This is a draft; email me with comments, typos, clarifications, etc.

An auction with finite types

Let's consider the simplest case of an auction with finite types: there are two players $i \in \{1, 2\}$ with types θ_i in $\Theta = \{1, \beta\}$. The probability that a player is of the *low* type is p (this is in contrast to Riley's notation, where the probability that a player is of the *high* type is p; we'll work to make this more consistent as the quarter goes on). If bidders tie, the object is to be allocated by a fair coin toss. What are bidding strategies in a symmetric Bayesian Nash equilibrium?

Since we have two types in the system, we'll consider them each separately. To start, consider the equilibrium strategy of the low-type bidder; suppose bidders of type 1 play a pure strategy and bid b_1.

- Suppose $b_1 > 1$. Then each low-type bidder's expected payoff is

$$V_i(1) = \left(\frac{p}{2} + (1-p)\, G_\beta\left(b_1\right)\right)(1 - b_1)$$

But since $1 - b_1 < 0$, this expectation is negative! The bidder is better off bidding 1 and losing for sure, but incurring no cost.

- Suppose $b_1 < 1$. Then each low-type bidder's expected payoff is

$$V_i(1) = \left(\frac{p}{2} + (1-p)\, G_\beta\left(b_1\right)\right)(1 - b_1)$$

where G_β is the CDF of the equilibrium bid strategy of the high-type bidder (we will later see that this term should vanish). Now let $\varepsilon > 0$; suppose that the low-type bidder deviates and bids $1 + \varepsilon$. Her expected payoff is now

$$\hat{V}_i^\varepsilon(1) = \left(p + (1-p)\, G_\beta\left(b_1 + \varepsilon\right)\right)(1 - b_1 - \varepsilon)$$

That is, by increasing the bid slightly the low-type bidder now beats the other low-type bidder for sure and possibly beats a high-type bidder a little more often.

Although the payoff from winning has shrunk, the probability of winning has increased dramatically, by at least $\frac{p}{2}$. It is clear that we can choose ε small enough that such a deviation is profitable in expectation (let $\varepsilon \to 0$ and notice that there is a discontinuity in the expected payoff at the limit). So $b_1 < 1$ cannot be a pure strategy equilibrium.

- Suppose $b_1 = 1$. Then each low-type bidder's expected payoff is

$$V_i(1) = \left(\frac{p}{2} + (1-p)\, G_\beta\left(b_1\right)\right)(1 - b_1) = 0$$

Now suppose the bidder deviates to $b' > 1$; as above, this will result in a negative expected utility. If the bidder deviates to $b' < 1$, she loses for sure and again obtains 0 expected utility. So $b_1 = 1$ is the best she can do and may constitute the low type's strategy in Bayesian Nash equilibrium.

Can we rule out a mixed strategy in equilbrium for the low-type player? Let G_1 be the CDF of a mixed strategy for a bidder of type 1. Recall that to play a mixed strategy, an agent must be indifferent across [almost] all elements in the support of the mixture; let $\underline{b}$ represent the lower bound of the support of the mixture, and $\bar{b}$ the upper. Two cases arise:

- $G_1(\underline{b}) = 0$. Assume for now that the high-type bidders will bid at least as much as the low-type bidders (and, if they randomize, their support will be weakly higher) — we will prove this later. Then by bidding $\underline{b}$ an agent has 0 probability of winning the object; hence the expected utility from bidding this amount is 0. If the low-type agent is to be indifferent among all items in the support of the mixture, the expected payoff from any other bid in the support must also be 0. The only other bid at which the expectation is 0 is $b_1 = 1$, so the support will consist of two items, $\{\underline{b}, 1\}$. But since we have assumed $G_1(\underline{b}) = 0$, this is equivalent to bidding 1 for sure; this is a contradiction of the assumption that the agent is mixing.

- $G_1(\underline{b}) = g > 0$. Again assume that high-type bidders will weakly outbid low-type bidders. Then the expected payoff from bidding $\underline{b}$ is

$$\frac{1}{2}\left(pG_1\left(\underline{b}\right) + (1-p)\,G_\beta\left(\underline{b}\right)\right)(1-\underline{b})$$

  However, for $\varepsilon > 0$ sufficiently small the agent may bid $\underline{b} + \varepsilon$ and obtain an expected payoff of

$$\left(pG_1\left(\underline{b}+\varepsilon\right) + (1-p)\,G_\beta\left(\underline{b}+\varepsilon\right)\right)(1-\underline{b}-\varepsilon)$$

  As $\varepsilon \to 0$ this quantity becomes larger than the expected payoff within the mixture, so the agent is better off by deviating to a pure strategy of $\underline{b} + \varepsilon$. Hence this cannot be an equilibrium mixture.

