3 Alternating Offers Bargaining

Consider the following setup following Rubinstein's (1982) extension of a model by Stahl (1972):

- Two players bargain over a pie with size normalized to 1.
- Given a division (x_1, x_2) with $x_1 \geq 0, x_2 \geq 0$ and $x_1 + x_2 \leq 1$, the instantaneous utility for the players is x_1 and x_2 respectively. Hence, we assume that i) agents care only about their own slice of the pie, ii) risk neutrality. Payoffs are discounted geometrically at (common) rate δ, so, from the point of view of the beginning of the game, the utility from agreeing on (x_1, x_2) at time t is $\delta^t x_1$ and $\delta^t x_2$
- Game has n periods ($n = \infty$) will be considered.
- The extensive form is as follows. In every odd period $t = 1, 3, 5, ... < n$, if the players have still not agreed on a division, first player 1 makes an offer (x_1^t, x_2^t), which player 2 either accepts or rejects. If player 2 accepts, the game ends with (time t) payoffs (x_1^t, x_2^t) [in terms of time 0 utility, the payoffs are $\left(\delta^t x_1^t, \delta^t x_2^t\right)$]. If player 2 rejects, the game proceeds to the next period.
- In even periods $t = 2, 4, 6, ... < n$ everything is just like odd periods except that 2 makes the offer and 1 accepts or rejects.
- If $t = n$ the players receive some exogenous division (s_1, s_2) [in terms of time 0 utility, the payoffs are $(\delta^n s_1, \delta^n s_2)$].

It will be useful to translate payoffs into time t units. Informally:

Definition 1 *The continuation payoff of a strategy profile in a subgame starting at time t is the utility in time t units of the outcome induced by the strategy profile.*

Below, you will see examples where the continuation payoff is the same for all histories, but, in general, this is usually not the case.

3.1 The Backwards Induction Equilibrium with 3 Periods

The last non-terminal nodes in the game are the acceptance decisions by player 1 at time 2, where we see that:

1. Player 1 must accept if offered $x_1^2 > \delta s_1$
2. Player 1 is indifferent if offered $x_1^2 = \delta s_1$
3. Player 1 must reject if offered $x_1^2 < \delta s_1$.

However, we note that if player 1 would reject δs_1, then the best response problem for player 2 at time 2 would be ill-defined due to an openness issue. It follows:

Claim *In any backwards induction equilibrium it must be that player 1 accepts an offer x_1^2 at time 2 if and only if $x_1^2 \geq \delta s_1$.*

proofSuppose that 1 would reject δs_1 with probability $r \in (0,1)$. Then, the expected payoff for player 2 is

$$\begin{aligned}(1-r)(1-\delta s_1) + r\delta s_2 &\leq (1-r)(1-\delta s_1) + r\delta(1-s_1) \\ &= (1-\delta s_1) - r[1-\delta].\end{aligned}$$

If 2 offers $\delta s_1 + \varepsilon$, player 1 accepts for sure and the expected payoff for 2 is

$$1 - \delta s_1 - \varepsilon > (1-\delta s_1) - r[1-\delta]$$

whenever $\varepsilon < r[1-\delta]$. Since $r(1-\delta) > 0$ we may for example let $\varepsilon = \frac{r(1-\delta)}{2} > 0$. ■

Hence, player 2 will offer δs_1 and keep $1-\delta s_1$. It follows that:

Claim *In any backwards induction equilibrium it must be that player 2 accepts an offer x_2^1 at time 1 if and only if $x_2^1 \geq \delta[1-\delta s_1]$*

The proof is the same as above. Hence,

Claim *There is a unique backwards induction equilibrium in which the equilibrium outcome is that player 1 offers division*

$$\left(x_1^1, x_2^2\right) = \left(1 - \delta\left[1 - \delta s_1\right], \delta\left[1 - \delta s_1\right]\right)$$

in the first period and player 2 accepts.

