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

Perfect Bayesian Equilibrium

We are by now familiar with the concept of Bayesian Nash equilibrium: agents are best responding given their beliefs, and behavior must be optimal along the equilibrium path. As we have seen, this leads to undesirable equilibrium outcomes in which players may engage in incredible threats (consider the incumbent-entrant game in the case with a strong entrant).

In simpler games, we got around these issues by applying subgame perfection to the standard notion of Nash equilibrium; however, in games of incomplete information it is often impossible to locate any proper subgames. Perfect Bayesian equilibrium allows us, in some sense, to extend the notion of subgame perfection to arbitrary games without proper subgames, and especially games of incomplete information.

Definition

A belief system μ is *consistent* according to some strategy profile σ if the probability assigned by μ to reaching any node in the game tree follows from Bayes' rule applied to the strategy profile σ (where this probability is well-defined).

The definition of consistency is fairly intuitive in most cases: where there is positive probability of reaching a node, the probability must be determined by Bayes' rule. Nodes which are reached with zero probability — or more specifically, those nodes which follow those which are reached with zero probability — may have their conditional probabilities arbitrarily defined.

Definition

A strategy profile σ is *sequentially rational* according to a belief system μ if at every information set, the player whose action is being taken is behaving optimally according to the beliefs generated by μ at this information set.

The definition of sequential rationality is also fairly intuitive: at every information set, players are best-responding conditional on the beliefs they maintain. Notice that this is significantly more general than requiring best-responding only along the equilibrium path! Here, it is assumed to be common knowledge that players are behaving rationally *everywhere*.

Definition

A Bayesian Nash equilibrium (σ, μ) is a *perfect Bayesian equilibrium* if σ is sequentially rational given μ, and μ is consistent.

The single-crossing property

It can be a real pain to pin down a precise definition of the single-crossing property. We will therefore begin with a vague mathematical statement and translate it slowly into Riley's definition (as found in *Essential Microeconomics*, chapter 11).

Intuitively, two curves satisfy the single-crossing property if they cross each other once and only once; in particular, two functions f_1 and f_2 satisfy the single-crossing property if $\exists! x, f_1(x) - f_2(x) = 0$. For our purposes, we must distinguish between crossing at a unique point and crossing once over a continuum of points. Wikipedia says that economics cares about crossing over a continuum of points; I would claim that Riley believes otherwise (and we'll see why in application).

Beyond the claim of a unique intersection, we need to know that the curves *actually cross*. That is, we do not consider two curves which simply touch at a single point to satisfy the single-crossing property. Therefore we can make the mathematical definition for our uses more precise,

$$\exists x f_1(x) = f_2(x) \wedge \forall y < x, f_1(x) < f_2(x) \wedge \forall y > x, f_1(x) > f_2(x)$$

Of course, f_1 and f_2 here are general, so the particular statement of, "f_1 is smaller when $y < x$," is not quite correct. There is also some question as to whether we should simply require that there exists some point below which the functions are ordered one way and above which they are ordered another; since we will see the single-crossing property almost uniquely applied to well-behaved problems and utility functions, it's safe to ignore the implications of issues of discontinuity.

How does this notion apply to economics? The ways may be many and varied, but for now we will consider single-crossing only in reference to indifference curves. In particular, if two agents' indifference curves go through the same point, they satisfy single-crossing if they never meet again; since for any allocation (and suitably well-behaved problem, although the constraints on this are not many) we can form an indifference curve, the claim that indifference curves satisfy the single-crossing property is always made in reference to a particular allocation. If not, then we may assume it holds for all allocations.

What does it mean for indifference curves to cross once and only once? Let x be the reference point at which the indifference curves intersect; then for all $y > x$, one agent's indifference curve lies about the other's, and for all $y < x$ the other agent's indifference curve lies above the one's. Depending on the orientation of preferences, this tells us that for all $y > x$ one agent prefers the other's indifferent-equivalent allocation (to coin an imprecise phrase) to her own indifferent allocation, and for all $y < x$ the other agent prefers the one's indifferent-equivalent allocation to his own allocation. For a decent graphical representation of this story, check out the answer to the 2009 Spring comp, question 3(a) below.

With this background, we can complete the transition to economics by copying out of *Essential Microeconomics*,

Definition

Type θ_t has a *stronger preference* for q than type θ_s if for all $(q, r), (q', r') \in Q \times R$ with $q > q'$, $(q, r) \succeq_s (q', r') \implies (q, r) \succ_t (q', r')$. The problem satisfies the *single-crossing property* if higher-indexed types have a stronger preference for q.

That is, if (q', r') is such that a lower type θ_s weakly prefers (q, r) to (q', r') (where (q, r) offers more q than (q', r')), then a problem satisfies the single-crossing property if a higher type θ_t must then strictly prefer (q, r) to (q', r'). Essentially, θ_t wants q more badly than does θ_s; if it helps, we can view this (in a way) as θ_t having a smaller elasticity of demand for q with respect to r — as the conditions on r become less favorable, the changes in q necessary to sustain indifference are lower for type θ_t than for type θ_s.

Essential Microeconomics, exercise 10.2-2

Buyers A and B each have a valuation of a single item for sale that is either 2 or 4. Valuations are independently distributed. A buyer's value is high with probability p. The item is to be sold in a sealedbid auction to the highest bidder at the second-highest price. If there is a tie, the winner will be selected randomly from the high bidders. Buyers must bid some $b \in B = \{2, 3, 4\}$ for the item.

(a) Explain why it is a Bayesian Nash equilibrium for buyer A to bid 2 regardless of his valuation and for buyer B to bid 4 regardless of her valuation.

Solution: suppose that these are equilibrium strategies. Then buyer B's ex ante expected utility is

$$p(4-2) + (1-p)(2-2) = 2p$$

Buyer A's ex ante expected utility is

$$p(0) + (1-p)0 = 0$$

That is, regardless of types buyer A loses the auction and buyer B wins. Given buyer B's strategy, buyer A can never obtain positive utility from bidding any value: he either continues to lose the auction, or wins and gains 0 utility. Buyer B can also do no better, since she is already winning with probability 1; further, by lowering her bid she cannot reduce his expected payment. Therefore this pair of strategies, while strange, constitutes a valid Bayesian Nash equilibrium.

