1 Basic models and notions: game, strategy, best response, equilibrium

1a A very simple auction: quite informal introduction

An *auctioneer* wants to sell an *object;*[1] two *players*[2] want to buy the object. They do not cooperate. Each player submits to the auctioneer a bid.[3] The maximal bid determines the *winner.*[4] The auctioneer sells the object to the winner, the price being equal to his (maximal) bid. That is the 'visible' side of the process.

In order to construct a mathematical model, we assume the following 'invisible' side. The first player believes that the object is worth of S_1 dollars. The second — S_2 dollars.[5] These S_1, S_2, called *signals* or *valuations,* are independent random variables, distributed uniformly on $(0, 1)$.[6]

Assume for a moment that the first player knows the action A_2 (that is, the bid) of the second player (which is an oversimplification, of course).[7] Then his payoff Π_1 depends on his action (bid) A_1 as follows:

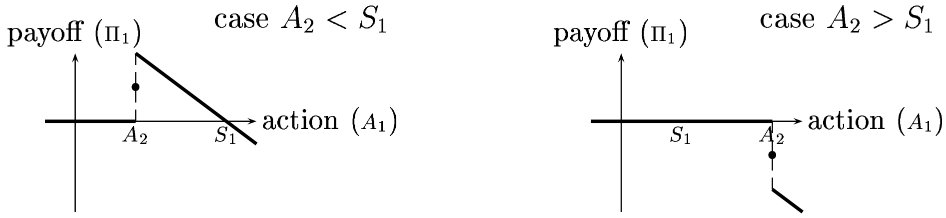

$$\Pi_1 = \begin{cases} 0 & \text{if } A_1 \in (-\infty, A_2), \\ \frac{1}{2}(S_1 - A_1) & \text{if } A_1 = A_2, \\ S_1 - A_1 & \text{if } A_1 \in (A_2, \infty). \end{cases}$$

Being discontinuous, that function does not reach its supremum, but anyway, for any $\varepsilon > 0$ the first player can get ε-close to the optimum by acting

$$A_1 = \begin{cases} A_2 + \varepsilon & \text{if } S_1 > A_2, \\ 0 & \text{otherwise.} \end{cases}$$

That strategy is an ε-best response to A_2. The problem with it is that the first player does not know A_2.

Assume now that the first player, not knowing A_2, knows[8] that A_2 is distributed uniformly on $(0, 1)$ (still an oversimplification). Now the payoff function is unknown, but its expectation

[1] A single indivisible object.

[2] Called also *bidders.* The auctioneer is not a player, she acts according to prescribed rules.

[3] *Sealed* bid; that is, not disclosed to the other player.

[4] If the two bids are equal, the winner is chosen at random (so-called *ties* breaking).

[5] Why X, Y are different? Maybe, the object is more useful to one player that to the other. Maybe, players have different private information (and/or intuition) about the object (and their ability to use it).

[6] You see, we strive to a *simple* model. In general, signals and valuations are different random variables, neither uniform nor independent.

[7] Or maybe he just believes that he knows A_2. In that case, following arguments are also his beliefs.

[8] Or just believes.

(average over all possible values of A_2) is known (to the first player):

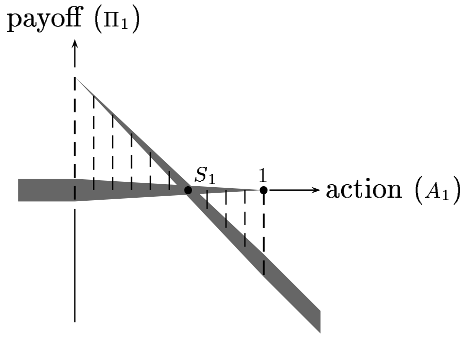

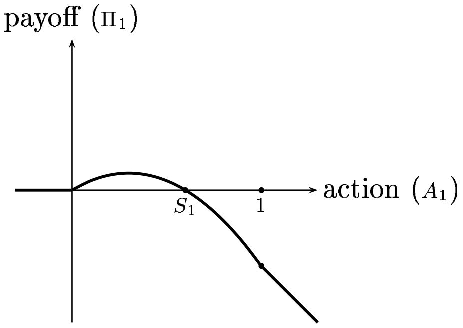

$$\text{Expected payoff} = \begin{cases} 0 & \text{if } A_1 \in (-\infty, 0), \\ (S_1 - A_1)A_1 & \text{if } A_1 \in (0, 1), \\ S_1 - A_1 & \text{if } A_1 \in (1, \infty). \end{cases}$$

