true value, so both the unknown proportion and the true value can be modeled as random variables.

Thus, Bonnie's opinion about the distribution of Alvin's bid is found by solving

$$A_B^* = \operatorname{argmax}_{a \in \mathbb{R}^+} (A_0 - a)F_{AB}(a)$$

$$A_C^* = \operatorname{argmax}_{a \in \mathbb{R}^+} (A_0 - a)F_{AC}(a)$$

and then assuming that Alvin bids the larger of those two random variables. So his bid is $A^* = \max\{A_B^*, A_C^*\}$.

Similarly, Bonnie belief about Clyde's bid C^* is that it has the distribution of $\max\{C_A^*, C_B^*\}$, where

$$C_A^* = \operatorname{argmax}_{c \in \mathbb{R}^+} (C_0 - c)F_{CA}(c)$$

$$C_B^* = \operatorname{argmax}_{c \in \mathbb{R}^+} (C_0 - c)F_{CB}(c)$$

and C_0 is Clyde's true value, with distribution G_{BC} , since it is unknown to Bonnie.

Just as before, Bonnie uses her beliefs about what Clyde thinks about Alvin's non-strategy and her non-strategy to identify F_{CA} and F_{CB} , respectively, and thus finds the distribution of C^* .

Bonnie has calculated her distribution for Alvin's bid A^* and Clyde's bid C^* . Now she should place the bid


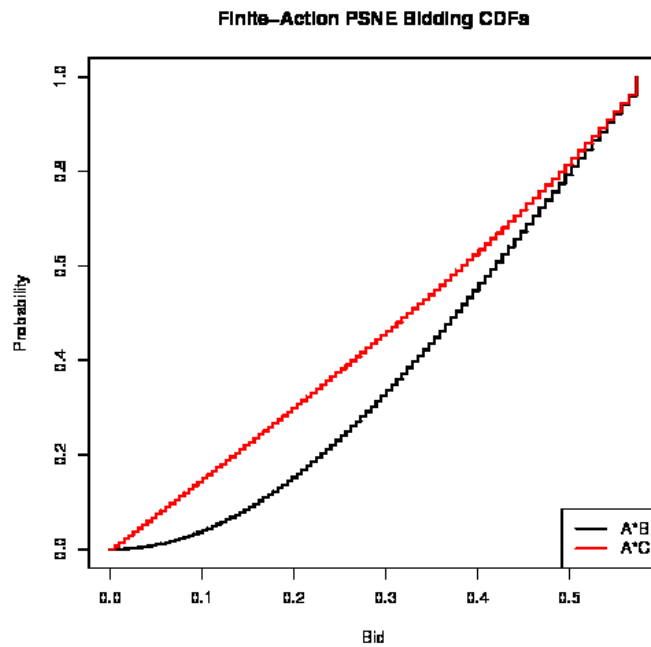
$$b^* = \operatorname{argmax}_{b \in \mathbb{R}^+} (b_0 - b)\mathbf{IP}[\max\{A^*, C^*\} < b].$$

For example, suppose Bonnie believes that Alvin thinks her value for the first edition is $\text{Beta}(1,1)$, and that Clyde's value is $\text{Beta}(2, 1)$.


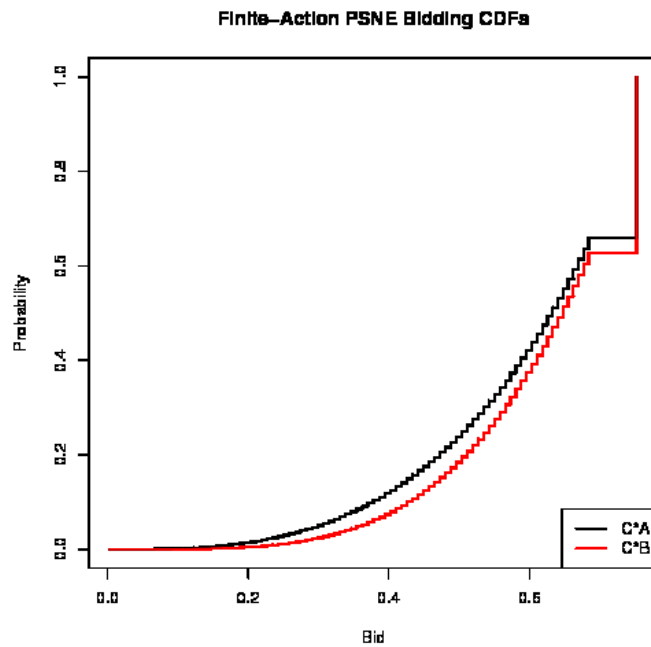
Similarly, she believes that Clyde thinks her value for the first edition is $\text{Beta}(4, 1)$, and she thinks Clyde thinks Alvin's value is $\text{Beta}(3,1)$.

