Economics 203: Section 4

Jeffrey Naecker

Februrary 6, 2014

1 Logistics

1.1 Office Hours This Week

Please note that my office hours today (the 6th) will be from 4:00 to 6:00 pm (one hour later than usual). They will still be in Landau 360.

1.2 Mid-Quarter Section Review

In the next few days, you will be getting an email with a survey about this section. Please fill this out; it should only take a few minutes, but it will help me immensely in making section better for you.

2 Concepts

2.1 Mixed Strategy Nash Equilibrium with Continuous Strategy Sets

Last week we talked about methods for finding MSNE that work very well for games where each player has only a few strategies. We need to employ some different techniques when the strategy sets become continuous, as the following example shows.

Example 1. Consider the Betrand model with 2 firms and capacity constraints, as in Doug's lecture notes. Firms simultaneously choose quantities, but they can't produce more than some limit K. Consumers buy up to K units from the lower price firm; if this does not satisfy demand, they buy the remainder from the other firm, as long as the price is not above their value. We assume demand is a step function: demand equals Q if $p \leq v$, 0 otherwise.

Let's assume without proof that a symmetric MSNE exists, where each player chooses their price according to a probability distribution whose support is $[a, v]$. Let F be the CDF of this distribution. Let's also assume that the CDF is atomless; this tells us there will be no ties (or more accurately, the event that the two firms naming the same price has zero probability of occurring).

Given these assumptions, let's find the CDF of this MSNE. To do so, imagine yourself as one of the firms, and assume your opponent is naming a price according to F. You are naming a price p in the support of F. If your opponent's price is higher than p, you sell K units for profit Kp. Note that this happens with probability $1 - F(p)$. Similarly, if your opponent's price is lower than p, you sell $(Q - K)$ units for a profit of $(Q - K)p$. This happens with probability $F(p)$. Thus your expected profit for naming p when your opponent is playing according to F is

$$(Q - K)pF(p) + Kp\,(1 - F(p)) = C,$$

where the constant C reminds us that for any p, your expected payoff must be the same. We can then solve for F:

$$F(p) = \frac{C - Kp}{(Q - 2K)p}.$$

Note that we have two more terms to nail down, however: the constant C and the lower bound a. We will find the two with two tricks that will be very useful for these types of problems.

First, we note that $F(v) = 1$ by definition. This allows us to find that $C = (Q - K)v$. Note that this can be interpreted as your profits as a firm in this game playing the MSNE. This makes a lot of sense, because if you deviate to $p = v$, you will sell only $Q - K$ units, since the other firm is undercutting you with probability 1; thus your profit is just $(Q - K)v$. (Note that this is effectively another way to find C.) Thus form of the CDF is completely nailed down:

$$F(p) = \frac{K}{2K - Q} - \frac{(Q - K)v}{(2K - Q)p}.$$

Lastly, we note that $F(a) = 0$ by definition. This allows us to find $a = \left(\frac{Q}{K} - 1\right) v$. Note that $a > 0$ since $K < Q$, so all named prices are positive. Note also that $a < v$ since $Q < 2K$, so the distribution is not degenerate.

2.2 Correlated Equilibrium

We noted in lecture that Nash equilibrium is a convincing solution concept at least partly because it is self-enforcing, in the sense that deviations cannot benefit you if all of your opponents are following the agreement. Yet this property is not unique to NE, and in fact can be found to hold in a much broader set of outcomes if we allow for players to condition on random events. These outcomes, if they have have same self-enforcing attribute, are know as correlated equilibria.

Definition 1. Consider a finite game (S, g). A probability density δ over S is a **correlated equilibrium (CE)** iff for all $i \in I$ and for all s_i played with strictly positive probability

$$s_i \in \arg\max_{s_i' \in S_i} E_{S_{-i}}[g(s_i', s_{-i})|s_i, \delta].$$

That is, s_i must be a best response to the distribution of outcomes implied by all players following δ and player i playing s_i. This may seem a bit circular, but you can think of the logic in the following way: Players agree to coordinate their actions on some publicly observable random event, which implies the joint distribution of outcomes δ. Suppose that player i has agreed to play s_i in a given state or states of the world. In this case, the possible outcomes are now distributed conditional on s_i and δ, since her signal that she is supposed to play s_i has given her some information about the state of the world. (She might know the exact outcome with certainty, or she may just have a distribution over some subset of the outcomes.) She then considers all possible deviations, i.e. maximizes over the s_i''s. Essentially, in any state of the world where i finds herself having agreed to play s_i, that strategy had better be a best response given i's knowledge about the state.[1]

Note that in MSNE, any mixtures had to be independent. Yet in a CE, the mixtures of players can be dependent, since the are conditioned on a publicly observable random event. This is what allows us to reach equilibrium payoffs outside the convex hull of the possible NE payoffs.