(b) Explain why the equilibrium is not trembling-hand perfect.

Solution: suppose that players $i \in \{A, B\}$ tremble across the possible bids $b \in \{2, 3, 4\}$ with probability $\Pr(b_i = b) = \varepsilon_i^b$. For the sake of normalization, represent $\varepsilon_A^2 = 1 - \varepsilon_A^3 - \varepsilon_A^4$ and $\varepsilon_B^4 = 1 - \varepsilon_B^2 - \varepsilon_B^3$ (note the lack of symmetry in this definition; this is to coincide with the strategies listed above).

To show that this equilibrium is not trembling-hand perfect, we will not directly show which equilibrium is (although hopefully it's fairly implicit after this explanation). Let's consider A's choice when he is type 4; his interim expected utility from bidding 2 in this Bayesian Nash equilibrium, subject to trembles, is

$$\frac{1}{2}\varepsilon_B^2\,(4-2) = \varepsilon_B^2 > 0$$

However, suppose he bids 4, his valuation. His expected utility is now

$$\varepsilon_B^2(4-2) + \varepsilon_B^3(4-3) = 2\varepsilon_B^2 + \varepsilon_B^3 > \varepsilon_B^2$$

So he prefers bidding his type.

It follows that the proposed Bayesian Nash equilibrium is not trembling-hand perfect. We could also prove this from player B's viewpoint, but showing that one agent sees fit to deviate is sufficient to show that this is not an equilibrium of the specified type.

(c) Explain why the Bayesian Nash equilibrium is a sequential perfect Bayesian equilibrium.

Solution: because players' information sets are never updated, the "beliefs in the limit of the trembles" correspond roughly to the limit of the trembles, which is the Bayesian Nash strategy profile we are testing.

(d) Show that for a Bayesian Nash equilibrium to trembling-hand perfect, both buyers must bid 4 with probability 1 when their values are high.

Solution: suppose otherwise. Let bidder A play 4 with probability less than 1 when he is a high type; then his expected utility is

$$\frac{1}{2}p_2\varepsilon_B^2(2) + \varepsilon_B^2 p_3(2) + \frac{1}{2}\varepsilon_B^3 p_3(1) + \varepsilon_B^2 p_4(2) + \varepsilon_B^3 p_4(1)$$

where p_b is the probability with which he mixes over his bids. Now suppose he plays 4 for sure; his expected utility is now

$$\varepsilon_B^2(2) + \varepsilon_B^3$$

We then look to check that

$$\varepsilon_B^2(2) + \varepsilon_B^3 > \frac{1}{2}p_2\varepsilon_B^2(2) + \varepsilon_B^2 p_3(2) + \frac{1}{2}\varepsilon_B^3 p_3(1) + \varepsilon_B^2 p_4(2) + \varepsilon_B^3 p_4(1)$$

$$\iff \quad 0 < \varepsilon_B^2\left(2 - p_2 - 2p_3 - 2p_4\right) + \varepsilon_B^3\left(1 - \frac{1}{2}p_3 - p_4\right)$$

Notice that $(1 - \frac{1}{2}p_3 - p_4) > 0$, and $(2 - p_2 - 2p_3 - 2p_4) > 0$. Therefore the inequality above must hold, and playing 4 is strictly preferred to placing any probability on bidding 2 or 3. Since we know that bidding your type is a weakly dominant strategy in a second-price auction (that is, bidding 4 is supported in equilibrium), it follows that for an equilibrium to be trembling-hand perfect high-type bidders must bid their type with probability 1.

(e) Is bidding 2 the unique trembling-hand perfect Bayesian equilibrium bidding strategy for a buyer with a low value?

Solution: yes. To see this, we need to consider Bayesian Nash equilibrium strategies in which high-type players bid their type. Since we want to show that bidding their valuation is the unique trembling-hand perfect Bayesian equilibrium bidding strategy, suppose that at least one of the low-type bidders does not follow this strategy; n particular, let bidder A mix over strategies with probability $\Pr(b_A = b) \equiv p_b$.

For trembles, let a low-type bidder B mix according to probabilities ε_B^b and a high-type bidder B mix according to probabilities ν_B^b. Bidder A's expected utility from the his mixture is then

$$p_3\left(\frac{1}{2}(1-p)\varepsilon_B^3(-1) + \frac{1}{2}p\nu_B^3(-1)\right) + p_4\left(\varepsilon_B^3(1-p)(-1) + \nu_B^3 p(-1) + \frac{1}{2}(1-p)\varepsilon_B^4(-2) + \frac{1}{2}p\nu_B^4(-2)\right)$$

Notice that this quantity is strictly negative (since trembles should be totally mixed). It follows that bidder A strictly prefers bidding 2 for sure — regardless of the ε tremble of bidder B — and so bidding 2 with probability 1 is the only strategy which is trembling-hand perfect.

2009 Fall comp, question 3

There is one firm in a market and a potential entrant. The incumbent firm can have marginal cost of production c_L or c_H with probabilities q and $(1-q)$, and the entrant has marginal cost c, where $c_H > c > c_L$. The cost of the entrant is known by the incumbent. Furthermore, the entrant faces an entry cost $f > 0$. There are two periods, and in each period, there is a single consumer with unkown reservation value $p > c_H$. In the first period, the incumbent is alone and chooses a price p_1 at which he offers to sell a unit to the consumer. In the second period, after observing this price, the entrant decides whether to enter or not. After entry, it observes the cost of the incumbent. Both firms then choose simultaneously their prices, and the consumer chooses to purchase from the cheaper of the two (or the lower-cost firm, in case of a tie).

Additionally, assume $f < c_H - c$, and that utility is over the sum of intertemporal profits.

(a) Model this problem as a Bayesian game.

Solution: see Figure 1. Notice by the * that we have assumed that high-cost incumbents cannot win against the entrant if it enters, and low-cost incumbents always beat the entrant. A more proper specification would make these the payoffs of Bertrand competition, but it is evident enough that this is going on; to this end, it may even be acceptable to simply say, "Bertrand competition," following 2's decision to enter or not.