3.2 A Stationary Equilibrium in the Infinite Horizon Model

Suppose that s_1 solves

$$s_1 = 1 - \delta\left[1 - \delta s_1\right]$$

or

$$s_1 = \frac{1-\delta}{1+\delta^2} = \frac{1}{1+\delta}.$$

Using on of the claims above we see that player 2 will (in the 3 period game) offer division

$$\left(\delta s_1, 1 - \delta s_1\right) = \left(\frac{\delta}{1+\delta}, \frac{1}{1+\delta}\right)$$

in period 2 and that player 1 will offer

$$\left(s_1, 1 - s_1\right) = \left(\frac{1}{1+\delta}, \frac{\delta}{1+\delta}\right)$$

in period 1. The critical observation is that the equilibrium division/payoffs is identical with the "exogenous" third period payoffs. Hence, we can add two, four, six or any even number of period and use the same argument recursively to conclude that the unique backwards induction equilibrium with (ad hoc) last period payoffs (odd number of periods) given by $\left(\frac{1}{1+\delta}, \frac{\delta}{1+\delta}\right)$ will have player 1 proposing $\left(\frac{1}{1+\delta}, \frac{\delta}{1+\delta}\right)$ in every odd period and player 2 proposing $\left(\frac{\delta}{1+\delta}, \frac{1}{1+\delta}\right)$ in every even period and each player accepting if and only if they are offered at least $\frac{\delta}{1+\delta}$.

In the infinite horizon game, we cannot backwards induct starting at the final period, but we can easily use the recursive structure to construct a stationary equilibrium where,

1. In every odd period t after any sequence of rejected offers player 1 offers a division $\left(s_1^t, s_2^t\right) = \left(\frac{1}{1+\delta}, \frac{\delta}{1+\delta}\right)$ and player 2 accepts any offer such that $x_2^t \geq \frac{\delta}{1+\delta}$.

2. In every even period t after any sequence of rejected offers player 2 offers a division $(s_1^t, s_2^t) = \left(\frac{\delta}{1+\delta}, \frac{1}{1+\delta}\right)$ and player 1 accepts any offer such that $x_2^t \geq \frac{\delta}{1+\delta}$.

To verify that this is an equilibrium, we note that the value of the game for player i is

$$V_i^t = \frac{1}{1+\delta}$$

in the beginning of every period t in which i makes a proposal (odd for 1 and even for 2). Hence, given these continuation payoffs, the unique acceptance rule is to accept a period t proposal if and only if

$$x_i^t \geq \frac{\delta}{1+\delta}.$$

Finally, given that in every period and after any history of play a proposal is accepted by the agent not making a proposal if and only if $x_i^t \geq \frac{\delta}{1+\delta}$, the optimal proposal in every period is to keep $\frac{1}{1+\delta}$ and give $\frac{\delta}{1+\delta}$ to the other agent. This verifies that the strategies specified are consistent with a "backwards induction" equilibrium [strictly speaking we would have to say that they are subgame perfect].

4 Subgame Perfection

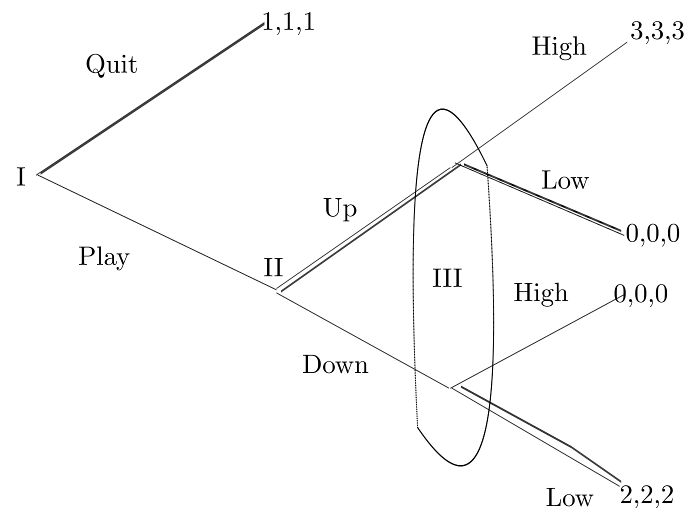

Figure 1: An example where backwards induction cannot be applie, but where the idea of sequential rationality generalizes