The best response (to the uniform distribution of A_2) is

$$A_1 = \frac{1}{2} S_1 \,.$$

If the first player uses (for any reason) the strategy $A_1 = \frac{1}{2} S_1$ as his guide to action, then A_1 is distributed uniformly on $(0, 1/2)$ rather than $(0, 1)$. Then, why assume $A_2 \sim \mathrm{U}(0, 1)$? The assumption $A_2 \sim \mathrm{U}(0, 1/2)$ should be just more clever. Let us find the best response to it:

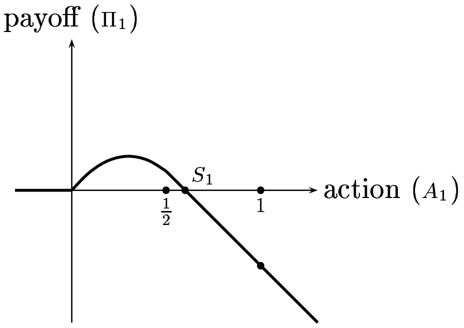

$$\text{Expected payoff} = \begin{cases} 0 & \text{if } A_1 \in (-\infty, 0), \\ (S_1 - A_1) \cdot 2A_1 & \text{if } A_1 \in (0, 1/2), \\ S_1 - A_1 & \text{if } A_1 \in (1/2, \infty). \end{cases}$$

Surprisingly, the best response is $A_1 = \frac{1}{2} S_1$ again. Good luck! We get a closed loop:

$$\text{bid} = \frac{1}{2}(\text{valuation}) \rightleftarrows \text{bid} \sim \mathrm{U}(0, 0.5)$$

If each player uses the strategy $A = \frac{1}{2} S$, then each player uses the best response to the strategy of the other player. (In other words, the strategy $A = \frac{1}{2} S$ is the best response to itself.) Such a behavior of players is called an equilibrium.

The game described above is radically different from, say, the chess game. The latter is deterministic. Being finite, it must have an ultimate solution. Fortunately, we are unable to find it in practice; but in principle, only one of the three possible results can appear if both players are clever enough: (a) white wins always; (b) black wins always; (c) the game ends in a draw, always.

In contrast, the auction game is inherently random. No strategy is clever enough for winning always; the equilibrium strategy is optimal (in some sense) rather than winning. A player must choose his action without knowing the signal of the other player. The game belongs to so-called games with incomplete information.

1b The very simple auction: more formal description

The game is specified by[9]

- signal spaces,
- action spaces,
- the distribution of signals,
- payoff functions.

The signal S_1 of the first player is a random variable whose values belong to a set $\mathcal{S}_1$ called the signal space (of the first player). It may be the real line $\mathbb{R} = (-\infty, +\infty)$, or the half-line $[0, \infty)$, a finite interval, say, $[0, 1]$, etc. Anyway, the distribution P_{S_1} of S_1 must be specified (see below), $P_{S_1}(A) = \mathbb{P}\big(S_1 \in A\big)$. Every set $\mathcal{S}_1 \subset \mathbb{R}$ such that $P_{S_1}(\mathcal{S}_1) = 1$ may serve as a signal space.[10] Any localization of signals, say, "$S_1 \geq 0$ always", is expressed by the distribution, say, $P_{S_1}\big([0, \infty)\big) = 1$. It may be expressed *again* by the signal space, say, $\mathcal{S}_1 = [0, \infty)$, but there is nothing in it. Equally well, we may take $\mathcal{S}_1 = \mathcal{S}_2 = \mathbb{R}$.

The first player chooses his action from his action space $\mathcal{A}_1$. The action must be a number,[11] which does not guarantee that $\mathcal{A}_1 = \mathbb{R}$ fits. If the auctioneer requires A_1 to belong, say, to the set $\{100, 110, 120, \dots\}$, the requirement must be expresses by $\mathcal{A}_1$.[12] For now we simply take $\mathcal{A}_1 = \mathcal{A}_2 = [0, \infty)$.[13]

The distribution of signals is, in general, their joint distribution, which means that the signals may be (inter-)dependent. For now, however, we assume that S_1, S_2 are independent. Thus, it is enough to specify one-dimensional distributions P_1, P_2. It can be made via cumulative distribution functions

$$F_{S_1}(s) = \mathbb{P}\big(S_1 \le s\big) = P_1\big((-\infty, s]\big)\,, \quad F_{S_2}(s) = \mathbb{P}\big(S_2 \le s\big) = P_2\big((-\infty, s]\big)\,.$$

The cumulative distribution function of the joint (two-dimensional) distribution is, by independence, the product

$$F_{S_1,S_2}(s_1, s_2) = \mathbb{P}\big(S_1 \le s_1,\, S_2 \le s_2\big) = F_{S_1}(s_1) F_{S_2}(s_2)\,.$$

---

[9] That is, conditions under which players operate are specified, but not their behavior. The description is rather simple since the game is one-stage.