2.2.1 Finding Correlated Equilibria

There is a nice trick for quickly verifying that a distribution δ is a CE. Consider all possible pairs $(s_i, s_i') \in S_i \times S_i$. Draw a matrix where the rows are labelled by s_i and the columns by s_i' (so the matrix is K_i by K_i). In each cell, put down the value of $E_{S_{-i}}[g(s_i', s_{-i})|s_i, \delta]$. Then for each row, circle the maximum payoff in that row. If all the payoffs along the diagonal are circled, then player i's strategies could be part of a CE. If we verify this diagonal property for each player, then we know δ is a CE.[2] We can see this at work in the example from lecture.

Example 2. Consider the game in Figure 1. Consider the distribution δ such that $P((A, a)) = \gamma$, $P((A, b)) = P((B, a)) = \frac{1-\gamma}{2}$, and $P(B, b) = 0$. For what values of γ is this a CE?

[1] For more discussion about this interpretation of CE, see Osbourne and Rubinstein.

[2] Note that technically we only need to consider rows that correspond to strategies that player i actually might play according to δ.

	a	b
A	$(9,9)$	$(6,10)$
B	$(10,6)$	$(0,0)$

Figure 1: The normal form of the game in Example 2.

We begin by construction the CE confirmation matrix as I described in the previous section. Without loss of generality, we consider 1's possible strategies, $s_i \in \{A, B\}$. We need to then calculate the four entries of the matrix; that is, the four values of $E_{S_2}[g(s_1', s_2)|s_1, \delta]$ for each pair (s_1', s_1):

- $(s_1', s_1) = (A, A)$: $E_{S_2}[g(A, s_2)|A, \delta] = \frac{2\gamma}{1+\gamma}9 + \frac{1-\gamma}{1+\gamma}6 = \frac{12\gamma+6}{1+\gamma}$.
- $(s_1', s_1) = (B, A)$: $E_{S_2}[g(B, s_2)|A, \delta] = \frac{2\gamma}{1+\gamma}10 + \frac{1-\gamma}{1+\gamma}0 = \frac{20\gamma}{1+\gamma}$.
- $(s_1', s_1) = (A, B)$: $E_{S_2}[g(A, s_2)|B, \delta] = 9$.
- $(s_1', s_1) = (B, B)$: $E_{S_2}[g(B, s_2)|B, \delta] = 10$.

To calculate the first entry, for example, we noted that 1 plays A with probability $\gamma + \frac{1-\gamma}{2} = \frac{1+\gamma}{2}$. Thus, *given 1 is playing* A, the outcome is (A, a) with probability $\frac{\gamma}{\frac{1+\gamma}{2}} = \frac{2\gamma}{1+\gamma}$, and the outcome is (A, b) with probability $\frac{1-\gamma}{2\frac{1+\gamma}{2}} = \frac{1-\gamma}{1+\gamma}$. These calculations allow us to fill in the matrix as in Figure 2. Clearly the maximum

	A	B
A	$\frac{12\gamma+6}{1+\gamma}$	$\frac{20\gamma}{1+\gamma}$
B	9	10

Figure 2: The CE confirmation matrix for player 1 in the game in Example 2.

in the second row is along the diagonal. The maximum in the first row will be on the diagonal as long as $\gamma \leq \frac{3}{4}$. The case for player 2 is symmetric, so we can conclude that the given distribution δ is a CE if $\gamma \leq \frac{3}{4}$.

2.3 Games of Incomplete Information: Bayesian Nash Equilibrium

Recall that a game is said to have **incomplete information** if some players do not know the payoffs of the other players with certainty. Analyzing such a game could in theory be very difficult, as we would have to reason about each player's beliefs about the others' payoffs, and the others' beliefs about those beliefs, and so on.

However, Harsanyi recognized that we already have the machinery to solve this problem, however, if we interpret this uncertainty about *payoffs* as uncertainty about *histories*. That is, we can interpret any game of incomplete information as one of imperfect information, in which nature selects some "types" for the players, and players observe their own types but not the types of others. We already have the tools to analyze such games of imperfect information.