Therefore, in equilibrium the low-type bidder can only bid 1. What now will a high-type bidder do? We establish a few guiding principles prior to developing the particular bidding strategy.

Principle i: high-type bidders cannot play a pure strategy.

Suppose high-type bidders play a pure strategy b_β in equilibrium. There are two cases:

- Let $b_\beta \geq \beta$. Then expected utility is weakly negative. By bidding $\beta - \varepsilon$ for some $\varepsilon > 0$, the bidder is sacrificing win probability but gaining payoff conditional on winning. Thus bidding $b_\beta \geq \beta$ cannot occur in a pure-strategy equilibrium.

- Let $b_\beta < \beta$; suppose that $b_\beta > b_1$ (this assumption is innocuous). Expected utility is now

$$V_i(\beta) = \left(p + \frac{1}{2}(1-p)\right)(\beta - b_\beta)$$

  However, by deviating up by some small $\varepsilon > 0$ the player can obtain

$$\hat{V}_i(\beta) = \beta - b_\beta - \varepsilon$$

  Clearly, for ε small enough this quantity is larger than that expected from bidding b_β; so bidding b_β cannot occur in a pure-strategy equilibrium.

Thus the high-type player must played a mixed strategy in equilibrium.

Principle ii: there cannot be a mass point in the high type's mixture at her valuation β.

Since players must be indifferent across their mixture and expected utility upon bidding their valuation is 0, if there is a mass point at β the expected utility from any bid in the distribution must be 0. This cannot be part of an equilibrium, as shown above.

Principle iii: there can be no mass points in the high type's mixture.

We assume that the high type's distribution does not extend beyond their valuation β; technically this is another principle, but it follows from the first above. Now suppose there is a mass point $b > 1$ in the CDF G_β. Given some $\varepsilon > 0$ define a new CDF H so that

$$H_\varepsilon(x;b) = \begin{cases} G(x) & \text{if } x < b \\ G(x) - G(b) + \lim_{y\to b} G(y) & \text{if } b \le x < b+\varepsilon \\ G(x) & \text{if } b+\varepsilon \le x \end{cases}$$

Essentially, H_ε represents a distribution in which the mass point in G has been "shifted rightward" by ε. Importantly, the probability density functions are identical except at b and $b+\varepsilon$. For ease of notation, for a CDF F let

$$\overrightarrow{F}(x) = \frac{1}{2}\left(F(x) + \lim_{t\to x} F(t)\right)$$

Essentially, $\overrightarrow{F}(x)$ captures the probability of winning given a bid x, accounting properly for the probability of tying.

We can represent expected utility in each bid distribution as

$$\begin{aligned} V_i^G(\beta) = &\int_{\underline{b}}^{b} \left(p + (1-p)\overrightarrow{G}(x)\right)(\beta - x)G'(x)dx \\ &+ \left(p + (1-p)\overrightarrow{G}(b)\right)\Pr{}_G(x=b)(\beta-b) \\ &+ \int_{b}^{b+\varepsilon} \left(p + (1-p)\overrightarrow{G}(x)\right)(\beta - x)G'(x)dx \\ &+ \left(p + (1-p)\overrightarrow{G}(b+\varepsilon)\right)\Pr{}_G(x=b+\varepsilon)(\beta-b-\varepsilon) \\ &+ \int_{b+\varepsilon}^{\bar{b}} \left(p + (1-p)\overrightarrow{G}(x)\right)(\beta - x)G'(x)dx \end{aligned}$$

$$\begin{aligned} V_i^H(\beta) = &\int_{\underline{b}}^{b} \left(p + (1-p)\overrightarrow{G}(x)\right)(\beta - x)H'(x)dx \\ &+ \left(p + (1-p)\overrightarrow{G}(b)\right)\Pr{}_H(x=b)(\beta-b) \\ &+ \int_{b}^{b+\varepsilon} \left(p + (1-p)\overrightarrow{G}(x)\right)(\beta - x)H'(x)dx \\ &+ \left(p + (1-p)\overrightarrow{G}(b+\varepsilon)\right)\Pr{}_H(x=b+\varepsilon)(\beta-b-\varepsilon) \\ &+ \int_{b+\varepsilon}^{\bar{b}} \left(p + (1-p)\overrightarrow{G}(x)\right)(\beta - x)H'(x)dx \end{aligned}$$

We now look at the difference between expected payoffs. Fortunately, most terms cancel or are 0, so we are left with

$$V_i^H(\beta) - V_i^G(\beta) = \left(p + (1-p)\overrightarrow{G}(b+\varepsilon)\right)\Pr{}_H(x=b+\varepsilon)(\beta-b-\varepsilon) - \left(p + (1-p)\overrightarrow{G}(b)\right)\Pr{}_G(x=b)(\beta-b)$$