[10] For more general games, a signal need not be a number; it may be a vector, a function, etc. In full generality, a signal space is just a measurable space.

[11] For more general games, it need not be a number.

[12] The distribution of the action is not a part of the specification.

[13] Not a realistic assumption, of course. First, the auctioneer probably does not want to sell the object for arbitrarily low (even zero) price. Second, the auctioneer probably does not want such a bid as, say, $\pi = 3.14159265\dots$ dollars.

The payoff Π_1 of the first player may depend on signals and actions,

$$\Pi_1 = \mathbf{\Pi}_1(A_1, S_1; A_2, S_2)\,; \quad \Pi_2 = \mathbf{\Pi}_2(A_2, S_2; A_1, S_1)\,;$$

functions $\mathbf{\Pi}_1, \mathbf{\Pi}_2$ are called payoff functions. However, our case is simpler; Π_1 does not depend on S_2, and Π_2 does not depend on S_1. Namely,

$$\mathbf{\Pi}_1(A_1, S_1; A_2, S_2) = \begin{cases} 0 & \text{if } A_1 < A_2, \\ \frac{1}{2}(S_1 - A_1) & \text{if } A_1 = A_2, \\ S_1 - A_1 & \text{if } A_1 > A_2. \end{cases} \tag{1b1}$$

Note that $\frac{1}{2}(S_1 - A_1)$ is not really a payoff; the payoff is either 0 or $S_1 - A_1$ and depends (in the case of a tie, $A_1 = A_2$) on an additional randomizer. The payoff function returns the conditional expectation of the payoff, given actions and signals. The formula for Π_2,

$$\mathbf{\Pi}_2(A_2, S_2; A_1, S_1) = \begin{cases} 0 & \text{if } A_2 < A_1, \\ \frac{1}{2}(S_2 - A_2) & \text{if } A_2 = A_1, \\ S_2 - A_2 & \text{if } A_2 > A_1, \end{cases}$$

is identical to the formula for Π_1 with (A_1, S_1) and (A_2, S_2) swapped; that is, $\mathbf{\Pi}_2(A_2, S_2; A_1, S_1) = \mathbf{\Pi}_1(A_2, S_2; A_1, S_1)$, which means just $\mathbf{\Pi}_1 = \mathbf{\Pi}_2$. Different values of Π_1, Π_2 result from different orders of the arguments, not from different functions.[14]

So, a game is described by

$$\big(\mathcal{S}_1, \mathcal{S}_2; \mathcal{A}_1, \mathcal{A}_2; P_1, P_2; \mathbf{\Pi}_1, \mathbf{\Pi}_2\big)\,; \tag{1b2}$$

that is a framework, more general than our specific game (the auction), but of course more specific than a general game.[15]

The game is called symmetric, if[16]

$$\mathcal{S}_1 = \mathcal{S}_2\,; \quad \mathcal{A}_1 = \mathcal{A}_2\,; \quad P_1 = P_2\,; \quad \mathbf{\Pi}_1 = \mathbf{\Pi}_2\,. \tag{1b3}$$

Our 'very simple auction' is the symmetric game described by

$$\begin{aligned} &\mathcal{S}_1 = \mathcal{S}_2 = \mathbb{R}\,; \quad \mathcal{A}_1 = \mathcal{A}_2 = [0, \infty)\,; \\ &P_1 = P_2 = \mathrm{U}(0,1)\,, \quad \text{the uniform distribution on } (0,1)\,; \\ &\mathbf{\Pi}_1 = \mathbf{\Pi}_2 \quad \text{is the function defined by (1b1)}\,. \end{aligned} \tag{1b4}$$

---

[14] For example, $\mathbf{\Pi}_1(9, 10; 3, 5) = \mathbf{\Pi}_2(9, 10; 3, 5) = 1$, which means that $A_1 = 9$, $S_1 = 10$, $A_2 = 3$, $S_2 = 5$ implies $\Pi_1 = 1$; equally well, it means that $A_2 = 9$, $S_2 = 10$, $A_1 = 3$, $S_1 = 5$ implies $\Pi_2 = 1$.

