UNIVERSITY OF OULU | ABSTRACT OF THE MASTER'S THESIS
Oulu Business School

Unit
Oulu Business School

Author
Alaste Tomi

Supervisor
Svento R. Professor

Title
Models of Two-Person Bargaining

Subject	Type of the degree	Time of publication	Number of pages
Economics	Master’s thesis	May 2017	103

Abstract

The purpose of the thesis is to present some theoretical studies on bargaining situations. The thesis concentrates on existing research on the subject, and we do not present any new results. A central aim of the thesis is in describing how different factors present in real life bargaining situations can be taken into account in theoretical models of bargaining. It is our intention to use these models to obtain some information about effects of these factors in bargaining situations. In addition, another central element of the thesis is in describing relationships between different models. The studied models are kept deliberately elementary to provide clear analyses.

We use axiomatic and strategic models to study bargaining situations. Axiomatic method is very general as it does not deal with any particular bargaining process. Instead, these models concentrate solely on the outcome of a bargaining situation. A weakness with axiomatic models is that it might be difficult to evaluate how accurately the assumptions made describe some particular bargaining process. In the thesis, we concentrate mostly on Nash bargaining solution and its properties. Strategic models studied in the thesis include the game of alternating offers and several variations of this game. The main purpose of these variations is to describe what kind of effects different factors present in bargaining situations have on bargaining. The models studied deal with such factors as uncertainty, inside options, outside options, commitment tactics, and incomplete information. We determine the subgame perfect equilibria of models of complete information and sequential equilibria of models of incomplete information.

The most important results of the work are Nash bargaining solution and Rubinstein’s solution to an infinite horizon model of alternating offers. Nash bargaining solution has been extremely important for the development of theoretical research on bargaining, since it is the first systematic study of bargaining situations and has had a significant impact on all later research. This solution is also closely connected to strategic models, as certain limits of solutions of several strategic models are given by Nash solution. This interplay between strategic and axiomatic models is of considerable importance: these models provide information about applicability and limitations of different models and one methods helps to understand the other. Another key result in the thesis is deriving Rubinstein’s solution to an infinite horizon model of the game of alternating offers. This result is also of considerable importance, since this proof involves first determination of subgame perfect equilibrium of an infinite horizon model. Similar arguments have been applied to solve a large number of different models, some of which we also consider in the thesis.

The thesis concentrates on theoretical models of bargaining. We have chosen to deal with elementary models, with an intention to provide clear analysis of these models, instead of general or complex models. In particular, we have chosen the models such that each one of them illustrates some particular element in bargaining situations. This means that all the models treated in the thesis can be generalized, and many interesting models can be generated by combining elements from the models studied. We present some applications of bargaining situations in economics, but the studied models can be applied to a much larger variety of economics problems.

Keywords
Nash bargaining solution, sequential games, subgame perfect equilibrium

Additional information

UNIVERSITY of OULU
OULUN YLIOPISTO

OULU BUSINESS SCHOOL

Tomi Alaste

MODELS OF TWO-PERSON BARGAINING

Master's thesis

Economics

May 2017

CONTENTS

LIST OF FIGURES

1 INTRODUCTION

The central aim of this thesis is to present some results on the theory of bargaining and some applications of bargaining situations in economics. As stated in Nash (1950), the term *bargaining* refers to a situation where

(i) agents have the possibility of concluding a *mutual beneficial* agreement;

(ii) there is a *conflict of interest* about which agreement to conclude;

(iii) no agreement may be imposed on any individual without his approval.

So, a bargaining situation includes some number of (economic) agents and some outcomes that benefit all the parties involved, but the outcome is determined by conflicting interests of the parties. Many real life bargaining situations, such as wage or loan negotiations, involve deciding how certain amount of some good (money, say) is divided between the parties involved. A conflict naturally arises from the fact that all the agents wish to have as much as possible for themselves. Depending on the situation in question, the agents might have different options while negotiating: employees might go to strike to support their demand for a higher wage, or loan negotiations might break up when the applicant receives a better offer from some other party. The purpose of this thesis is to describe how such bargaining situations can be analyzed mathematically and how different options during negotiations can be taken into account in this analysis.

We apply two methods, strategic and axiomatic, to study bargaining situations. In strategic approach, some bargaining situation is studied by modeling it as some kind of game between the agents. When the model is built, one needs to consider how the actual negotiation proceeds, such as who makes an offer and when. After this, some game theoretical (equilibrium) concepts are applied to the resulting model to single out some outcomes as solutions. Axiomatic approach, on the other hand, does not deal with any particular bargaining process. In a strategic model the outcome is determined by players' equilibrium behavior, but in axiomatic approach one concentrates solely on the outcome: some properties of bargaining outcome are taken as granted, and then one studies whether a solution satisfying these properties exists. Choice of these properties is a delicate process: small number is desirable to keep the solution traceable, but too few might yield only vague information.

We do not seek to present most general models or results: all the models we treat can be extended. Instead, we have chosen the models with an intention to demonstrate what kind of effects different options during bargaining have on the outcome. We also consider only bargaining situations of two agents, and large part of the thesis is devoted to partitioning a cake of size one between the agents. We have chosen to concentrate on two–player bargaining situations as there is a sharp contrast between two–player and general bargaining situations. Ultimately, bargaining is about maximizing agent's outcome under the constraints created by the other bargainers. In the case of three or more negotiators, the agents might be able to form *coalitions* to support their personal interests. Full analysis of such situations requires considering questions like what kind of coalitions we expect to occur, what kind of payments are made inside coalitions, etc. These problems do not exist in bargaining between two players where, roughly, any improvement for one of the players comes at the expense of the other.

In the thesis, we concentrate on theoretical models of bargaining. We do not consider such questions as what is a fair agreement or how accurately these models describe real life bargaining situations. The reader may refer to Binmore (1994, 1998, 2007a) or Korth (2009) for such discussions, and Chatterjee (2014) for a discussion about aims, strengths, and weaknesses of bargaining theory in general.

The thesis is organized as follows. We begin in the next chapter by introducing some notation and recalling some concepts that we use throughout the thesis. We also provide a detailed discussion about properties of von Neumann–Morgenstern utility functions. We begin our study of bargaining situations in Chapter 3. This chapter concentrates on the Nash bargaining solution, originally given in Nash (1950). The assumptions behind the Nash solution can be questioned, but the importance of this solution cannot be overestimated since it is a first systematic study of bargaining and it has had a significant impact on later research on this subject. In Chapter 4, we concentrate on the game of Alternating Offers, where the players make offers and counteroffers each in turn. We describe the famous solution of this model, originally presented in Rubinstein (1982), and several variations of the model to illustrate how different factors during bargaining affect on the outcome. In the second to last chapter, we discuss briefly Folk Theorem and Coase Theorem, and in the last chapter we present some conclusion about the material presented in the thesis.

2 PRELIMINARIES

We gather briefly some notation and definitions that we use throughout the thesis.

Definition 2.1. Let $f : X \to \mathbb{R}$ be a function. A point $x^* \in X$ is a *maximizer* of f on X if and only if $f(x^*)$ is the largest value f attains on X, that is, $f(x) \leq f(x^*)$ for every $x \in X$. Define

$$\arg\max_{x \in X} f = \{x^* \in X : \ x^* \text{ is a maximizer of } f \text{ on } X\}.$$

In general, a function defined on some set X need not attain its maximum, and the above defined set might be empty. In all the cases when we use the above notation, the set X is a non–empty compact subset of $\mathbb{R}^2$ and f is a continuous function on X. In this case, f attains its maximum on X and the above defined set contains at least one element.

Definition 2.2. Given vectors $x = (x_1, \ldots, x_n)$ and $y = (y_1, \ldots, y_n)$ of $\mathbb{R}^n$, we denote

(i) $x \leq y$ if and only if $x_i \leq y_i$ for every $i = 1, \ldots, n$.

(ii) $x < y$ if and only if $x \leq y$ and $x \neq y$.

(iii) $x \ll y$ if and only if $x_i < y_i$ for every $i = 1, \ldots, n$.

Note that $x < y$ if and only if $x_i \leq y_i$ for every $i = 1, \ldots, n$ with strict inequality for at least one i. So, $x \ll y$ implies $x < y$, and $x < y$ implies $x \leq y$, but reverse implications need not hold.

Definition 2.3. A subset C of $\mathbb{R}^n$ is *convex* if and only if $x + t(y - x) \in C$ for all $x, y \in C$ and for every $t \in [0, 1]$.

Note that, as t attains all the values from the interval $[0, 1]$, the points $x + t(y - x)$ trace out the line segment (straight line) between x and y. In other words, a subset of $\mathbb{R}^n$ is convex if and only if this set contains the line segment between any two of its points.

Remark 2.4. A *convex combination* of some finite number of points $x_1, \ldots, x_m$ of $\mathbb{R}^n$ is a sum $\alpha_1 x_1 + \alpha_2 x_2 + \ldots + \alpha_m x_m$, where $\alpha_i \geq 0$ for every $i = 1, \ldots, m$ and

$\alpha_1+\alpha_2+\ldots+\alpha_m = 1$. Note here that $x+t(y-x) = (1-t)x+ty$ for all $x, y \in \mathbb{R}^n$ and for every $t \in [0, 1]$, and so the line segment between two points of $\mathbb{R}^n$ consists of convex combinations of these points. The above given definition for a convex set uses only a pair of points of C. However, a subset of $\mathbb{R}^n$ is convex if and only if this set contains all the convex combinations of its points.

Definition 2.5. A function $T : \mathbb{R}^n \to \mathbb{R}^n$ is *affine* if and only if there exist a linear function $L : \mathbb{R}^n \to \mathbb{R}^n$ and $v \in \mathbb{R}^n$ such that $T(x) = L(x) + v$ for every $x \in \mathbb{R}^n$.

Remark 2.6. In this thesis, we apply the following easily verified property of an affine transformation: an affine function preserves convex combinations. That is, if $T : \mathbb{R}^n \to \mathbb{R}^n$ is an affine transformation and $x = \alpha_1 x_1 + \alpha_2 x_2 + \ldots + \alpha_m x_m$ is a convex combination, then we have $T(x) = \alpha_1 T(x_1) + \alpha_2 T(x_2) + \ldots + \alpha_m T(x_m)$. In particular, if C is a convex subset of $\mathbb{R}^n$, then the set $T(C) = \{T(x) : x \in C\}$ is also convex.

2.1 Some game theory

A central theme of this thesis is to study bargaining situations between two players (bargainers). We often model a bargaining situation as some kind of game between the players, and then apply some equilibrium concepts to analyze these games. Since our approach relies on game theory, we recall some necessary concepts in this section. We consider only bargaining situations between two players which, with the lack of imagination, we call Player A and Player B. Therefore, we give the necessary definitions in this context and not in their full generality. Also, we hope that the verbal explanations given below are sufficient to follow the material presented in this thesis. The reader may consult, for example, Fudenberg and Tirole (2004), Maschler, Solan, and Zamir (2013), Osborne (2009), Owen (2008), Rasmusen (2008), or Tadelis (2013) for a more precise and general treatment of the concepts given below.

Definition 2.7. A *strategy* of a player in some game is a list of rules which describe an action of the player in every possible situation that might occur during the game. Let σ_A and σ_B be strategies of Player A and Player B, respectively. The strategy σ_A is a *best response* to the strategy σ_B if and only if Player A cannot increase his (or her) outcome from the game by choosing any other strategy, knowing that Player B will play according to σ_B. Strategy pair (σ_A, σ_B) is a *Nash equilibrium* if and only if σ_A is a best response to σ_B and *vice versa.*

Note that the concept of strategy requires that the players are prepared to act at any point of the game. If the players choose some strategies, it might be that the game ends after the first move. However, the concept of strategy requires that we have described players actions also at those points of the game which were never reached. In other words, a strategy describes what a player will do *if* the game reaches some point, but every point of the game need not be reached during a particular play.

A Nash equilibrium (σ_A, σ_B) is a pair of strategies such that the only way the players could (possibly) increase their outcome is if *both* of the players chose some other strategies. Indeed, as long as Player B plays according to σ_B, Player A has no reason to switch to any other strategy, since σ_A is a best response to σ_B. (Player A might get the *same* outcome with some other strategy, knowing that Player B plays according to σ_B. However, this other strategy might drive the players out from equilibrium, since σ_B need not be a best response anymore.) In other words, a Nash equilibrium is a pair of "no–regret strategies" in a sense that, having played these strategies, neither of the players has a reason to think that he (or she) should have chosen some other strategy. See, for example, McCain (2010: 63–88), Montet & Serra (2003: 63–80), or Osborne (2009: 13–52) for examples and discussion about Nash equilibrium.

When we apply the concept of Nash equilibrium, we assume that the players are *rational*. For one thing, this means that the players are trying to maximize their outcome from the game. In addition, we need to assume that rationality is *common knowledge*. This means that both of the players know that the other player is rational, both know that both know that both of the players are rational, and so on. Indeed, when Player A chooses his (or her) equilibrium strategy, he (or she) needs to know that Player B is also rational and that Player B will also play his (or her) best response. See Hargreaves Heap and Varoufakis (1995: 51–62) for a further discussion about rationality assumptions lying behind the concept of Nash equilibrium.

One should be somewhat careful in interpreting the concept of Nash equilibrium. In some relatively simple games, real players might reason the Nash equilibrium (should it exist) rather quickly and play the equilibrium strategies. However, in more complex games the players might need some practice to reach equilibrium. In the games treated in this thesis, we assume that the players are experienced enough to reach a Nash equilibrium.

Nash equilibrium is one of the most fundamental concepts in game theory. This concept also has its shortcomings, such as it does not always exist, or there can be many Nash equilibria and one cannot deduce without further information which equilibrium one should expect to occur. For the purposes of this thesis, a major shortcoming is that Nash equilibrium does not deal with threats presented by players. For example, the player second in turn might say "If you don't play the strategy I like, I will do something that causes you harm". We need a more delicate concept to analyze whether such threats are actually credible, and this is given as follows.

Definition 2.8. A *subgame* begins at some point of the original game, where both of the players have all the relevant information concerning the game so far, and continues thereafter just like the original game would continue from this point onwards. A pair of strategies (σ_A, σ_B) is a *subgame perfect equilibrium* if and only if this pair of strategies determines a Nash equilibrium in every subgame.

Note that the whole game is a subgame of itself, and so every subgame perfect equilibrium is also a Nash equilibrium. But every Nash equilibrium need not qualify as a subgame perfect equilibrium (we will give examples of this later), and so the concept of subgame perfect equilibrium is more restrictive than the concept of Nash equilibrium.

We need to be little cautious while defining the concept of subgame so that it also applies to games of imperfect information. In these games, at least one of the players has some additional information (concerning the game) compared to the other player at some point of the game. This means that one of the players might not know at which point of the game he (or she) is or will be later. Rather, the less–informed player faces some probabilities about different possibilities he (or she) might be at. A subgame does not begin at such a point of the game: it can only begin at a point where the actions taken so far do not cause any uncertainty on present or later choices. (There can be uncertainty later in the subgame, but not at the moment when the subgame begins.)

This concept is indeed efficient in analyzing different kinds of threats presented by the players. We will give proper applications of this concept later, but note that if Player B makes a threat to do something, and Player A makes a move to a point where Player B should execute his (or her) threat, the requirement of Nash equilibrium in the resulting subgame guarantees that Player B will execute his (or her) threat only in the case that this is part of his (or her) equilibrium strategy. Otherwise, the threat presented by Player B is not credible. (Executing the threat anyway would mean that Player B is not rational, and we do not consider such players in this thesis.)

There is a useful result concerning subgame perfect equilibrium which simplifies some of our reasonings. For a proof of the following statement, see Fudenberg and Tirole (2004: 108–110).

Theorem 2.9. *A strategy pair* (σ_A, σ_B) *is a subgame perfect equilibrium if and only if neither of the players benefits by deviating from his (or her) strategy in a single stage of the game.*

A player may deviate from a given strategy a number of times during the game. However, the previous statement asserts that if any number of deviations increases player's payoff, then there exists also a single deviation which benefits the player.

2.2 Preferences and utility functions

Uncertainty and time are important factors present in bargaining situations. Time is naturally involved in bargaining, since bargaining usually takes time and the time spent on bargaining has an opportunity cost. Uncertainty is also involved in bargaining situations as there is usually a possibility that bargaining ends for some reason before an agreement is reached. In this section, we treat in some detail player's preferences over uncertain outcomes. We show that, under certain conditions on player's preferences over uncertain outcomes, these preferences can be represented by utility function with some desirable properties, known as *von Neumann–Morgenstern utility function*. As usual, these utility functions depend on player's preferences, so we begin with some terminology related to preferences.

Definition 2.10. Let O be some set of outcomes. A *preference relation* of Player i over O is a relation $\succsim_i$ on O. We use the following notation with respect to a preference relation:

(i) $x \approx_i y$ means that Player i is indifferent between x and y.

(ii) $x \succsim_i y$ means that Player i prefers x to y or is indifferent between x and y.

(iii) $x \succ_i y$ means that Player i strictly prefers x to y.

If we consider only one player in some context, we denote a preference relation of this player simply by $\succsim$.

If not explicitly otherwise stated, we assume that every preference relation $\succsim$ treated in this thesis satisfies the following properties.

Completeness: Given $x, y \in O$, we have either $x \succsim y$ or $y \succsim x$.

Reflexiveness: For every $x \in O$, we have $x \succsim x$.

Transitiveness: If $x \succsim y$ and $y \succsim z$ for some $x, y, z \in O$, then $x \succsim z$.

Completeness means that a player can compare any two outcomes and choose the more attractive of these (or conclude that the outcomes are equally good). Transitiveness means that if a player prefers an outcome x to y and the outcome y to z, then we assume that he (or she) also prefers x to z. For example, if a player prefers Coke to Pepsi and Pepsi to Dr. Pepper, then we assume that he (or she) will choose Coke if Coke and Dr. Pepper are offered. Reflexiveness is perhaps so easy to accept that this assumption might even seem vain. Indeed, this assumption states that a player considers any outcome at least as good as the outcome itself. Why do we need to assume that this is the case? At this point, we consider preferences at a very general level, and one *can* construct preferences which satisfy completeness and transitiveness but not reflexiveness. However, our interpretation about preferences would then mean that a player is not capable of comparing some outcome to itself, and this seems to be rather difficult to accept. Preferences not satisfying reflexiveness are more like a mathematical curiosity, and they do not seem to correspond to any real world situations. Furthermore, as the reader might recall from basic microeconomic courses, the indifference curves of a consumer play an important role in consumer decision problems. Recall that an indifference curve consists of all those outcomes which the consumer finds equally attractive. Reflexivity assumption guarantees that, for every outcome, there exists at least one point on the indifference curve, namely the outcome itself.