Consider Figure 1. Clearly, (Quit, Up, Low) is a Nash equilibrium as a deviation by 1 would give a payoff of 0 instead of 1. Deviations by players 2 and 3 don't change the outcome, so we conclude that the strategy profile indeed is a Nash equilibrium. Still, there is something implausible about this equilibrium in that if the game started at the node where player 2 decides between "Up" and "Down", then (Up, Low) would not be a Nash equilibrium.

That is, there is something "non-credible" about assuming that players 2 and 3 would play (Up, Low) following play. If the option "Quit" would be eliminated for player 1, then this play would not be consistent with Nash equilibrium in the reduced game. However, we cannot strictly speaking apply backwards induction as this game has imperfect information.

The generalization of backwards induction for games with imperfect information is called *subgame perfection*. Intuitively, the equilibrium concept simply requires play to be Nash equilibrium play in any part of the game that can be thought of as a game in itself.

Definition 2 *A subgame of an extensive form game K is a subset K′ of the game that could be analyzed as a stand alone game with the properties:*

1. *There exists a node t' such that $t' \prec t$ for all t in the subgame K' and an information set h_j such that $h_j = \{t'\}$.(root is singleton information set)*
2. *For every node t in the subgame K', if $t \in h_j$ and $\widetilde{t} \in h_j$, then $\widetilde{t}$ is also a node in the subgame (no broken information sets).*

In a finite game of perfect information, every node is the beginning of a subgame.

Definition 3 *A strategy profile s^* is called subgame perfect if it induces Nash equilibrium play in every subgame K' of K.*

4.1 A Slight Problem With Subgame Perfection

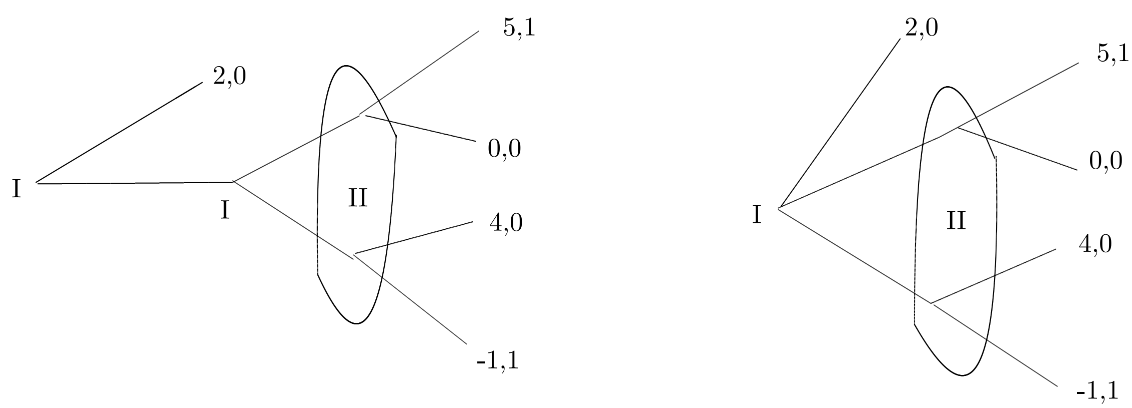

Figure 2: Equivalent Games with Different Predictions of Subgame Perfection

Consider the two extensive forms in Figure 2. We see that:

- Both games have the same reduced normal form

	L	R
U	$2,0$	$2,0^*$
T	$5,1^*$	$0,0$
B	$4,0$	$-1,1$

- The reduced normal form has two Nash equilibria (U,R) and (T,L)

- In the game to the left, we note that T strictly dominates B in the subgame. The unique Nash in the subgame is thus (T,L). Hence, the unique subgame prefect equilibrium is (AT,L) .

