Binding Agreements I: The Bargaining Framework

Debraj Ray, October 2006

1. Basics

The following model of bargaining, based on Rubinstein (1982) and Ståhl (1977) represents possibly one of the simplest examples of an infinite game which isn't repeated. It *looks* "repeated", as you will see in a minute, but it is not formally a repeated game. It is also a beautiful example of the bite created by subgame perfection. In addition, Rubinstein-Ståhl bargaining is of interest in itself and has found many applications in the literature.

Suppose there are two persons, call them 1 and 2. They are dividing a cake of size 1. They take turns in proposing divisions of the cake; at each round, person i proposes a division and person j must accept or reject. If there is an acceptance, the game ends and the proposed division is implemented. If there is a rejection, we move on to the next round, and proposer and responder switch roles. If a period passes, the next period is discounted. The discount factor of player i is $\delta_i \in (0,1)$. Thus, if a division $(x, 1-x)$ is settled on at date t, the two payoffs are $\delta_1^t x$ and $\delta_2^t(1-x)$, and if no division is ever settled on at all, then the payoffs are zero.

This is a special case of the infinite extensive form already studied, and the definitions of Nash and subgame perfect equilibrium apply without any change.

2. Nash, Without Subgame Perfection

There are lots of Nash equilibria of this game. Specifically, fix *any* division of the cake; call it $(x_1, x_2) = (y, 1-y)$. Now think of the following strategies. Each proposer proposes this division initially and following any history. Each responder i uses the response rule following any history and proposal: yes if and only if she is given at least x_i. It is easy to check that this strategy profile is Nash.

However, such an equilibrium is not subgame perfect. Under the strategies described above, nodes in which player i is offered a bit less than x_i will never be visited. But subgame perfection requires that these nodes be checked for equilibrium behavior.

In particular, if person i feels that she will get x_i tomorrow she should certainly be willing to accept $\delta_i x_i + \epsilon$ today. But for small ϵ this number is in fact smaller than x_i, and the going strategy is telling her to refuse such an offer. So the equilibrium isn't subgame perfect.

3. Perfect Equilibrium

A remarkable property of this two-person bargaining model is that subgame perfection wipes out all but one of these multiple Nash equilibria:

Theorem **1.** *There is a unique subgame perfect equilibrium payoff vector in the two-person bargaining model.*

Proof. Let M_i be the *maximum* equilibrium payoff and m_i be the *minimum* equilibrium payoff to person i. Then it is obvious that i (as proposer) can always get at least $1 - \delta_j M_j$ in equilibrium. This proves that

$$m_i \geq 1 - \delta_j M_j. \tag{1}$$

Now examine M_j. Suppose that j wants to try and clinch an agreement today. She cannot get more than $1 - \delta_i m_i$. On the other hand, if j makes an unacceptable offer, the max she can get from tomorrow (discounted to today) is $\delta_j M_j$. It follows that

$$M_j \leq \max\{1 - \delta_i m_i, \delta_j M_j\}.$$

It is easy to see that $M_j > 0$ (why?). Therefore, the second term on the RHS above cannot be the one that attains the max. Consequently,

$$M_j \leq 1 - \delta_i m_i. \tag{2}$$

Combining (1) and (2), it is easy to see that

$$m_i \geq 1 - \delta_j M_j \geq 1 - \delta_j(1 - \delta_i m_i),$$

or

$$m_i \geq \frac{1 - \delta_j}{1 - \delta_i \delta_j}. \tag{3}$$

Now combining (1) and (2) in a slightly different way and using the same logic,

$$M_j \leq 1 - \delta_i m_i \leq 1 - \delta_i(1 - \delta_j M_j),$$

so that

$$M_j \leq \frac{1 - \delta_i}{1 - \delta_i \delta_j}.$$

Flipping the indices i and j,

$$M_i \leq \frac{1 - \delta_j}{1 - \delta_i \delta_j}. \tag{4}$$

Combining (3) and (4), we conclude that

$$m_i = M_i = \frac{1 - \delta_j}{1 - \delta_i \delta_j}. \tag{5}$$

□

We can now unpack this to figure out supporting strategies. Player 1 as proposer will always propose the division $(x^*, 1 - x^*)$, where

$$x^* = \frac{1 - \delta_2}{1 - \delta_1 \delta_2},$$

and accept any proposal that gives her at least $\delta_1 x^*$. Likewise, Player 1 as proposer will always propose the division $(1 - y^*, y^*)$, where

$$y^* = \frac{1 - \delta_1}{1 - \delta_1 \delta_2},$$

and accept any proposal that gives her at least $\delta_2 y^*$. If the two are equally patient with common discount factor δ, then the proposer picks up $1/(1 + \delta)$ and the responder picks up

$\delta/(1+\delta)$. The difference arises from the first-mover advantage that the proposer has but in any case will wash out as the discount factor converges to one.

4. More Than Two Players

Additional issues arise in Rubinstein bargaining when there are three players or more. Of these, the most basic concerns issues of *protocol*: who proposes, who responds, etc. Two standard examples:

1. First rejector of going proposal proposes next.

2. New proposer drawn at random

We can also create a general class of protocols that encompass these two.

In what follows we abstract from questions of heterogeneous patience; assume that everyone has the same discount factor δ. As before, there is a cake of unit size to divide.

4.1. Equilibrium in Stationary Strategies.

A *stationary strategy* (sometimes called a Markovian strategy) calls upon the agent to *take the same action* always when it is her turn to propose or respond, provided that the ambient situation is the same.[1] Neither history nor calendar time matters. [Of course, when we require such strategy profiles to be equilibria, they must be full-blown equilibria in the class of *all* strategies, stationary or not.