One can now use the BID algorithm developed by Tim Au to solve this three-person game. In this application, of course, we are supporting Bonnie.

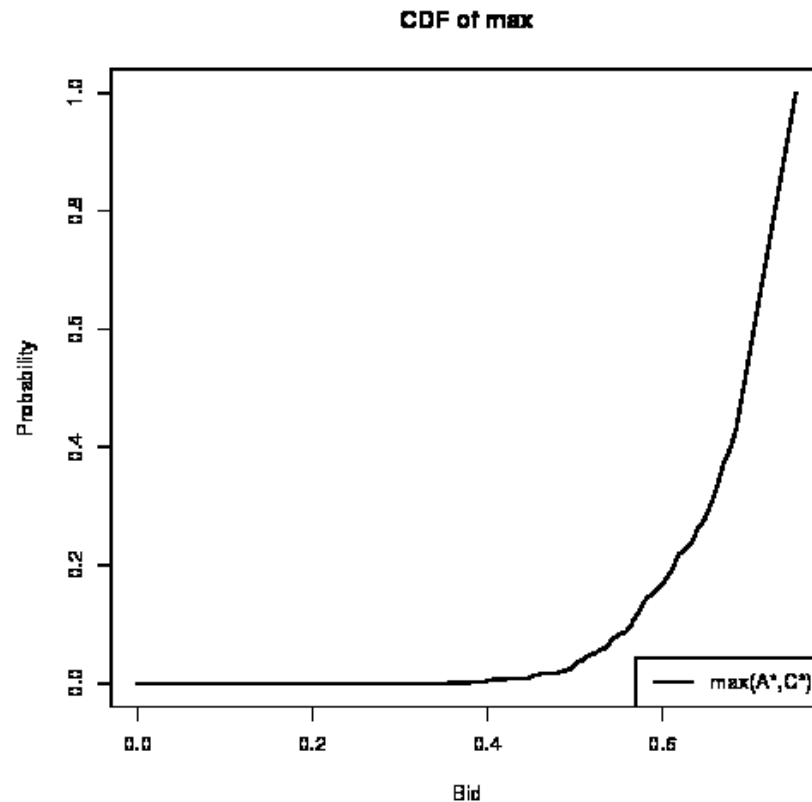
Alvin's bid is $A^* = \max\{A_B^*, A_C^*\}$. The left panel shows the cdfs of A_B^* and A_C^* , and the right shows the cdf of A^* .



Clyde's bid is $C^* = \max\{C_A^*, C_B^*\}$. The left panel shows the cdfs of C_A^* and C_B^* , and the right shows the cdf of C^* .



This figure shows the distribution of the maximum of the optimal bids for Alvin and Clyde.



Under these assumptions about the beliefs of Alvin and Clyde, if Bonnie's true value for the book is 0.95, then her optimal bid is 0.7523.

Now consider the use of the mirror equilibrium solution concept when there are three bidders. This assumes that all bidders are solving the problem in the same way, but with possibly different subjective distributions over all unknown quantities.