[15] Or rather, that is the idea of a framework, since we do not formulate mathematical requirements for these objects. We do not answer such questions as: Must the spaces be one-dimensional? Must the distribution be nonatomic? Must the profit functions be monotone? And many others.

[16] It means that the two players operate under equal conditions. Still, they may behave differently. Maybe, they just have different temperaments.

1c Strategies

A strategy[17] of (say) the first player is a joint (2-dimensional) distribution $\mu = P_{S_1,A_1}$ of random variables S_1 and A_1. That is, a probability distribution on the two-dimensional space $\mathcal{S}_1 \times \mathcal{A}_1 = \mathbb{R} \times [0,\infty)$ whose projection to $\mathcal{S}_1$ is P_1. In other words, the marginal distribution of S_1 must be equal to P_1. Of course, the same holds for every player.

1c1. Example. Let $A = \frac{1}{2}S$. Here, μ is concentrated on a straight segment and uniformly distributed on the segment.

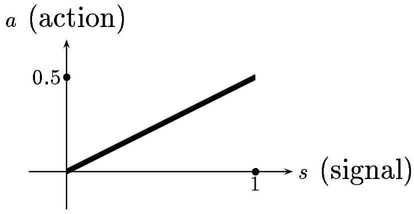

This time, A is uniquely determined by S. The conditional distribution of A given $S = s$ is degenerate, concentrated at $\frac{1}{2}s$. That is, $\mathbb{P}\big(A = \frac{1}{2}s \,\big|\, S = s \big) = 1$. The strategy is optimal in some sense, as was seen in Sect. 1a.

Other examples do not pretend to be optimal. They are just some strategies. 'Optimal' means: better than *all* others. The notion 'optimal strategy'[18] is based on the notion 'arbitrary strategy'.

1c2. Example. Let $A = 4S(1-S)$. Here, μ is concentrated on a curvilinear segment, non-uniformly.

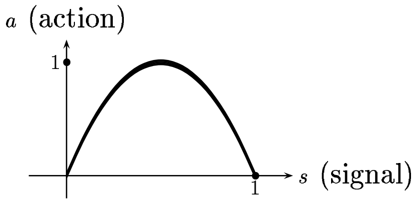

Still, conditional distributions are degenerate. However, a higher signal can cause a lower action.[19]

1c3. Example. Let $\dfrac{A}{S}$ be a random variable independent of S and taking on two values 0, 1 only, equiprobably. In other words, $A = S\xi$, where ξ is another random variable independent of S, and $\mathbb{P}\big(\xi = 0\big) = 0.5 = \mathbb{P}\big(\xi = 1\big)$. Here, μ is concentrated on (the union of) two straight segments.

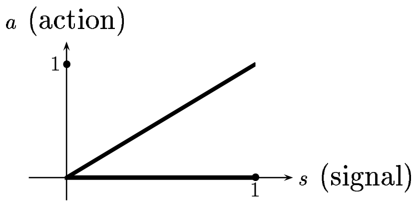

---

[17] Several kinds of strategies are well-known. I start with so-called distributional strategies. In general they are mixed (randomized). Pure strategies will be treated as a special case of distributional strategies.

[18] Or rather, 'equilibrium strategy', see 1e.

[19] Is it a good idea for the player? In fact, it is not. We'll see that such a strategy cannot be optimal for such a game.

The conditional distribution of A, given $S = s$, is concentrated at two equiprobable atoms located at 0 and s. That is, $\mathbb{P}\left(A = 0 \,\middle|\, S = s \right) = 0.5 = \mathbb{P}\left(A = s \,\middle|\, S = s \right)$. We may imagine that the player tosses a (fair) coin for choosing either $A = 0$ or $A = S$.[20]

1c4. Example. The player chooses either $A = 0$ or $A = S$ but with different probabilities:

$$\mathbb{P}\left(A = 0 \,\middle|\, S = s \right) = 1 - s\,; \quad \mathbb{P}\left(A = s \,\middle|\, S = s \right) = s\,.$$