- In the game to the right, there is no other subgame than the full game. Hence, both (T,L) and (U,R) are subgame perfect.

- Disturbing example as we think that the (reduced) normal form and the extensive form are just different ways of modelling the same strategic situation.

4.2 Repeated Games

Example 1:

Consider the following Normal form game, where the Prisoner's dilemma has been extended with a "punishment strategy" P,

	C	D	P
C	$1,1$	$-1,\underline{2}$	$\underline{-2},-2$
D	$\underline{2},-1$	$\underline{0},\underline{0}$	$\underline{-2},-2$
P	$-2,\underline{-2}$	$-2,\underline{-2}$	$-3,-3$

The unique Nash equilibrium is (D,D).

Consider the case where this repeated twice. Consider the following strategy:

1. In period 1, play $s_i^1 = C$

2. In period 2, play

$$s_i^2 = \begin{cases} D & \text{if } (C,C) \text{ in period 1} \\ P & \text{if } (a_1^1, a_2^1) \neq (C,C) \text{ in period 1} \end{cases}$$

If both firms play in accordance to these strategies they each get

$$1 + 0 = 1$$

3. Suppose that a player deviates in the first period, then the best deviation is D and the payoff is at most

$$2 + (-2) = 0.$$

4. Also, no second period deviation is profitable as equilibrium play prescribes (C,C) in the second period, which is a static Nash equilibrium. Hence, we have specified a Nash equilibrium.

5. Not subgame prefect as the construction relies on play different from (D,D) in final period.

Example 2

Now consider

	C	D	A	B
C	$5, 5$	$-1, \underline{6}$	$-1, -1$	$-1, -1$
D	$\underline{6}, -1$	$\underline{0}, \underline{0}$	$-1, -1$	$-1, -1$
A	$-1, -1$	$-1, -1$	$\underline{2}, \underline{2}$	$0, 0$
B	$-1, -1$	$-1, -1$	$0, 0$	$\underline{1}, \underline{1}$

We see that (D,D) , (A,A) ,and (B,B) are all equilibria.

Consider the following strategy:

1. In period 1, play $s_i^1 = C$

2. In period 2, play A of $(s_1^1, s_2^1) = (C, C)$

3. In period 2, play D of $(s_1^1, s_2^1) \neq (C, C)$

If player i follows the candidate equilibrium strategy the payoff is

$$5 + 2 = 7$$

If player i wold deviate in the first period, the best deviation is D and the payoff is

$$6 + 0 = 6 < 7.$$

Since play in the second period is Nash after any history we conclude that we've constructed a subgame perfect equilibrium

Remark 1 *The key idea, which is used in many contexts, is that multiplicity of equilibria allows us to create "credible punishments". If players "behave" then play a good equilibrium in later stages. I players don't behave, then play a bad equilibrium in the later stages.*

5 Infinitely Repeated Games

Consider a stage game:

- $G = (I, A, u)$, where $A = \times_{i=1}^n A_i$ are the action spaces (strategy spaces in one-shot game) and $u_i : A \to R$. Also suppose that payoffs are discounted by $\delta \in (0, 1)$.
- Let $G(\delta, \infty)$ denote the infinitely repetition of G given discount factor δ.

We now note that:

1. Under the assumption that we have "perfect monitoring" so that all actions in previous rounds are observable we have that a history is the collection of all actions up to the last period. That is

$$H_t = A^t.$$

That is, $H_0 = \{\varnothing\}, H_1 = A, H_2 = A \times A, ...$

An *information set* is thus identified by the history leading up to that information set.