It is rather tedious to analyze decision problems using only preference relations. A more useful tool is to use utility functions which capture player's preferences.

Definition 2.11. Let O be a set of outcomes and let $\succsim$ be a preference relation over O. A function $u : O \to \mathbb{R}$ is a *utility function representing* these preferences if and only if, for all outcomes x and y, we have

$$x \succsim y \qquad \text{if and only if} \qquad u(x) \geq u(y).$$

Note that a utility function attaches a real number to every outcome in such a way that player's preferences are captured by the order of real numbers. It is a simple consequence that, if u represents player's preferences, then the player is indifferent between x and y if and only if $u(x) = u(y)$.

It is not obvious that one can actually represent player's preferences using utility functions. Above we gave three conditions for preferences of a player, and it is actually true that if we wish to represent player's preferences using utility functions, these conditions must hold. That is, if any of these three conditions does not hold, then we cannot find a utility function to represent these preferences. For example, suppose that a utility function exists and take any two outcomes x and y. Since $u(x)$ and $u(y)$ are some real numbers, either $u(x) \geq u(y)$ or $u(y) \geq u(x)$ must hold. But if u represents player's preferences, we must have either $x \succsim y$ or $y \succsim x$, and so the given preference relation is complete. Similarly, reflexiveness and transitiveness of preferences follow from corresponding properties of real numbers.

The above given three conditions, however, are not sufficient to guarantee that player's preferences can be represented by utility function. An example is given by *lexicographical ordering*. Let $x = (x_1, x_2)$ and $y = (y_1, y_2)$ be consumption bundles and suppose that player's preferences are such that the player strictly prefers x to y whenever $x_1 > y_1$ or $x_1 = y_1$ and $x_2 > y_2$. It is straightforward to verify that this preference relation satisfies the above given three conditions. (Note that every indifference curve contains only one point.) However, this preference relation cannot be represented by a utility function; see Gravelle and Rees (2004: 41–42) for further details.

Next, we turn our attention to preferences over uncertain outcomes. We proceed to describe sufficient conditions under which player's preferences over uncertain outcomes can be represented by utility functions with some desirable properties. For these purposes, we need to consider lotteries. Presentation that follows is adapted from Maschler *et al.* (2013).

Definition 2.12. A *simple lottery* is a finite sequence

$$[p_1(A_1), p_2(A_2), \ldots, p_N(A_N)]$$

such that $p_1 + p_2 + \ldots + p_N = 1$ and $A_1, A_2, \ldots, A_N$ are some outcomes. A *compound lottery* is a finite sequence

$$[q_1(L_1), q_2(L_2), \ldots, q_M(L_M)],$$

such that $q_1 + q_2 + \ldots + q_M = 1$ and $L_1, L_2, \ldots, L_M$ are simple lotteries.

An interpretation of a simple lottery is that the outcome A_1 (monetary payoff, for example) occurs with probability p_1, the outcome A_2 occurs with probability p_2, and so on, and similar interpretation applies to compound lotteries. Note that the outcome of a compound lottery is some simple lottery L_m, and the outcome of this simple lottery is some (physical) outcome A_n. Note also that some of the probabilities above might be zero. In particular, we may regard every outcome A as a simple lottery $[1(A)]$ which yields the outcome A with probability 1 (that is, certainly), and so we may assume that every outcome is also a lottery. Similarly, we may regard every simple lottery as a compound lottery, with this particular simple lottery occurring with probability 1. (These identifications will be important later. We wish to construct a utility function on the set of compound lotteries. During this process, we like to think that this set also contains all the outcomes and simple lotteries over these outcomes.)

We proceed to show that, under certain conditions on player's preferences over lotteries, these preferences can be represented by certain utility functions. The property on utility function we are after is the following.

Definition 2.13. A utility function u on simple lotteries is *linear* if and only if

$$u(L) = p_1 u(A_1) + p_2 u(A_2) + \ldots + p_N u(A_N)$$

for every simple lottery $L = [p_1(A_1), p_2(A_2), \ldots, p_N(A_N)]$.

A linear utility function is also called *von Neumann–Morgenstern utility function.* Note that the term linear means that the utility function is linear with respect to probabilities and, under such a utility function, player's utility of a simple lottery is determined by the utilities obtained from different outcomes weighted by the probabilities of these outcomes.

Next, we discuss some sufficient conditions (axioms) on player's preferences over lotteries under which these preferences can be represented by a linear utility function. We use the same notation $\succsim$ both for the player's preferences over lotteries and outcomes. (Recall that we may regard every outcome as a lottery.) The reader may consult Binmore (2007b: 111–142), Luce and Raiffa (1957: 12–38), or Maschler *et al.* (2013: 9–38) for a further discussion on these assumptions and whether individuals actually behave according to these assumptions. (Perhaps not surprisingly, the answer is not always positive.)

Continuity: If A, B, and C are outcomes such that $A \succsim B \succsim C$, then there exists some $\theta \in [0,1]$ such that $B \approx [\theta(A), (1-\theta)(C)]$.

Note that the player prefers A to B and B to C. We may interpret this assumption as follows. If the player is indifferent between B and the lottery $[\theta(A), (1-\theta)(C)]$, the player is willing to accept this lottery instead of certain outcome B despite the fact that there is the probability $(1-\theta)$ that a worse outcome C appears. The number θ depends on how risk averse the player is. (A risk averse player requires a high probability on A in change for certain outcome B because of the possibility of the worse outcome C.)

Monotonicity: Let $\alpha, \beta \in [0,1]$ and let A and B be outcomes such that $A \succ B$. Then

$$[\alpha(A), (1-\alpha)(B)] \succsim [\beta(A), (1-\beta)(B)]$$

if and only if $\alpha \geq \beta$.

Note that, in this case, a player strictly prefers an outcome A to B. We assume that the player prefers the lottery where the probability for the more preferred outcome A is higher.

Remark 2.14. It would also seem reasonable that a player facing two lotteries with the same outcomes chooses the lottery with higher probability on the more preferred outcome. That is, if outcomes A and B satisfy $A \succ B$ and $\alpha, \beta \in [0,1]$ satisfy $\alpha > \beta$, then we would expect that the player strictly prefers the lottery $[\alpha(A), (1-\alpha)(B)]$ to the lottery $[\beta(A), (1-\beta)(B)]$. This statement is actually included in the previous assumption. To see this, suppose that $\alpha, \beta \in [0,1]$ satisfy $[\alpha(A), (1-\alpha)(B)] \approx [\beta(A), (1-\beta)(B)]$. Monotonicity Axiom implies $\alpha \geq \beta$, since $\succsim$ includes the possibility that the player is indifferent between the lotteries.

But we may change the order of these lotteries and the player is still indifferent between them. Therefore, we must also have $\beta \geq \alpha$, and so $\alpha = \beta$, thereby verifying our claim. (To clarify our train of thought, if $\alpha, \beta \in [0, 1]$ satisfy $\alpha > \beta$, we have $[\alpha(A), (1-\alpha)(B)] \succsim [\beta(A), (1-\beta)(B)]$ by Monotonicity Axiom. Our claim above implies that the player cannot be indifferent between these lotteries, for this would imply $\alpha = \beta$.)

The next assumption might first seem somewhat tedious, as we are dealing with outcomes, simple lotteries over these outcomes, and compound lotteries over these simple lotteries. We give an explanation of this assumption below.

Simplification: Let $A_1, A_2, \ldots, A_N$ be outcomes, let

$$L_m = [p_1^m(A_1), p_2^m(A_2), \ldots, p_N^m(A_N)]$$

be a simple lottery over these outcomes for every $m = 1, \ldots, M$, and let

$$\widehat{L} = [q_1(L_1), q_2(L_2), \ldots, q_M(L_M)]$$

be a compound lottery over these simple lotteries. Define

$$r_n = q_1 p_n^1 + q_2 p_n^2 + \ldots + q_M p_n^M$$

for every $n = 1, \ldots, N$. Then $\widehat{L} \approx [r_1(A_1), r_2(A_2), \ldots, r_N(A_N)]$.

To clarify the previous assumption, note that there are M simple lotteries over the outcomes $A_1, A_2, \ldots, A_N$, and so the compound lottery $\widehat{L}$ depends, ultimately, on these outcomes. The above assumption can be justified by assuming that it is only the probabilities of different outcomes which matter for the player, and not the actual form of the lottery in question. To explain this, consider some outcome A_n. Let us show that the probability of this outcome is the same in the compound lottery $\widehat{L}$ given above and in the lottery $[r_1(A_1), r_2(A_2), \ldots, r_N(A_N)]$. Of course, the probability of A_n in this simple lottery is r_n, so let us show that this is also the probability of A_n in the compound lottery $\widehat{L}$. First, $\widehat{L}$ yields the simple lottery L_1 with probability q_1, and L_1 yields the outcome A_n with probability p_n^1. Therefore, the probability that the outcome A_n appears from this combination is $q_1 p_n^1$. Similarly, the outcome A_n could appear by obtaining L_2 from the compound lottery (with probability q_2), and then obtaining A_n from L_2 (with probability p_n^2); the probability of this combination is $q_2 p_n^2$. Continuing this way, we see that the probability of the outcome A_n in the compound lottery is

$$q_1 p_n^1 + q_2 p_n^2 + \ldots + q_M p_n^M,$$

which is precisely r_n, as required.

Finally, we introduce the last of our assumptions.

Independence: Let $\widehat{L} = [q_1(L_1), q_2(L_2), \ldots, q_M(L_M)]$ be a compound lottery and let L be a simple lottery such that the player is indifferent between L and L_m for some $m = 1, \ldots, M$. Then

$$\widehat{L} \approx [q_1(L_1), \ldots, q_{m-1}(L_{m-1}), q_m(L), q_{m+1}(L_{m+1}), \ldots, q_M(L_M)].$$

We assume that we may replace one component of a compound lottery with something that the player considers equally attractive, and this change does not affect on player's valuation on the lottery. Just like the previous assumption, we may interpret Independence Axiom in a way that the player's preferences are determined only by the probabilities of different outcomes, and not by the actual form of the lottery.

Remark 2.15. In the previous assumption, we assume that we may replace one component of a compound lottery. Applying this assumption repeatedly, we see that we may replace any number of simple lotteries in a compound lottery and we obtain a lottery which the player considers precisely as good as the original lottery.

Now, let us show that these four axioms (with completeness, reflexiveness, and transitiveness) are sufficient for the existence of a linear utility function. The proof of the next theorem given below is adapted from Maschler *et al.* (2013), where the statement is proved in the case that there exists a finite number of outcomes. Note that, in this case, completeness and transitiveness of preferences imply that the player is capable of choosing the best (most preferred) outcome, the second best, and so on, and the assumption on the existence of best and worst outcomes given below becomes redundant.

Theorem 2.16. *Suppose that the set O of outcomes contains elements A_b (best) and A_w (worst) such that $A_b \succsim A \succsim A_w$ for every outcome A. Suppose also that the players preferences are determined over all the finite lotteries (both simple* *lotteries and compound lotteries), and that these preferences satisfy the above axioms. Then there exists a linear utility function on the set $\widehat{\mathcal{L}}$ of compound lotteries which represents these preferences.*

Proof. The proof is somewhat lengthy, so we divide it into a number of steps. Let us begin by dealing with an extreme case where the player finds all the outcomes equally good. In this case, the player also finds all the lotteries equally attractive. Indeed, take any simple lottery $L = [p_1(A_1), p_2(A_2), \ldots, p_N(A_N)]$. Independence Axiom[1] implies that $L \approx [p_1(A_w), p_2(A_w), \ldots, p_N(A_w)]$. Since $p_1 + \ldots + p_N = 1$, this lottery yields the certain outcome A_w, and the Axiom of Simplification implies that $L \approx [1(A_w)]$. Since every simple lottery satisfies this, another use of Axiom of Simplification implies that $\widehat{L} \approx [1(A_w)]$ for every compound lottery $\widehat{L}$. In conclusion, the player is indifferent between any two compound lotteries, and so any constant function will represent his (or her) preferences.

Next, let us deal with the more interesting case where the player finds some of the outcomes more attractive than the others. Since A_b is the most attractive outcome (or equivalent to it) and A_w is the least attractive outcome (or equivalent to it), we have $A_b \succ A_w$. Let us begin with a preliminary result for our proof.

Step 1: For every outcome A, there exists a unique number $\theta_A \in [0, 1]$ such that $A \approx [\theta_A(A_b), (1 - \theta_A)(A_w)]$.

To prove this statement, let A be any outcome. By our assumption on A_b and A_w, we have $A_b \succsim A \succsim A_w$. Such a number θ_A exists by the Continuity Axiom, so we need only to show that θ_A is unique. If $\alpha, \beta \in [0, 1]$ satisfy $A \approx [\alpha(A_b), (1-\alpha)(A_w)]$ and $A \approx [\beta(A_b), (1 - \beta)(A_w)]$, then the player is indifferent between these two lotteries. But then, $\alpha = \beta$ by Monotonicity Axiom (see Remark 2.14), as required. In particular, note that the relations $A_b \approx [1(A_b), 0(A_w)]$ and $A_w \approx [0(A_b), 1(A_w)]$ imply that the numbers θ_b and θ_w corresponding to the outcomes A_b and A_w, respectively, satisfy

$$\theta_b = 1 \qquad \text{and} \qquad \theta_w = 0. \tag{2.1}$$

Step 2: Let us show how to define the function u. (We then need to prove that this is the function we are after.) Let $\widehat{L} = [q_1(L_1), q_2(L_2), \ldots, q_M(L_M)]$ be any compound lottery, where $L_m = [p_1^m(A_1), p_2^m(A_2), \ldots, p_N^m(A_N)]$ is a simple lottery

[1] Recall that we may regard every outcome as a simple lottery.

for every $m = 1, \ldots, M$.

There is a small trick involved we should mention. Every lottery L_m depends on a finite number of outcomes, but these outcomes (and their number) might be different between these lotteries. However, here we have M simple lotteries L_m, each one of which depends on a finite number of outcomes, so the total number of outcomes ultimately appearing in $\widehat{L}$ is some finite number N. Now, we may assume that every simple lottery L_m depends on all these outcomes: simply put a probability zero on those outcomes which did not exist in the lottery in the first place. This is justified by the Simplification Axiom. In conclusion, we may assume that all the simple lotteries L_m depend on the same number N of outcomes.

For every $n = 1, \ldots, N$, define

$$r_n = q_1 p_n^1 + q_2 p_n^2 + \ldots + q_M p_n^M. \tag{2.2}$$

Recall that r_n is the probability of obtaining the outcome A_n from the compound lottery $\widehat{L}$. Pick a unique number $\theta_{A_n} \in [0, 1]$ for every $n = 1, \ldots, N$ such that $A_n \approx [\theta_{A_n}(A_b), (1 - \theta_{A_n})(A_w)]$ (by Step 1). Define a function[2] $u : \widehat{\mathcal{L}} \to \mathbb{R}$ by

$$u(\widehat{L}) = r_1\theta_{A_1} + r_2\theta_{A_2} + \ldots + r_N\theta_{A_N}. \tag{2.3}$$

This value is given as follows. For every outcome A_n appearing in the compound lottery $\widehat{L}$, find the probability r_n of obtaining A_n from $\widehat{L}$ and the probability θ_{A_n} such that $A_n \approx [\theta_{A_n}(A_b), (1-\theta_{A_n})(A_w)]$. Then $u(\widehat{L})$ is the sum of products of these probabilities over all the outcomes $A_1, A_2, \ldots, A_N$ appearing in $\widehat{L}$. In particular, if $L = [p_1(A_1), p_2(A_2), \ldots, p_N(A_N)]$ is a simple lottery, then the probability of an outcome A_n, $n = 1, \ldots, N$, in L is p_n. In conclusion,

$$u(L) = p_1\theta_{A_1} + p_2\theta_{A_2} + \ldots + p_N\theta_{A_N} \tag{2.4}$$

for every simple lottery L.

Step 3: Next, we claim that $u(A) = \theta_A$ for every outcome A. Recall that we have identified an outcome A with the lottery $L = [1(A)]$. (To be precise, we still should identify L with a compound lottery, but we already know how to calculate $u(L)$ from Step 2.) For this simple lottery, Equation (2.4) gives $u(L) = \theta_A$, as required.

[2] At this point, we need to know that the numbers θ_{A_n} are unique.

Step 4: It is now straightforward to show that u is linear. Indeed, given any simple lottery $L = [p_1(A_1), p_2(A_2), \ldots, p_N(A_N)]$, Equation (2.4) and Step 3 give

$$u(L) = p_1\theta_{A_1} + \ldots + p_N\theta_{A_N} = p_1 u(A_1) + \ldots + p_N u(A_N),$$

as required.

Step 5: We still need to show that u represents player's preferences. We establish the following result to finish our proof. Let $\widehat{L}$ be a compound lottery and let $\alpha = u(\widehat{L})$ (where $u(\widehat{L})$ is defined in Equation (2.3)). We claim that the player is indifferent between $\widehat{L}$ and the simple lottery $[\alpha(A_b), (1-\alpha)(A_w)]$.

Let us verify the previous claim. Let $\widehat{L} = [q_1(L_1), q_2(L_2), \ldots, q_M(L_M)]$ be some compound lottery, where $L_m = [p_1^m(A_1), p_2^m(A_2), \ldots, p_N^m(A_N)]$ is a simple lottery for every $m = 1, \ldots, M$. First, $\widehat{L} \approx [r_1(A_1), r_2(A_2), \ldots, r_N(A_N)]$ by Simplification Axiom, where r_n is given by Equation (2.2). Let $M_n = [\theta_{A_n}(A_b), (1-\theta_{A_n})(A_w)]$ for every $n = 1, \ldots, N$ and note that $A_n \approx M_n$ by the choice of θ_{A_n}. By Independence Axiom, we have $\widehat{L} \approx [r_1(M_1), r_2(M_2), \ldots, r_N(M_N)]$. Therefore, $\widehat{L}$ is equivalent to a lottery on just two outcomes, namely A_b and A_w. As in the explanation given after the Simplification Axiom, the probability p of obtaining the outcome A_b from this lottery is

$$p = r_1\theta_{A_1} + r_2\theta_{A_2} + \ldots + r_N\theta_{A_N} = u(\widehat{L}) = \alpha.$$

Simplification Axiom gives $\widehat{L} \approx [p(A_b), (1-p)(A_w)]$ (since $\widehat{L}$ depends only on these two outcomes), which together with the above equality verifies our claim.