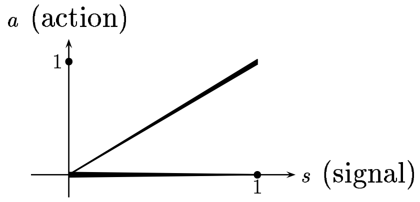

Here, μ is concentrated on the same segments as in the previous example, but distributed differently. We may imagine that the player generates (say, by a roulette) a random number $\xi \sim \mathrm{U}(0,1)$ (that is, distributed uniformly on $(0,1)$), and chooses A as follows:[21]

$$A = \begin{cases} 0 & \text{if } \xi < 1 - s, \\ S & \text{if } \xi > 1 - s. \end{cases}$$

1c5. Example. Let $A = S\xi$ where $\xi \sim \mathrm{U}(0,1)$ is independent of S. Here, μ is distributed in a triangle, non-uniformly. The distribution μ has a (two-dimensional) density f_μ,

$$f_\mu(s,a) = \begin{cases} \frac{1}{s} & \text{if } 0 < a < s < 1, \\ 0 & \text{otherwise.} \end{cases}$$

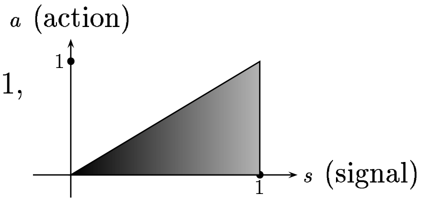

The conditional distribution of A given $S = s$ is $\mathrm{U}(0,s)$. Note that the same μ may be obtained from $A = S(1-\xi)$ or, say,

$$A = \begin{cases} S\xi & \text{if } S < 0.5, \\ S(1-\xi) & \text{otherwise.} \end{cases}$$

These are different implementations of *the same* strategy μ.

As was said, a strategy is a joint (2-dim) distribution $\mu = P_{S,A}$ of random variables S and A that conforms to the given (1-dim) distribution P_S of the signal S. In other words, the projection of μ to the signal space $\mathcal{S}$ must be equal to P_S. What about the action A and its (1-dim) distribution P_A? It is the projection of μ to the action space, and it is just an arbitrary distribution on the action space.

[20] Is it a good idea for the player, to toss a coin? In fact, it is not. We'll see that such a strategy cannot be optimal for such a game.

[21] No matter what happens for $\xi = 1 - S$, since that is of probability 0.

Here are (cumulative) distribution functions F_A and densities f_A (if exist) for all our examples.