Notice that $\Pr_G(x=b) = \Pr_H(x=b+\varepsilon)$; since we are concerned only with the sign of the difference, these terms may be removed. We are then left with

$$V_i^H(\beta) - V_i^G(\beta) = -p\varepsilon + (1-p)\left(\overrightarrow{G}(b+\varepsilon)(\beta-b-\varepsilon) - \overrightarrow{G}(b)(\beta-b)\right)$$

Letting $\varepsilon \to 0$, we find $V_i^H(\beta) - V_i^G(\beta) \to \frac{1}{2}(1-p)\Pr_G(x=b)(\beta-b) > 0$. That is, for ε small enough shifting the CDF rightward is profitable. Thus there cannot be a mass point above the low type's equilibrium bid (more generally: the two high-type bidders cannot have a common mass point; hence by the symmetry in this example there can be no mass points).

Principle iv: there are no "flat" ranges in the mixture distribution.

Suppose that there is a range $[h_\ell, h_u]$ over which G is constant (note that we can WLOG assume this range is closed since we have already demonstrated that G has no mass points). For some $\varepsilon_\ell \geq 0$, $\varepsilon_u \geq 0$ the densities $G'(h_\ell - \varepsilon_\ell)$ and $G'(h_u + \varepsilon_u)$ are positive; moreover, ε_ℓ and ε_u may be arbitrarily small. Since agents must be indifferent over their mixture, we need

$$(p + (1-p)\,G\,(h_\ell - \varepsilon_\ell))\,(\beta - h_\ell + \varepsilon_\ell) = (p + (1-p)\,G\,(h_u + \varepsilon_u))\,(\beta - h_u - \varepsilon_u)$$

Letting ε_ℓ, ε_u tend toward 0, we retain the equivalence

$$(p + (1-p)G\,(h_\ell))\,(\beta - h_\ell) = (p + (1-p)G\,(h_u))\,(\beta - h_u)$$

Or $h_\ell = h_u$, a contradiction. Thus there are no flat ranges in G.

Principle v: the lower bound of the support of the high type's mixture must coincide with the strategy of a low-type bidder.

Here, we will consider the low-type bidder as playing a pure strategy; the discussion will generalize to the case where the low-type bidder mixes (which will not happen, but is necessary to investigate to maintain logical consistency with our previous discussion of the low type's strategy) if we consider the upper bound of the mixture rather than the bid in the pure strategy equilibrium. As this is directly analogous and introduces headaches in notation we'll ignore it here.

Suppose the high-type players are mixing over some support $[\underline{b}, \overline{b}]$ with $\underline{b} > 1$. Recalling that players must be indifferent among all elements in the support of their mixtures, we see that expected utility for the high-type player is

$$p\,(\beta - \underline{b})$$

Let $\varepsilon > 0$ be such that $\underline{b} - \varepsilon > 1$. Then the expected utility from bidding $\underline{b} - \varepsilon$ is

$$p\,(\beta - \underline{b} + \varepsilon) > p\,(\beta - \underline{b})$$

So deviation is preferred.

Now suppose that $\underline{b} < 1$. Expected utility from bidding $\underline{b} < 1$ is 0, so this must be the expected utility across the support of the mixture; but the only points at which this is possible are $\underline{b}$ and β. This directly implies a mass point and a flat range in the distribution, both contradictions. So this cannot be an equilibrium.

To recap: a high-type bidder must be playing a mixed strategy with continuous support on $[1, \overline{b}]$ for some $\overline{b} \leq \beta$. The CDF representing the mixture is continuous and, if differentiable, has a strictly positive first derivative on its interior.

How can we uncover this CDF? We need to construct indifference conditions for the high-type bidder; from above, we know that the low type's bid of 1 will be the lower bound of the mixture. Naïvely, it seems that utility from bidding 1 should be

$$\frac{1}{2}p(\beta - 1)$$

That is, the low-type bidder is beaten half of the time and the high-type bidder is never beaten (recalling that the bidding CDFs have no mass points). However, consider an ε deviation; expected utility becomes

$$(1 + (1-p)\,G(1+\varepsilon))\,(\beta - 1 - \varepsilon)$$

Letting $\varepsilon \to 0$, we see that expected utility goes to

$$p(\beta - 1)$$

So long as 1 is bid with 0 probability, we will run into no real mathematical issues if we assume this as our indifference condition (if we assume the naïve indifference condition we will obtain a mass point at 1, a contradiction). We use this to construct an indifference equality and subsequently verify that there is 0 probability on bidding 1.

For indifference, we have

$$p(\beta - 1) = (p + (1-p)\, G(b))\, (\beta - b)$$

By algebraic rearrangement, this gives us

$$G(b) = \left(\frac{p}{1-p}\right)\left(\frac{b-1}{\beta - b}\right)$$

This CDF satisfies the properties of non-flatness, no mass points, and a lower bound at 1. To obtain its upper bound, we set $G(b) = 1$; this gives us

$$\overline{b} = (1-p)\beta + p$$

With $\beta > 1$, $\overline{b} < \beta$ satisfying our remaining restriction. Thus we are all good.