The two-person system extends so that the basic problem is to solve

$$\begin{aligned}
 A^* &= \mathbf{argmax}_{a \in \mathbb{R}^+} (A_0 - a)F_A^*(a) \\
 B^* &= \mathbf{argmax}_{b \in \mathbb{R}^+} (B_0 - b)F_B^*(b) \\
 C^* &= \mathbf{argmax}_{c \in \mathbb{R}^+} (C_0 - c)F_C^*(c)
 \end{aligned} \tag{1}$$

from the perspective of each of the players, where $F_I^*(x)$ is what bidder I thinks is the chance that a bid of x will win. Bonnie does not know F_I^* , but she can use ARA to find F_I , which is her belief about what each opponent thinks is the chance that a given bid is successful.

The figure shows the notation that describes what Bonnie thinks each person believes about the distributions for each of the other bidders' true values. The G_{IJ} is what Bonnie thinks bidder I believes is distribution of the true value for bidder J , and G_{IJK} is the distribution that Bonnie thinks bidder I thinks bidder J has for the true value of the book to bidder K .

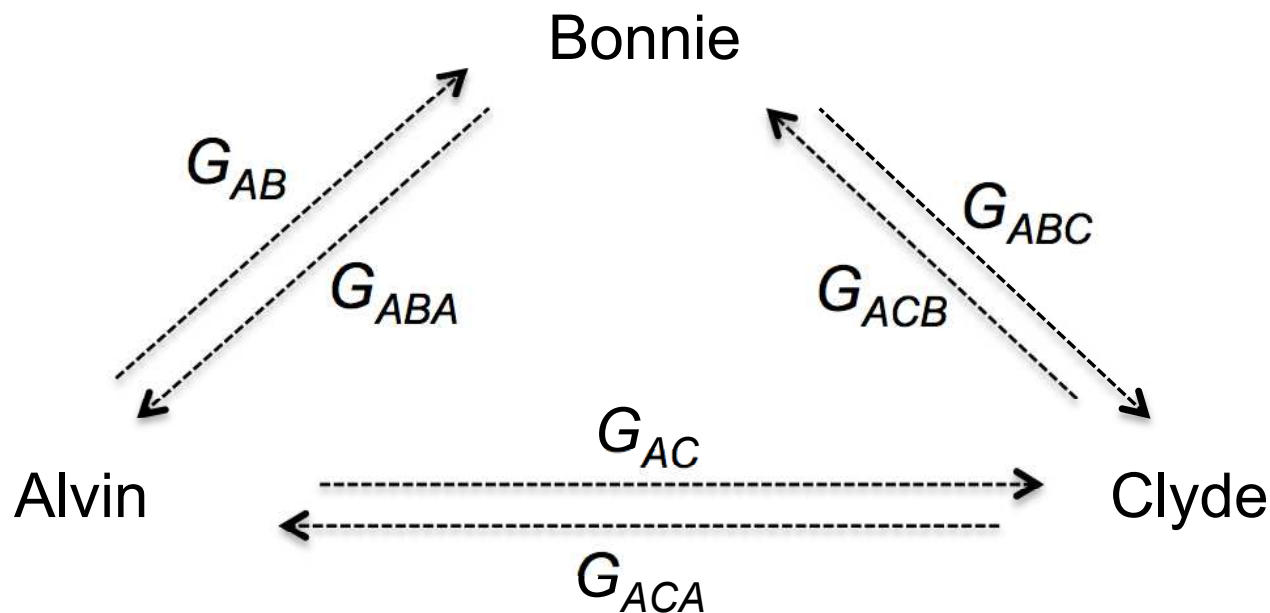


Figure 1: A representation of what Bonnie believes about the opinions held by each of the bidders regarding the value of the book to each the other bidders.

First, she models Alvin's logic. Bonnie thinks he obtains his distribution for her bid by solving (1) with $A_0 \sim G_{ABA}$, $B_0 \sim G_{AB}$, and $C_0 \sim G_{ABC}$. Since he, like Bonnie, does not know the true F_I^* , he must develop his own beliefs about them.

Here, his F_A is the distribution of the maximum of B^* and C^* , F_B is the distribution of the maximum of A^* and C^* , and F_C is the distribution of the maximum of B^* and C^* . After numerical computation to find the equilibrium solution, he obtains F_{AB} , his belief about the distribution of Bonnie's bid.