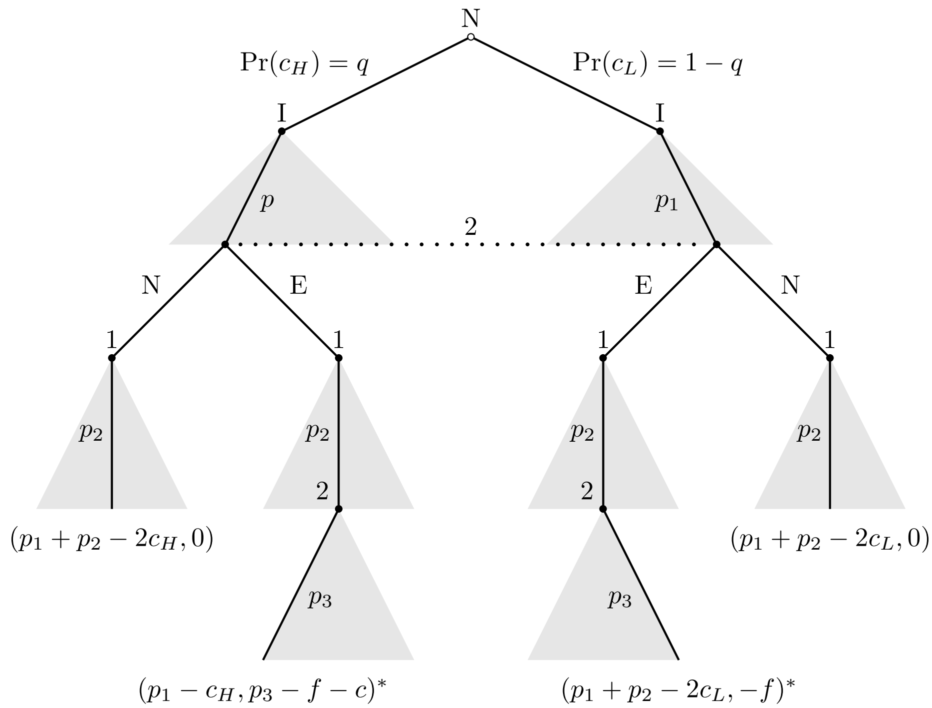

Figure 1: extensive-form Bayesian game for 2009 Fall comp, question 3

(b) Find all the separating ($p_H \neq p_L$) Perfect Bayesian equilibria of this game.

Solution: first, consider the second-round dynamics. In a separating equilibrium, if the entrant believes the incumbent is of low type the expected utility from entering is $-f$: the low type can always undercut the entrant on cost, so the consumer will buy from the incumbent firm. Thus the entrant will not enter. If the entrant believes the incumbed is of high type the expected utility from entering is $c_H - c - f > 0$: the entrant can now undercut the high-cost incumbent and Bertrand pricing gives us the equilibrium price level.

Considering that prices are set *after* observing whether or not entry has occurred, in the second round if there is no entry the price will be the maximum possible, which is the consumer's reservation value p. Otherwise, price is determined by Bertrand competition.

Now, notice that in any separating equilibrium if the incumbent's type is c_H its first-period price will be p. This follows from the two possible beliefs that the entrant can have in a separating equilibrium:

- If the entrant believes that a first-period price of p indicates an incumbent type of c_L, it will not enter in the second round. Thus an incumbent of type c_H obtains the highest possible price and keeps the entrant out.

- If the entrant believes that a first-period price of p indicates an incumbent type of c_H, it will enter in the second round; however, it will also enter from any other price p which indicates an incumbent type of c_H, so if the high-cost incumbent is to reveal its type it should charge the highest price possible. If there is incentive to deviate to the low-cost's strategy, we are not in equilibrium, so we may safely ignore this possibility.

What will the low-cost incumbent do? To begin, we compute the price at which the high-cost incumbent is indifferent between "reporting" a high type and a low type. We find this as

$$(p_L - c_H) + (p - c_H) = (p - c_H) + 0$$

where p_L is the price of the low-cost incumbent. Rearranging, we find that the threshold is

$$p_L = c_H$$

So any price $p'_L \leq p_L$ is such that the high-cost incumbent weakly prefers revealing its type by stating a price of p.

Now, we pose the opposite question: what is the lowest value of p_L such that a low-cost incumbent is indifferent between reporting a low type and a high type? We find this as

$$(p_L - c_L) + (p - c_L) = (p - c_L) + (c - c_L)$$

Rearraning, the lower threshold is

$$p_L = c$$

So we now know that we can support a separating equilibrium with a price range of $p_L \in [c, c_H]$.

What beliefs are necessary to support these equilibria? There is a large family for any particular p_L, but it must have the property that p_L is the *highest* price such that the entrant believes the incumbent faces a low marginal cost. Otherwise, depending on the particular belief structure, a low-cost incumbent might deviate to a higher price point or a high-cost incumbent might deviate to a lower price point (certainly one or the other will occur).

This set of strategies fully characterizes Perfect Bayesian equilibrium.

(c) Characterize the pooling ($p_H = p_L$) equilibrium of this game. Does a pooling equilibrium always exist?

Solution: in a pooling equilibrium, the entrant's beliefs upon observed play are that the incumbent faces a low marginal cost with probability q and a high marginal cost with probability $(1-q)$. Applying the standard results of third-stage Bertrand competition, the entrant witnesses expected utility of

$$q(0) + (1-q)(c_H - c) - f = (1-q)(c_H - c) - f$$

Then the entrant will enter if and only if $f < (1-q)(c_H - c)$.