Step 6: It is time to finish our proof. Let $\widehat{L}_1$ and $\widehat{L}_2$ be two compound lotteries such that $\widehat{L}_1 \succsim \widehat{L}_2$. The proof is completed once we show that $u(\widehat{L}_1) \geq u(\widehat{L}_2)$. By Step 5, we have

$$\widehat{L}_1 \approx [\alpha(A_b), (1-\alpha)(A_w)] \quad \text{and} \quad \widehat{L}_2 \approx [\beta(A_b), (1-\beta)(A_w)],$$

where $\alpha = u(\widehat{L}_1)$ and $\beta = u(\widehat{L}_2)$. Since the player prefers $\widehat{L}_1$ to $\widehat{L}_2$, we must have $\alpha \geq \beta$ by Monotonicity Axiom, and we are finished. □

Remark 2.17. Step 5 of the previous proof contains a point we consider worth of mentioning. First, it is our intention to construct a utility function which represents player's preferences over compound lotteries. As already noted, every compound lottery depends ultimately on some outcomes $A_1, \ldots, A_N$, so we seek for a utility function based on these outcomes. However, these outcomes are also related to each others via Continuity Axiom, since the player is willing to trade any outcome A to the lottery $[\theta_A(A_b),(1-\theta_A)(A_w)]$ for suitable $\theta_A \in [0,1]$. In conclusion, the value of a compound lottery to the player is determined ultimately by these extreme, best and worst, outcomes, and the probabilities θ_A. This fact is highlighted in Step 5, as it states that the value $u(\widehat{L})$ is chosen in such a way that the player accepts the simple lottery $[\theta(A_b),(1-\theta)(A_w)]$ in change for $\widehat{L}$, where $\theta = u(\widehat{L})$. In other words, we may interpret the value $u(\widehat{L})$ as the probability the player requires on the best outcome so that he (or she) would trade $\widehat{L}$ to a simple lottery on best and worse outcomes.

Let us finish this section by showing that the above constructed utility function is unique up to a positive affine transformation (defined below).

Definition 2.18. Let X be a non–empty set and let $u, v : X \to \mathbb{R}$ be functions. The function v is a *positive affine transformation* of u if and only if there exist real numbers $\alpha > 0$ and β such that $v(x) = \alpha u(x) + \beta$ for every $x \in X$.

Theorem 2.19. *If u is a linear utility function representing player's preferences over compound lotteries, then every positive affine transformation of u is also a linear utility function representing player's preferences. On the other hand, if u and v are linear utility functions representing player's preferences over compound lotteries, then there exist real numbers $\alpha > 0$ and β such that $v = \alpha u + \beta$.*

Proof. First, let u be a linear utility function over compound lotteries which represents player's preferences. Let $\alpha > 0$ and β be real numbers and consider the function $v = \alpha u + \beta$. To see that v also represents player's preferences, let $\widehat{L}_1$ and $\widehat{L}_2$ be compound lotteries. Since

$$v(\widehat{L}_1) = \alpha u(\widehat{L}_1) + \beta \quad \text{and} \quad v(\widehat{L}_2) = \alpha u(\widehat{L}_2) + \beta$$

and $\alpha > 0$, we see that $v(\widehat{L}_1) \geq v(\widehat{L}_2)$ is equivalent to $u(\widehat{L}_1) \geq u(\widehat{L}_2)$. Since u represents player's preferences, so does v.

To show that v is also linear, let $L = [p_1(A_1), p_2(A_2), \ldots, p_N(A_N)]$ be any simple lottery. The linearity of u and $p_1 + p_2 + \ldots + p_N = 1$ give

$$\begin{aligned} v(L) &= \alpha u(L) + \beta \\ &= \alpha[p_1 u(A_1) + p_2 u(A_2) + \ldots + p_N u(A_N)] + (p_1 + p_2 + \ldots + p_N)\beta \\ &= p_1(\alpha u(A_1) + \beta) + p_2(\alpha u(A_2) + \beta) + \ldots + p_N(\alpha u(A_N) + \beta) \\ &= p_1 v(A_1) + p_2 v(A_2) + \ldots + p_N v(A_N), \end{aligned}$$

thus verifying the first statement.

To prove the second statement, we deal first with the case where u is the function constructed in the proof of Theorem 2.16. Since u and v are linear and every compound lottery depends ultimately on some outcomes $A_1, \ldots, A_N$, it is enough to show that there exist real numbers $\alpha > 0$ and β such that $v(A) = \alpha u(A) + \beta$ for every outcome A. If the player is indifferent between all the outcomes, then u and v are both constant functions, and we may choose $\alpha = 1$ and $\beta = v - u$. As in the previous proof, the more interesting case is when $A_b \succ A_w$, and we assume now that this is the case.

To motivate our upcoming choices for α and β, note that *if* α and β exist, then Equation (2.1) gives

$$\begin{aligned} v(A_b) &= \alpha u(A_b) + \beta = \alpha + \beta \\ v(A_w) &= \alpha u(A_w) + \beta = \beta. \end{aligned}$$

Solving for α and β, we obtain $\beta = v(A_w)$ and $\alpha = v(A_b) - v(A_w)$. Let us show that these choices satisfy the statement.

Let A be any outcome. Recall that the number θ_A was chosen under the condition that $A \approx L = [\theta_A(A_b), (1 - \theta_A)(A_w)]$. Since v is linear and represents player's preferences, we have

$$v(A) = v(L) = \theta_A v(A_b) + (1 - \theta_A) v(A_w).$$

Our choices for α and β and Step 3 in the proof of Theorem 2.16 give

$$\alpha u(A) + \beta = (v(A_b) - v(A_w))\theta_A + v(A_w) = \theta_A v(A_b) + (1 - \theta_A) v(A_w) = v(A).$$

So, the statement is proved in the case that u is the function constructed in the proof of Theorem 2.16.

Finally, suppose that v_1 and v_2 are any two linear utility functions on compound lotteries both of which represent player's preferences. By what we just proved, there exist real numbers $\alpha_1 > 0$, β_1, $\alpha_2 > 0$, and β_2 such that $v_1 = \alpha_1 u + \beta_1$ and $v_2 = \alpha_2 u + \beta_2$. The second equation gives

$$u = \frac{1}{\alpha_2} v_2 - \frac{\beta_2}{\alpha_2}.$$

Substituting this into the first equation gives

$$v_1 = \frac{\alpha_1}{\alpha_2} v_2 + (\beta_1 - \frac{\alpha_1 \beta_2}{\alpha_2}),$$

thus finishing the proof. □

2.3 Risk aversion

In this section, we treat briefly player's attitudes towards risk. The reader may find a more comprehensive discussion in Cowell (2006: 190–197) or Maschler *et al.* (2013: 23–26). We consider only one player, so we denote his (or her) utility function by U. We assume that the player's preferences satisfy the axioms given in the previous section. We also assume that the outcomes O are some real numbers, for example, monetary payoffs.

Players attitudes towards risk are captured by certain equalities and inequalities describing player's willingness to change between lotteries and certain outcomes. For these purposes, we need the following definition. Recall that the outcomes $A_1, A_2, \ldots, A_N$ are now real numbers, so the definition given below is reasonable.

Definition 2.20. Let $L = [p_1(A_1), p_2(A_2), \ldots p_N(A_N)]$ be a simple lottery. The *expected value* of L is the number

$$\mu = \sum_{k=1}^{N} p_k A_k = p_1 A_1 + p_2 A_2 + \ldots + p_N A_N$$

Recall that the expected value of a lottery is the *long term average payout.*

Definition 2.21. A player is *risk neutral* if and only if

$$pU(A_1) + (1 - p)U(A_2) = U(pA_1 + (1 - p)A_2).$$

for any two outcomes A_1 and A_2 and any $0 \leq p \leq 1$. A player is *risk averse* if and only if

$$pU(A_1) + (1-p)U(A_2) \leq U(pA_1 + (1-p)A_2)$$

for any two outcomes A_1 and A_2 and any $0 \leq p \leq 1$. A player is *risk loving* (or *risk seeking*) if and only if

$$pU(A_1) + (1-p)U(A_2) \geq U(pA_1 + (1-p)A_2)$$

for any two outcomes A_1 and A_2 and any $0 \leq p \leq 1$.

Note that, in all these equations, the left hand side is the expected utility of the player and the term inside the parenthesis on the right side is the expected value μ of a lottery $[p(A_1), (1-p)(A_2)]$. So, players attitudes towards risk are captured by whether he (or she) is willing to trade a lottery to the certain outcome of expected value μ of the lottery. Risk averse player rather takes the certain outcome μ, but risk loving player rather takes the lottery than takes the certain outcome μ.

The condition for risk averse player defines a concave function. Indeed, note that this condition says that for any two outcomes the line segment joining the points on the curve is below the curve. Therefore, preferences of a risk averse player are captured by a concave utility function, for example, by x^α for some $0 < \alpha \leq 1$. Similarly, preferences of a risk loving player are captured by a convex function. Most of the time in this thesis, we consider risk neutral or risk averse players; we make only a few comments concerning risk loving players.

3 BARGAINING PROBLEMS AND NASH'S SOLUTION

The purpose of this thesis is to study bargaining situations between two players and to present some applications of bargaining situations in the field of economics. We begin in this chapter by introducing general (abstract) bargaining problems and presenting the famous Nash's bargaining solution, originally presented in Nash (1950). As we shall see, this is a very general bargaining solution, since the structure of a bargaining situation (such as who makes an offer and when) is not described in any way. Instead, the solution is axiomatic in the sense that it is derived from a list of assumptions about the bargaining outcome that one could consider reasonable. However, this method also has its limitations, since it can be questioned whether these axioms actually describe real life bargaining processes. We study strategic models of bargaining in the next chapter.

We begin in the first section by briefly describing a bargaining situation where the players are trying to find an agreement over a partition of a cake between the players. We only describe the underlying situation briefly, and we will return to this problem a number of times as we proceed. In the next section, we introduce general (abstract) bargaining problems and the concept of bargaining solution. The most important results of this chapter are given in Section 3. Here, we present Nash's bargaining solution and some of its basic properties. As we shall see, Nash's solution does not take into account possible differences between the bargaining abilities of the players. In the fourth section, we present a generalization of Nash's solution, originally given in Harsanyi and Selten (1971), which takes these possible differences into account. In the second to last section, we present an application of Nash bargaining solution to working in teams. Finally, in the last section we present some further comments on the results given in this chapter.

We make the following convention for the rest of this chapter: *We assume that the preferences of Player A and Player B are represented by linear utility functions by* U_A *and* U_B*, respectively.*

3.1 Partition of a cake

Let us begin by treating briefly a situation where Player A and Player B are trying to find an agreement over how to divide a cake between the players. Although the underlying situation is relatively simple, this partition problem has drawn a lot of attention and it has been studied using a number of different approaches. For one thing, this kind of partition problems arise in a number of economic applications as, in many cases, economic activity creates surplus and it is natural to ask how this surplus is divided between the agents. We also provide some examples of division of surplus between economic agents in later sections.

In this section, we do not consider the actual bargaining process at all, or even any outcomes (divisions of a cake) that we might expect to result. Instead, the purpose of introducing this bargaining problem at this point is to clarify the definition of a general bargaining problem given in the next section. We return to this particular bargaining problem a number of times as we proceed.

Bargaining situation over the partition of a cake is given as follows. Player A and Player B are trying to find an agreement on how to divide a cake of size 1 between the players. At this point, we assume that no cake wasted, that is, if Player A obtains a share x_A, $0 \leq x_A \leq 1$, of the cake, then Player B will obtain $1 - x_A$, the rest of the cake.[3] Under this assumption, the *possible agreements* are represented by the set

$$X = \{(x_A, 1 - x_A) : 0 \leq x_A \leq 1\}.$$

The first and the second coordinates are the shares of cake obtained by Player A and Player B, respectively. If the players reach an agreement $(x_A, 1 - x_A)$, then the resulting utilities are $U_A(x_A)$ and $U_B(1 - x_A)$, respectively. However, if the players fail to reach an agreement, then neither of the players receives any cake. We assume that the players have some utility levels d_A and d_B, respectively, before they start bargaining. So, in the case of disagreement, the resulting utilities of Player A and Player B are d_A and d_B, respectively.[4]

Even without any information about how the bargaining process proceeds, there are some assumptions concerning this problem we find reasonable. First of all, we assume that $d_A \geq U_A(0)$ and $d_B \geq U_B(0)$, that is, having no cake will make neither of the players better off. To make sure that we have a reason to study this bargaining situation, we assume that there exists some $0 \leq x_A \leq 1$ such that $U(x_A) > d_A$ and $U_B(1 - x_A) > d_B$. For otherwise, one (or both) of the players

[3] This means that the partition is *efficient*. We return to this concept later.

[4] We are implicitly assuming that, in the case of disagreement, time spent on bargaining does not matter for the players: d_A and d_B do not depend on time.

would not benefit anything from bargaining, and we have no reason to expect that the bargaining process would take place between the players. Having this particular bargaining situation in mind, we now turn our attention to general bargaining problems.

3.2 Bargaining problems

In this section, we focus on general bargaining problems between two players. Let us begin by considering a bargaining situation between these players in some detail. Whatever the players are bargaining over, the possible (physical) outcomes in the case of agreement are given (as above) by some set of *possible agreements*

$$X = \{(x_A, x_B)\}.$$

Here, x_A and x_B might be something relatively simple, like a share of a cake, or something much more complex, like a highly specified contract.[5] It would perhaps seem reasonable to base a study of bargaining situations on these sets. However, following J. Nash, we choose a different approach. Instead of concentrating on the actual (physical) outcomes, we turn our attention to players' utilities determined by these outcomes. So, if the players reach an agreement (x_A, x_B), the resulting utilities are $U_A(x_A)$ and $U_B(x_B)$, respectively. We base our treatment on the *attainable set*

$$\{(U_A(x_A), U_B(x_B)) : (x_A, x_B) \in X\}$$

of possible utilities available to players from bargaining. Given some reasonable assumptions about players' utility functions (such as they are strictly increasing), we can always deduce the physical outcome from the resulting utilities. Next, let us turn our attention to general bargaining problems.

Definition 3.1. A *bargaining problem* is a pair (K, d) such that

(i) K is a non–empty compact and convex subset of $\mathbb{R}^2$;

(ii) d is some point in K;

(iii) there exists some $u \in K$ such that $d \ll u$.

[5] One can argue which one is more complex: ink on a paper or a delicious combination of ingredients tied together by several chemical reactions.

Given a bargaining problem (K, d), the set K is called the *attainable set* and the point d is called the *disagreement point* (*threat point*). We denote by $\mathcal{B}$ the set of all bargaining problems.

As roughly explained above, an interpretation of the above definition is given as follows. Every element (u_A, u_B) of K represents possible *utilities* of Player A and Player B, respectively, that they could achieve from bargaining situation. The disagreement point $d = (d_A, d_B)$ represents the utilities of the players in the case of disagreement ("status quo", as put in Roth (1977)). The third condition guarantees that there exists some outcome that benefits both of the players, that is, the players have an incentive to bargain.

Since we base our study of bargaining situations to utilities instead of physical outcomes, this means that should we begin with physical outcomes the players could achieve by finding an agreement, we first transform these outcomes into utility pairs by using players' utility functions. For example, in the case of the partition of a cake of size 1, the physical outcomes are described Figure 1a. (A point inside the triangle represents a division where some part of the cake is wasted. Those partitions where no cake is wasted are represented by the line $x_B = 1 - x_A$.)

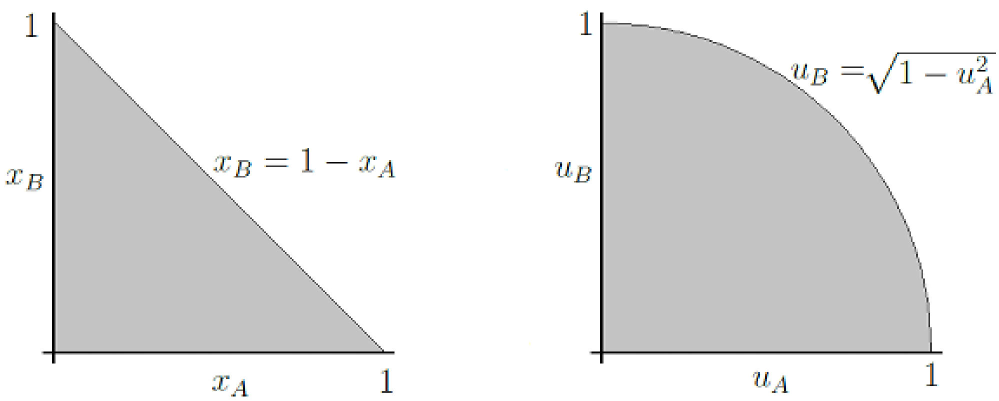

(a) Possible agreements (b) Attainable set

Figure 1: Partition of a cake – possible agreements and utilities

For the sake of an example, let us assume that the utility functions of Player A and Player B are given by $U_A(x) = U_B(x) = \sqrt{x}$. Fix some possible utility u_A of Player A, so that $u_A = \sqrt{x_A}$ for some $0 \leq x_A \leq 1$. Then $x_A = u_A^2$, so the share

x_B obtained by Player B is at most $1 - x_A = 1 - u_A^2$, and so the utility of Player B is at most $\sqrt{1-u_A^2}$. Therefore, the attainable set is given as in Figure 1b.

Next, we need to describe what we mean by a solution to a bargaining problem. Keeping in mind our interpretation of a bargaining problem (K, d), that is, the members of K represent possible utilities of rational negotiators, we would not expect the players to reach an agreement which would make (at least) one of the players worse off. Therefore, any possible agreement should be a member of the set (see Figure 2)

$$K(d) = \{u \in K : d \leq u\}.$$

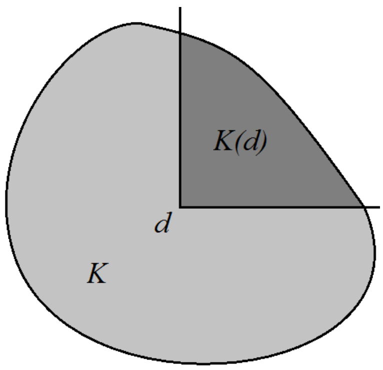

Figure 2: Attainable set K and the set $K(d)$

Remark 3.2. Given any bargaining problem (K, d), the pair $(K(d), d)$ is also a bargaining problem. Indeed, the set $K(d)$ is non–empty since $d \in K(d)$, and it is compact and convex as the intersection of K with the closed and convex set $\{x \in \mathbb{R}^2 : d \leq x\}$ (See Figure 2). Also, by property (iii) of a bargaining problem, there exists some $u \in K$ such that $d \ll u$. Clearly, $u \in K(d)$, thus showing that condition (iii) of a bargaining problem also holds for $(K(d), d)$.