Formally, we have a **Bayesian game** (S, g, Θ, F), where the components are defined as follows:

- $S = \{S_i\}_{i=1}^I$ is the strategy profile set.
- Each player has a **type** $\theta_i \in \Theta_i$. Let $\Theta = \Theta_1 \times \ldots \times \Theta_I$.
- The types are jointly distributed according to CDF $F(\theta)$.
- Payoffs for player i depend on i's type: $g = (g_1(s, \theta_i), \ldots, g_I(s, \theta_i))$.

We can then define an **extended game** (Σ, u) as follows:

- A pure strategy for player i is a function $\sigma_i : \theta_i \to s_i$ for $\theta_i \in \Theta_i$ and $s_i \in S_i$. Let $\sigma = (\sigma_1, \ldots, \sigma_I) \in \Sigma$. So Σ_i, i's strategy set, is all possible mappings from types to strategies.
- Let $u(\sigma) = (u_1(\sigma), \ldots, u_I(\sigma))$, where $u_i = E_\theta[g_i(\sigma_1(\theta_1), \ldots, \sigma_I(\theta_I), \theta_i)]$ is i's expected payoff given strategy profile σ.

Definition 2. A pure strategy **Bayesian Nash equilibrium (BNE)** of the Bayesian game (S, g, Θ, F) is a pure strategy Nash equilibrium of the extended game (Σ, u). That is, a BNE is a strategy profile σ^* such that $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*)$ for all $\sigma_i \in \Sigma_i$ and all $i \in I$.

Formally, we are modeling the decision process as follows: The agents pick a strategy that determines what they will do for each possible type they could draw. Then types are drawn by nature and the corresponding strategies executed.

However, there is another way to think about such games: Each type of a particular player is in fact a separate player. All of these players pick their strategy, and then nature randomly chooses which players will actually get to participate. Those players' strategies are then implemented. This leads to an alternative definition of a BNE:

Definition 3. A collection of decision rules $\sigma = (\sigma_1, \ldots, \sigma_I)$ is pure strategy **Bayesian Nash equilibrium** of the Bayesian game (S, g, Θ, F) iff for all $i \in I$ and for all $\theta_i \in \Theta_i$ occurring with positive probability, we have

$$\sigma(\theta_i) \in \arg\max_{s_i \in S_i} E_{\theta_{-i}}[g_i(s_i, \sigma_{-i}(\theta_{-i}), \theta_i) | \theta_i].$$

That is, given our second interpretation, each player who has a positive chance of entering the game is choosing her best (pure strategy) response.

2.3.1 Finding Bayesian Nash Equilibria

The first definition of a BNE suggests an exceedingly simple way to deal with games of incomplete information:

1. Write down the extended game in normal for if possible. Drawing the extensive form is usually helpful here.
2. Find the NE of this extended game using the methods of the previous section of this course. That's it!

This method is the best way to go for games with a small number of strategies and types, where the extensive and normal forms are easy to write down.

The second definition suggests another way to find BNE, which is preferred if we are dealing with continuous strategies and/or types.

1. Propose a BNE set of decision rules $(\sigma_1(\theta_1), \ldots, \sigma_I(\theta_I))$.
2. Write down a general formula for $E_{\theta_{-i}}[g_i(s_i, \sigma_{-i}(\theta_{-i}), \theta_i)|\theta_i]$; that is, write down an expression for i's expected payoffs from some pure strategy s_i given the other players are playing according to σ.
3. Maximize this expression to find player i's best response function, and confirm that it is the proposed decision rule.
4. Repeat for all players if necessary.

Example 3. Let's see both of these approaches at work in an example from the notes. Consider the game of incomplete information *Friend or Foe*, shown in Figure 3.

	h	t
H	$(3,1)$	$(0,0)$
T	$(2,1)$	$(1,0)$

Friend

	h	t
H	$(3,0)$	$(0,1)$
T	$(2,0)$	$(1,1)$

Foe

Figure 3: The normal form of the game in Example 3.

Player 1's type space just a singleton, and 2's type is drawn from $\Theta_2 = \{\text{Friend}, \text{Foe}\}$. Nature selects $\theta_2 = \text{Friend}$ with probability p. We can draw the extensive form of the extended game, which then allows us to write down the normal form as in Figure 4.[3] We can then easily verify that (H, ht) is a NE if $p \geq \frac{1}{2}$, and (T, ht) is a NE if $p \leq \frac{1}{2}$.