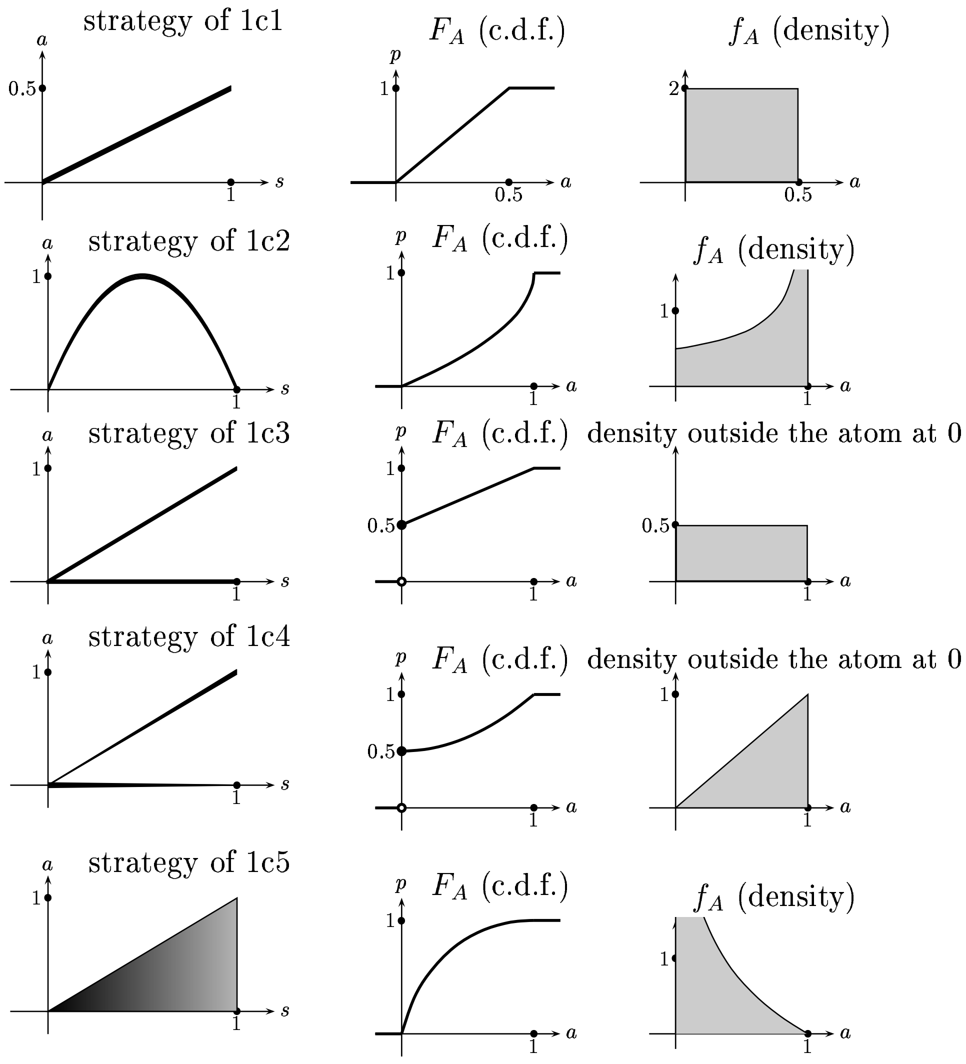

You may calculate these F_A, f_A in different ways. I prefer the universal formula

$$\mathbb{P}\big(\, A \le a \,\big) = \mathbb{E}\Big(\mathbb{P}\big(\, A \le a \,\big|\, S \,\big)\Big)\,;$$

the probability is the expectation of the conditional probability. Thus, for the strategy of 1c2,

$$\begin{aligned}
&\mathbb{P}\big(\, A \le a \,\big|\, S = s \,\big) \;=\; \begin{cases} 0 & \text{if } 4s(1-s) > a, \\ 1 & \text{if } 4s(1-s) \le a; \end{cases} \\
&\mathbb{P}\big(\, A \le a \,\big) = \mathbb{P}\big(\, 4S(1-S) \le a \,\big) = \mathbb{P}\big(\, 1 - 4(S - \tfrac12)^2 \le a \,\big) = \\
&= \mathbb{P}\big(\, S \notin (\tfrac12(1-\sqrt{1-a}), \tfrac12(1+\sqrt{1-a}))\,\big) = 1 - \sqrt{1-a}\,; \\
&F_A(a) = 1 - \sqrt{1-a}\,; \qquad f_A(a) = \frac{d}{da}F_A(a) = \frac{1}{2\sqrt{1-a}}
\end{aligned}$$

(assuming $a \in (0,1)$ and $s \in (0,1)$, of course). For the strategy of 1c4,

$$\mathbb{P}\big(A \le a \,\big|\, S = s \big) = \begin{cases} 1 - s & \text{if } s > a, \\ 1 & \text{if } s \le a; \end{cases}$$

$$\mathbb{P}\big(A \le a \big) = \int_0^a 1\,ds + \int_a^1 (1-s)\,ds = a + \frac{(1-a)^2}{2} = \frac{1}{2}(1+a^2)\,;$$

$$F_A(a) = \frac{1}{2}(1+a^2)\,; \qquad F'_A(a) = a \quad (\text{except for the atom})\,.$$

For the strategy of 1c5,

$$\mathbb{P}\big(A \le a \,\big|\, S = s \big) = \begin{cases} a/s & \text{if } s \ge a, \\ 1 & \text{if } s \le a; \end{cases}$$

$$\mathbb{P}\big(A \le a \big) = \int_0^a 1\,ds + \int_a^1 \frac{a}{s}\,ds = a - a \ln a\,;$$

$$F_A(a) = a - a \ln a\,; \qquad f_A(a) = -\ln a\,.$$

1d Strategies in use

A pair (μ_1, μ_2) of strategies (μ_1 for the first player, μ_2 for the second) determines uniquely all probabilities and expectations. Especially, the winning probability of the first player is

$$w_1 = \mathbb{P}\big(A_1 > A_2 \big) + \frac{1}{2}\mathbb{P}\big(A_1 = A_2 \big)\,;$$

actions A_1, A_2 are independent[22] random variables; their (one-dimensional) distributions are determined by μ_1, μ_2 respectively (as explained in 1c). The payoff Π_1 of the first player is a random variable,

$$\Pi_1 = \mathbf{\Pi}_1(A_1, S_1; A_2, S_2)\,.$$

It has an expectation $\mathbb{E}\,\Pi_1$, just a (non-random) number. That is, *before* getting his signal, the first player expects $\mathbb{E}\,\Pi_1$. *After* getting his signal S_1 he expects $\mathbb{E}\big(\Pi_1 \,\big|\, S_1 \big)$, the *conditional* expectation of Π_1 given A_1. After choosing his action A_1 he expects $\mathbb{E}\big(\Pi_1 \,\big|\, S_1, A_1 \big)$.[23]