Our discussion above leads us to the following definition. Note that we require bargaining solution to be a member of $K(d)$, as these points represent outcomes which benefit both of the players. The members of K outside $K(d)$ (the light gray area in Figure 2) do not represent reasonable outcomes to a bargaining problem, since these points represent situations where at least one of the players is worse off than in *status quo*. We may interpret these points in such a way that, although the players are aware of these possibilities, as rational bargainers they know that these outcomes will never be reached.

Definition 3.3. A *bargaining solution* is a function F which chooses a unique member of $K(d)$ for every bargaining problem (K, d).

For a bargaining solution F and a bargaining problem (K, d), we may interpret the value $F(K, d) = (u_A^*, u_B^*)$ such that, according to this solution, we expect the players to reach an agreement where their resulting utilities are u_A^* and u_B^*, respectively. Next, we turn our attention to a particular bargaining solution, namely to the one introduced by J. Nash.

3.3 Nash bargaining solution

A somewhat natural method to study a bargaining situation is to model the bargaining process as some kind of game between the players and to use tools of game theory to analyze the resulting game. This approach, however, requires some knowledge (or assumptions) about the actual bargaining process, such as who makes an offer and when. In his seminal paper Nash (1950), J. Nash chose a somewhat opposite approach, as he did not treat any particular bargaining process. Instead, he started with some assumptions (axioms) about the *outcome* of a bargaining situation that one could consider reasonable. He was able to show that there exists a unique bargaining solution satisfying these axioms and, furthermore, that this solution is given by a certain optimization problem. In this section, we describe Nash's solution. This solution is treated in a number of literature; the presentation given here follows Osborne and Rubinstein (1990: 13–17).

To describe Nash's solution, we first introduce the four axioms Nash used to derive his bargaining solution. Whether these axioms represent actual bargaining situations can be questioned: we discuss these axioms only briefly, but the reader may find a more comprehensive treatment in Hargreaves Heap and Varoufakis (1995: 118–128).

In what follows, let F be a bargaining solution. We denote by F_A and F_B the components of F, that is, in $F(K, d) = (F_A(K, d), F_B(K, d))$, these components give the utilities of Player A and Player B, respectively, in the case of agreement. The first axiom is as follows.

EFF (*Pareto Efficiency*) If (K, d) is a bargaining problem and $u, v \in K$ satisfy $u \ll v$, then $F(K, d) \neq u$.

This assumption states that the players will never agree on an outcome if there is a possibility of making both of the players better off. We may interpret this assumption as follows: any outcome u such that $u \ll v$ for some $v \in K$ would leave room for "renegotiation". Indeed, suppose that $u, v \in K$ satisfy $u \ll v$ and consider the set $K(u) = (w \in K : u \leq w)$. Now, $(K(u), u)$ is another bargaining problem (see Remark 3.2 and Figure 3a below). Every member of $K(u)$ represents an agreement in the original bargaining problem. In addition, all these outcomes make both of the players at least as well off as agreeing to u. Therefore, the players would have an incentive to continue bargaining to reach an agreement in the set $K(u)$. The efficient outcomes are described by the curve between P and Q in Figure 3b.

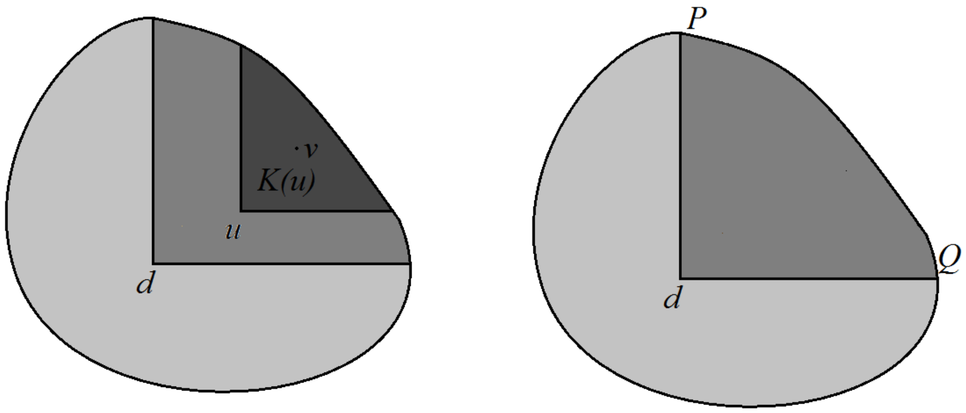

(a) Renegotiation set $K(u)$ (b) Efficient outcomes

Figure 3: Pareto efficiency

Note here that if F satisfies **EFF**, then the players will never disagree by property (iii) of a bargaining problem. Furthermore, if bargaining process takes time and some resources are used during this process, then **EFF** implies that an agreement is reached immediately. To introduce the second axiom, we first need a definition.

Definition 3.4. A bargaining problem (K, d) is *symmetric* if and only if $d_A = d_B$ and $(u_A, u_B) \in K$ if and only if $(u_B, u_A) \in K$.

Note that a bargaining problem (K, d) is symmetric if and only if the set K is symmetric (as a subset of $\mathbb{R}^2$) with respect to the line $y = x$ and the disagreement point is on this line. (The mapping $(x_1, x_2) \mapsto (x_2, x_1)$ reflects a point (x_1, x_2) of $\mathbb{R}^2$ symmetrically with respect to the line $y = x$.) The second axiom concerns symmetric bargaining problems.

SYM (*Symmetry*) If (K, d) is symmetric, then $F_A(K, d) = F_B(K, d)$.

Suppose that F is a bargaining solution satisfying **SYM** and let (K, d) be a symmetric bargaining problem. According to the solution F, the players obtain the same outcome (in terms of their utilities), which means that the solution $F(K, d) = (u_A^*, u_B^*)$ is on the line $y = x$. Axiom **SYM** is usually taken to mean that the players are equally "tough" bargainers, and this is the interpretation given in Nash (1950). However, Nash abstracts from this viewpoint in Nash (1953): "With people who are sufficiently intelligent and rational there should not be any question of "bargaining ability", a term which suggests something like skill in duping the other fellow. The usual haggling process is based on imperfect information, the hagglers trying to propagandize each other into misconceptions of the utilities involved. Our assumption of complete information makes such an attempt meaningless." Instead, Nash takes Axiom **SYM** simply to mean that the players are intelligent and rational beings

Example 3.5. These two axioms, **EFF** and **SYM**, are already enough to give a solution to the problem of partition of a cake presented in Figure 1b. (Recall that we assumed that the players have the same utility functions.) This problem is symmetric, so the solution lies on the line $y = x$ (or $u_B = u_A$, see Figure 4). Then **EFF** implies that the solution must lie on the boundary circle (point P in Figure 4). The players have equal utilities ($u_A^* = u_B^* = 1/\sqrt{2}$ to be precise), which means that both of the players get half of the cake.

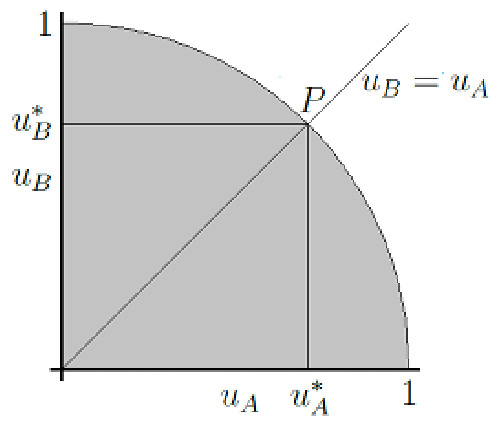

Figure 4: Partition of a cake – symmetric case

To introduce the next axiom, let α_A, $\alpha_B > 0$ and β_A, β_B be real numbers. Define $T : \mathbb{R}^2 \to \mathbb{R}^2$ by $T(x, y) = (\alpha_A x + \beta_A, \alpha_B y + \beta_B)$ and put

$$d' = T(d) = (\alpha_A d_A + \beta_A, \alpha_B d_B + \beta_B) \tag{3.1}$$

and

$$K' = T(K) = \{(\alpha_A u_A + \beta_A, \alpha_B u_B + \beta_B) : (u_A, u_B) \in K\}. \tag{3.2}$$

Note here that T is an affine transformation, and so (K', d') is also a bargaining problem (see Remark 2.6).

INV (*Invariance to Equivalent Utility Representations*) Take any bargaining problem (K, d). If K' and d' are as above, then

$$F_A(K', d') = \alpha_A F_A(K, d) + \beta_A \quad \text{and} \quad F_B(K', d') = \alpha_B F_B(K, d) + \beta_B.$$

Remark 3.6. A short way to formulate **INV** is $F(T(K), T(d)) = T(F(K, d))$.

This axiom can be justified by assuming that the players' preferences are the basic for the bargaining problem, and not the particular utility functions used. This is perhaps best understood if we consider the actual (physical) outcomes of bargaining. Suppose that $F(K, d) = (u_A^*, u_B^*)$. These utilities correspond to some physical outcome (x_A, x_B), and so $u_A^* = U_A(x_A)$ and $u_B^* = U_B(x_B)$. Let α_A, α_B, β_A, and β_B be as above and consider the utility functions

$$U_A' = \alpha_A U_A + \beta_A \quad \text{and} \quad U_B' = \alpha_B U_B + \beta_B.$$

Recall that these functions also represent players' preferences (see Theorem 2.19). The physical outcomes are turned into the bargaining problem (K', d') under these utility functions. Since the players' preferences have not changed, we assume that the players reach the same physical agreement (x_A, x_B) with resulting utilities $U_A'(x_A)$ and $U_B'(x_B)$. Now,

$$U_A'(x_A) = \alpha_A U_A(x_A) + \beta_A = \alpha_A u_A^* + \beta_A = \alpha_A F_A(K, d) + \beta_A,$$

and similar equality holds for $U_B'(x_B)$. This is precisely Axiom **INV**. Finally, we introduce the last of Nash's axioms.

IIA (*Independence of Irrelevant Alternatives*) Suppose that (K', d) and (K, d) are bargaining problems such that K contains K' and $F(K, d)$ is a member of K'. Then $F(K', d) = F(K, d)$.

An interpretation of the previous axiom is given as follows. First, note that the disagreement point in the previous definition is the same for both of the bargaining problems. In the larger attainable set K, the players reach an agreement $F(K, d)$.

If $F(K,d)$ is an element of the smaller attainable set K' (the light gray area in Figure 5), then we assume that the players reach the same agreement in K'. In other words, we consider all the other members of K' except $F(K,d)$ as irrelevant, since the players did not find these members attractive when they had more options on K. Note here that **IIA** is automatically satisfied if F is defined in such a way that $F(K,d)$ is the maximizer of some function on K and this maximizer is also a member of K'.

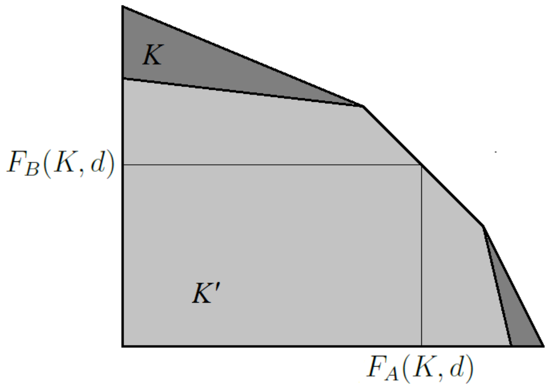

Figure 5: Axiom **IIA**

Whether axiom **IIA** is reasonable depends on the actual bargaining process. If the bargaining process eliminates outcomes until just one is left, then **IIA** seems to be a reasonable assumption, but in general one should be rather cautious on accepting this axiom. In fact, **IIA** is the most criticized of the Nash's axioms, see, for example, Luce and Raiffa (1957: 132–134) and Kalai and Smorodinsky (1975: 514–515).

Now, we are ready to present the famous Nash's bargaining solution. We give a proof for this theorem below; except for the uniqueness of the solution, the proof is rather straightforward. Note carefully the content of the next theorem: if it is reasonable that the previous axioms describe some bargaining situation, then the bargaining solution is given by the maximization problem given below (as this is the *only* solution satisfying the axioms). Also, since Nash solution satisfies these axioms, we should not apply Nash solution if we have some reason to believe that one of the previous assumptions (**IIA**, for example) does not hold.

Theorem 3.7. *There exists a unique bargaining solution F^N satisfying **EFF**, **SYM**, **INV**, and **IIA**. This* Nash solution *is given by*

$$F^N(K,d) = \operatorname*{arg\,max}_{(u_A,u_B)\in K(d)} (u_A - d_A)(u_B - d_B).$$

Proof. We divide the proof into a number of steps. First, let us show that the above equality is well–defined, that is, the given maximization problem admits a unique solution. Let (K, d) be a bargaining problem and let

$$f(x, y) = (x - d_A)(y - d_B).$$

Since f is continuous on $\mathbb{R}^2$, it attains its maximum value on the compact set $K(d)$. Furthermore, f is strictly quasi–concave on the set $\{v \in K(d) : d \ll v\}$. (This is the dark gray area in Figure 2 with horizontal and vertical lines removed.) Since this set is not empty by assumption (iii) of a bargaining problem, function f attains its maximum value at a unique point of $K(d)$ (see Chiang & Wainwright (2005: 364–371)). Next, let us show that F^N satisfies the four axioms given above.

EFF: Let (K, d) be a bargaining problem and let $F^N(K, d) = (u_A^*, u_B^*)$, that is, (u_A^*, u_B^*) is the maximizer of f on $K(d)$.[6] Consider any $(x, y) \in \mathbb{R}^2$ such that $(u_A^*, u_B^*) \ll (x, y)$. Now,

$$(x - d_A) > (u_A^* - d_A) \geq 0 \quad \text{and} \quad (y - d_B) > (u_B^* - d_B) \geq 0,$$

and so $f(x, y) > f(u_A^*, u_B^*)$. Since $f(u_A^*, u_B^*)$ is the maximum value of f on $K(d)$, the point (x, y) does not belong to the set $K(d)$. Therefore, there does not exist $(u_A, u_B) \in K(d)$ satisfying $(u_A^*, u_B^*) \ll (u_A, u_B)$, thus showing that F^N is efficient.

SYM: Let (K, d) be a symmetric bargaining problem. Recall that $d_A = d_B$. Put $F^N(K, d) = (u_A^*, u_B^*)$, that is, (u_A^*, u_B^*) is the maximizer of the function $f(x, y) = (x - d_A)(y - d_A)$ on $K(d)$. Now, $f(u_A^*, u_B^*) = f(u_B^*, u_A^*)$, which means that (u_B^*, u_A^*) also maximizes f on $K(d)$. Since f attains its maximum on $K(d)$ at a unique point, we must have $u_A^* = u_B^*$, as required.

INV: We use the same notation as in Axiom **INV** and the discussion preceding it. (See Equations (3.1) and (3.2).) First, note that $(u'_A, u'_B) \in K'$ if and only if $u'_A = \alpha_A u_A + \beta_A$ and $u'_B = \alpha_B u_B + \beta_B$ for some $(u_A, u_B) \in K$. If $(u'_A, u'_B) \in K'$, then the definition of d' gives

$$u'_A - d'_A = \alpha_A u_A + \beta_A - (\alpha_A d_A + \beta_A) = \alpha_A(u_A - d_A),$$

[6] The function f depends on the disagreement point $d = (d_A, d_B)$, but we use the same notation throughout this proof.

and similarly $u'_B - d'_B = \alpha_B(u_B - d_B)$. Therefore,

$$(u'_A - d'_A)(u'_B - d'_B) = \alpha_A\alpha_B(u_A - d_A)(u_B - d_B).$$

Since $\alpha_A, \alpha_B > 0$, we see that (u^*_A, u^*_B) maximizes $(u_A - d_A)(u_B - d_B)$ on $K(d)$ if and only if $(\alpha_A u^*_A + \beta_A, \alpha_B u^*_B + \beta_B)$ maximizes $(u'_A - d'_A)(u'_B - d'_B)$ on $K'(d')$. In conclusion,

$$F^N_A(K', d') = \alpha_A u^*_A + \beta_A = \alpha F^N_A(K, d) + \beta.$$

A similar equality holds for $F^N_B(K', d')$, and so F^N satisfies **INV**.

IIA: Let (K, d) and (K', d) be bargaining problems such that $K' \subseteq K$ and $F^N(K, d) \in K'$. By definition, the point $F^N(K, d)$ maximizes f on $K(d)$. Since $F^N(K, d) \in K'$ and $K' \subseteq K$, the point $F^N(K, d)$ must also be the maximizer of f on $K'(d)$. Therefore, $F^N(K, d) = F^N(K', d)$.