[3] It is important to note that player 2 has two information sets: one after seeing that he is of type Friend, and the other after seeing that he is type Foe. He does *not* observe the strategy of player 1.

	hh	ht	th	tt
H	$(3, p)$	$(3p, 1)$	$(3-3p, 0)$	$(0, 1-p)$
T	$(2, p)$	$(1+p, 1)$	$(2-p, 0)$	$(1, 1-p)$

Figure 4: The normal form of the extended game in Example 3.

Applying the second paradigm is a bit redundant, but it is important to see alongside the other method. For this approach, I propose that the following decision rules constitute at BNE if $p > \frac{1}{2}$:

$$\sigma_1 = H$$
$$\sigma_2(\theta_2) = \begin{cases} h & \text{if } \theta_2 = \text{Friend} \\ t & \text{if } \theta_2 = \text{Foe} \end{cases}$$

Let's check that this is in fact a BNE. First, check player 1's response. Note that

$$E[g_1(s_1, \sigma_2)] = \begin{cases} 3p + 0(1-p) = 3p & \text{if } s_1 = H \\ 2p + 1(1-p) = 1 + p & \text{if } s_1 = T. \end{cases}$$

We see that H is a best-response if $p > \frac{1}{2}$, as desired. Since player 1 has only one type, we are done with her. Next we check player 2's responses. If $\theta_2 =$ Friend, then

$$E[g_2(s_2, \sigma_1, \theta_2)|\theta_2] = \begin{cases} 1 & \text{if } s_2 = h \\ 0 & \text{if } s_2 = t. \end{cases}$$

So, h is a best response for this type for any p. Similarly,

$$E[g_2(s_2, \sigma_1, \theta_2)|\theta_2] = \begin{cases} 0 & \text{if } s_2 = h \\ 1 & \text{if } s_2 = t. \end{cases}$$

when $\theta_2 =$ Foe, so this type is best-responding as well. This covers all types for player 2, so we are done with him. Since all types of all players are best-responding, we have a BNE. The cases for $p < \frac{1}{2}$ and $p = \frac{1}{2}$ can be done similarly.

Of course, I am being a bit pedantic with this example, but the point is that the approach is the same when the game is made much more complex: propose a set of decision rules, and then confirm that each type of each player is best-responding.

3 Problems

Problem 1. *Asymmetric all-pay auction (Problem Bank #23).*

Consider an asymmetric all-pay auction between two bidders whose values are common knowledge. The object is worth v to Bidder 1 and v' to Bidder 2 where $v' > v$. All bidders simultaneously submit sealed bids b_i which must be non-negative; the highest bidder wins the object and all bidders pay their bids. In the event of a tie, the good is given to Bidder 2.

(a) What bidding strategies are strictly dominated? Can you iteratively eliminate any more strictly dominated strategies?

(b) Are there any pure-strategy Nash equilibria? If so, characterize the set of PSNE. If not, explain why not.

We now consider a mixed-strategy Nash equilibrium where both bidders randomize.

(c) At a MSNE, would Bidder 1 ever bid strictly higher than v with positive probability? Would Bidder 2 ever bid strictly higher than v with positive probability?

(d) Consider a MSNE where Bidder i randomly selects a bid based on the CDF F_i. You can assume (without proof) that the support of F_i is $[0, v]$.

 (i) Write an expression for 2's payoff from bidding x in $[0, v]$, fixing F_1.

 (ii) Use the fact that $x = v$ is in the support of Bidder 2's strategy to solve for 2's payoff.

 (iii) Use answers from the last two parts to solve for $F_1(x)$. What is $F_1(0)$?

 (iv) Write an expression for 1's payoff from bidding x in $[0, v]$, fixing F_2. Use the fact that $x = v$ is in the support of Bidder 1's strategy to solve for 1's payoff.

 (v) Use answers from the previous part to solve for $F_2(x)$. What is $F_2(0)$?

(e) What is the expected revenue? As v' varies in the set (v, ∞), is expected revenue increasing, constant, or decreasing in v'?

Problem 2. *A three-player CE. (Based on Osbourne and Rubinstein 48.1)*

Consider the 3-player game in Figure 5. Player 1 chooses rows, 2 chooses columns, and 3 chooses matrices.

(a) What are all the possible PSNE payoffs in this game?

(b) Show that there is a correlated equilibrium where 3 chooses B for sure, and 1 and 2 play (T, L) and (B, R) with equal probabilities.