A proposal now is a vector $\mathbf{x} \in \mathbb{R}^n$ such that $\sum_i x_i \leq 1$.

Focus on responses. With stationary strategy profiles they must look like this: say yes if the amount given to you, person i, is above some threshold m_i, otherwise say no. Actually it is a bit more complicated than that because if the proposal gives you more (than your threshold) but gives some responder who responds after you less (than her threshold), you don't want to accept because you want to grab the initiative, at least in the rejector-proposes protocol. So the technically correct description of a response vector is a collection $(m_1^*, \ldots m_n^*)$ such that each responder i says yes if the proposal gives $x_j^* \geq m_j$ to every responder j who responds after i and says no if the proposal gives $x_j^* < m_j$ to some responder j who comes after i.[2]

Now go back to proposals made by i. He will simply try to get

$$z_i \equiv 1 - \sum_{j \neq i} m_j,$$

assuming this is nonnegative, and zero otherwise.

---

[1] In the present context, "the ambient situation is the same" simply means that she is responding to the same proposal.

[2] Strictly speaking, i is perfectly allowed to randomize if $x_i = m_i$ but we simply presume that he will accept then. It is easy to see that if he rejects with some probability in the indifference case then this cannot form part of any stationary equilibrium, because the corresponding proposer will not have a well-defined best response.

Variant 1. First rejector proposes. By rejecting i takes the initiative and gets z_i tomorrow, valued today at δz_i. Therefore

$$m_i = \delta z_i = \delta(1 - \sum_{j \neq i} m_j) \tag{6}$$

It is easy to check (do so!) that there is a unique solution to (6) with $m_i = m_j \equiv m^*$ for all i and j, so

$$m^* = \frac{\delta}{1 + (n-1)\delta},$$

and indeed, this is what each responder gets. The proposer picks up the remainder, which is easily seen to be

$$\frac{1}{1 + (n-1)\delta}.$$

All of this goes to equal division as $\delta \to 1$.

Variant 2. Random choice of proposers. If i rejects a proposal, two things can happen. One (with probability $1/n$), i gets the initiative and therefore $(1 - \sum_{j \neq i} m_j)$ tomorrow. Two (with probability $(n-1)/n$), i remains a responder and then gets m_i tomorrow. Discounting and then taking expected values, we must conclude that Therefore

$$m_i = \delta \left[\frac{1}{n}(1 - \sum_{j \neq i} m_j) + \frac{n-1}{n} m_i \right]. \tag{7}$$

Again, it is easy to see that (please verify) that there is a unique solution to (7) with $m_i = m_j \equiv \hat{m}$ for all i and j, so

$$\hat{m} = \frac{\delta}{n},$$

and this is what each responder gets. [The proposer picks up the remainder as in Variant 1.]

Notice that $\hat{m}$ in Variant 2 is smaller than m^* in Variant 1 (check this). This is as it should be, because in Variant 1 the rejector has "more power". Notice, however, that even in this case the outcome goes to equal division as the discount factor converges to one.

4.2. **Other Equilibria.** Unfortunately, the uniqueness result for two-person Rubinstein bargaining no longer survives with three or more players. The argument, due to Herrero and Shaked (see Herrero (1985)) is good practice and it is worth writing down. We do so for the rejector-proposes protocol.

Following the language of repeated games, let us try to write down a penal code. "Phase i" is described as follows:

Player i asks for the entire cake, giving zero to all other players.

Now "connect" these phases as follows. Notice that the description above can also apply in situations in which player i has deviated and is now asking for the whole cake. This *looks* like phase i but it isn't a "valid" phase, it is something that needs to be punished. So we will develop the notion of a "valid" and "invalid" phases recursively.

First the recursive strategy specification, then the initial condition:

[A] Suppose that we are currently in a valid phase i. Then all responders are required to accept the proposal.

[B] Otherwise, the phase is invalid. In this case, the responder k (among the remaining responders) is required to reject the proposal if and only of she gets no more than $1/(n-1)$, and start phase k.

[C] A phase is *valid* if it is a going reaction to an invalid phase as specified by [A] or [B], or if it is the phase specified at the very start of the game. Otherwise, it is *invalid*.

And the initial condition: begin with the specification that some (valid) phase i is started up.

This entire construction yields a penal code which we shall employ to support various outcomes as equilibria. Let us check conditions under which it is an equilibrium. We only need to check the unprofitability of one-shot deviations. Consider [A]. If a responder does not accept in a valid phase, then by [C], he starts an invalid phase and using [B], he will subsequently get 0. So it is optimal to accept in [A], given that everyone else follows the prescription.

Now for [B]. Consider a responder who gets more than $1/(n-1)$. Suppose she rejects. Then by [C], she starts an invalid phase and will surely get 0 if we follow the prescription thereafter. If the responder gets less than $1/(n-1)$, then by accepting she can get no more than what she is being currently offered.[3] By rejecting, she gets 1 after a lag discounted by δ. So she rejects whenever

$$\delta > \frac{1}{n-1}. \tag{8}$$

Note the crucial important point about this construction: it only works when $n \geq 3$, because the inequality (8) must hold.

This set of punishments is so strong (because you use it to give any deviator 0) that it can be used to support all sorts of equilibria, including those that are inefficient. The formalities of this are left to you.

[3] For if she accepts, the proposal will either be implemented or someone else is asked to reject, and she will get 0.