What can we say about strategies and beliefs in the first round? Suppose first that $f < (1-q)(c_H - c)$. Then we need beliefs to be such that the low-cost incumbent cannot deviate and prevent the entrant from entering; this alone will be sufficient to keep the high-cost incumbent from deviating in the same way, by single-crossing. Further, we can ignore the possibility that the low-cost incumbent deviates upward in price as the high-cost incumbent would also want to do this. Obviously, the prevailing first-stage equilibrium price will be the highest price which supports pooling; let p_L be the highest price such that the entrant believes the incumbent faces a low cost. Then we need

$$(p_L - c_L) + (p - c_L) \leq (p_1 - c_L) + (c - c_L)$$

That is,

$$p_L + p \leq p_1 + c$$

Since the entrant is entering, consider the payoff from deviating upward. The entrant will either believe the types are still mixed, or that the deviation indicates a high marginal cost; regardless, the entrant will enter. So both types of incumbents strictly prefer deviating upward if possible. It follows that when $f < (1-q)(c_H - c)$ the first-stage price $p_1 = p$. Then appealing to the above equation, beliefs must satisfy $p_L \leq c$, or the highest possible price at which a low marginal cost is indicated is no larger than c, and that the distribution of types given an observed price of p is $\{q, (1-q)\}$.

Suppose now that $f > (1-q)(c_H - c)$. Similar to the previous, we now need that the high-cost incumbent does not want to deviate upward and reveal its price. It must be that a first-stage price of p indicates pooling or a high marginal cost (from the low type's incentive constraints), so we consider deviating to this price. We need

$$(p_1 - c_H) + (p - c_H) \geq (p - c_H)$$

That is,

$$p_1 \geq c_H$$

Piecing these two together, we can support a pooling equilibrium at price p_1 as long as p_1 is the greatest value at which the entrant believes pooling is occurring, with $f > (1-q)(c_H - c)$ and $p_1 \geq c_H$, or $f < (1-q)(c_H - c)$, $p_1 = p$, and the highest price at which the entrant believes the incumbent faces a low marginal cost is $p_L \leq c$.

It follows that a pooling equilibrium should always exist, given particular beliefs on the part of the entrant. This can be seen fairly simply: suppose that the entrant believes that prices convey no information. Then both incumbent types will set a price of p and the entrant will enter or not based on expected profits.

(d) Do all the pooling and separating equilibria you found in (b) and (c) satisfy the Intuitive Criterion?

Solution: no; in fact, most of the equilibria above do not satisfy the intuitive criterion. First, look at the separating equilibrium; so long as $p_L < c_H$, the low-cost firm should be able to renegotiate a better initial price level while still being convincing that it is the low type.

In the pooling equilibrium with entry, it is possible for the low-marginal-cost incumbent to persuade the entrant that only it is willing to play some price lower than p, namely c_H. Thus even this pooling equilibrium does not withstand the intuitive criterion.

2009 Spring comp, question 3

Each consumer purchases either 0 or 1 unit of a commodity. A type t buyer's utility gain from paying a price r for a unit of quality q, $u_i(q, r)$, is a strictly increasing function of q and strictly decreasing function of r. Moreover the single-crossing property holds, that is

$$-\frac{\partial u_t}{\partial q} \Big/ \frac{\partial u_t}{\partial r} > -\frac{\partial u_s}{\partial q} \Big/ \frac{\partial u_s}{\partial s} \qquad \forall s, t > s$$

The fraction of the population of type t is f_t, $t \in \{1, \ldots, T\}$. The cost of producing a unit of quality q is cq. Any offer made to one customer must be made to all customers.

Let $r = R(q)$ be the indifference curve for type s through (q', r') and (q'', r'') where $q'' > q'$.

(a) With the help of a graph, explain why any type $t > s$ strictly prefers (q'', r'') and any type $t < s$ strictly prefers (q', r').

Solution: intuitively, the single-crossing ratio above tells us that the change in r necessary to keep a type-t agent indifferent between levels of q is greater than that for type s (assuming $t > s$; when $t < s$ the symmetrically opposite case holds); that is, in terms of derivatives, for type t, "r is changing utility relatively less than q," when compared with type s. This gives us a clue as to how indifference curves should appear when they cross one another.

What can we say about the slope of indifference? Since u is strictly increasing in q and strictly decreasing in r, indifference curves should be monotonic, and in particular they should be increasing. Suppose that there is a decreasing segment of an indifference curve; then we may sustain indifference by increasing q and decreasing r which by our assumptions should leave the any agent strictly better off.

Lastly, we appeal to the nature of u to claim that more q is better and more r is worse; this tells us that the area below-and-to-the-right-of the indifference curve is preferred, while the area above-and-to-the-left-of the indifference curve is not. A graphical proof of the claim in this question appeals to this fact and is seen in Figure 1.

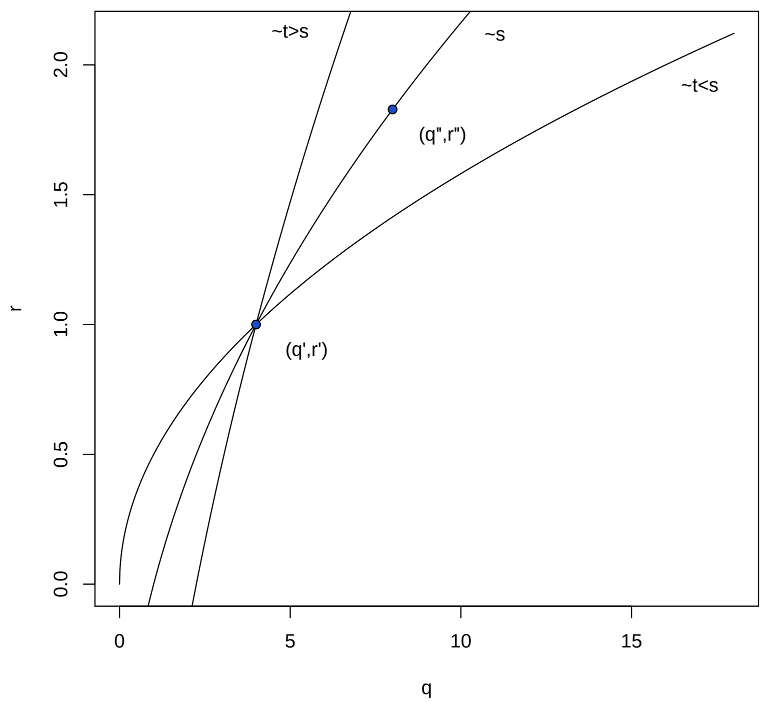

Figure 2: graphical evidence of the claim in the 2009 Spring comp, question 3(a) $(u = \sqrt{\theta q} - r)$

(b) Explain why the slope of $R(q)$ satisfies $R'(q) = -\frac{\partial u_s}{\partial q}(q, R(q))/\frac{\partial u_s}{\partial r}(q, R(q))$. Hence prove that the statement in part (a) is true.