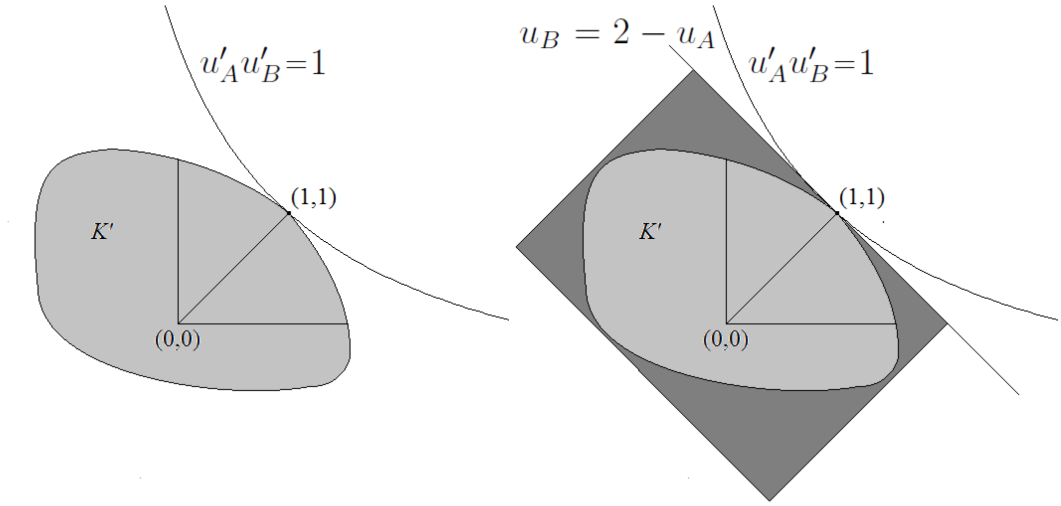

(a) Nash solution of K' (b) Enlarged bargaining problem

Figure 6: Proof of Nash's Theorem

We still need to show that F^N is the only bargaining solution which satisfies the given axioms, so suppose that F is another such a solution. We need to show that F and F^N give the same solution to every bargaining problem, so let (K, d) be some bargaining problem. The proof involves several steps, but the underlying idea is quite simple, and is illustrated in Figure 6. First, transform (K, d) into a bargaining problem K' whose disagreement point and Nash solution are $(0, 0)$ and $(1, 1)$, respectively (Figure 6a). Then "enlarge" this problem into a symmetric bargaining problem R *without changing* the Nash solution (Figure 6b). (This is an important part of the proof: the transformed bargaining problem should be on the left side of the line $u_B = 2 - u_A$, so that we can "enlarge" it properly.) Then the axioms imply that F gives the same solution $(1,1)$ to the transformed bargaining problem. The details are given as follows.

Put $F^N(K,d) = (u_A^*, u_B^*)$. Since f attains its maximum on $K(d)$ at this point and f attains a strictly positive value on $K(d)$ by property (iii) of a bargaining problem, we must have $u_A^* > d_A$ and $u_B^* > d_B$. Define

$$\alpha_A = \frac{1}{u_A^* - d_A}, \quad \beta_A = \frac{d_A}{u_A^* - d_A}, \quad \alpha_B = \frac{1}{u_B^* - d_B}, \quad \beta_B = \frac{d_B}{u_B^* - d_B},$$

and $T : \mathbb{R}^2 \to \mathbb{R}^2$ by $T(x,y) = (\alpha_A x - \beta_A, \alpha_B y - \beta_B)$. Then $T(d_A, d_B) = (0,0)$ and $T(u_A^*, u_B^*) = (1,1)$. In what follows, we denote the origin $(0,0)$ of $\mathbb{R}^2$ simply by 0. Put $K' = T(K)$ and note that $(K', 0)$ is the bargaining problem obtained from (K,d) under the transformation T. Since both F and F^N satisfy **INV**, we have

$$F_i(K',0) = \alpha_i F_i(K,d) + \beta_i \quad \text{and} \quad F_i^N(K',0) = \alpha_i F_i^N(K,d) + \beta_i.$$

for $i = A, B$. Since $\alpha_A, \alpha_B > 0$, we see that $F(K,d) = F^N(K,d)$ if and only if $F(K',0) = F^N(K',0)$. By Remark 3.6, we have

$$F^N(K',0) = F^N(T(K,d)) = T(F^N(K,d)) = T(u_A^*, u_B^*) = (1,1).$$

Therefore, it is enough to show that $F(K',0) = (1,1)$ to finish the proof. Note here that, by the definition of the Nash solution, $F^N(K',0) = (1,1)$ means that the maximum value of $(u'_A - 0)(u'_B - 0) = u'_A u'_B$ on $K'(0)$ (the light gray area in Figure 6 bounded by heavy lines) is 1.

Here comes the crux of the proof. We claim that there does not exist $(u'_A, u'_B) \in K'$ such that $u'_A + u'_B > 2$. (This is clear from Figure 6. Since the product $u'_A u'_B$ attains the maximum value 1 on K' at the point $(1,1)$, the curve $u'_A u'_B = 1$ is tangent to K' at $(1,1)$. We also provide an algebraic verification.) Suppose that such a point exists and define

$$v'_A = (1-\varepsilon) \cdot 1 + \varepsilon \cdot u'_A \quad \text{and} \quad v'_B = (1-\varepsilon) \cdot 1 + \varepsilon \cdot u'_B,$$

where $0 < \varepsilon < 1$. Since K' is convex, we have $(v'_A, v'_B) \in K'$ for every $0 < \varepsilon < 1$.

By multiplying and gathering the terms, we see that

$$v'_A v'_B = 1 + (u'_A + u'_B - 2)\varepsilon + (1 - u'_A - u'_B + u'_A u'_B)\varepsilon^2.$$

Since $u'_A + u'_B - 2 > 0$, we see that $v'_A v'_B > 1$ for $\varepsilon > 0$ small enough.[7] This contradiction with the conclusion at the end of the previous paragraph verifies the claim.

The rest of the proof is straightforward. We know by the previous claim that K' is contained in the half–plane $\{(u_A, u_B) \in \mathbb{R}^2 : u_B \leq 2 - u_A\}$ (this is the area to the left from the line $u_B = 2 - u_A$ in Figure 6). Since K' is compact, there exists a compact rectangle R in $\mathbb{R}^2$ such that R is contained in the half–plane $\{(x, y) \in \mathbb{R}^2 : y \leq 2 - x\}$, R is symmetric with respect the line $u_B = u_A$, and R contains the set K'. Therefore, $(R, 0)$ is a symmetric bargaining problem. Since F satisfies **SYM**, the solution $F(R, 0)$ must be on the line $u_B = u_A$. Since F satisfies **EFF**, $(1, 1) \in R$, and $x + y \leq 2$ for every $(x, y) \in R$, we must have $F(R, 0) = (1, 1)$. Finally, **IIA** implies $F(K', 0) = F(R, 0) = (1, 1)$, thus finishing the proof. □

Remark 3.8. Suppose that F is a bargaining solution satisfying **SYM** and **EFF**. The last part of the previous proof shows that F and F^N give the same solution for every symmetric bargaining problem.

The product $(u_A - d_A)(u_B - d_B)$ is known as the *Nash product*. We may interpret the terms $(u_A - d_A)$ and $(u_B - d_B)$ as *excess utilities*, since these terms represent the additional utility (compared to *status quo*) obtained by the players in the case of an agreement (u_A, u_B). Therefore, the Nash solution is given by the condition that the product of the players' excess utilities is maximized.

Example 3.9. Let us apply Nash solution to the problem of partitioning a cake when the disagreement point $d = (d_A, d_B)$ and the utility functions of the players are $U_A(x) = U_B(x) = x$, respectively. In order to apply Nash solution, we need to assume that $d_A + d_B < 1$. (Otherwise, property (iii) of a bargaining problem does not hold, that is, there are no mutually beneficial agreements.) Efficiency requires that no cake is wasted. Therefore, we need to find the maximum of

$$(x - d_A)(1 - x - d_B).$$

[7]When $\varepsilon > 0$ is small, the term ε^2 is neglible compared to ε. Therefore, for small ε the given term is very close to $1 + (u'_A + u'_B - 2)\varepsilon$, which is strictly greater than 1.

Equating the derivative of this function with zero, we find that the maximum value is obtained at $x^* = (1 + d_A - d_B)/2$. We may write the shares of cake obtained by Player A and Player B, respectively, by

$$x_A = d_A + \frac{1}{2}(1 - d_A - d_B) \quad \text{and} \quad x_B = d_B + \frac{1}{2}(1 - d_A - d_B).$$

These expressions admit a nice interpretation. Namely, first give both of the players their utilities in the case of disagreement, d_A and d_B, and then split the rest of the cake equally between the players.

Example 3.10. Let us consider again the problem of partitioning a cake. This time, let us assume that the utility functions are given by $U_A(x_A) = x_A^\gamma$ and $U_B(x_B) = x_B^\delta$ for some $\gamma, \delta \in]0, 1]$, and let $d = (0, 0)$. In this case, the Nash product $(u_A - d_A)(u_B - d_B) = u_A u_B$. Efficiency requires that the whole cake is used, so $x_B = 1 - x_A$. Since utilities are functions of x_A, we can solve the problem by maximizing the function $x_A^\gamma(1 - x_A)^\delta$. The first order condition is

$$\gamma x_A^{\gamma-1}(1 - x_A)^\delta - \delta x_A^\gamma(1 - x_A)^{\delta-1} = 0.$$

Write this equation as

$$x_A^{\gamma-1}(1 - x_A)^{\delta-1}[\gamma(1 - x_A) - \delta x_A] = 0.$$

This equation has three solutions. Two solutions are $x_A = 0$ and $x_A = 1$, but we notice that these values *minimize* our target function (giving value zero). The third solution, which is obtained by equating zero with the term in the parenthesis above, is the one we are after. Together with the equation $x_B = 1 - x_A$, we obtain

$$x_A = \frac{\gamma}{\gamma + \delta} \qquad \text{and} \qquad x_B = \frac{\delta}{\gamma + \delta}.$$

Note that if $\gamma = \delta$, then $x_A = x_B = 1/2$ as we would expect since, in this case, the bargaining problem is symmetric. Note also that

$$\frac{x_A}{x_B} = \frac{\gamma}{\delta}.$$

Therefore, it is the *ratio* of γ and δ, and not the particular values of γ and δ, which determines the partition. For example, both $\gamma = 0.3$, $\delta = 0.9$ and $\gamma = 0.1$, $\delta = 0.3$ yield a division where Player B obtains 75 % of the cake.

The last equality also shows that more risk averse player will get a smaller share of cake. That is, if $\gamma < \delta$, then Player A obtains less than half of the cake. Furthermore, if γ decreases, that is, Player A becomes more risk averse, then his (or her) share of cake also decreases. (This statement is true for more general utility functions than the ones treated here, see Osborne and Rubinstein (1990: 17–19).) We may interpret these results in a way that a risk averse player will rather settle for a small share of cake than take a risk that negotiations break down. Also, the more risk averse the player is compared to the other player, the smaller share he (or she) will settle.

The Nash solution admits also a nice geometric interpretation, which we proceed to describe. For these purposes, we need the following concept.

Definition 3.11. Let (K, d) be a bargaining problem. The *strong Pareto frontier* of K is the set

$$\{u \in K : \text{there does not exist } v \in K \text{ such that } u < v\}.$$

The *weak Pareto frontier* of K the set

$$\{u \in K : \text{there does not exist } v \in K \text{ such that } u \ll v\}.$$

Note in Figure 7 below that x and y belong to the weak Pareto frontier of K but they do not belong to the strong Pareto frontier of K. There does not exist points of K with *both* of the coordinates larger than in x or y, and so x and y are in the weak Pareto frontier of K. However, u and v satisfy $x < u$ and $y < v$, and so x and y are not in the strong Pareto frontier of K. The strong Pareto frontier consists of the curve between u and v, whereas the weak Pareto frontier also contains the horizontal and vertical lines containing the points x and y.

Remark 3.12. Axiom **EFF** means that the Nash solution $F^N(K, d)$ of a bargaining problem (K, d) belongs to the weak Pareto frontier of K. However, a stronger statement is actually true, namely that the Nash solution is always a member of the strong Pareto frontier of K. If $u = (u_A, u_B)$ and $v = (v_A, v_B)$ satisfy $d \ll u, v$ and $u < v$, then either $u_A - d_A < v_A - d_A$ or $u_B - d_B < v_B - d_B$. But then v gives a higher value to the Nash product than u, and so u is not the Nash solution.

Theorem 3.13. *Let (K, d) be a bargaining problem and let $u_2 = g(u_1)$ be the equation for the strong Pareto frontier of K.*

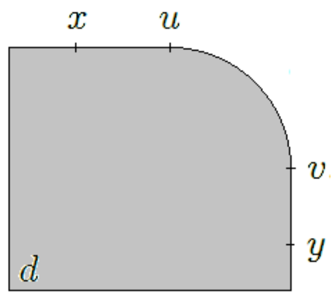

Figure 7: Pareto frontiers

(i) The Nash solution $F^N(K,d)$ is the pair (u_1^*, u_2^*) if and only if $u_2^* = g(u_1^*)$ and u_1^* maximizes the function $f(u_1) = (u_1 - d_1)(g(u_1) - d_2)$.

(ii) If g is differentiable, then $F^N(K,d) = (u_1^*, u_2^*)$ if and only if

$$u_2^* = g(u_1^*) \quad \textit{and} \quad |g'(u_1^*)| = \frac{u_2^* - d_2}{u_1^* - d_1}.$$

Proof. The first statement is clear, since we argued above that the Nash solution belongs to the strong Pareto frontier of K. The second statement is easy to prove by writing the first order condition for the function f and solving for $g'(u_1)$. (See the proof of Theorem 3.16.) However, maximizing the Nash product means that we are maximizing the area of a certain rectangle inside K, so let us also provide another reasoning for statement (ii).

Note that the Nash product $(u_A - d_A)(u_B - d_B)$ is the area of the rectangle with corners at (d_A, d_B) and (u_A, u_B) (see Figure 8a). For simplicity, we assume that $d = (0,0)$.[8] So, we seek for a point P from the strong Pareto frontier of K which maximizes the area of this rectangle. Let us show that if $(u_A, g(u_A))$ is a point on the strong Pareto frontier of K such that

$$\frac{g(u_A)}{u_A} > -g'(u_A), \tag{3.3}$$

then the area of the rectangle (Nash product) is not maximized. (Note that the left hand side is the fraction given in the statement, since $d_A = d_B = 0$), Similar reasoning shows that the reverse inequality cannot hold at the Nash solution, and so we must have equality at the Nash solution.

[8] There is no loss of generality. We apply the transformation $T(x,y) = (x - d_A, y - d_B)$ to K. This transformation moves K to a new position, but it does not affect on areas.

So, suppose that the above inequality holds. Using the definition of a derivative, we see that for a small $\varepsilon > 0$ we have

$$\frac{g(u_A)}{u_A + \varepsilon} > -\frac{g(u_A + \varepsilon) - g(u_A)}{\varepsilon} = \frac{g(u_A) - g(u_A + \varepsilon)}{\varepsilon}.$$

Multiplying we get

$$\varepsilon g(u_A) > u_A g(u_A) - u_A g(u_A + \varepsilon) + \varepsilon g(u_A) - \varepsilon g(u_A + \varepsilon).$$

Canceling $\varepsilon g(u_A)$ and rearranging we have

$$(u_A + \varepsilon) g(u_A + \varepsilon) > u_A g(u_A).$$

In terms of areas given in Figure 8b, this means $A + B > A + C$, and so $B > C$. When we move from u_A to $u_A + \varepsilon$, the area decreases by C and increases by B. Therefore, the net change is $B - C$, which we have just shown to be positive. In conclusion, when $\varepsilon > 0$ is small enough, we can increase the area of the rectangle by moving from u_A (slightly) to the right, which proves the statement. □

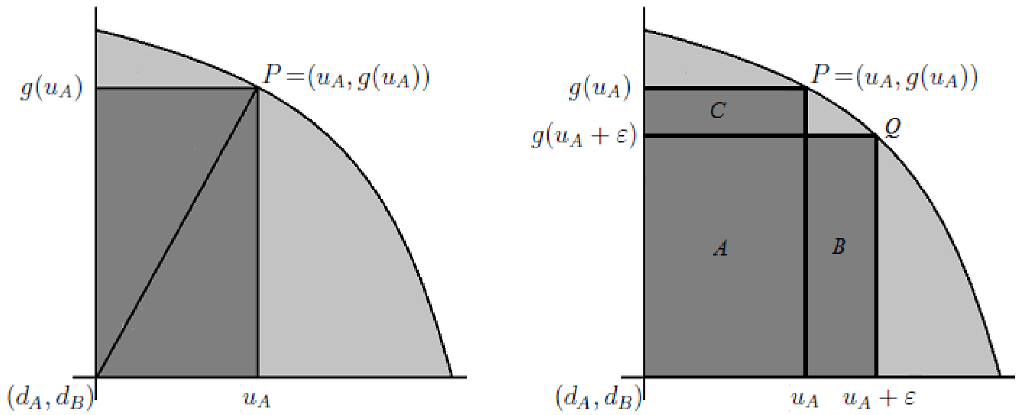

(a) Nash product as area

(b) Change in area

Figure 8: Nash solution

The fraction given in statement (ii) of the previous theorem is the slope of the line from the disagreement point to the point (u_A^*, u_B^*). Therefore, the previous theorem states that the Nash solution is given as that point P on the strong Pareto frontier of K where the line from d to P and the tangent line at P have the same absolute value of their slopes. In Figure 9 below, the point P is the Nash solution. At point Q, the absolute value of the slope of the tangent line is less than the slope of the line from d to Q, and at R we have the reverse order. At point P, the absolute values of the slopes are the same.

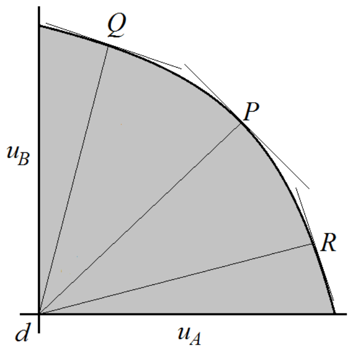

Figure 9: Nash solution

If the strong Pareto frontier of K is not differentiable, the Nash solution might be a corner point of K. One cannot use derivative to describe this point, but a similar condition as given above is still valid, see Muthoo (1999: 14–15).

Nash bargaining solution is convenient in applications as it is determined by a maximization problem. However, it is rather difficult to provide an interpretation for the maximization of the product of excess utilities. We finish this section by providing such an interpretation. In the next theorem, we concentrate on actual (physical) outcomes instead of utilities and we denote by D the disagreement event. Next statement is given in Muthoo (1999) and Osborne and Rubinstein (1990).

Theorem 3.14. *Suppose that x^* is some (physical) outcome with the following property: if x is another outcome such that one of the players strictly prefers the lottery $[\theta(x), (1-\theta)(D)]$ to x^* for some $\theta \in [0, 1]$, then the other player prefers the lottery $[\theta(x^*), (1-\theta)(D)]$ to x. Then the Nash solution to this bargaining problem yields the outcome x^*.*

Proof. Let x^* be as above and let the utility functions U_A and U_B be such that $U_A(D) = U_B(D) = 0$. To show that the Nash solution yields the outcome x^*, we need to show that x^* maximizes the Nash product $U_A(x)U_B(x)$.