Problem 3. *Cournot with uncertain costs. (Problem Bank # 32)*

	L	R
T	$(0,0,3)$	$(0,0,0)$
B	$(1,0,0)$	$(0,0,0)$

A

	L	R
T	$(2,2,2)$	$(0,0,0)$
B	$(0,0,0)$	$(2,2,2)$

B

	L	R
T	$(0,0,0)$	$(0,0,0)$
B	$(0,1,0)$	$(0,0,3)$

C

Figure 5: The normal form of the game in Problem 1.

Consider the linear Cournot model. Now, however, suppose that each firm has probability μ of having unit costs of c and $(1-\mu)$ of having costs d, where $d > c$. Solve for the Bayesian Nash equilibrium.

4 Solutions

Solution 1.

(a) For player i, bidding $b_i > v_i$ gives a negative payoff for sure, but bidding $b_i = 0$ guarantees a non-negative payoff. Thus $b_1 > v$ and $b_2 > v'$ are dominated.

Given that 1 will not bid more than v, $b_2 > v$ is dominated by $b_2 = v$. (Note the tiebreaking rule here.) Thus we eliminate $b_2 > v$ with one round of iterative deletion.

Now, given that both bidders are bidding in $[0, v]$, can we do any further rounds of deletion? It turns out no. Given $b_1 \in [0, v]$, 2's best response is to set $b_2 = b_1$. Thus all of 2's strategies are best responses for some strategy of 1, and can't be deleted.

Showing that none of 1's remaining strategies is dominated is more difficult. Suppose, by way of contradiction, that some $\widehat{b}_1 \in [0, v]$ is dominated by a pure strategy. That is, there exists some $\tilde{b}_1$ such that $g_1(\tilde{b}_1, b_2) > g_1(\widehat{b}_1, b_2)$ for all $b_2 \in [0, v]$. But the consider two cases:

- Consider $\tilde{b}_1 < \widehat{b}_1$. Suppose that $b_2 \in (\tilde{b}_1, \widehat{b}_1)$. Then $g_1(\tilde{b}_1, b_2) = -\tilde{b}_1 \leq 0 \leq v - \widehat{b}_1 = g_1(\widehat{b}_1, b_2)$, a contradiction.
- Consider $\tilde{b}_1 > \widehat{b}_1$. Suppose that $b_2 > \tilde{b}_1$. Then $g_1(\tilde{b}_1, b_2) = -\tilde{b}_1 < -\widehat{b}_1 = g_1(\widehat{b}_1, b_2)$, a contradiction.

So, we can't have $b_1 \in [0, v]$ dominated by a pure strategy. One can show by similar methods and a lot more algebra that the same is true for a mixed strategy. Thus no $b_1 \in [0, v]$ is dominated.

(b) There are no PSNE. We proceed by cases:

- Suppose $0 < b_i < b_j$. Then j could lower his bid slightly but still win the auction, thus improving payoffs.
- Suppose $0 < b_1 = b_2$. In this case, 1 will lose the auction, giving him a negative payoff. He can do better by bidding 0 for a payoff of 0.
- Suppose $0 = b_1 = b_2$. Then player 1 can deviate to a bid greater than 0 to win the object and guarantee positive payoffs.

(c) We know that any strategy that is eliminated by IDDS will not have positive weight in a MSNE. Thus neither bidder will bid above v with positive probabillty.

(d) (i) If 2 bids $b_2 = x$, she wins the object with probability $F_1(x)$ for a payoff of $v' - x$, and loses with probabillty $1 - F_1(x)$ for a payoff of $-x$. Thus 2's expected payoff from bidding $b_2 = x$ is $(v' - x)F_1(x) + (-x)(1 - F_1(x)) = v'F_1(x) - x$.

(ii) If 2 bids v, she wins the object for sure and gets payoff $v' - v$. This must be her payoff in the MSNE, since she must get the same payoff from each pure strategy in her support.

(ii) We know 2's payoffs are $v'F_1(x) - x = v' - v$. Solving for $F_1(x)$, we find that
$$F_1(x) = \frac{v' - v + x}{v'}.$$
Note that $F_1(0) = \frac{v'-v}{v'}$.

(iv) If 1 bids x, he wins the object with probability $F_2(x)$ for a payoff of $v - x$, and loses with probabililty $1 - F_2(x)$ for a payoff of $-x$. Thus 1's expected payoff from bidding $b_2 = x$ is $(v - x)F_2(x) + (-x)(1 - F_2(x)) = vF_2(x) - x = C$. Since $F(v) = 1$, we find that $C = 0$. Alternatively, note that if 1 bids v he gets payoff 0 with probability 1.