Solution: along an indifference curve, we have

$$\mathbf{u} = u\left(\bar{q}, R\left(\bar{q}\right)\right)$$

To sustain indifference, we take the derivative with respect to q and set it equal to 0,

$$0 = \frac{\partial}{\partial q} u_s\left(\bar{q}, R\left(\bar{q}\right)\right) + \frac{\partial}{\partial r} u_s\left(\bar{q}, R\left(\bar{q}\right)\right) R'\left(\bar{q}\right)$$

Algebraically, this comes to

$$R'(q) = -\frac{\partial u_s}{\partial q}\left(q, R\left(q\right)\right) \Bigg/ \frac{\partial u_s}{\partial r}\left(q, R\left(q\right)\right)$$

Now, we know an agent of type s is indifferent between (q', r') and (q'', r''). Let $t > s$; from the single-crossing property, we know that

$$R'_s(q) < R'_t(q) \quad \forall\,(q, R_s(q)) \sim_s (q', r')$$

It follows that $(q'', R_s(q''))$ falls below t's indifference point (relative to (q', r')) at q'', $(q'', R_t(q''))$. Since paying less is preferred to paying more, the player of type $t > s$ will strictly prefer $(q'', r'') \succ_t (q'', R_t(q''))$ to (q', r').

There was demand in section for a more mathematical proof of this conjecture; the proof given was correct in spirit but fundamentally flawed (in particular, we do not know $R'_s(q) < R'_t(q) \forall q$, since single-crossing applies only per (q, r)-tuples). Moving from (q', r') to (q'', r''_i), to support indifference we must have

$$r''_s = r' + \int_{q'}^{q''} R'_s(q)dq$$
$$r''_t = r' + \int_{q'}^{q''} R'_t(q)dq$$

Since $R'_s(q') < R'_t(q')$, we know that for q'' sufficiently close to q' we must have $r''_s < r''_t$. How can we show this more generally? Suppose $r''_t < r''_s$; then by continuity properties of all involved functions there must be some $q^* \in (q', q'')$ such that $\int_{q'}^{q^*} R'_s(q)dq = \int_{q'}^{q^*} R'_t(q)dq$. By the previous argument, for all $q = q^* + \varepsilon$ for ε positive and sufficiently small we must then have $\int_{q'}^{q} R'_s(q)dq < \int_{q'}^{q} R'_t(q)dq$. It follows that along this indifference curve we can never reach $r''_t < r''_s$.

What about $r''_t = r''_s$? Appealing to the previous argument we choose some (q^*, r^*), $q' < q^* < q''$ such that $(q', r') \prec_s (q^*, r^*)$ and $(q', r') \succ_t (q^*, r^*)$; that is, (q^*, r^*) lies below type-t's (q', r')-indifference curve and above type-s's. By assumption, we know that $(q'', r'') \succ_s (q^*, r^*)$, hence the type-$s$ indifference curve through (q^*, r^*) must cross the type-t indifference curve through (q', r') *from below*, violating the assumption that $R'_s(q) < R'_t(q)$ along indifference curves (note that this technically requires us to make the indifference-slope argument from above for $t < s$, but this is a nonissue).

It follows that $r''_t > r''_s$ — the price for quality necessary to keep a type-t agent indifferent between (q', r') and (q'', r''_t) is greater than that necessary to keep a type-s agent indifferent between (q', r') and (q'', r''_s). Since a lower price is certainly preferred (u is strictly decreasing in r) it follows that $(q', r') \sim_s (q'', r'') \prec_t (q'', r''_t)$. Then the type-$t$ agent sees $(q', r') \prec_t (q'', r'')$. Note that the $r''_t \neq r''_s$ proof above could be made more rigorous, but all necessary steps are there; graphing it out may help to clarify the argument.

The proof of the opposite for $t < s$ will follow identically.

(c) Explain briefly why, for the direct mechanism $\{(q_t, r_t)\}_{t=1}^T$ to be incentive-compatible, $\{q_t\}_{t=1}^T$ must be increasing.

Solution: this essentially states that agents with a stronger preference for quality must receive goods of higher quality in equilibrium. Otherwise, agents with a weaker preference for quality will prefer the stronger type's allocation. Consider parts (a) and (b) above; to retain incentive compatibility, if $q_t < q_s$ we need $(q_t, r_t) \preceq_s (q_s, r_s)$ and $(q_t, r_t) \succeq_t (q_s, r_s)$. But if $(q_t, r_t) \sim_s (q_s, r_s)$, then $(q_t, r_t) \prec_t (q_s, r_s)$; thus $(q_t, r_t) \prec_s (q_s, r_s)$.

There must then be some $r'_t < r_t$ such that $(q_t, r'_t) \sim_s (q_s, r_s)$. Since u is strictly decreasing in r, $(q_t, r'_t) \succ_t (q_t, r_t)$. But from (b) above, $(q_t, r'_t) \succ_t (q_s, r_s)$; appealing to transitivity, we then have $(q_t, r_t) \succ_t (q_s, r_s)$ regardless of s's preference. Then it cannot be that $q_s > q_t$.

(d) Consider any direct mechanism $\{(q_t, r_t)\}_{t=1}^T$. Suppose that $\{q_t\}_{t=1}^T$ is increasing and the local downward constraints are binding. Show that this mechanism is incentive-compatible.

Solution: we show first that a bidder of type t will not want to deviate upward, then that they will not want to deviate downward.