A final statement of Bayesian Nash equilibrium is that bidders of type 1 bid 1 for sure, while bidders of type β bid according to the mixed strategy $G(b) = (\frac{p}{1-p})(\frac{b-1}{\beta-b})$.

In general, we don't need to go through all of this complexity to answer these questions; it is sufficient to set up the lower-bound indifference condition and then directly solve for the distribution. Still it is useful to have a catalogue of the necessary properties and keep them in mind while working through other questions.

A slight generalization of finite types

There are several ways to generalize auctions with finite types — more players, more types, etc. — but we're going to leave the general concepts of the auction the same and change the specification only so that players may have different high types, $\theta_i \in \{1, \beta_i\}$, with different probabilities $(1 - p_i)$.

Most of the arguments for the simpler case hold just as well here: low types will still bid 1, high types must mix, there should be no flat ranges in the distribution, and high types should have no overlapping mass points. Naïvely, then, we can setup indifference conditions for the high-type bidders:

$$\begin{aligned} & p_{-i}\,(\beta_i - 1) = (p_{-i} + (1-p_{-i})\, G_{-i}(b))\,(\beta_i - b) \\ \iff \quad & G_{-i}(b) = \left(\frac{p_{-i}}{1-p_{-i}}\right)\left(\frac{b-1}{\beta_i - b}\right) \end{aligned}$$

What is the upper bound of the bid distribution for the high-type bidders? Setting $G_{-i}(b) = 1$, we find

$$\overline{b}_{-i} = p_{-i} + (1-p_{-i})\,\beta_i$$

The upper bounds of the mixtures will then match if and only if

$$\begin{aligned} & p_1 + (1-p_1)\,\beta_2 = p_2 + (1-p_2)\,\beta_1 \\ \iff \quad & p_1 - 1 + (1-p_1)\,\beta_2 = p_2 - 1 + (1-p_2)\,\beta_1 \\ \iff \quad & \frac{\beta_2}{1-p_2} = \frac{\beta_1}{1-p_1} \end{aligned}$$

If this holds, we are done.

What if the upper bounds of the mixtures do not coincide? Suppose that bidder 1's upper bound $\bar{b}_1$ lies above bidder 2's upper bound $\bar{b}_2$. The expected payoffs from these actions are

$$\begin{aligned}\left(p_2 + (1-p_2)\, G_2\left(\bar{b}_2\right)\right)\left(\beta_1 - \bar{b}_2\right) &= \beta_1 - \bar{b}_2 \\ \left(p_2 + (1-p_2)\, G_2\left(\bar{b}_1\right)\right)\left(\beta_1 - \bar{b}_1\right) &= \beta_1 - \bar{b}_1\end{aligned}$$

But since $\bar{b}_1 > \bar{b}_2$, the expected payoff from bidding $\bar{b}_2$ is higher! This contradicts the nature of the mixed strategy (to be fair, we should be certain to examine this within ε-neighborhoods). How can the strategies be altered so as to share an upper bound?

Intuitively, we have two options: we can add a mass point to the higher distribution (the one with the higher upper bound) at the upper bound of the lower distribution. Assuming $\bar{b}_1 > \bar{b}_2$, this means placing a mass point in G_1 at $\bar{b}_2$; in particular, we let $\hat{G}_2(\bar{b}_2) = 1$. But if bidder 2 is of the high type, she can now win the auction for sure (making a profit) when bidding $\bar{b}_2 + \varepsilon$ for any $\varepsilon > 0$. Her expected payoff from doing so is $\beta_2 - \bar{b}_2 - \varepsilon$. Now suppose she bids $\bar{b}_2$; her expected payoff is

$$\left(p_1 + (1-p_1)\, \overrightarrow{G}_1\left(\bar{b}_2\right)\right)\left(\beta_2 - \bar{b}_2\right)$$

As usual, for ε small enough the expected utility from deviating is greater than that from playing the existing equilibrium strategy. So this cannot be an equilibrium.

What if, instead, we add a mass point to the lower bidder's mixture at the lower (common) bound of the support? Without performing the calculations, we'll appeal to the proof above and existing ε deviation arguments and state that this cannot be an equilibrium either.

That is to say, given the existing setup, when $\frac{\beta_2}{1-p_2} \neq \frac{\beta_1}{1-p_1}$ *there is no pseudo-symmetric Bayesian Nash equilibrium* (pseudo-symmetric is by no means a technical term, it merely reflects the fact that we expect symmetric behavior from low-type bidders). Is there a way around this? We have the option to rewrite the tiebreaking rule in a manner similar to how we might address existence issues in Bertrand competition with differing marginal costs (if this is not familiar to you, consider what the Nash equilibrium would be in a Bertrand competition game with differing marginal cost structures): in the event of a tie, the object should be allocated to the bidder with the *higher* valuation. Note that this makes an assumptions that we paper over, namely that the seller has the ability to know a bidder's true valuation — which might make holding an auction in the first place seem ridiculous — but let's leave that be for now.