Next, Alvin considers Clyde. Bonnie thinks he solves (1) with $A_0 \sim G_{ACA}$, $B_0 \sim G_{ACB}$, and $C_0 \sim G_{AC}$. He proceeds as before, and obtains F_{AC} , his belief about the distribution of Clyde's bid. From this, Bonnie thinks his distribution for the probability of winning with a bid of a is F_A , where F_A is the distribution of the maximum of $B \sim F_{AB}$ and $C \sim F_{AC}$.

Bonnie's analysis for Clyde is analogous. To find Clyde's distribution for Bonnie's bid, she thinks he solves (1) with $A_0 \sim G_{CBA}$, $B_0 \sim G_{CB}$, and $C_0 \sim G_{CBC}$ to obtain F_{CB} . Similarly, to find Clyde's distribution for Alvin's bid, he uses $A_0 \sim G_{CA}$, $B_0 \sim G_{CAB}$, and $C_0 \sim G_{CAC}$ to obtain F_{CA} . Putting these together, Bonnie thinks that Clyde thinks the probability that a bid of c will win is $F_C(c)$, which is the distribution of the maximum of $A \sim F_{CA}$ and $B \sim F_{CB}$.

Based on this reasoning, Bonnie thinks that Alvin's bid will be

$$A^* = \operatorname{argmax}_{a \in \mathbb{R}^+} (A_0 - a)F_A(a) \sim F_{BA},$$

where $A_0 \sim G_{BA}$.

Bonnie thinks Clyde's bid will be

$$C^* = \operatorname{argmax}_{c \in \mathbb{R}^+} (C_0 - c)F_C(c) \sim F_{BC},$$

where $C_0 \sim G_{BC}$. From this, the chance that a bid of b will win is $F_B(b)$, where F_B is the distribution of the maximum of $A^* \sim F_{BA}$ and $C^* \sim F_{BC}$. Now Bonnie uses her known value b_0 and solves

$$b^* = \operatorname{argmax}_{b \in \mathbb{R}^+} (b_0 - b)F_B(b)$$

to obtain her best bid under the mirror equilibrium solution concept.

Lebrun (1999, 2006) shows that an equilibrium solution always exists, and that, under a mild log concavity condition, the equilibrium is unique.

3. ARA in General

In ARA one takes the side of one agent, using only her beliefs and knowledge, rather than assuming common knowledge and trying to solve all of the agents' problems simultaneously. The selected agent must have

- a subjective probability about the actions of each opponent,
- subjective conditional probabilities about the outcome for every set of possible choices, and
- perfect knowledge of her own utility function.

Thus Daphne believes Apollo has probability $\pi_D(a)$ of choosing action $a \in \mathcal{A}$. She has a subjective probability $p_D(s | d, a)$ for each possible outcome $s \in \mathcal{S}$ given every choice $(d, a) \in \mathcal{D} \times \mathcal{A}$. And she knows her own utility $u_D(d, a, s)$ for each combination of outcome and pair of choices.

Daphne maximizes her expected utility by choosing the action d^* such that

$$\begin{aligned} d^* &= \mathbf{argmax}_{d \in \mathcal{D}} \mathbf{IE}_{\pi_D, p_D} [u_D(d, A, S)] \\ &= \mathbf{argmax}_{d \in \mathcal{D}} \int_{s \in \mathcal{S}} \int_{a \in \mathcal{A}} u_D(d, a, s) p_D(s | d, a) \pi_D(a) da ds \end{aligned}$$

where A is the random action chosen by Apollo and S is the random outcome that results from choosing A and d .

In practice, the most difficult quantity to obtain is $\pi_D(a)$. The $p_D(s | d, a)$ is found by conventional risk analysis and $u_D(d, a, s)$ is a personal utility.