It is convenient to abuse the symbol $\mathbf{\Pi}$ by writing

(1d1)
$$\mathbb{E}\,\Pi_1 = \mathbf{\Pi}_1(\mu_1; \mu_2)\,,$$
$$\mathbb{E}\big(\Pi_1 \,\big|\, A_1, S_1 \big) = \mathbf{\Pi}_1(A_1, S_1; \mu_2)\,,$$

[22] Actions are independent, since they result from independent signals. Each player may toss a coin; naturally, we assume that players use *independent* coins.

[23] It is rather stupid, to choose at random among actions of different profitability. We'll see that a strategy cannot be optimal unless $\mathbb{E}\big(\Pi_1 \,\big|\, S_1 \big) = \mathbb{E}\big(\Pi_1 \,\big|\, S_1, A_1 \big)$. Recall however that the notion 'optimal strategy' is based on the notion 'arbitrary strategy'.

etc. More formally,

$$
\text{(1d2)}\qquad
\begin{aligned}
\mathbb{E}\left(\Pi_1 \,\middle|\, A_1 = a_1, S_1 = s_1 \right) &= \iint_{\mathcal{A}_2 \times \mathcal{S}_2} \mathbf{\Pi}_1(a_1, s_1; a_2, s_2)\, d\mu_2(a_2, s_2) = \mathbf{\Pi}_1(a_1, s_1; \mu_2)\,, \\
\mathbb{E}\,\Pi_1 &= \iint_{\mathcal{A}_1 \times \mathcal{S}_1} \iint_{\mathcal{A}_2 \times \mathcal{S}_2} \mathbf{\Pi}_1(a_1, s_1; a_2, s_2)\, d\mu_1(a_1, s_1)\, d\mu_2(a_2, s_2) = \mathbf{\Pi}_1(\mu_1; \mu_2)\,.
\end{aligned}
$$

The same for the second player. Note that $\mathbb{E}\,\Pi_1 = \mathbb{E}\big(\mathbb{E}\left(\Pi_1 \,\middle|\, A_1, S_1 \right)\big) = \mathbb{E}\big(\mathbb{E}\left(\Pi_1 \,\middle|\, A_2, S_2 \right)\big)$, that is,

$$
\text{(1d3)}\qquad
\begin{aligned}
\mathbf{\Pi}_1(\mu_1; \mu_2) &= \iint_{\mathcal{A}_1 \times \mathcal{S}_1} \mathbf{\Pi}_1(a_1, s_1; \mu_2)\, d\mu_1(a_1, s_1)\,, \\
\mathbf{\Pi}_1(\mu_1; \mu_2) &= \iint_{\mathcal{A}_2 \times \mathcal{S}_2} \mathbf{\Pi}_1(\mu_1; a_2, s_2)\, d\mu_2(a_2, s_2)\,.
\end{aligned}
$$

1d4. Example. Let μ_2 be the strategy of 1c3:

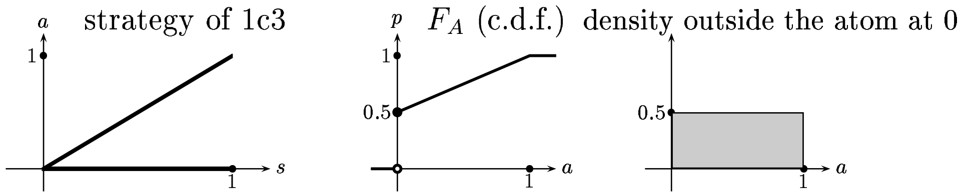

Let us calculate $\mathbf{\Pi}_1(a_1, s_1; \mu_2)$ for all a_1, s_1. In principle it is $\iint \mathbf{\Pi}_1(a_1, s_1; a_2, s_2)\, d\mu_2(a_2, s_2)$. However, our case is simpler: the payoff does not depend on s_2. Thus, not the whole 2-dim distribution μ_2 on $\mathcal{A}_2 \times \mathcal{S}_2$ matters, but only its projection P_{A_2} to $\mathcal{A}_2$. That is,

$$
\mathbf{\Pi}_1(a_1, s_1; \mu_2) = \int \mathbf{\Pi}_1(a_1, s_1; a_2)\, dP_{A_2}(a_2)\,;
$$