Since the local downward constraints bind, we have $(q_s, r_s) \sim_s (q_{s-1}, r_{s-1})$; by the argument in part (c) we know that $(q_s, r_s) \prec_{s-1} (q_{s-1}, r_{s-1})$. Suppose that $(q_s, r_s) \succeq_{s-2} (q_{s-1}, r_{s-1})$. Then there is some $r'_{s-1} > r_{s-1}$ such that $(q_s, r_s) \sim_{s-2} (q_{s-1}, r'_{s-1})$ and hence — with q_t increasing — $(q_s, r_s) \succ_{s-1} (q_{s-1}, r'_{s-1})$. This contradicts $(q_s, r_s) \prec_{s-1} (q_{s-1}, r_{s-1})$ with $r_{s-1} < r'_{s-1}$; so it must be that $(q_s, r_s) \prec_{s-2} (q_{s-1}, r_{s-1})$. Since $(q_{s-2}, r_{s-2}) \succ_{s-2} (q_{s-1}, r_{s-1})$ (again by appeal to part (c) above), it follows that $(q_{s-2}, r_{s-2}) \succ_{s-2} (q_s, r_w)$. This argument may be iterated arbitrarily to show that

$$(q_{s-k}, r_{s-k}) \succ_{s-k} (q_s, r_s) \qquad \forall k < s$$

To show the other direction, apply the binding local downward constraints to see $(q_{s-2}, r_{s-2}) \sim_{s-1} (q_{s-1}, r_{s-1})$; hence $(q_{s-2}, r_{s-2}) \prec_s (q_{s-1}, r_{s-1})$. Since $(q_{s-1}, r_{s-1}) \sim_s (q_s, r_s)$ it follows that $(q_{s-2}, r_{s-2}) \prec_s (q_s, r_s)$. Again, this argument may be generalized forward for all $s-k$, $k < s$.

Therefore this system of allocations is incentive compatible.

(e) Solve for the profit-maximizing quality choices if $(f_1, f_2) = (\frac{3}{4}, \frac{1}{4})$, $u_t(q, r) = \theta_t q - \frac{1}{2}q^2 - r$, where $(\theta_1, \theta_2) = (10, 30)$ and $c = 2$.

Solution: in a separating equilibrium, the low type's participation constraint will bind (this is simple to see graphically). So we have

$$r_1 = 10q_1 - \frac{1}{2}q_1^2$$

The high type's local downward constraint will bind,

$$30q_2 - \frac{1}{2}q_2^2 - r_2 = 30q_1 - \frac{1}{2}q_1^2 - r_1$$

Rearranging, this gives us

$$r_2 = 30\left(q_2 - q_1\right) - \frac{1}{2}\left(q_2^2 - q_1^2\right) + 10q_1 - \frac{1}{2}q_1^2$$

The firm's optimization is then

$$\max_{q_i} \frac{3}{4}\left(10q_1 - \frac{1}{2}q_1^2 - 2q_1\right) + \frac{1}{4}\left(30\left(q_2 - q_1\right) - \frac{1}{2}\left(q_2^2 - q_1^2\right) + 10q_1 - \frac{1}{2}q_1^2 - 2q_2\right)$$

First-order conditions yield

$$\begin{aligned}
\frac{\partial}{\partial q_1}: \quad & \frac{15}{2} - \frac{3}{4}q_1 - \frac{3}{2} - \frac{15}{2} + \frac{1}{4}q_1 + \frac{5}{2} - \frac{1}{4}q_1 = 0 \\
\iff \quad & \frac{3}{4}q_1 = 1 \\
\iff \quad & q_1 = \frac{4}{3} \quad \Longrightarrow \quad r_1 = \frac{112}{9}
\end{aligned}$$

$$\begin{aligned}
\frac{\partial}{\partial q_2}: \quad & \frac{15}{2} - \frac{1}{4}q_2 - \frac{1}{2} = 0 \\
\iff \quad & \frac{1}{4}q_2 = 7 \\
\iff \quad & q_2 = 28 \quad \Longrightarrow \quad r_2 = \frac{964}{3}
\end{aligned}$$

Note that we can see that the low-type agent receives 0 utility graphically, or we can solve the question under the opposite assumption (that the high-type agent receives 0 utility) and see that this implies that the low-type agent receives negative utility.

2008 Spring comp, question 4

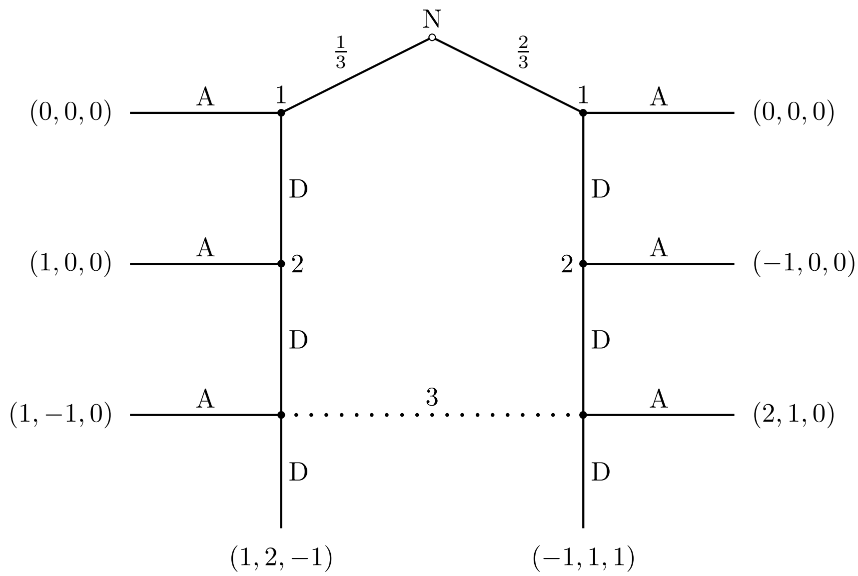

Figure 3: the extensive-form game for the 2008 Spring comp, question 4

Answer the following questions for the game in Figure 3.

(a) Consider a PBE in which player 3 plays $zA + (1 - z)D$. For what values of z will player 1 always play D? For what values of z will player 2 always play D? Use your answer to show that for all z player 3's information set will always be reached in equilibrium.