Previously, we laid out ARA methods for obtaining $\pi_D(a)$, in the cases of the the non-strategic opponent, the Nash equilibrium seeking opponent, the opponent whose analysis mirrors that of the decision-maker, and the opponent who is a level- k thinker. Implementing these approaches imposes different cognitive loads upon the analyst.

The following shows how the cognitive load depends upon the kind of ARA. Each row corresponds to a different level of reasoning in level- k thinking.

The table displays the quantities that Daphne must assess in order to implement a level- k analysis. Row 0 corresponds to the utilities and beliefs of Daphne and Apollo, as perceived by themselves. Subsequently, row k contains the additional utilities and probabilities that Daphne would have to assess in order to perform a level- k analysis.

- The first column contains what Daphne believes are the utility functions that Apollo ascribes to her.
- The second column contains the probabilities of the outcome, conditional on both her action and Apollo's, that she believes Apollo ascribes to her.
- The third column contains her opinion of what Apollo thinks is her distribution for he will do.
- The fourth column contains the utility functions she ascribes to Apollo.
- The fifth column contains the conditional probabilities of the outcome, given her choice and Apollo's, that she ascribes to Apollo.
- The sixth column contains what she thinks is Apollo's distribution over her choice.

	1	2	3	4	5	6
0	u_D	$p_D(\cdot d, a)$	$\pi_D(a)$	u_A	$p_A(\cdot d, a)$	$\pi_A(d)$
1	U_D^1	$P_D^1(\cdot d, a)$	$\Pi_D^1(a)$	U_A^1	$P_A^1(\cdot d, a)$	$\Pi_A^1(d)$
2	U_D^2	$P_D^2(\cdot d, a)$	$\Pi_D^2(a)$	U_A^2	$P_A^2(\cdot d, a)$	$\Pi_A^2(d)$
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots

The upper case characters in rows 1 and higher indicate that these quantities are all random variables to Daphne.


In terms of the table, different solution concepts require information in different cells:

- Traditional game theory requires cells (0,1), (0,2), (0,4), (0,5) and assumes that these are common knowledge.
- The non-strategic adversary analysis requires cells (0,1), (0,2) and (0,3), where the (0,3) cell is assessed from historical data and/or expert opinion.
- When the adversary seeks a Nash equilibrium solution, the analysis requires cells (0,1), (0,2) and (1,1), (1,2), (1,4) and (1,5). It uses these last four cells to infer cell (0,3).
- The level- k adversary approach requires cells (0,1), (0,2) and:
 - for a level-1 analysis, cells (1,4), (1,5) and (1,6) can produce (0,3);
 - for a level-2 analysis, cells (1,1), (1,2) and (1,3) produce (1,6), which, with (1,4), (1,5) can then produce (0,3);
 - and so forth for larger k .
- The mirror equilibrium approach requires cells (0,1), (0,2) and uses a consistency condition between (1,4), (1,5), (1,6) and (1,1), (1,2) and (1,3) to produce (0,3).

The main message is that all of these methods entail significant effort.

4. Routing Games

The most famous routing game is Nash. But a simpler game with practical importance is to route a convoy through a city street network when an adversary may place IEDs.



The Defender has imperfect information about the placement of the IEDs, and the Attacker's resources and utilities. Symmetrically, the Attacker has probabilistic knowledge about route choice, convoy value, and Defender utilities. But Harsanyi's common-knowledge analysis is untenable.

As with auctions, this leads to coupled probability equations. But here the stochastic payoffs for the Attacker and the Defender have additive structure.

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

$\tilde{\mathbf{X}}$: the random variable which the Defender uses to model the Attacker's gain matrix—it has probability distribution F ;

$\tilde{\mathbf{Y}}$: the random variable that the Defender uses to describe the Attacker's beliefs about the Defender's loss matrix—it has probability distribution G ;

$\tilde{\mathbf{a}}$: the random vector that the Defender uses to model the Attacker's decision—it has distribution P with support \mathcal{A} ;

$\tilde{\mathbf{r}}$: the random vector that the Defender uses to model the Attacker's belief about the Defender's decision—it has distribution Q with support \mathcal{D} .