Consider an arbitrary outcome x. If neither of the players prefers x to x^*, then $U_A(x^*)U_B(x^*) \geq U_A(x)U_B(x)$, and we are done. Suppose that Player A, say, strictly prefers x to x^*. Then $U_A(x) > U_A(x^*)$, and so $U_A(x^*)/U_A(x) < 1$. Pick some $0 < \theta < 1$ such that $U_A(x^*)/U_A(x) < \theta$, that is, $U_A(x^*) < \theta U_A(x)$. Since $U_A(D) = 0$, the value $\theta U_A(x)$ is the expected utility of Player A over the lottery $[\theta(x), (1-\theta)D]$. By assumption, Player B prefers the lottery $[\theta(x^*), (1-\theta)D]$ to x, so $U_B(x) \leq \theta U_B(x^*)$, and so $U_B(x)/U_B(x^*) \leq \theta$.

So, for any $0 < \theta < 1$ satisfying $U_A(x^*)/U_A(x) < \theta$, we have $U_B(x)/U_B(x^*) \leq \theta$. This means that we must have $U_B(x)/U_B(x^*) \leq U_A(x^*)/U_A(x)$. This yields the inequality $U_A(x)U_B(x) \leq U_A(x^*)U_B(x^*)$, which finishes the proof. □

The previous theorem may be explained as follows. Suppose that the offer x^* is "on the table". If Player A, say, is willing to object this outcome by offering another outcome x despite the risk that the negotiation breaks down with probability $1-\theta$, then Player B is willing to take a similar risk to get the outcome x^* instead of x. In other words, the Nash solution x^* represents a stable bargaining convention.

3.4 Asymmetric Nash solution

Nash bargaining solution established in the previous section puts the players in a symmetric position, as highlighted by Axiom **SYM**. According to this approach, any possible differences between the players must be captured in the bargaining problem (K, d), for example, by differences in their utility functions. To consider a more general approach which takes into account possible differences between bargaining abilities of the players, we treat in this short section an approach taken in Harsanyi and Selten (1971). The material in this section has similarities to the one presented in the previous section, so we only state the main results.

Definition 3.15. Let $0 < \alpha < 1$ and let (K, d) be a bargaining problem. The *generalized Nash bargaining solution* is given by

$$F^\alpha(K, d) = \operatorname*{arg\,max}_{(u_A, u_B) \in K(d)} (u_A - d_A)^\alpha (u_B - d_B)^{1-\alpha}.$$

The exponents α and $1 - \alpha$ are the *bargaining powers* of Player A and Player B, respectively.

We should mention that the term "bargaining power" does not refer to bargaining *skills*. As rational beings, both of the players will bargain as well as it is possible to bargain. Any differences in bargaining powers mean that one of the players has an *advantage* in the bargaining game. This could result, for example, from the fact that one of the players is very patient and the other one is anxious, but not, for example, from the fact that one of the players is more experienced negotiator.

Let $0 < \alpha < 1$ and let (K, d) be a bargaining problem. Consider the function

$f(x, y) = (x - d_A)^\alpha (y - d_B)^{1-\alpha}$, the *generalized Nash product.* This function is strictly quasi–concave on the set $\{u \in \mathbb{R}^2 : d \ll u\}$, and so f attains its maximum on $K(d)$ at a unique point. Also, the solution F^α satisfies axioms **EFF**, **INV**, and **IIA**; the proof of Theorem 3.7 applies from word to word. (As in Remark 3.12, we may also conclude that $F^\alpha(K, d)$ is actually in the strong Pareto frontier of K.) However, F^α satisfies **SYM** if and only if $\alpha = 1/2$ (in which case $\alpha = 1-\alpha$). Note that if $\alpha = 1/2$, then the generalized Nash product is nothing but the square root of the usual Nash product, and so they attain the maximum value at the same point. In conclusion, $F^{1/2}$ is the Nash solution described in the previous section.

Again, we obtain a geometric condition for the generalized Nash solution.

Theorem 3.16. *Let* $0 < \alpha < 1$*, let* (K, d) *be a bargaining problem, and let the strong Pareto frontier of* K *be given by the equation* $u_2 = g(u_1)$*. If* g *is differentiable, then* $F^\alpha(K, d) = (u_A^*, u_B^*)$ *if and only if*

$$g(u_A^*) = u_B^* \quad \textit{and} \quad |g'(u_A^*)| = \frac{\alpha}{1-\alpha} \cdot \frac{u_B^* - d_B}{u_A^* - d_A}. \tag{3.4}$$

Proof. The first equation merely says that the outcome (u_A^*, u_B^*) is on the strong Pareto frontier of K. To verify the second equation, note that the solution is, then, obtained by finding the maximizer of $[u_A - d_A]^\alpha [g(u_A) - d_B]^{1-\alpha}$. Since logarithm is an increasing function, we may solve this problem by finding the maximizer of $\alpha \cdot \ln(u_A - d_A) + (1-\alpha) \cdot \ln(g(u_A) - d_B)$. The maximizer u_A^* satisfies the first order condition

$$\alpha \cdot \frac{1}{u_A^* - d_A} + (1-\alpha) \cdot \frac{g'(u_A^*)}{g(u_A^*) - d_B} = 0.$$

Solving $g'(u_A^*)$ from this equation, we obtain the second condition in Equation (3.4), as required. □

As the Nash solution presented in the previous section, the generalized Nash solutions admits also a characterization through the introduced axioms.

Theorem 3.17. *A bargaining solution* F *satisfies* ***EFF****,* ***INV****,* ***IIA****, if and only if there exists* $0 < \alpha < 1$ *such that* $F = F^\alpha$*.*

Proof. See Napel (2002: 13–16) or Binmore (2007b: 471–479). (In these texts, efficiency is defined using the strong Pareto frontier of a bargaining problem. As noted above, $F^\alpha(K, d)$ is on the strong Pareto frontier of K.) □

Example 3.18. Let us consider once again the problem of partitioning a cake. Let the utility functions be as in Example 3.10, but this time let us also consider the bargaining power α. Now, we need to find the maximizer of the function

$$(x_A^{\gamma})^{\alpha}[(1-x_A)^{\delta}]^{1-\alpha} = x_A^{\gamma\alpha}(1-x_A)^{\delta(1-\alpha)}.$$

This is the function maximized in Example 3.10 with γ replaced by $\gamma\alpha$ and δ replaced by $\delta(1-\alpha)$. Therefore, the solution is

$$x_A = \frac{\gamma\alpha}{\gamma\alpha + \delta(1-\alpha)} \qquad \text{and} \qquad x_B = \frac{\delta(1-\alpha)}{\gamma\alpha + \delta(1-\alpha)}.$$

Note again that

$$\frac{x_A}{x_B} = \frac{\gamma\alpha}{\delta(1-\alpha)} = \frac{\gamma}{\delta} \cdot \frac{\alpha}{1-\alpha}.$$

If $\alpha = 1/2$, the term $\alpha/(1-\alpha)$ equals to 1, and we have the same solution as in Example 3.10, as we should. Otherwise, the division favors (as compared to the division obtained in Example 3.10) the player with higher bargaining power. For example, let $\gamma = 0.3$, $\delta = 0.9$, and $\alpha = 0.75$. Then $\gamma/\delta = 1/3$, and with equal bargaining powers Player A will get only 1/4 of the cake. When we substitute $\alpha = 0.75$ to the above expression, we see that when we take the bargaining power into account, Player A and Player B split the cake equally.

Note that the term $\alpha/(1-\alpha)$ approaches to infinity as α tends to 1. This means that even a highly risk averse player can obtain a major share of cake if he (or she) has high enough bargaining power, for example, being much more patient than the other player.

Example 3.19. Let us also find the Nash bargaining solution for the problem of partitioning a cake with general disagreement point $d = (d_A, d_B)$ in the case that the utility functions are $U_A(x) = U_B(x) = x$. We could solve a similar maximization problem as above, but we may also apply Theorem 3.16. Note that, in this case, the equation of the strong Pareto frontier is $g(u_A) = 1 - u_A$, and so $g'(u_A) = -1$. Then it is easy to verify that

$$x_A = d_A + \alpha \cdot (1 - d_A - d_B) \quad \text{and} \quad x_B = d_B + (1-\alpha) \cdot (1 - d_A - d_B)$$

satisfy Equation (3.4). (Note that $u_A = x_A$.) This partition admits a natural interpretation, namely that both of the players first obtain their utilities in *status quo*, and the rest of the cake is split in the ratio of bargaining powers.

3.5 Moral hazard

In this section, we present a simple application of Nash bargaining solution to working in teams. This example is given in Muthoo (1999). For other uses of Nash bargaining solution, see Muthoo (1999: 16–29) for applications on crime control, optimal asset ownership, and union–firm negotiations, or Montet & Serra (2003: 242–246) for applications on bilateral monopoly and union–firm negotiations.

Let us treat a situation where Player A and Player B are working on a joint project. The players may decide how much effort they use. We assume that if the players choose effort levels e_A and e_B, respectively, then the output is given by $Q = 2e_A^{1/2}e_B^{1/2}$. Furthermore, we assume that the cost of an effort level e_i to Player i, $i = A, B$, is $C_i = \alpha_i e_i^2/2$ for some constant $\alpha_i > 0$.

What follows is a two–stage game. In the first stage, the players bargain over how to divide the output they produce. If they find an agreement, then, in the second stage the players choose how much effort they put in the project, and they divide the resulting output according to what they agreed in the first stage. So, the players choose their effort levels *after* they have agreed how to divide the output, and so they might work hard or shirk off depending on the share of output they receive. (The players are unable to make a binding contract on how much effort they use in the project.)

We proceed using backward induction, that is, we "solve" the second part of the game first. Suppose that the players have agreed to divide the output in such a way that Player A obtains the share x and Player B obtains the share $1 - x$, where $0 < x < 1$. Let us determine how much effort the players are willing to put in the project. If Player B chooses effort level e_B, Player A wishes to find e_A such that his (or her) profit (output received minus cost)

$$x \cdot 2e_A^{1/2}e_B^{1/2} - \alpha_A \cdot \frac{e_A^2}{2}$$

is maximized. The first order condition gives the equation

$$xe_A^{-1/2}e_B^{1/2} = \alpha_A e_A,$$

which determines the optimal effort level e_A^* for Player A as

$$e_A^* = \frac{x^{2/3} e_B^{1/3}}{\alpha_A^{2/3}}. \tag{3.5}$$

Note that this is the best response of Player A to the effort level e_B. Similarly, we find that the best response function of Player B is given by

$$e_B^* = \frac{(1-x)^{2/3} e_A^{1/3}}{\alpha_B^{2/3}}.$$

In equilibrium, these values are best responses to each others. Substituting e_B^* into Equation (3.5) and solving for e_A^*, we find that

$$e_A^* = \frac{x^{3/4}(1-x)^{1/4}}{\alpha_A^{3/4}\alpha_B^{1/4}} \quad \text{and} \quad e_B^* = \frac{x^{1/4}(1-x)^{3/4}}{\alpha_A^{1/4}\alpha_B^{3/4}}. \tag{3.6}$$

With these values, the output is given as

$$Q = 2 \cdot \frac{x^{3/8}(1-x)^{1/8}}{\alpha_A^{3/8}\alpha_B^{1/8}} \cdot \frac{x^{1/8}(1-x)^{3/8}}{\alpha_A^{1/8}\alpha_B^{3/8}} = 2 \cdot \frac{x^{1/2}(1-x)^{1/2}}{\alpha_A^{1/2}\alpha_B^{1/2}},$$

and the costs for the players are

$$C_A = \frac{x^{3/2}(1-x)^{1/2}}{2\alpha_A^{1/2}\alpha_B^{1/2}} \quad \text{and} \quad C_B = \frac{x^{1/2}(1-x)^{3/2}}{2\alpha_A^{1/2}\alpha_B^{1/2}}, \tag{3.7}$$

respectively. These values imply that the players' profits are

$$U_A(x) = \beta \cdot x^{3/2}(1-x)^{1/2} \quad \text{and} \quad U_B(x) = \beta \cdot x^{1/2}(1-x)^{3/2},$$

respectively, where $\beta = 3/(2\alpha_A^{1/2}\alpha_B^{1/2})$.

Let us apply the Nash solution to determine the number x above. Calculating the derivatives $U_A(x)$ and $U_B(x)$, we find that $U_A(x)$ is increasing in the interval $[0, 3/4]$, decreasing in the interval $[3/4, 1]$, and so attains its maximum at $x = 3/4$, and that $U_B(x)$ is increasing in the interval $[0, 1/4]$, decreasing in the interval $[1/4, 1]$, and so attains its maximum at $x = 1/4$. Therefore, efficiency requires that the bargaining solution is in the interval $[1/4, 3/4]$. Since disagreement in bargaining produces no output, the Nash product is

$$U_A(x)U_B(x) = \beta^2 x^2 (1-x)^2.$$

The first order condition is $4x^3 - 6x^2 + 2x = 0$. This equation has three solutions, namely $x = 0$, $x = 1/2$, and $x = 1$, but the only solution satisfying the efficiency criterion is $x = 1/2$. In conclusion, the output is divided equally between the players according to the Nash bargaining solution.

Note that, with $x = 1/2$, Equation (3.6) gives

$$e_A^* = \frac{1}{2} \cdot \frac{1}{\alpha_A^{3/4} \alpha_B^{1/4}} \quad \text{and} \quad e_B^* = \frac{1}{2} \cdot \frac{1}{\alpha_A^{1/4} \alpha_B^{3/4}},$$

and Equation (3.7) shows that the costs for the players are the same. Since the players divide the output equally and choose their effort levels in such a way that both bear the same cost, the players obtain equal profits. Note also that the ratio of effort levels is

$$\frac{e_A^*}{e_B^*} = \frac{\alpha_A^{1/4} \alpha_B^{3/4}}{\alpha_A^{3/4} \alpha_B^{1/4}} = (\frac{\alpha_B}{\alpha_A})^{1/2}.$$

If, for example, Player A faces a higher cost of participating, that is, $\alpha_A > \alpha_B$, then Player A puts less effort in the project than Player B, but Player A still obtains half of the output. We may interpret this in a way that both of the players know that without his (or her) contribution to the project, the output will be zero. Therefore, the players are able to require half of the output, not by supplying equal amount of work, but by providing an amount of work which equalizes players' profits.

3.6 Some additional comments

In this chapter, we have focused on bargaining problems of two players. Nash bargaining solution can be extended to bargaining problems of N players using a straightforward extension of the Nash product as follows:

$$\prod_{k=1}^{N} (u_i - d_i) = (u_1 - d_1) \cdot (u_2 - c_2) \cdot \ldots \cdot (u_N - d_N).$$

In the case of $N > 2$ players, there is a possibility that some of the players form a *coalition*. The above given Nash product is a reasonable candidate for a bargaining solution only in the case that such coalition forming does not occur, see Napel (2002: 21)

As noted, some of the axioms used by Nash are troublesome, and there are a number of other axiomatic solutions to bargaining problems. The texts Peters (1992) and Roth (1979) contain a detailed discussions about different axiomatic models of bargaining, and Binmore (2007b: 482–485) and Napel (2002: 22–25) provide short discussions of some axiomatic bargaining models.

In the proof of Theorem 3.7, we used all the four axioms to deduce that the Nash solution is the unique solution satisfying these axioms. It is natural to ask whether the uniqueness of the solution actually requires all these axioms. That is, if we exclude any of the four axioms, is the Nash solution still the unique solution satisfying the rest of the axioms. In all the cases, the answer is negative, see Osborne and Rubinstein (1990: 20–23) for details. Of particular interest is the solution which maximizes the *sum* of the players' excess utilities instead of product as in Nash solution. (A further condition is required to guarantee that a unique member of K is obtained.) This solution satisfies Axioms **EFF**, **IIA**, and **SYM**, but not **INV**.

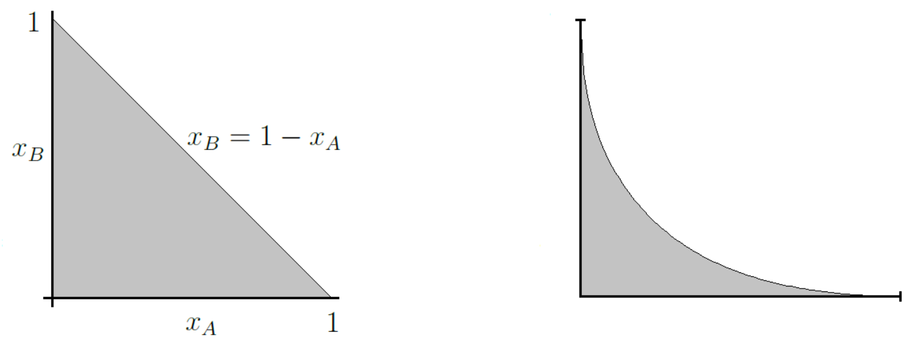

(a) Possible agreements (b) Attainable set

Figure 10: Partition of a cake for risk loving players

In the given examples on the partition of a cake, we have assumed that the players are risk averse. (This is shown by their utility functions.) What about if the players are risk loving? For example, let the utility functions be given by $U_A(x) = U_B(x) = x^2$. We could do our calculations again to find when the Nash product $u_A u_B = x^2(1-x)^2$ is maximized, but the problem in here is that this approach is completely wrong from the very beginning. When the players are risk loving, the set of possible outcomes is not turned into a convex set under the utility functions, and so we have no hope of applying Nash's solution. Indeed, if $u_A = x_A^2$ for some $0 \leq x_A \leq 1$, then the share of cake obtained by Player B is at most $x_B = 1 - x_A = 1 - \sqrt{u_A}$, and so his (or her) utility is at most $(1 - \sqrt{u_A})^2$.

The attainable set is given as in Figure 10b, and it clearly is not convex.

We can take the convex hull of the attainable set (this is the smallest convex set containing the given set) as in Figure 11 and apply Nash solution to this set. This set is symmetric, so the Nash solution is given by $(1/2, 1/2)$. The solution, however, is not anymore a member of K, so how do we interpret this solution? The answer is that this is a lottery where both of the players have probability 1/2 of obtaining the whole cake. That is, risk loving players rather gamble for the whole cake than settle for a smaller share of the cake. See Binmore (2007b: 194-195) for further details.

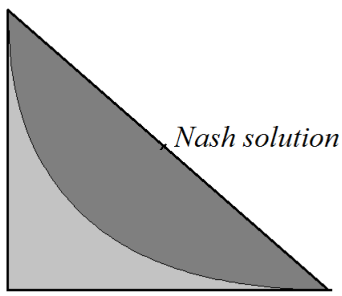

Figure 11: Nash solution

There is one more property of the Nash solution we consider worth of mentioning. Axiom **EFF** represents *collective rationality* in the sense that the players will not agree on an outcome if there is a possibility of making *both* of the players better off. However, A. Roth has demonstrated that one may replace this assumption by certain assumption on *individual rationality* and the content of Nash's theorem still remains valid. This is an interesting extension of Nash's result, since one usually considers individual rationality as a more elementary requirement than collective rationality. The assumption Roth used to prove his claim is given as follows. (He used the term *individual rationality* for the next axiom; the following term is used in Osborne and Rubinstein (1990).)

