Bayesian Nash Equilibrium and Spence Signaling Notes

Jacob Kohlhepp

October 29, 2021

1 Bayesian Nash Equilibrium (BNE)

Big idea: BNE is just Nash Equilibrium with expectations when we have types of players, and players usually only know their own type for sure.

Definition 1. *Type t_i's expected utility from action a_i given strategies $s_{-i}(t_{-i})$ of other players is given by:*

$$\mathbb{E}[u_i(a_i, s_{-i}(t_{-i}), \theta)|t_i] = \int_{t_{-i}} u_i(a_i, s_{-i}(t_{-i}), \theta) dF_{T_{-i}}(t_{-i})$$

Definition 2. *Action a_i is a best response for type t_i to strategy $s_{-i}(t_{-i})$ of other players if:*

$$\mathbb{E}[u_i(a_i, s_{-i}(t_{-i}), \theta)|t_i] \geq \mathbb{E}[u_i(a_i', s_{-i}(t_{-i}), \theta)|t_i] \forall a_i'$$

Definition 3. *A strategy profile $s^* = (s_1^*, s_2^*, ..., s_N^*)$ is a Bayesian Nash Equilibrium (BNE) if each type t_i of each player i is playing a best response in expectation to the strategies of others $s_j^*(\theta_j)$*

These definitions apply very generally. In this class we will often have that people have i.i.d. types/values. In this case we usually write that true type is just θ_i and the reported type is $\tilde{\theta}_i$. In this case, we can re-write the best response definition as:

Definition 4. *Action a_i is a best response for type θ_i to strategy $s_{-i}(\theta_{-i})$ of other players if:*

$$\mathbb{E}[u_i(a_i, s_{-i}(\theta_{-i}), \theta_i)|\theta_i] \geq \mathbb{E}[u_i(a_i', s_{-i}(\theta_{-i}), \theta_i)|t_i] \forall a_i'$$

Notice that since θ_i's are i.i.d. the above objects are much simpler, in that $\theta_{-i}|\theta_i = \theta_{-i}$ in distribution. If this confuses you, ignore it.

1.1 Example: First-Price Auction

Consider the first-price auction. The action space is a bid. So to find the BNE of the auction we need to find $a(\theta_i)$ which maps values to bids. We know utility from a bid is:

$$u_i(a(\theta_i), s_{-i}(\theta_{-i})) = (\theta_i - a(\theta_i))\mathbb{I}\{a(\theta_i) > \max\{a(\theta_j)\}_{j \neq i}\}$$

Denote the symmetric bid function of others b. Assuming symmetric strategies and i.i.d. values, we can take the conditional expectation of this to get that:

$$\mathbb{E}[u_i(a(\theta_i), s_{-i}(\theta_{-i}))] = (\theta_i - a(\theta_i))[F(b^{-1}(a(\theta_i)))]$$

The BNE maximizes this w.r.t a holding fixed bids of others. Doing this and a bunch of tricks gets the answer:

$$a(\theta_i) = \theta_i - \frac{\int_{\underline{\theta_i}}^{\theta_i} F(x)^{N-1} dx}{F(\theta_i)^{N-1}}$$

This can be derived in an easier way using mechanism design.

2 Bonus: Spence Signaling Model

Source: NYU Lecture Notes [1]

2.1 Model

- Two types of workers: H, L which have different productivities, θ. Specifically, assume $\theta_H > \theta_L$.
- The worker is high type p fraction of the time.
- Education is costly, and the unit cost is based on productivity: $c(e) = e/\theta$, so that high types find acquiring education less costly.
- There is one firm who can hire one worker.
- Timing: nature moves, and a worker is born with private information about their type. The worker then chooses publicly observed education level e. The firm then offers a salary, s, conditional on e.
- There is perfect competition for workers, so that the wage will be $E[\theta|e, s^*]$, that is the expected productivity given the education choice and equilibrium strategies.
- Final payoffs to the worker are just $s(e) - e/\theta$. Payoffs to the firm are: $\theta - s(e)$.

Note: The key condition for the below equilibria (particularly separating) is the single-crossing of expected utility given type, $U(e, \theta)$. This comes from the fact that the cost of education is higher for low types.

2.2 Solution

I consider all Perfect Bayesian Equilibrium. It turns out there are three types.

1. *Separating Equilibrium.* This is the most obvious one, where education is a good signal of ability, and the firm pays two wages: a high wage for high education and a low for low education.

   To solve for this equilibrium, we must make the crucial observation that given that education is revealing and has no productive value, it must be that the low type gets education of 0. Next, we need two conditions to be met.

   First, the high type must prefer to make the education investment, call it e^*, over no education and the low wage. That is:

$$s_H - e^*/\theta_H \geq s_L$$

[1] https://www.econ.nyu.edu/user/debraj/Courses/05UGGameLSE/Handouts/05uggl10.pdf

Second, the low type must prefer the low wage and no education to exerting effort for education and the high wage. That is:

$$s_L \geq s_H - e^*/\theta_L$$

Solving both for e^* yields:

$$\theta_H(s_H - s_L) \geq e^* \qquad e^* \geq \theta_L(s_H - s_L)$$

Given $\theta_H > \theta_L$ there exists many such e^*. We select one of them.

Finally, we must specify both on and off-equilibrium beliefs for the firm. Given the strategies of the workers, the only restriction on beliefs is that $e = e^*$ means the worker is high type for sure, while $e = 0$ means the worker is low type for sure. For all other values of e, there are many choices. An easy specification is just that $e \geq e^* \implies \theta_H$ while $e < e^* \implies \theta_L$.

This means that $s_H = \theta_H$ and $s_L = \theta_L$. Optimal education for high types is any e^* in the range$[\theta_L(\theta_H - \theta_L), \theta_H(\theta_H - \theta_L)]$.

2. *Pooling equilibrium.* As the name states, in this equilibrium high types and low types choose the same levels of education, and blend in with one another. In this equilibria we posit that both types choose the same education level $\tilde{e}$. There are actually many variants of this equilibrium. I consider one. Suppose the firm believes that any $e \neq \tilde{e}$ represents the low type for sure, while any $e = \tilde{e}$ represents either high or low, with the same prior belief. Clearly this is compatible with on-path strategies of workers (who only play $\tilde{e}$).

   Then the firm pays a medium salary of $s_M = p\theta_H + (1-p)\theta_L$ for education level $\tilde{e}$ and then a low salary θ_L for any other education choice. Clearly high types have no reason to deviate: any other education choice lower the wage. For low types, it must be that the education cost is worth the benefit of the salary bump, that is:

$$s_M - \tilde{e}/\theta_L \geq s_L \implies \theta_L(s_M - s_L) \geq \tilde{e} \implies \theta_L(p\theta_H - p\theta_L) \geq \tilde{e}$$

   As long as the chosen education level satisfies this inequality (and is weakly greater than 0) we have equilibrium.

3. *Hybrid Equilibria.* There are rather uninteresting hybrid equilibria where types mix over education level.

Note: Because there is so much flexibility over off-path beliefs, people generally consider refining the equilibria by imposing what is called the intuitive criterion. This criterion says that the firm cannot believe the worker is a type after a signal that is dominated for that type is received. This rules out all the equilibria above except the separating equilibria where $e_L = 0$ and $e_H = \theta_L(\theta_H - \theta_L)$.