2. A *strategy* is a full contingent plan of action. One way to think of this is as thinking of a strategy as a sequence

$$\begin{aligned} s_i &= \left(s_i^1, s_i^2, ..., s_i^t, ...\right) \\ s_i^1 &\in A_i \\ s_i^t &: H_t \to A_i \end{aligned}$$

Sometimes it is useful/more elegant to note that we may write

$$\begin{aligned} H &= \cup_{t=1}^{\infty} H_t \\ s_i &: H \to A_i \end{aligned}$$

3. Given a strategy $s = (s_1, .., s_n)$ the outcome path is given by

$$\begin{aligned} a(s) &= \left(a^1, a^2, ..., a^t, ...\right) \\ a^1 &= \left(s_1^1, s_2^1, ..., s_n^1\right) \in A \\ a^2 &= \left(s_1^2\left(a^1\right), s_2^2\left(a^1\right), ..., s_n^2\left(a^1\right)\right) \\ &.... \\ a^t &= \left(s_1^t\left(a^1, ..., a^{t-1}\right), s_2^t\left(a^1, ..., a^{t-1}\right), ..., s_n^t\left(a^1, ..., a^{t-1}\right)\right) \\ &= \left(s_1^t(h_t), s_2^t(h_t), ..., s_n^t(h_t)\right) \end{aligned}$$

after introducing the convenient notation

$$h_t = \left(a^1, ..., a^{t-1}\right)$$

4. Payoffs

$$u_i(s) = \sum_{t=1}^{\infty} \delta^t u_i\left(a^t(s)\right)$$

A Nash equilibrium is defined in the obvious/standard way:

Definition 4 *s^* is a Nash equilibrium if $u_i(s^*) \geq u_i\left(s_i, s_{-i}^*\right)$ or every $i \in I$ and $s_i : H \to A_i$.*

For subgame perfection we have to introduce some notation to be able to say that starting after any sequence of realized actions all players are playing a Nash equilibrium.

- Given strategy s_i write $s_i|h_t$ to denote the continuation strategy following history h_t. Just like s_i we have that

$$s_i|h_t : H \to A_i.$$

- Stack strategies h_τ and h_t into $h_{\tau+t} = (h_\tau, h_t)$ and note that

$$s_i^t|h_\tau (h_t) = s_i^{t+\tau} (h_\tau, h_t)$$

Definition 5 *s^* is a subgame perfect equilibrium in $G(\delta, \infty)$ if for all histories h_t we have that*

$$s^*|h_t = (s_1^*|h_t, ..., s_n^*|h_t)$$

is a Nash equilibrium in subgame after history h_t.

5.1 Cooperation in an Infinitely Repeated Prisoner's Dilemma

Consider the Prisoner dilemma stage game

	D	C
D	$1,1$	$4,0$
C	$0,4$	$3,3$

5.1.1 GRIM-TRIGGER

In any finite repetition we know that (D, D) in every period is the unique (Nash) equilibrium outcome. However, the infinite repetition allows us to consider (for example) the following *GRIM-TRIGGER* strategies where each players starts by cooperating and cooperates until somebody defects. In case somebody has defected, then all players defect for ever. That is,

$$\begin{aligned} s_i^1 &= C \\ s_i^t &= \begin{cases} C & \text{if } h_t = (CC, CC, ..., CC) \\ D & \text{otherwise} \end{cases} \end{aligned}$$

Obviously, the outcome from both players playing *GRIM-TRIGGER* is CC in every period. Now, consider a deviation *on the equilibrium path.* If player i follows the recommended play, the payoff is

$$u_i(s) = \sum_{t=0}^{\infty} \delta^t 3 = \frac{3}{1-\delta},$$

whereas a deviation gives

$$u_i(s \backslash D) = 4 + \frac{\delta}{1-\delta}.$$

Hence,

$$\begin{aligned}
\frac{3}{1-\delta} &\geq 4 + \frac{\delta}{1-\delta} \Leftrightarrow \\
\frac{3-\delta}{1-\delta} &\geq 4 \Leftrightarrow 3-\delta \geq 4(1-\delta) \Leftrightarrow \\
3\delta &\geq 1 \Leftrightarrow \delta \geq \frac{1}{3}
\end{aligned}$$