SIR (*Strong Individual Rationality*) $F(K, d) > d$ for every bargaining problem.

Note that this axiom means that *one* of the players benefits from the bargaining process. The following statement was proved in Roth (1977).

Theorem 3.20. *The Nash solution is the unique bargaining solution satisfying*

***SIR**, **SYM**, **INV**, and **IIA**.*

When we compare this result with Theorem 3.7, it follows that the four axioms given above together imply Axiom **EFF**. (This is actually what Roth proves in his paper.)

Finally, we mention that the disagreement point in the Nash solution is fixed throughout bargaining. In Nash (1953), Nash considered a variable threat game where the players first choose some threat points, and then start bargaining. If they cannot find agreement via bargaining, then the outcome is given by the Nash solution of bargaining problem where disagreement point is given by the chosen threats. See Binmore (2007b: 500–502) or Owen (2008: 198–201) for a further discussion of this game.

4 ALTERNATING OFFERS

We began our study of bargaining in the previous chapter by describing Nash's bargaining solution. As noted, the axiomatic approach to bargaining situations is very general, since the actual bargaining procedure is not taken into account in any way in the solution, but the drawback of this approach is that the assumptions on which the solution is based might be too restrictive to describe real bargaining situations. For one thing, axiom **EFF** implies that if bargaining entails some costs, we should expect an immediate agreement. This does not seem to describe very well, for example, bargaining situations between employees and labor unions, as many strikes have occurred during these negotiations.

In this chapter, we turn our attention to strategic models of bargaining. This means that we model the bargaining process between the players as some kind of game. A central theme of this chapter is the model of *Alternating Offers* which, roughly, means that the players make offers and counteroffers about the division of a cake until agreement is reached. Alternating Offers has drawn a lot of attention and a large number of different variations of this model have been studied; we concentrate only on a small number of these.

In analyzing the resulting games, we are mostly interested in the subgame perfect equilibria. As we shall see, the concept of Nash equilibrium is not very powerful tool in analyzing these games (or bargaining games in general), as there are simply far too many Nash equilibria. In the game of Alternating Offers, Nash equilibrium tells that we can expect any outcome to occur. This is not an illustrative result, and therefore we need to apply the more delicate concept of subgame perfect equilibrium to restrict the set of outcomes. This is, indeed, a successful approach, as in many cases we find that there exists a *unique* subgame perfect equilibrium.

This chapter serves two purposes. First of all, we wish to present various models of bargaining to describe how different factors affecting on bargaining situations can be taken into account in models of bargaining. Second, we wish to find some support from these models to how and when we may apply Nash bargaining solution. This interplay between axiomatic and strategic models of bargaining is known as the *Nash program*. As described in Nash (1953): "The two approaches to the problem, via the negotiation model or via the axioms, are complementary; each helps to justify and clarify the other."

In the first section, we describe the Ultimatum game, which is the simplest form of Alternating Offers with only one offer made, and determine the unique subgame perfect equilibrium of this game. In the Ultimatum game, it is assumed that the time spent on bargaining does not matter. In the next two sections, we remove this assumption by studying a model of Alternating Offers between impatient players. The second section contains a finite horizon model, which means that there is some number N of offers and counteroffers. The most important results of this chapter are given in Section 3, where we determine the subgame perfect equilibrium of an infinite horizon model of Alternating Offers. This solution, originally given in Rubinstein (1982), is one of the cornerstones of game theory, since this is a first description of subgame perfect equilibria in an infinite horizon model where backward induction does not work.

Rubinstein's model implicitly assumes that the only option the players have is to bargain with each others and that it is impatience which drives the players towards agreement. In the next sections, we relax these assumptions and add some other factors into the model. We deal with uncertainty (Section 4), outside options (Section 5), inside options (Section 6), and commitment tactics (Section 7). In many cases we find that the original Rubinstein's solution somehow affects the equilibrium partition of cake. The models up to Section 7 assume complete information, and we find that the players always reach immediate agreement. In Section 8, we study a model of incomplete information, which somehow explains a delay in reaching an agreement.

In all the models considered in this chapter, the underlying idea is division of a cake of size one. We apply the following assumptions throughout this chapter.

Assumption 1: Cake is desirable, that is, the players always prefer a larger share of cake to smaller share.

Assumption 2: No cake is wasted, that is, if Player A obtains a share x of cake, then Player B obtains the share $1 - x$.

4.1 Ultimatum game

The simplest form of Alternating Offers is the *Ultimatum game*, where only one offer is made. In what follows, we interpret every offer made by either of the players as a demand on a share of cake. So, either of the players demanding a share x of cake is as this player saying: "I'll take share x and You get the rest". When Player A demands a share x, Player B either accepts or rejects this offer. If Player B rejects the offer, then neither of the players receives anything. In extensive form, this game is given as follows.

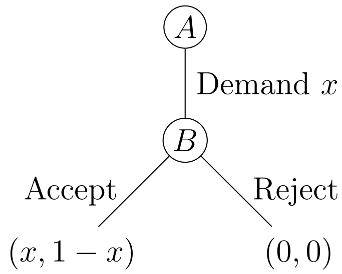

Figure 12: Ultimatum game

In any outcome (u, v) (like $(x, 1-x)$ above), the first and the second coordinates are always the payoffs to Player A and Player B, respectively.

The following statement, given in Tadelis (2013: 223), shows that the concept of Nash equilibrium is not very powerful tool in analyzing bargaining problems. Indeed, using only Nash equilibrium, we may expect any outcome to occur in the Ultimatum game.

Proposition 4.1. *Let $0 \leq s \leq 1$. Then there exists a Nash equilibrium in the Ultimatum game such that the outcome is $(s, 1-s)$.*

Proof. Let $0 \leq s \leq 1$ and consider the following strategies:
A: Demand the share s.
B: Accept a demand x if and only if $1 - x \geq 1 - s$.
So, Player B will accept an offer if and only if his (or her) share of cake is at least $1-s$. These strategies form a Nash equilibrium. If Player A demands the share s, then Player B is playing optimally since he (or she) gets the share $1-s$ instead of 0 given by rejecting. (Note that if $s = 1$, the strategy given above says that Player B will also accept this demand.) Player B would also accept any demand less than s, but this would decrease the payoff of Player A, so it is not optimal for Player A to demand less than s. If Player A demands more than s, Player B will reject the offer and the outcome is 0. In conclusion, s is the optimal demand for Player A. When Player A demands the share s and Player B accepts, the outcome is $(s, 1-s)$. □

The reasoning given above is valid in the sense that it verifies the given statement, but it has a serious drawback which is related to the subgame perfect equilibrium. To see what this means, let $s = 1/2$. The above given strategy for Player B is like Player B saying before the game: "I will only accept an offer which gives me at least half of the cake." Now, suppose that Player A demands the share $3/4$. It is not rational for Player B to execute the threat of rejection, since this gives the outcome 0 instead of $1/4$. (This is where the concept of subgame perfect equilibrium emerges: the strategy of Player B is not optimal in the *subgame* which begins when Player B makes his (or her) decision.) The same reasoning applies to any $s > 0$, which means that (rational) Player B will accept any positive share of cake.[9] This fact puts Player B in a very poor position, as is illustrated by the following theorem. Note that using subgame perfect equilibrium instead of Nash equilibrium, we narrow the set of possible outcomes significantly.

Theorem 4.2. *In the Ultimatum game, the following pair of strategies*
A: Demand the share** 1 **(the whole cake).
B: Accept any offer.
is the unique subgame perfect equilibrium. The outcome of this equilibrium is that Player A obtains the whole cake.

Proof. It is easy to verify that the above given strategies form a subgame perfect equilibrium. Indeed, both of the players are playing optimally. Since Player B accepts any offer, it is optimal for Player A to demand the whole cake. Also, when Player A demands the whole cake, Player B cannot increase his (or her) outcome by rejecting some offers.

To show that the given pair of strategies is the only subgame perfect equilibrium, it is enough to show that Player B will accept the demand 1 in any subgame perfect equilibrium. We already argued that Player B will accept any positive share in a subgame perfect equilibrium. This, in fact, assures that Player B will also accept the demand 1 in any subgame perfect equilibrium. To see why, suppose that Player B rejects this offer. Then the outcome is zero for both of the players, and Player A would benefit by making an offer $(1-\varepsilon, \varepsilon)$ for some $\varepsilon > 0$ instead of $(1, 0)$. However, since Player B will accept any positive share, Player A does not

[9] Here, theory and experience are in a serious conflict. It seems that people usually expect a "fair" division of cake, and they reject small shares. Also, the offers made by the first player are significant shares instead of 0, see Tadelis (2013: 223–224) or Binmore (2007a: 103–117). Binmore also finds that it is the one–stage Ultimatum game where the difference between theory and experience is strict.

have a best response to the strategy of Player B. For any $\varepsilon > 0$, Player A would do better by offering $\varepsilon/2$ to Player B instead. The lack of best response contradicts the definition of a subgame perfect equilibrium, and finishes our proof. □

Ultimatum game does not seem to have much of "bargaining", since Player B can only accept or reject an offer made by Player A. Let us see what happens if we try to include "more of bargaining" by allowing Player B to make a counteroffer. In this case, the extensive form of the game is given in Figure 13. (This game is called *two–stage Ultimatum Game.*)

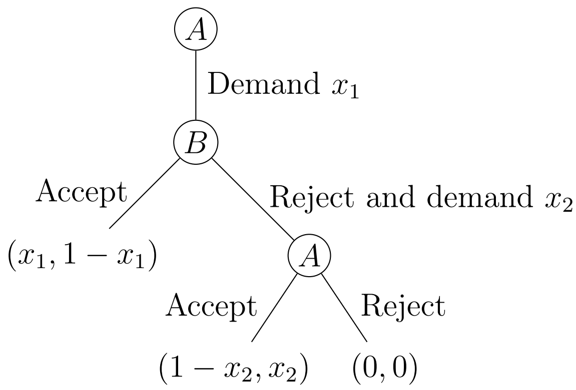

Figure 13: Two–stage ultimatum game

The possibility for Player B to make a counteroffer puts all the bargaining power on Player B. Indeed, if Player B rejects the offer made by Player A, then the resulting game is nothing but the Ultimatum game with Player B as the first player. Therefore, in a subgame perfect equilibrium, Player B will obtain the whole cake.

Now, consider an Ultimatum game with N rounds for some integer N, that is, the players make a total of N offers each in turn. What we just noticed about the case $N = 2$, it should be obvious that the player who makes the *last* offer has all the bargaining power. This game has several (depending on N) subgame perfect equilibria, but the outcome is always the same, namely that the last player to make an offer takes the whole cake. (The player to make the last offer certainly obtains the whole cake in the last round, but he (or she) might also obtain the whole cake on some earlier round.)

Let us think what happens if we allow the players to negotiate over undefined (infinite) number of rounds. This game has no last player advantage, but we might guess that every time the player in turn to make an offer demands the whole cake.

Furthermore, in this game the players do not have a proper incentive to reach an agreement: the time of agreement does not matter for the players, which seems somewhat disturbing. As put in Cross (1965): "If it did not matter when people agreed, it would not matter whether or not they agreed at all." Therefore, we need to add some "friction" between the players to give them an incentive to reach an agreement. In the next section, we create this friction by assuming that the players are impatient.

4.2 Alternating offers over finite horizon

In this and the next section, we consider a model of Alternating Offers where we assume that the players are impatient. We divide this model into two parts, the one where the players have some fixed total number N of offers to use, and another where the bargaining process could, at least in principle, continue for an arbitrary long time. We consider the first model in this section and the second one in the next section.

To include the time used in bargaining into the model, we assume that each player makes one offer in one unit of time. The game begins at time $t = 0$ with an offer made by Player A. If Player B accepts this offer, the game ends. Otherwise, the game proceeds to the next round with time $t = 1$ where Player B makes an offer. If Player A accepts this, the game ends. Otherwise, the game proceeds to the next round with time $t = 2$ where Player A again makes an offer, and so on. (Note that an offer and an acceptance occur during the same unit of time; time t increases only after a rejection.)

We assume that the players have discount factors (per unit of time) $\delta_A < 1$ and $\delta_B < 1$, respectively. Therefore, the payoffs to Player A and Player B after an agreement $(x, 1-x)$ on round n are given by

$$u_A = \delta_A^n x \qquad \text{and} \qquad u_B = \delta_B^n (1-x),$$

respectively. Furthermore, if the players fail to reach an agreement during the N rounds, then we assume that the outcome is $(0, 0)$. The first three rounds of this game are described in Figure 14 below.

Since $\delta_A, \delta_B < 1$, we see that the payoffs get smaller (size of the cake decreases)

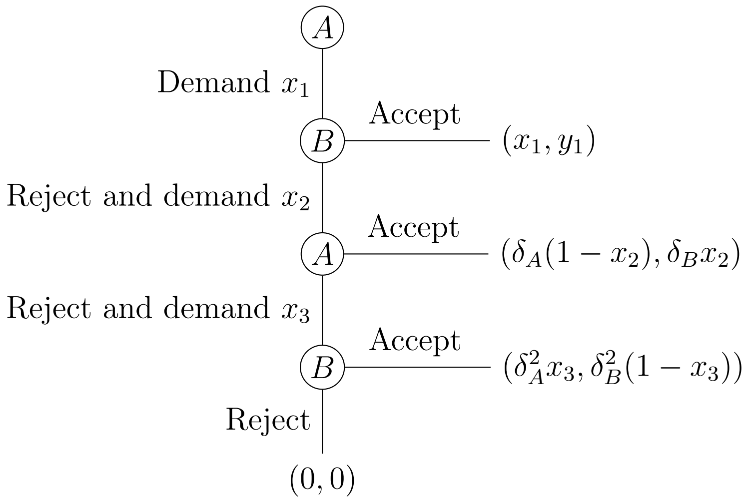

Figure 14: Alternating Offers with three rounds

as n increases. Therefore, the players have an incentive to reach an agreement as soon as possible, and this is what happens in a subgame perfect equilibrium.

The next theorem was proved in Ståhl (1972). We should emphasize that this is a model of perfect information, which means that both of the players know both of the discount factors δ_A and δ_B and the number N of rounds.

Theorem 4.3. *The game of Alternating Offers with N rounds has a unique subgame perfect equilibrium where the players reach an immediate agreement. If N is even, so that $N = 2k$ for some integer k, then the share obtained by Player A is*

$$x_A^* = 1 - \delta_B \cdot \frac{1-(\delta_A\delta_B)^k}{1-\delta_A\delta_B} + \delta_A\delta_B \cdot \frac{1-(\delta_A\delta_B)^{k-1}}{1-\delta_A\delta_B} \tag{4.1}$$

If N is odd, so that $N = 2k+1$ for some integer k, then the share obtained by Player A is

$$x_A^* = 1 - \delta_B \cdot \frac{1-(\delta_A\delta_B)^k}{1-\delta_A\delta_B} + \delta_A\delta_B \cdot \frac{1-(\delta_A\delta_B)^k}{1-\delta_A\delta_B}. \tag{4.2}$$

Proof. See Napel (2002: 28–29) for a full proof. We prove the statement in the case that $N = 3$, so that $k = 1$. This illustrates the use of *backward induction*. In this case, the extensive form is given as in Figure 14 above.

Let us calculate the optimal offers made by each player. In a subgame perfect equilibrium, both of the players adjust their demands in such a way that the other player is indifferent between accepting and rejecting. (The same reasoning as in the proof of Theorem 4.2 applies to show that, in a subgame perfect equilibrium, a player will accept an offer if he (or she) is indifferent between accepting and rejecting.)

First, note that if the play reaches the point where Player A makes his (or her) last demand x_3, the resulting game is the Ultimatum game, and it is optimal for Player A to demand the share 1. This gives the outcome $(\delta_A^2, 0)$.

Next, Player B knows what will happen if Player A gets to make the last offer, so Player B will adjust the demand x_2 in such a way that Player A is indifferent between accepting and rejecting this offer. Rejecting would give payoff δ_A^2 to Player A by what we have just reasoned. Therefore, Player B will demand the share x_2 which satisfies

$$\delta_A(1 - x_2) = \delta_A^2.$$

The solution is $x_2 = 1 - \delta_A$, which gives Player B the payoff $\delta_B(1 - \delta_A)$.

Finally, Player A knows that Player B will make the above demand at the second round of play (if reached), and so Player A will adjust the demand x_1 in such a way that Player B is indifferent between accepting and rejecting. This gives the equation

$$1 - x_1 = \delta_B(1 - \delta_A).$$

The solution is $x_1 = 1 - \delta_B + \delta_A\delta_B$, which is Equation 4.2 for $k = 1$. Since the players are indifferent between accepting and rejecting at every point of the game, Player B accepts this demand in the first round. □

Note that the term $(\delta_A\delta_B)^k$ tends to zero when k increases, which means that the share obtained by Player A decreases as k (and N) increases. In fact, whether N is even or odd, we have

$$\lim_{k\to\infty} x_A^* = 1 - \delta_B \cdot \frac{1}{1 - \delta_A\delta_B} + \delta_A\delta_B \cdot \frac{1}{1 - \delta_A\delta_B} = \frac{1 - \delta_B}{1 - \delta_A\delta_B},$$

which means that Player A obtains approximately this share for large number of rounds. For example, if both of the players use 5 % of discounting, so that $\delta_A = \delta_B = 0.9524$, the above value $x_A^* = 0.51$, and so Player A obtains only slightly over half of the cake if N is large.

This game still has the last player advantage, but the fact that time now matters for the players gives some bargaining power also to the other player. Suppose, for example, that Player A makes the last offer. When N is small, the threat of Player B to delay the agreement (by rejecting or making high demands) is not credible, since Player A can obtain a high utility by "stealing" the cake in the last round. However, when N (and k) increases, this threat becomes more credible, since a long delay in the game would also decrease the payoff to Player A substantially. Therefore, Player A may just as well offer a proper share to Player B in order to reach a quick agreement instead of waiting until the last round to get the whole cake.

We have just described what happens in the game of Alternating Offers when the number of rounds becomes large. A problem with this model is that the number of offers is predetermined in the bargaining situation. This raises the question who determines this number. It is difficult to believe that this number is set by the bargainers. Next, we turn our attention to a version of Alternating Offers where the number of rounds is not (somewhat artificially) determined as above.