namely,

$$
\mathbf{\Pi}_1(a_1, s_1; \mu_2) = \frac{1}{2}\mathbf{\Pi}_1(a_1, s_1; 0) + \frac{1}{2}\int_0^1 \mathbf{\Pi}_1(a_1, s_1; a_2)\, da_2\,;
$$

using (1b1),

$$
\begin{aligned}
\mathbf{\Pi}_1(a_1, s_1; \mu_2) &= \frac{1}{2}(s_1 - a_1) + \frac{1}{2}\int_0^{a_1} (s_1 - a_1)\, da_2 + \frac{1}{2}\int_{a_1}^1 0\, da_2 = \\
&= \frac{1}{2}(s_1 - a_1) + \frac{1}{2}(s_1 - a_1)a_1 = \frac{1}{2}(s_1 - a_1)(1 + a_1)\,,
\end{aligned}
$$

provided that $a_1 > 0$. For $a_1 = 0$ we have

$$
\mathbf{\Pi}_1(0, s_1; \mu_2) = \frac{1}{2} \cdot \frac{1}{2}(s_1 - 0) + \frac{1}{2}\int_0^1 0\, da_2 = \frac{1}{4} s_1\,.
$$

1d5. Example. Let μ_1 be the strategy of 1c1,

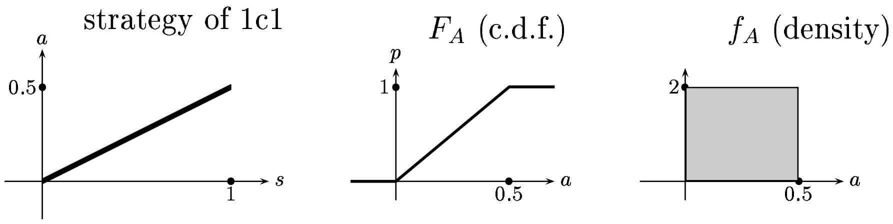

while μ_2 still be the strategy of 1c3. Let us calculate $\mathbf{\Pi}_1(\mu_1;\mu_2)$. We have

$$\mathbf{\Pi}_1(\mu_1;\mu_2) = \iint \mathbf{\Pi}_1(a_1,s_1;\mu_2)\,d\mu_1(a_1,s_1) = \int_0^1 \mathbf{\Pi}_1(\tfrac12 s_1, s_1;\mu_2)\,ds_1 = \\ = \int_0^1 \tfrac12(s_1 - \tfrac12 s_1)(1+\tfrac12 s_1)\,ds_1 = \frac14 \int_0^1 (s_1 + \tfrac12 s_1^2)\,ds_1 = \frac14 \cdot \Bigl(\frac12 + \frac16\Bigr) = \frac16\,.$$

Both a_1 and s_1 matter.

1e Best response and equilibrium

The first player wants to maximize his expected profit $\mathbf{\Pi}_1(\mu_1;\mu_2)$ over all strategies μ_1. However, he wants too much, if he seeks a single μ_1 optimal against every μ_2.

A strategy μ_1 is called a *best response*[24] (of the first player) to a given strategy μ_2, if

(1e1) $$\mathbf{\Pi}_1(\mu_1,\mu_2) = \sup_{\mu_1'} \mathbf{\Pi}_1(\mu_1',\mu_2)\,;$$

the supremum is taken over all strategies μ_1'. Similarly, a strategy μ_2 is called a best response (of the second player) to a given strategy μ_1, if $\mathbf{\Pi}_2(\mu_2,\mu_1) = \sup_{\mu_2'} \mathbf{\Pi}_2(\mu_2',\mu_1)$. Of course, μ_1 and μ_1' are (possible) strategies of the first player, while μ_2 and μ_2' — of the second. For a symmetric game (recall (1b3)) it is the same, and we may say 'best response' without specifying, of which player.

A pair (μ_1,μ_2) of strategies is called an *equilibrium,* if both μ_1 is a best response (of the first player) to μ_2, and μ_2 is a best response (of the second player) to μ_1.

For a symmetric game, an equilibrium (μ_1,μ_2) is called *symmetric*, if $\mu_1 = \mu_2$. In other words, a symmetric equilibrium is a strategy μ that is a best response to itself.

We'll see that the strategy of 1c1 is a symmetric equilibrium of the 'very simple auction' game (1b4). Moreover, the game has no other equilibrium.

[24]Neither existence nor uniqueness (of a best response) is implied by the definition.