If ties go to the bidder with the higher type, we still cannot support a mass point at the upper bound of the lower bidder's distribution (check for yourself why this is the case). However, we are now able to let high-type bidders play the low-type's bid with positive probability. So again assume $\bar{b}_1 > \bar{b}_2$; we add a mass point for bidder 1 at 1. Bidder 2's indifference statement is now

$$\left(p_1 + (1-p_1)\, G_1(1)\right)\left(\beta_2 - 1\right) = \left(p_1 + (1-p_1)\, G_1(b)\right)\left(\beta_2 - b\right)$$

This gives us a form for bidder 1's mixture,

$$G_1(b) = \frac{\left(p_1 + (1-p_1)\, G_1(1)\right)\left(\beta_2 - 1\right) - p_1\left(\beta_2 - b\right)}{\left(\beta_2 - b\right)\left(1-p_1\right)} = \frac{p_1(b-1) + (1-p_1)\, G_1(1)\left(\beta_2 - 1\right)}{\left(\beta_2 - b\right)\left(1-p_1\right)}$$

Bidder 2's indifference conditions are unchanged. It may seem strange that the bidder with the "higher" distribution is the bidder with the mass point at the lower bound; but this ignores the truth about the bidder with the higher distribution: agents' mixtures are set to keep the *opposing* player indifferent between options. So the bidder with the "higher" distribution is the one with the lower overall valuation (here we are hand-waving prodigiously, but the intuition is solid), net of the opponent's type distribution. For a simple example of this, let $p_1 = p_2 = \frac{1}{2}$, $\beta_1 = 2$, and $\beta_2 = 3$. Then bidder 1's expectation from his lowest possible bid is $\frac{1}{2}$ while bidder 2's is 1; the upper bound on bidder 2's distribution must then be $\frac{3}{2}$ and the upper bound on bidder 1's must be 2 (set up indifference conditions on both ends of the distribution to check this).

Thus even though bidder 1 values the object less, he must be willing to bid more to keep bidder 2 willing to mix. In this light, it makes sense that the bidder with the higher mixing distribution must put a mass point at 1.

The only quantity remaining to calculate is $G_1(1)$. This we can achieve by setting $G_1(\bar{b}_2) = 1$, or $G_1(p_2 + (1-p_2)\beta_1) = 1$. Plugging in, we find

$$\begin{aligned} 1 &= \frac{p_1\left(p_2 + \left(1-p_2\right)\beta_1 - 1\right) + \left(1-p_1\right)G_1(1)\left(\beta_2 - 1\right)}{\left(\beta_2 - p_2 - \left(1-p_2\right)\beta_1\right)\left(1-p_1\right)} \\ \iff \quad \left(\beta_2 - p_2 - \left(1-p_2\right)\beta_1\right)\left(1-p_1\right) &= p_1\left(1-p_2\right)\left(\beta_1 - 1\right) + \left(1-p_1\right)G_1(1)\left(\beta_2 - 1\right) \\ \iff \quad G_1(1) &= \frac{\left(\beta_2 - p_2 - \left(1-p_2\right)\beta_1\right)\left(1-p_1\right) - p_1\left(1-p_2\right)\left(\beta_1 - 1\right)}{\left(1-p_1\right)\left(\beta_2 - 1\right)} \end{aligned}$$

All-pay auctions

All-pay auctions happen precisely as the name would suggest: each agent submits a bid; the high bidder wins the object up for auction and *every agent* pays their bid. As an exercise in working through the steps of solving auction problems, it is useful to solve for the equilibrium bid function in an all-pay auction.

To begin, we need to define the value function. A bidder of type θ_i will win the object with probability $w(\theta_i)$ by placing bid $b(\theta_i)$. But where in the standard auction the expected payment is $w(\theta_i)b(\theta_i)$, here the bid is paid regardless so expected payment is $b(\theta_i)$. The value function is then

$$V\left(\theta_i\right) = \theta_i w\left(\theta_i\right) - b\left(\theta_i\right)$$

Taking the derivative, we find

$$V'\left(\theta_i\right) = w\left(\theta_i\right) + \theta_i w'\left(\theta_i\right) - b'\left(\theta_i\right)$$

But from the envelope theorem, we know the latter two terms sum to 0. We can see this by looking at an agent's choice over possible type-reports rather than the more abstract value function. Let $u(\theta;\theta_i)$ be the utility that an agent of type θ_i receives from reporting type θ; this is

$$u(\theta;\theta_i) = \theta_i w(\theta) - b(\theta)$$

The agent will choose to report a type which maximizes expected utility, or $\frac{\partial u}{\partial \theta} = 0$. We then see

$$\frac{\partial u}{\partial \theta}\left(\theta;\theta_i\right) = \theta_i w'(\theta) - b'(\theta) = 0$$

Since these are precisely the latter two terms in $V'(\theta_i)$, we can see that they reduce to 0.