ARA Algorithm: Assume $\tilde{\mathbf{X}} \sim F$, $\tilde{\mathbf{Y}} \sim G$, $\mathbf{a} \in \mathcal{A}$ and $\mathbf{r} \in \mathcal{D}$.

1. **Initialize.** The Defender starts with a pair of probability distributions (P_0, Q_0) , where P_0 is a distribution for $\tilde{\mathbf{a}}$ and Q_0 is a distribution for $\tilde{\mathbf{r}}$.

2. **Iterate.** Given (P_k, Q_k) , iterate to convergence.

2.A Simulate many realizations of $\tilde{\mathbf{X}}$. For each, the Defender mimics the Attacker's analysis and solves $\mathbf{a}^* = \max_{\mathbf{a} \in \mathcal{A}} \mathbf{a}' \tilde{\mathbf{X}} \mathbb{E}_{Q_k}[\tilde{\mathbf{r}}]$. Since $\tilde{\mathbf{X}} \sim F$, the resulting maximizer is a random variable, and the distribution of the solutions is an estimate of P_{k+1} , the updated distribution of $\tilde{\mathbf{a}}$; i.e.,

$$\operatorname{argmax}_{\mathbf{a} \in \mathcal{A}} \mathbf{a}' \tilde{\mathbf{X}} \mathbb{E}_{Q_k}[\tilde{\mathbf{r}}] = \tilde{\mathbf{a}} \sim P_{k+1}.$$

2.B Update Q_k of $\tilde{\mathbf{r}}$ using P_{k+1} by generating realizations of $\tilde{\mathbf{Y}}$ and solving

$$\operatorname{argmin}_{\mathbf{r} \in \mathcal{D}} \mathbb{E}_{P_{k+1}}[\tilde{\mathbf{a}}]' \tilde{\mathbf{Y}} \mathbf{r} = \tilde{\mathbf{r}} \sim Q_{k+1}.$$

3. **Terminate.** The Defender chooses $\mathbf{r}^* = \operatorname{argmin}_{\mathbf{r} \in \mathcal{D}} \mathbb{E}_{P^*}[\tilde{\mathbf{a}}]' \mathbf{Y} \mathbf{r}$. In this final step the Defender uses the true loss matrix \mathbf{Y} .

For this situation, we can obtain two theorems:

Theorem 1: A mirroring fixed-point for the system of equations exists.

Theorem 2: If there exists a total order $\succeq_{\mathcal{A}}$ on \mathcal{A} and $\succeq_{\mathcal{D}}$ on \mathcal{D} such that

1. $\tilde{V}_{a,r} := \mathbf{a}\tilde{\mathbf{X}}\mathbf{r}$ has increasing difference in (\mathbf{a}, \mathbf{r}) ,
2. $\tilde{W}_{a,r} := \mathbf{a}\tilde{\mathbf{Y}}\mathbf{r}$ has decreasing difference in (\mathbf{a}, \mathbf{r}) ,

then the ARA Algorithm converges to the mirroring fixed point.

The second theorem is technical, using submodularity, but is satisfied if most (some) of the losses or gains are of opposite sign. For example, it holds for zero-sum games.

In general, finding fixed-point solutions in game theory is hard. For the special structure of the routing game, we know that equilibria exist and can provide conditions under which a reasonable algorithm converges to the mirroring method solution.

The mirroring argument provides an explicit mechanism for modeling the reasoning of one's opponents. Previously, the decision-theoretic Bayesians who did game theory simply declared a distribution over the actions of their opponents.

Kadane (2009) points to a passage in Poe's The Purloined Letter that illustrates the naturalness of the ARA approach, in contrast to the minimax solution. Dupin recalls:

I knew one [school-boy] about eight years of age, whose success at guessing in the game of “even and odd” attracted universal admiration. This game is simple, and is played with marbles. One player holds in his hand a number of these toys and demands of another whether that number is even or odd. If the guess is right, the guesser wins one; if wrong, he loses one. The boy to whom I allude won all the marbles of the school. Of course he had some principle of guessing; and this lay in mere observation and admeasurement of the astuteness of his opponents.