Solution: notice that in the left branch of the game tree, player 1 is better off from playing D in every case than from playing A; so down this branch player 1 will play D regardless of z. Further, player 2 is better off (in expectation) from playing D than from playing A so long as $z < \frac{2}{3}$.

In the right branch of the game tree, player 2 is better off from playing D in every case than from playing A; so down this branch player 2 will play D regardless of z. Further, player 1 is better off playing D than A so long as

$$p_D^2(-1) + \left(1 - p_D^2\right) z(2) + \left(1 - p_D^2\right)(1 - z)(-1) = -1 + 3\left(1 - p_D^2\right) z > 0$$

When player 2 plays D for sure, player 1 is better off when $z > \frac{1}{3}$. In a perfect Bayesian equilibrium, players must be properly responding at all information sets — even those off the equilibrium path — so we know that player 2 will play D for sure in perfect Bayesian equilibrium.

We now check several cases:

- When $z \leq \frac{2}{3}$ and the left branch is played, players 1 and 2 both choose D and player 3's information set is reached.
- When $z > \frac{2}{3}$ and the left branch is played, player 1 will play D but player 2 will play A; player 3's information set is not reached.
- When $z \geq \frac{1}{3}$ and the right branch is played, players 1 and 2 both choose D and player 3's information set is reached.

- When $z < \frac{1}{3}$ and the right branch is played, player 1 will play A; neither player 2's nor player 3's information set is reached.

Now, what can we say about z? Provisional on his information set being reached, the following choices arise:

- When $z < \frac{1}{3}$, if it is his turn player 3 knows that he is in the left branch. Then his best response is to play A, implying $z = 1$, a contradiction.

- When $z > \frac{2}{3}$, if it is his turn player 3 knows that he is in the right branch. Then his best response is to play D, implying $z = 0$, a contradiction.

- When $z \in [\frac{1}{3}, \frac{2}{3}]$, player 3 does not know which branch he is in. His expected payoff from playing A is 0 while his expected payoff from playing D is $\frac{1}{3}(-1) + \frac{2}{3}(1) = \frac{1}{3}$. Thus he should play D, implying $z = 0$, a contradiction.

How can we make player 3 willing to randomize? We need to reduce the posterior belief that he has that he is in the right branch. We have seen that player 1 has a dominant strategy in the left branch while player 2 has a dominant strategy in the right branch. Moreover, if player 2 is mixing in the left branch (possible when $z = \frac{2}{3}$) then player 1 is *not* mixing in the right branch. Therefore we cannot reduce the posterior probability that 3 believes he is in the right branch by having 2 mix in the left branch.

The remaining option is to have player 1 mix in the right branch. To support this, we need $z = \frac{1}{3}$. How can we support 3 mixing? We compute indifference, recalling that 1 and 2 both have pure strategy best responses in the left branch.

$$
\begin{aligned}
& E\left[u_3(A)\right] = E\left[u_3(D)\right] \\
\iff \quad & 0 = \left(\frac{1}{3}\right)(-1) + \left(\frac{2}{3}\right)\Pr(a_1 = D)(1) \\
\iff \quad & \frac{1}{2} = \Pr(a_1 = D)
\end{aligned}
$$

Then player 3 is indifferent between A and D when player 1 mixes in the right branch, playing A with probability $\frac{1}{2}$ and D with probability $\frac{1}{2}$.

This fully categorizes the desired perfect Bayesian equilibrium. Notice that 3's information set is not reached with certainty! That is, there is some positive probability that player 3's information set is not reached (in particular, this probability is $\frac{1}{3}$). To the extent that this contradicts the phrasing of the question, we are left with two options:

- The payoffs to player 3 from playing D should be adjusted so that they are 0 in expectation; here, we could set $u_3 = -2$ in the left branch.
- We may be applying the wrong definition of beliefs. It's possible that 3 could believe that 1 is mixing according to this strategy, while in reality 1 is playing a pure strategy. This of course means that 3's beliefs are incorrect on the equilibrium path — a requirement for perfect Bayesian equilibrium — but there are equilibrium notions which revolve around *falsifiability*. Since 3 cannot falsify that this is the strategy that 1 is playing, he's free to respond by randomizing in equilibrium. I should note here that this does not appear to be consistent with the definition of perfect Bayesian equilibrium, but it seems to be worth considering on its own merit.

(b) Find all the perfect Bayesian equilibria in which player 3 gets to move.

Solution: the argument above is sufficient to establish that there is a unique perfect Bayesian equilibrium. That is, we have sufficiently established that there is a unique possibility for player 3 to randomize; in the course of this argument, we also established that pure strategies on the part of 3 lead to contradictions of optimal behavior subject to Bayesian updating and proper beliefs. Therefore there are no remaining perfect Bayesian equilibria.

A good test of comprehension may be to apply the "change the payoffs" strategy mentioned in the bullets above to see what pure strategies (if any) now arise for player 3.

(c) Is there a Bayesian Nash equilibrium which is *not* perfect? Explain.

Solution: yes. Consider the equilibrium in which 1 plays D in the left branch but plays A in the right due to a contingent belief that 2 will play A in the next stage. In this equilibrium, player 2 will play A in the left branch since player 3 will know which branch is being played. This equilibrium is consistent with players' beliefs, but players' beliefs are not consistent with subsequent optimal behavior; in particular, player 1's belief about player 2's action in the right branch is inconsistent with what 2's behavior would actually be. It follows that this is not a perfect Bayesian equilibrium.

It was raised in section that equilibrium should be specified more fully here. This is a good point, although the above seems sufficient to indicate why a non-perfect Bayesian Nash equilibrium exists. To this end, let 1 have the beliefs above, let 2 play A in both branches, and let 3 play A regardless; the beliefs supporting this equilibrium emphasize the fact that a player's view of future outcomes has no need to respect reality in any way in a Bayesian Nash equilibrium.

2007 Spring comp, question 4

Three sellers each have one car to sell to a single buyer. All the cars come from one of two plants: at plant G, fraction p of the cars produced are of high quality and the rest are of low quality; at plant B none of the cars produced are of high quality and all are of low quality. Each seller observes the quality of his own car, but not which plant it comes from. The probability that a car comes from a given plant is $\frac{1}{2}$.