We then have $V'(\theta_i) = w(\theta_i)$. Expressing the value function by integrating up its derivative, we have

$$\begin{aligned} V\left(\theta_i\right) &= \int_{\theta_0}^{\theta_i} V'(\theta)d\theta \\ &= \int_{\theta_0}^{\theta_i} w(\theta)d\theta \end{aligned}$$

But from our initial definition we also know that $V(\theta_i) = \theta_i w(\theta_i) - b(\theta_i)$. This gives us

$$\theta_i w\left(\theta_i\right) - b\left(\theta_i\right) = \int_{\theta_0}^{\theta_i} w(\theta)d\theta$$

Rearranging, we may find an explicit form for the [symmetric] equilibrium bid function,

$$b(\theta_i) = \theta_i w(\theta_i) - \int_{\theta_0}^{\theta_i} w(\theta) d\theta$$

Generally, we take this at least one step further; integrating by parts, we obtain

$$\begin{aligned} b(\theta_i) &= \theta_i w(\theta_i) - \theta w(\theta)\Big|_{\theta=\theta_0}^{\theta_i} + \int_{\theta_0}^{\theta_i} \theta w'(\theta) d\theta \\ &= \theta_0 w(\theta_0) + \int_{\theta_0}^{\theta_i} \theta w'(\theta) d\theta \\ &= \int_{\theta_0}^{\theta_i} \theta w'(\theta) d\theta \qquad \text{(the low type cannot win)} \\ b(\theta_i) &= \int_{\theta_0}^{\theta_i} \theta F'(\theta) d\theta \end{aligned}$$

where F is the CDF of bidder types. Notice that we crucially assumed that the low type cannot win the auction. This is sensible if the bid function is strictly increasing and the type distribution is everywhere-continuous (that is, there are no jumps or probability mass points). We can go ahead and verify this assumption in the solution,

$$b'(\theta_i) = \theta_i F'(\theta_i) > 0$$

So everything is alright.

Solved using revenue equivalence

Can we take this a step further? Recall that in the case of a continuous distribution revenue equivalence holds as long as allocation is efficient and the low-type buyer incurs no unnecessary costs. If the seller expects identical revenue across auctions, each bidder should expect to make identical payments across auctions. To this end, let's consider an agent's expected payment in a second-price auction and use this to derive the equilibrium bid strategy in the all-pay auction.

In the second-price auction, the agent expects to pay the value of the next-highest type if she wins the auction, and 0 otherwise. Let $r(\theta_i)$ represent the expected payment of a bidder of type θ_i. We find

$$\begin{aligned} r(\theta_i) &= \Pr(\text{win}|\theta_i) E[\max \theta_{-i} | \theta_{-i} < \theta_i] \\ &= w(\theta_i) \int_{\theta_0}^{\theta_i} \theta \left(\frac{w'(\theta)}{w(\theta_i)} \right) d\theta \\ &= \int_{\theta_0}^{\theta_i} \theta w'(\theta) d\theta \\ r(\theta_i) &= \int_{\theta_0}^{\theta_i} \theta F'(\theta) d\theta \end{aligned}$$

There is one step left to determine the equilibrium bid in the all-pay auction: we must consider how an agent's expected payment relates to what she bids. In the all-pay auction, expected payment is exactly the bid, so we find

$$b(\theta_i) = r(\theta_i) = \int_{\theta_0}^{\theta_i} \theta F'(\theta) d\theta$$

This is exactly the result from above! It is not necessarily shorter nor necessarily better to proceed this way, but we have obtained a useful check of our previous algebra.

2007 Fall comp, question 3

Lance Armstrong and Floyd Landis agree to an exhibition bike race. Whoever trains hardest will win a prize they each value at $V = 2$, but training is expensive: the cost is ct where $t \in [0, 1]$ is the level of training chosen and $c \in [1, 2]$ is individual-specific marginal cost. Lance and Floyd each know their own marginal cost of training and behave as if the other's marginal cost is uniformly distributed on $[1, 2]$.

(a) Show that, whatever strategy Lance follows, Floyd's best response is weakly decreasing in his own cost.