(v) Thus we find that
$$F_2(x) = \frac{x}{v},$$
noting that $F_2(0) = 0$.

(e) The expected revenue from player 1 is $\int_0^v x\frac{1}{v'}dx = \frac{v^2}{2v'}$, and the expected revenue from player 2 is $\int_0^v x\frac{1}{v}dx = \frac{v}{2}$. Thus total expected revenue is $\frac{v^2}{2v'} + \frac{v}{2}$, which is decreasing in v'.

Note. It is important to see in the above problem that player 2 has an atom in her distribution, at $b_1 = 0$. However, we didn't have to consider this explicitly in part (d) because their was no chance of the bidders exactly tying.

Solution 2.

(a) One can verify that the pure strategy Nash equilibria are (T, R, A), (B, L, A), (T, R, C), and (B, L, C). The possible payoffs are $(0, 0, 0)$, $(1, 0, 0)$, and $(0, 1, 0)$.

(b) First, we need to figure out what distribution δ will imply the choices described in the problem. Since 3 plays B for sure, all the probability mass must be in the middle matrix. Then we know that $P((T, L, B)) = P((B, R, B)) = \frac{1}{2}$, and the probability of all other outcomes is 0.

Next, let's verify that the appropriate conditions are satisfied for player 1. Note that if $P(s_2 = L, s_3 = B|s_1 = T, \delta) = P(s_2 = R, s_3 = B|s_1 = B, \delta) = 1$. Thus we find the the following:

$$E[g_1(s_1', s_{-1})|s_1, \delta] = \begin{cases} 2 & \text{if } s_1' = T \text{ and } s_1 = T \\ 0 & \text{if } s_1' = B \text{ and } s_1 = T \\ 0 & \text{if } s_1' = T \text{ and } s_1 = B \\ 2 & \text{if } s_1' = B \text{ and } s_1 = B \end{cases}$$

Thus the CE confirmation matrix is

	T	B
T	2	0
B	0	2

so clearly player 1's conditions are satisfied.

The equations for player 2 are identical, so it just remains to check player 3. Note that since δ implies that only B is played by 3, we only need to check deviations from $s_3 = B$. But note that

$$E[g_3(s_3', s_{-3})|s_3, \delta] = \begin{cases} 1.5 & \text{if } s_3' = A \text{ and } s_3 = B \\ 2 & \text{if } s_3' = B \text{ and } s_3 = B \\ 1.5 & \text{if } s_3' = C \text{ and } s_3 = B, \end{cases}$$

so clearly player 3's conditions check out as well. Thus the proposed profile δ is a correlated equilibrium.

Solution 3.

Note that each player has just two types, parameterized by the possible costs c and d. Thus I propose decision rules of the following form:

$$\sigma_i(\theta_i) = \begin{cases} q_i^c & \text{if } \theta_i = c \\ q_i^d & \text{if } \theta_i = d \end{cases}$$

for $i \in \{1, 2\}$. For this to be a BNE, we need both types of both players to be best-responding. Consider first type $\theta_1 = c$, and let us calculate his expected payoff given 2 is playing according to the proposed decision rule σ_2:

$$E_{\theta_2}[g_1(q_1^c, \sigma_2(\theta_2), \theta_1)|\theta_1 = c] = \mu\left[a - b(q_1^c + q_2^c) - c\right] q_1^c + (1-\mu)\left[a - b(q_1^c + q_2^d) - c\right] q_1^c.$$

We want to find the q_1^c that maximizes this payoff, so we take the FOC and solve for q_1^c, which gets us

$$q_1^c = \frac{a - c - \mu b q_2^c + (1-\mu) q_2^d}{2b}.$$

By a similar calculation, we find that the optimal response for 1 when $\theta_1 = d$ is

$$q_1^d = \frac{a - d - \mu b q_2^c + (1-\mu) q_2^d}{2b}.$$

So far we have been assuming that we know q_2^c and q_2^d. Yet note that by the symmetry of the problem, we must have $q_1^c = q_2^c = q^c$ and $q_1^d = q_2^d = q^d$. This, plus the two equations above, allows us to solve for the optimal responses q^c and q^d. With a bit of algebra, we can find that

$$q^c = \frac{a - \frac{1-\mu}{2}(c-d) - c}{3b}, \text{ and}$$
$$q^d = \frac{a - \frac{\mu}{2}(d-c) - d}{3b}.$$