The sellers value high-quality cars at 10 each and low-quality cars at 5 each.

The sellers and the buyer have agreed to use the following mechanism: the sellers each report the quality of their car. If two or more of them report high, sellers who report high will receive 12 and sellers (if any) who report low will receive 11; if two or more of them report low, sellers who report low will receive 10 and sellers (if any) who report high will receive 9. The buyer has agreed to accept all cars at these prices.

(a) Consider the event that the other sellers observe H. Explain why the conditional probabilities of this event are $\Pr(HH|H) = p^2$ and $\Pr(HH|L) = p^2(\frac{1-p}{2-p})^2$, where a seller is conditioning on his own observation. Solve for the conditional probabilities of the other events.

Solution: a large part of this question is in the interpretation. Here, *all* cars come from the same plant, otherwise the desired results here make no sense. So if a given seller's car is of high quality, she knows that the car was produced at plant G. Hence the remaining two cars also come from plant G, each of high quality with i.i.d. probability p. Then we have

$$\Pr(HH|H) = p^2$$

On the other hand, if the car is of low quality it might have come from either plant. We see

$$\Pr(HH|L) = \Pr(HH|L,G)\Pr(G|L) + \Pr(HH|L,B)\Pr(B|L)$$

We know $\Pr(HH|L,B) = 0$, since the bad plant cannot produce good cars. Further, $\Pr(HH|L,G) =$

$\Pr(HH|G) = p^2$. So we need only compute $\Pr(G|L)$; this will follow from Bayes' rule,

$$\begin{aligned}\Pr(G|L) &= \frac{\Pr(L|G)\Pr(G)}{\Pr(L)} \\ &= \frac{\frac{1}{2}(1-p)}{\Pr(L|G)\Pr(G) + \Pr(L|B)\Pr(B)} \\ &= \frac{\frac{1}{2}(1-p)}{\frac{1}{2}(1-p) + \frac{1}{2}} \\ &= \frac{1-p}{2-p}\end{aligned}$$

This gives us

$$\Pr(HH|L) = p^2\left(\frac{1-p}{2-p}\right)$$

A full table of probabilities is

$$\Pr(HH|H) = p^2 \qquad \Pr(HH|L) = p^2\left(\frac{1-p}{2-p}\right)$$

$$\Pr(HL|H) = 2p(1-p) \qquad \Pr(HL|L) = 2\left(\frac{(1-p)^2 p}{2-p}\right)$$

$$\Pr(LL|H) = (1-p)^2 \qquad \Pr(LL|L) = \frac{(1-p)^3+1}{2-p}$$

Notice that there is a critical typo in this comp question (presumably the question writer assumed conditional independence of $H|L$). This is rare, but people make mistakes.

(b) For $p = 0.9$, show that it is a Bayesian Nash equilibrium for all sellers to report truthfully.

Solution: suppose that all other sellers are reporting truthfully. A particular seller's payoff conditional on a low signal is

$$\begin{aligned}u(L|L) &= (11-5)p^2\left(\frac{1-p}{2-p}\right) + (10-5)(2)\left(\frac{(1-p)^2 p}{2-p}\right) + (10-5)\left(\frac{(1-p)^3+1}{2-p}\right) \\ u(H|L) &= (7-5)p^2\left(\frac{1-p}{2-p}\right) + (12-5)(2)\left(\frac{(1-p)^2 p}{2-p}\right) + (9-5)\left(\frac{(1-p)^3+1}{2-p}\right)\end{aligned}$$

Subtracting, we find

$$u(L|L) - u(H|L) = -p^2\left(\frac{1-p}{2-p}\right) - 4\left(\frac{(1-p)^2 p}{2-p}\right) + \left(\frac{(1-p)^3+1}{2-p}\right)$$

We concern ourselves then with the sign of

$$-p^2(1-p) - 4(1-p)^2 p + (1-p)^3 + 1 \tag{1}$$

Intuitively, when p is large the $-(1-p)^k$ terms are more than outweighed by the $+1$ term at the end. So truthful reporting is a dominant strategy conditional on the low type, given that other players are reporting truthfully.

Now consider the conditional utility of a seller with a high-type car.

$$\begin{aligned}u(L|H) &= (11-10)p^2 + (10-10)(2)p(1-p) + (10-10)(1-p)^2 \\ u(H|H) &= (12-10)p^2 + (12-10)(2)p(1-p) + (9-10)(1-p)^2\end{aligned}$$

Subtracting, we find

$$u(H|H) - u(L|H) = p^2 + 4p(1-p) - (1-p)^2 \tag{2}$$

Again, for p large the negative terms will be dominated by the positive terms.

Thus when $p = 0.9$, truthful reporting is a Bayesian Nash equilibrium.

(c) For $p = 0.5$, show that it is not a Bayesian Nash equilibrium for all sellers to report truthfully.

Solution: we now check (1) and (2) to see their values at $p = \frac{1}{2}$.

$$u(H|H) - u(L|H) = \frac{1}{4} + 4\left(\frac{1}{2}\right)\left(\frac{1}{2}\right) - \left(\frac{1}{2}\right)^2 > 0$$

So the high type is still willing to truthfully report.

We turn to the numerator of the low type,

$$-p^2(1-p) - 4(1-p)^2 p + (1-p)^3 + 1 = -\left(\frac{1}{2}\right)^2\left(\frac{1}{2}\right) - 4\left(\frac{1}{2}\right)^2 \frac{1}{2} + \left(\frac{1}{2}\right)^3 + 1 = \frac{1}{2}$$

Provided our algebra is correct, it seem that the desired answer to this question follows from the incorrect conditional probability given in part (a). So we cannot show what is asked for here, but we can see that for p sufficiently small a seller with a high-quality car is prefers to misreport.

(d) For $p = 0.5$, find a symmetric Bayesian Nash equilibrium in behavior strategies.

Solution: given the previous pieces of the question, suppose that p is sufficiently small that truthful reporting cannot arise in Bayesian Nash equilibrium. Then two natural strategies arise: always report high, and always report low. It is obvious that these are equilibria.