Solution: by the revelation principle, we can consider this as a game where Lance and Floyd report their types and are then told how much to train. If Floyd is of type c_F but reports type c, his expected utility is then

$$u_F(c; c_F) = 2w(c) - c_F t(c)$$

where $w(c)$ is the probability that an agent of type c wins the race. For truthful reporting to be a dominant strategy, it must be incentive-compatible; that is, an agent of type c must be willing to report type c. Let $c' > c$; we then require

$$2w(c) - ct(c) \geq 2w(c') - ct(c')$$
$$2w(c') - c't(c') \geq 2w(c) - c't(c)$$

By rearranging algebraically we obtain

$$2(w(c) - w(c')) \geq c(t(c) - t(c'))$$
$$2(w(c') - w(c)) \geq c'(t(c') - t(c))$$

These inequalities may be put together,

$$c'(t(c) - t(c')) \leq c(t(c) - t(c'))$$

Since we have assumed $c' > c$, it follows that $t(c) - t(c') \leq 0$ for all $c < c'$. Thus $t(\cdot)$, Floyd's best response, is decreasing.

(b) Find a symmetric Bayesian Nash equilibrium in smooth, strictly decreasing strategies.

Solution: notice that this situation looks vaguely like an all-pay auction. Either racer's value function will look like

$$V_i(c) = 2w(c) - ct(c)$$

According to our usual envelope theorem argument, we have

$$V_i'(c) = -t(c)$$

Integrating up from initial conditions and equating with our definition of $V(\cdot)$, we have

$$\int_2^c V_i'(c)dc + V_i(2) = 2w(c) - ct(c)$$

Since $V_i(2) = 0$ — the *high*-type racer has zero probability of winning — we may rewrite this as

$$\int_2^c -t(c)dc = 2w(c) - ct(c)$$

Where should we go from here? We can consider substituting in for t' by taking first-order conditions of $u_i(c; c_i)$ as above, but it seems that it will be just as quick to solve this system directly. By the revelation principle, $w(c) = 1 - F(c)$; since types are believed to be uniformly distributed on $[1, 2]$, this gives us $w(c) = 2 - c$. So we have

$$\int_2^c -t(c)dc = 4 - 2c - ct(c)$$

Taking the derivative with respect to c, we find

$$-t(c) = -2 - ct'(c) - t(c) \qquad \Longrightarrow \qquad ct'(c) = -2$$

It follows that

$$t(c) = -2\ln c + K$$

We may determine K by observing that $V_i(2) = 0$; with $w(2) = 0$, it must also be that $t(2) = 0$. So we have

$$-2\ln 2 + K = 0 \qquad \Longrightarrow \qquad K = 2\ln 2$$

Then the symmetric equilibrium training strategy is

$$t(c) = 2\ln 2 - 2\ln c$$

The key trick to note in this question is that the value of the *high*-type racer is 0, rather than the value of the *low*-type racer being 0 as in a more standard auction context. This flips most of our equations (we base calculations from 2 rather than 1) but the general spirit of the math remains the same.

Essential Microeconomics, exercise 10.1.1

Also, 2008 Fall comp question 3.

Batman has just learned that the Joker plans a big caper for tomorrow, but he does not know if the target of the Joker's attack will be the Museum of the Tower. Batman can guard one of these targets but not both. Tomorrow's outcome depends on the actions of Batman and the Joker and on the weather.

- If Batman guards the wrong target (that is, the target that the Joker does not attack) then the Joker will succeed regardless of the weather. Batman values this outcome at -4 and the Joker values it as $+4$.
- If Batman guards the Museum and the Joker attacks the Museum and the weather is bad, Batman will catch the Joker. Batman values this outcome at $+8$ and the Joker values it at -20.
- In every other circumstance the Joker will be foiled but will escape. Batman and the Joker value this outcome at 0.

Batman must make his choice today, before he knows the weather. The Joker can make his choice tomorrow when he sees the weather. Neither player sees the action of the other. It is common knowledge that the probability that the weather will be good is $\frac{3}{4}$ and the probability that the weather will be bad is $\frac{1}{4}$.

(a) Depict the Bayesian game in extensive (tree) form.

Solution: TBD.

(b) Explain carefully why there can be no equilibrium in pure strategies.

Solution: this is best handled in cases.

- Suppose that Batman always guards the Museum; then the Joker will always attack the Tower and Batman would rather guard the Tower.
- Suppose that Batman always guards the Tower; then the Joker will always attack the Museum and Batman would rather guard the Museum. Therefore Batman cannot play a pure strategy in equilibrium.
- Suppose that the Joker plays a the same pure strategy regardless of the state of the weather. Then just as in the previous two cases Batman will strictly prefer thwarting the Joker to not, leading the Joker to deviate.
- Suppose that the Joker attacks the Museum (for sure) when the weather is bad; Batman's expected payoff from guarding the Museum is at least $\frac{1}{4}(8)+\frac{3}{4}(-4)=-1$ while Batman's expected payoff from guarding the Tower is no greater than $\frac{1}{4}(-4)+\frac{3}{4}(0)=-1$. These payoffs are obtained assuming that the Joker attackes the Tower (for sure) when the weather is good. Since Batman is indifferent, he is willing to play a mixed strategy and we run into no contradiction of Batman playing a pure strategy.
- There is one remaining pure strategy for the Joker: attack the Tower when the weather is bad and attack the Museum when the weather is good. Then Batman sees probability $\frac{3}{4}$ of thwarting the Joker if he guards the Museum and probability $\frac{1}{4}$ of thwarting the Joker if he guards the Tower. Batman will then strictly prefer guarding the Museum, contradicting the fact that Batman cannot play a pure strategy in equilibrium.