For example, an arrant simpleton is his opponent, and, holding up his closed hand, asks, "Are they even or odd?" Our school-boy replies, "Odd," and loses; but upon the second trial he wins, for he then says to himself: "The simpleton had them even upon the first trial, and his amount of cunning is just sufficient to make him have them odd upon the second; I will therefore guess odd"; he guesses odd, and wins. Now, with a simpleton a degree above the first, he would have reasoned thus: "This fellow finds that in the first instance I guessed odd, and, in the second, he will propose to himself, upon the first impulse, a simple variation from even to odd, as did the first simpleton; but then a second thought will suggest that this is too simple a variation, and finally he will decide upon putting it even as before. I will therefore guess even"; he guesses even, and wins. Now this mode of reasoning in the schoolboy, whom his fellows termed "lucky," what, in its last analysis, is it?

It is merely, I said, an identification of the reasoner's intellect with that of his opponent.

5. 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:

$$\begin{aligned} \mathbf{IE}[V_x] &= -\mathbf{IP}[\mathbf{Bart\ folds}] + \mathbf{IP}[\mathbf{Bart\ bets\ and\ Lisa\ folds}] \\ &\quad + (1+b)\mathbf{IP}[\mathbf{Lisa\ bets\ and\ loses}] \\ &\quad - (1+b)\mathbf{IP}[\mathbf{Lisa\ bets\ and\ wins}]. \end{aligned}$$

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

$$\begin{aligned} \mathbf{IE}[V_x] &= -[1 - g(x)] + g(x)\mathbf{IP}[\text{Lisa folds} \mid \text{Bart bets}] \\ &\quad + (1 + b)g(x)x\mathbf{IP}[\text{Lisa bets} \mid \text{Bart bets}] \\ &\quad - (1 + b)g(x)(1 - x)\mathbf{IP}[\text{Lisa bets} \mid \text{Bart bets}]. \end{aligned}$$

For optimal play, Bart needs to find $\mathbf{IP}[\text{Lisa bets} \mid \text{Bart bets}]$.

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

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

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 $x \leq \tilde{y}$, 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

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

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

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

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 .

4.1 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 conditional on 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 $V_x = \mathbb{E}_{g(x),h(\cdot|x)} [V_B|X = x]$:

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

Bart's first-order ARA solution is

$$\{g^*(x), h^*(\cdot|x)\} \in \mathbf{argmax}_{g(x), h(\cdot|x)} \mathbf{IE}_{g(x), h(\cdot|x)} [V_B | X = x].$$

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

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

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

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

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

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

Now Bart is able to compute the following quantities:

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

Combining these expressions shows:

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

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

$$b^*(x) \in \underset{b \in [\epsilon, K]}{\operatorname{argmax}} \quad \tilde{y}^*(b) + 2(x - \tilde{y}^*(b))^+(1+b) - (1 - \tilde{y}^*(b))(1+b),$$

$$\Delta^*(x) \equiv \max_{b \in [\epsilon, K]} \tilde{y}^*(b) + 2(x - \tilde{y}^*(b))^+(1+b) - (1 - \tilde{y}^*(b))(1+b).$$

Then, Bart's first-order ARA solution is

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

$$h^*(b|x) = \delta(b - b^*(x)),$$

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:


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

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

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

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 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.



6. Conclusions

ARA allows the analyst to flexibly model the thought-process of the opponents. This fits naturally with a large body of modern work in behavioral game theory, and avoids awkward assumptions about rationality and common knowledge.

It also partitions the total uncertainty into usefully distinct parts (aleatory, epistemic, and concept uncertainty), which facilitates elicitation and calculation.

The talk described ARA perspectives in two game settings: convoy routing and *La Relance*. The examples find interestingly different results than one obtains under traditional solution concepts.

Also, the ARA formulation leads to new research questions, as with the n -person version of *La Relance* in which one can model all the pairwise beliefs that bidders have about each other's valuations, or the n -person auction.