Claim *s is a Nash equilibrium whenever $\delta \geq \frac{1}{3}$.*

We checked deviations on the equilibrium path above. However, off the equilibrium path play is (D, D) which is the unique stage game Nash. Hence,

Claim *s is a subgame perfect Nash equilibrium whenever $\delta \geq \frac{1}{3}$.*

5.1.2 Tit For Tat

Instead consider,

$$\begin{aligned}
s_i^1 &= C \\
s_i^t(h_t) &= \begin{cases} C & \text{if } a_j^{t-1} = C \\ D & \text{if } a_j^{t-1} = D \end{cases}
\end{aligned}$$

The outcome is (C, C) in every period, which gives a payoff of $u_i(s) = \frac{3}{1-\delta}$.

- Checking all possible deviations pretty tricky in this case.
- Will talk about one shot deviation principle next.

To see the need/use. Consider deviating once and then following the specified strategy. The play will then evolve according too

$$(D, C), (C, D), (D, C), ...$$

and the payoff from such a deviation is

$$\begin{aligned} u\left(s \backslash s_i^1 = D\right) &= 4 + \delta 0 + \delta^2 4 + \delta^3 0 + \delta^4 4... \\ u\left(s \backslash s_i^1 = D\right) &= \frac{1}{1-\delta^2} 4 \end{aligned}$$

Another possibility is to deviate once and then avoid punishing the opponent for the punishment in the next period in which case we get

$$(D, C), (C, D), (C, C), (C, C) ....$$

This deviation gives a payoff of

$$4 + \delta 0 + \delta^2 \frac{3}{1-\delta}.$$

We then see that

$$\begin{aligned} & 4 + \delta^2 \frac{3}{1-\delta} - \frac{1}{1-\delta^2} 4 \\ = & -4 \frac{\delta^2}{(1-\delta)(1+\delta)} + \delta^2 \frac{3}{1-\delta} \\ = & \frac{\delta^2}{1-\delta} \left[3 - \frac{4}{1+\delta} \right] = \frac{\delta^2}{1-\delta^2} [3\delta - 1] \end{aligned}$$

So:

- The second deviation is better than the first if $\delta \geq \frac{1}{3}$
- The equilibrium candidate is better than second deviation if

$$\begin{aligned} \frac{3}{1-\delta} &\geq 4 + \delta^2 \frac{3}{1-\delta} \\ 3 + \delta 3 + \delta^2 \frac{3}{1-\delta} &\geq 4 + \delta^2 \frac{3}{1-\delta} \\ \delta &\geq \frac{1}{3}. \end{aligned}$$

- We conclude that, under the assumption that $\delta \geq \frac{1}{3}$, then:

   1. $4 + \delta^2 \frac{3}{1-\delta}$ must be the best deviator payoff (since 0 always follows D). Hence, Tit-for-Tat is a Nash equilibrium given the restriction $\delta \geq \frac{1}{3}$.

   2. However, Tit-for-Tat is not subgame perfect. After (D, C) the strategy prescribes

   $$(C, D), (D, C), (D, C) ....$$

   and we concluded already that the strategy where 1 would play C (against C) after the punishment in the first period after the deviation (C, D) does better

One way to make the Tit-for-Tat subgame perfect is to amend it so that agents take part in the punishments of themselves. That is

$$\begin{aligned} s_i^1 &= C \\ s_i^t(h_t) &= \begin{cases} C & \text{if } \left(a_1^{t-1}, a_2^{t-1}\right) = (C, C) \\ D & \text{if } \left(a_1^{t-1}, a_2^{t-1}\right) \neq (C, C) \end{cases} . \end{aligned}$$

Here, any unilateral deviation prescribes (D, D) in the next period.