Since Batman cannot have an equilibrium in pure strategies, there is no pure strategy Nash equilibrium. However, it is *possible* that the Joker plays a pure strategy (per state of the world) while Batman mixes[1].

(c) Let y be the probability that Batman chooses the Museum. Show that there is a unique y such that the Joker is indifferent between M and T if the weather is bad, and a second y such that the Joker is indifferent if the weather is good.

Solution: for the Joker to be indifferent between the Museum and the Tower when the weather is bad we need

$$E\left[u_J\left(M\right)|\text{bad}\right]=E\left[u_J\left(T\right)|\text{bad}\right]$$

That is,

$$\begin{aligned} & (-20)y_b+(4)(1-y_b)=(4)y_b+(0)(1-y_b) \\ \iff \quad & y_b=\frac{1}{7} \end{aligned}$$

For the Joker to be indifferent between the Museum and the Tower when the weather is good we need

$$E\left[u_J\left(M\right)|\text{good}\right]=E\left[u_J\left(T\right)|\text{good}\right]$$

That is,

$$\begin{aligned} & (0)y_g+(4)(1-y_g)=(4)y_g+(0)(1-y_g) \\ \iff \quad & y_g=\frac{1}{2} \end{aligned}$$

(d) Let x be the probability that the Joker chooses M. For each y determined in (c) above, examine the payoff of Batman and show that there is a unique Bayesian Nash equilibrium of this game.

[1] This will be ruled out in subsequent discussion.

Solution: suppose $y = y_b$. Then when the weather is good the Joker will certainly attack the Museum. For Batman to be indifferent between guarding the Museum and guarding the Tower, we need

$$E\left[u_B\left(M\right)\right] = E\left[u_B\left(T\right)\right]$$

Explicitly, this is

$$\begin{aligned} & (8)\left(\frac{1}{4}\right)x + (-4)\left(\frac{1}{4}\right)(1-x) = (-4)\left(\frac{1}{4}\right)x + (-4)\left(\frac{3}{4}\right) \\ \iff & 2x - (1-x) = -x - 3 \\ \iff & x = -\frac{1}{2} \end{aligned}$$

This probability is negative! Essentially, since guarding the Museum is advantageous when the weather is good and also yields (potentially) a very large payoff when the weather is bad, Batman is very highly incentivized to guard the Museum.

Now suppose $y = y_g$. Then when the weather is bad the Joker will certainly attack the Tower. Batman's indifference conditions have not changed (as compared to those stated above), but the explicit form has. We now set

$$\begin{aligned} & (-4)\left(\frac{1}{4}\right) + (-4)\left(\frac{3}{4}\right)(1-x) = (-4)\left(\frac{3}{4}\right)x \\ \iff & -1 - 3(1-x) = -3x \\ \iff & x = \frac{2}{3} \end{aligned}$$

So we have a Bayesian Nash equilibrium in which Batman guards the Museum with probability $y_g = \frac{1}{2}$ and the Joker attacks the Tower when the weather is bad, and attacks the Museum with probability $x = \frac{2}{3}$ when the weather is good.

Is this equilibrium unique? The only other candidate for equilibrium was described above in part (a): the Joker attacks the Tower when the weather is good and the Museum when the weather is bad, while Batman mixes somehow. We have seen $y_b < y_g$, so we can handle this concept in two cases. Suppose $y < y_g$; then the Joker will strictly prefer attacking the Museum when the weather is good. Hence to obtain the desired result, we require $y > y_g$. But this implies $y > y_b$; when this is the case, the Joker strictly prefers attacking the Tower when the weather is bad. Thus we cannot find a mixed strategy for Batman which will induce the Joker to follow the desired contingent strategy. It follows that the Bayesian Nash equilibrium derived here is unique.

At this equilibrium, Batman's utility is

$$\begin{aligned} E\left[u_B\left(\sigma_B, \sigma_J\right)\right] &= E\left[u_B\left(T, \sigma_J\right)\right] \\ &= (-4)\left(\frac{3}{4}\right)x \\ &= -3\left(\frac{2}{3}\right) \\ &= -2 \end{aligned}$$

Poor Batman.