4.3 Alternating offers over infinite horizon

In this section, we consider the same game as in the previous section with the exception that the number of offers is not bounded (see Figure 15 below). This model was first solved in Rubinstein (1982). Although it might seem unrealistic to give the players an infinite number of offers to use, the model of this section is perhaps more realistic than in the previous section, since the fixed number N of rounds used in the previous section is somewhat hard to justify.

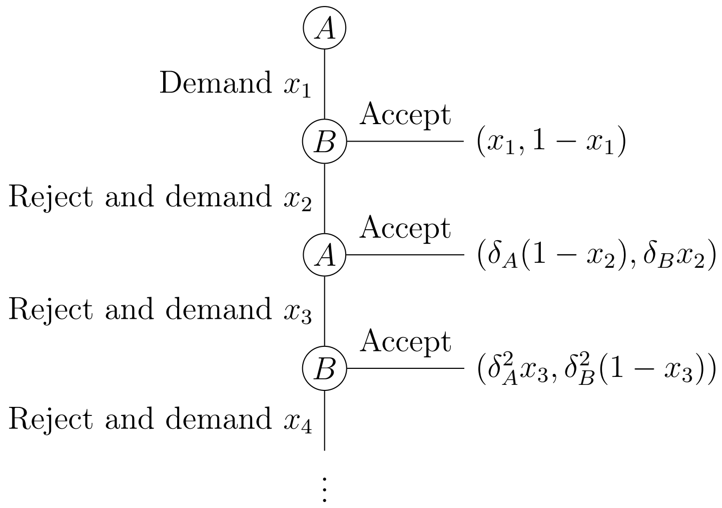

Figure 15: Alternating Offers

We turn right a way to the main result of this section. Note that, as in the previous section, the players still have a strong incentive to reach a quick agreement. Again, this is a game of complete information, that is, the players know the discount factors of each others.

Theorem 4.4. *The game of Alternating Offers over infinite horizon has a unique subgame perfect equilibrium where Player A and Player B always demand the following shares*

$$x_A^* = \frac{1-\delta_B}{1-\delta_A\delta_B} \qquad \textit{and} \qquad x_B^* = \frac{1-\delta_A}{1-\delta_A\delta_B},$$

respectively, and accept only offers giving at least the shares $\delta_A x_A^*$ *and* $\delta_B x_B^*$*, respectively. The outcome of this equilibrium is an agreement on the first round, and Player A and Player B obtain the shares* x_A^* *and* $\delta_B x_B^*$*, respectively.*

The division given above is the one obtained in the end of the previous section by taking the limit when N tends to infinity. However, this simple reasoning is not a valid proof for the statement above. The games of this and the previous section are of completely different nature since, this time, we have no last round from where to begin backward induction.

Let us first discuss the previous statement shortly before giving a proper proof. Of particular interest in the previous statement is the fact that the players' strategies in the unique subgame perfect equilibrium are *stationary*, that is, these strategies do not depend in any way on the round when a player makes a decision. The structure of the game allows players to condition their actions on the history of the game. However, in the unique subgame perfect equilibrium the players make no such conditioning. Let us give a heuristic reasoning why this is so.

Suppose that Player A has optimal first demand. The previous statement claims that, in the subgame perfect equilibrium, it is optimal for Player A to make the same demand on every round. Suppose that, for some reason, the play reaches round three, which means that both of the players have rejected one offer from each others. When Player A looks at the game beginning in round three, *nothing has changed from his (or her) perspective.* Since there is no last offer, Player A faces precisely the same game as in round one. (This is not true in the game studied in the previous section. If the game reaches round three, Player A has one less offer in the resulting game than in the beginning of the game.) The players' outcomes depend only on the share of cake they receive and on the round the agreement is reached. This means that Player A does not feel regret for having made a bad offer on round one or upset for having his (or her) offer rejected. As a rational agent, Player A does his (or her) best on round three to maximize his (or her) payoff from the resulting game. There is still the same amount of cake to be divided, and Player A wishes to make an offer that, first, Player B would accept, and second, gives Player B as small share as possible. But if Player A has such an offer on round one, then the same offer is what Player A should offer on round three. In conclusion, if Player A has an optimal offer on round one, then this is the optimal offer on every round.

We give a proof, following Fudenberg and Tirole (2004: 114–116), to demonstrate how the uniqueness of a subgame perfect equilibrium can be deduced in an infinite horizon game where backward induction is powerless. Similar reasoning applies also in later results presented. The reader may also consult Tadelis (2013: 228) for a proof in the case that $\delta_1 = \delta_2$.

Proof. First, let us show that the above given strategies form a subgame perfect equilibrium. Note that the above demands x_A^* and x_B^* satisfy

$$1 - x_A^* = \frac{1 - \delta_A\delta_B - (1 - \delta_B)}{1 - \delta_A\delta_B} = \delta_B \cdot \frac{1 - \delta_A}{1 - \delta_A\delta_B} = \delta_B x_B^* \tag{4.3}$$

and

$$1 - x_B^* = \frac{1 - \delta_A\delta_B - (1 - \delta_A)}{1 - \delta_A\delta_B} = \delta_A \cdot \frac{1 - \delta_B}{1 - \delta_A\delta_B} = \delta_A x_A^*. \tag{4.4}$$

Therefore, both of the demands leave the other player just the amount he (or she) is willing to accept, and so neither of the players has a reason to demand for a smaller share.

We need also to verify that neither of the players benefits by demanding a higher share of cake. At this point, we apply Theorem 2.9. Suppose that, for example, Player A makes a demand higher than x_A^* on some round n. Player B will reject and offer the share $1-x_B^*$ to Player A in the next round. Player A will accept this (since he (or she) deviates from the above strategy only on round n) and obtain the outcome $\delta_A^{n+1}(1 - x_B^*)$. But making demand x_A^* on round n, which Player B would accept, gives Player A the outcome $\delta_A^n x_A^*$. Equation (4.4) and $\delta_A < 1$ imply that

$$\delta_A^{n+1}(1 - x_B^*) = \delta_A^{n+1} \cdot \delta_A x_A^* = \delta_A^{n+2} x_A^* < \delta_A^n x_A^*.$$

Therefore, Player A does not benefit from making a higher demand on round n.

Similar reasoning applies to Player B.

Finally, we need to show that it is optimal for Player A and Player B to accept only shares at least $\delta_A x_A^*$ and $\delta_B x_B^*$, respectively. Suppose that, Player B this time, rejects some share x on round n. Player B will make the demand x_B^* on round $n+1$. Since Player A will accept this offer, Player B will get the payoff $\delta_B^{n+1} x_B^*$. Therefore, it is reasonable for Player B to reject the share x if and only if

$$\delta_B^n x < \delta_B^{n+1} x_B^*.$$

This gives $x < \delta_B x_B^*$, which is what we wanted to prove. Precisely the same reasoning applies to Player A.

Next, we turn our attention to the serious part of the proof, namely showing that there are no other subgame perfect equilibria. The proof is somewhat involved, and we break it into a number of steps.

To proceed, we define the *continuation payoffs of a strategy profile in a subgame starting at time* t to be the outcome *in time* t *units* determined by these strategy profiles. For example, if Player A gets the whole cake in time $t = 3$, then the value of this outcome is δ_A in time $t = 2$ units (discount one unit of time) and δ_A^3 in time $t = 0$ units (discount three units of time).

Denote by v_A and V_A the lowest and highest[10] continuation payoffs to Player A resulting from any subgame perfect equilibrium of any subgame that begins with Player A making the first offer, and let w_A and W_A be the lowest and highest continuation payoffs to Player A in any subgame perfect equilibrium of any subgame that begins with Player B making the first offer. Define v_B, V_B, (here, Player B begins) and w_B and W_B (here, Player A begins) for Player B similarly. Next, we establish some relationships between these terms. In what follows, we assume that the players are playing according to some subgame perfect equilibrium strategies.

Claim 1: Player A will accept on any round a share at least $\delta_A V_A$, and Player B will accept on any round a share at least $\delta_B V_B$.

Suppose that Player B makes some offer at time $t = n$. If Player A rejects this

[10] At this point, we do not actually know that the smallest and the greatest value exist. More precisely, these numbers are the *infimum* and *supremum* of possible outcomes.

offer, the maximum outcome available to Player A from the resulting game (which begins with an offer made by Player A at time $t = n + 1$) is V_A. On round $t = n$, Player A values this outcome as $\delta_A V_A$. Therefore, Player A will accept the offer if his (or her) share is at least $\delta_A V_A$. Similar reasoning applies to Player B.

Claim 2: Player A will reject any offer giving a share less than $\delta_A v_A$, and Player B will reject any offer giving a share less than $\delta_B v_B$.

This follows from similar reasoning as given above. Since Player A can guarantee at least the outcome v_A from the resulting game that follows by rejecting, he (or she) will reject any share less than $\delta_A v_A$.

Claim 3: $v_A \geq 1 - \delta_B V_B$ and $v_B \geq 1 - \delta_A V_A$.

Suppose that $v_A < 1 - \delta_B V_B$. Then there exists a subgame perfect equilibrium which gives Player A a continuation payoff $u_A < 1 - \delta_B V_B$ following his (or her) offer. Pick some u'_A such that $u_A < u'_A < 1 - \delta_B V_B$. If Player A makes the demand u'_A on the first round of the subgame, the share available to Player B is $1 - u'_A \geq \delta_B V_B$. By Claim 1, Player B will accept this offer and this gives Player A the outcome u'_A. But this means that the outcome u_A cannot result from a subgame perfect equilibrium and proves the claim.

Claim 4: $W_A \leq \delta_A V_A$

By Claim 1, Player B will never offer more than $\delta_A V_A$ to Player A, and so the maximum payoff available to Player A by accepting is $\delta_A V_A$. If Player A rejects an offer made by Player B, then the maximum payoff available to Player A from the resulting game (which begins with an offer by Player A) is at most $\delta_A V_A$. So, whether Player A accepts or rejects an offer made by Player B, the maximum payoff available to Player A is $\delta_A V_A$, as claimed.

Claim 5: $V_A \leq \max\{1 - \delta_B v_B, \delta_A W_A\} \leq \max\{1 - \delta_B v_B, \delta_A^2 V_A\}$

The second inequality follows from Claim 4. To prove the first inequality, let u_A be some demand made by Player A. If Player B rejects this offer, the maximum payoff available to Player A from the resulting game (which begins with an offer by Player B in the next round) is $\delta_A W_A$. If Player B accepts the demand, we must have $1 - u_A \geq \delta_B v_B$ by Claim 2. This gives $u_A \leq 1 - \delta_B v_B$, that is, the maximum payoff to Player A in the case of agreement is $1 - \delta_B v_B$. Therefore, the payoff to Player A following his (or her) offer is at most larger of the numbers $1-\delta_B v_B$ and $\delta_A V_A$, as claimed.

Claim 6: $\max\{1-\delta_B v_B, \delta_A^2 V_A\} = 1-\delta_B v_B$.

Note first that $1-\delta_B v_B > 0$ since $\delta_B < 1$ and $v_B \le 1$. (Recall that v_B denotes the smallest outcome to Player B.) Seeking for a contradiction, suppose that $\max\{1-\delta_B v_B, \delta_A^2 V_A\} = \delta_A^2 V_A$. Then the previous claim implies that $V_A \le \delta_A^2 V_A$. Since $\delta_A < 1$ and $V_A \ge 0$, we must have $V_A = 0$. But then, the above maximum would be zero. This is impossible since there exists also a positive element, and this contradiction proves the claim.

Now we are almost done. The previous two claims (with the roles of the players reversed) imply that

$$V_A \le 1-\delta_B v_B \quad \text{and} \quad V_B \le 1-\delta_A v_A. \tag{4.5}$$

Combining the first inequality above with Claim 1, we have

$$V_A \le 1-\delta_B v_B \le 1-\delta_B(1-\delta_A V_A).$$

Solving for V_A, we obtain

$$V_A \le \frac{1-\delta_B}{1-\delta_A\delta_B}. \tag{4.6}$$

Similarly, combining the second inequality of Equation (4.5) with Claim 1, we obtain

$$v_A \ge 1-\delta_B V_B \ge 1-\delta_B(1-\delta_A v_A).$$

Solving for v_A, we obtain

$$v_A \ge \frac{1-\delta_B}{1-\delta_A\delta_B}. \tag{4.7}$$

Since V_A and v_A are the largest and the smallest continuation payoffs, respectively, we must have $v_A \le V_A$. Then Equations (4.6) and (4.7) imply that $v_A = V_A$. Similarly, we may deduce that

$$v_B = V_B = \frac{1-\delta_A}{1-\delta_A\delta_B}, \quad w_A = W_A = \frac{\delta_A(1-\delta_B)}{1-\delta_A\delta_B}, \quad w_B = W_B = \frac{\delta_B(1-\delta_A)}{1-\delta_A\delta_B}.$$

These equalities show that the outcomes in all subgame perfect equilibria are the same. Then Claim 1 and Claim 2 imply that the players always offer the shares $\delta_A v_A$ and $\delta_B v_B$, respectively, to the other player. As in the proof of Theorem 4.2, we see that the players will accept these offers by the definition of a subgame perfect equilibrium. □

Note that the ratio of shares obtained by the players is

$$\frac{x_A^*}{\delta_B x_B^*} = \frac{1-\delta_B}{1-\delta_A\delta_B} \cdot \frac{1-\delta_A\delta_B}{\delta_B(1-\delta_A)} = \frac{1}{\delta_B} \cdot \frac{1-\delta_A}{1-\delta_B}.$$

If $\delta_A = \delta_B$, so that the players are equally patient, there is a first mover advantage. Indeed, since $\delta_B < 1$, we see that the above fraction is strictly greater than 1, which means that the share obtained by Player A is strictly larger than the share obtained by Player B.

Remark 4.5. In the previous model, we have used constant discount rates to capture players' time preferences. Another interesting case is provided by utility functions with *constant costs of delay*: if the players reach an agreement (x_A, x_B) in round t, then the resulting utilities to Player A and Player B are $x_A - c_A t$ and $x_B - c_B t$, respectively, where $c_A, c_B > 0$ are constants. Note that c_A and c_B are the "costs" of one round of negotiation to Player A and Player B, respectively. With these utility functions, the outcome is very different compared to the one obtained in Theorem 4.4, namely the player with smaller cost of bargaining has all the bargaining power. If $c_A < c_B$, then Player A obtains the whole cake in the first round, and if $c_A > c_B$, then Player A obtains the share c_B in the first round. The only reason that Player A obtains something in the second case is that Player B is indifferent between obtaining $1 - c_B$ in the first round and the whole cake in the second round. See Rubinstein (1982: 107–108).

Next, let us discuss what happens in the above model when the time interval approaches zero. This seems to be a rather reasonable idea, since we would expect that the players would like to make offers and counteroffers as quickly as possible.

Let us divide the given time interval t into n equally spaced intervals of length t/n. Since δ_A and δ_B represent players' time preferences over time interval t, we use the discount factors $\delta_A^{1/n}$ and $\delta_B^{1/n}$ on the time interval t/n. When the players make offers and counteroffers after time t/n, the shares of cake obtained by Player A and Player B are given, as above, by

$$\frac{1-\delta_B^{1/n}}{1-\delta_A^{1/n}\delta_B^{1/n}} \quad \text{and} \quad \delta_B^{1/n} \cdot \frac{1-\delta_A^{1/n}}{1-\delta_A^{1/n}\delta_B^{1/n}},$$

respectively. Let us find the limit of these expressions when n tends to infinity, that is, the time between an offer and a counteroffer approaches zero. Using l'Hospitals rule and the fact that the derivative of $a^{f(x)}$ is $f'(x) \cdot a^{f(x)} \cdot \ln a$, we find that

$$\begin{aligned}
\lim_{n\to\infty} \frac{1-\delta_B^{1/n}}{1-\delta_A^{1/n}\delta_B^{1/n}} &= \lim_{n\to\infty} \frac{\frac{1}{n^2}\cdot\delta_B^{1/n}\cdot\ln\delta_B}{\frac{1}{n^2}\cdot\delta_A^{1/n}\cdot\ln\delta_A\cdot\delta_B^{1/n}+\delta_A^{1/n}\cdot\frac{1}{n^2}\cdot\delta_B^{1/n}\cdot\ln\delta_B} \\
&= \lim_{n\to\infty} \frac{\delta_B^{1/n}\cdot\ln\delta_B}{\delta_A^{1/n}\cdot\ln\delta_A\cdot\delta_B^{1/n}+\delta_A^{1/n}\cdot\delta_B^{1/n}\cdot\ln\delta_B} \\
&= \frac{\ln\delta_B}{\ln\delta_A+\ln\delta_B}.
\end{aligned}$$

Similarly (or using the fact that Player B obtains the rest of the cake), we find

$$\lim_{n\to\infty} \delta_B^{1/n}\cdot\frac{1-\delta_A^{1/n}}{1-\delta_A^{1/n}\delta_B^{1/n}} = \frac{\ln\delta_A}{\ln\delta_A+\ln\delta_B}.$$

In conclusion, we obtain the following statement given in Muthoo (1999) and Osborne and Rubinstein (1990).

Theorem 4.6. *When the time interval in Alternating Offers approaches zero, the equilibrium shares of cake obtained by Player A and Player B are*

$$\frac{\ln\delta_B}{\ln\delta_A+\ln\delta_B} \quad \textit{and} \quad \frac{\ln\delta_A}{\ln\delta_A+\ln\delta_B},$$

respectively.

Note that if we take $\gamma = \delta = 1$ and $\alpha = \ln\delta_B/(\ln\delta_A + \ln\delta_B)$ in Example 3.18, the partition is the one obtained above. In conclusion, asymmetric Nash solution gives a good approximation for the equilibrium partition of the game of Alternating Offers when the time interval between offers is small and the bargaining power α is chosen as above. It is worth of noticing that the disagreement point in Example 3.18 is $(0, 0)$, which is the outcome in Alternating Offers in the case of perpetual disagreement (that is, no agreement is reached).

Note also that if $\delta_A = \delta_B$, then, when the time interval is small, the players split the cake equally. Therefore, the first player advantage disappears as the time interval becomes small. This is also the same division as obtained in Example 3.5, which it should since equal discount factors put the players in a symmetric position. In conclusion, symmetric Nash solution gives a good approximation for the equilibrium partition of Alternating Offers in the case that the players are equally patient.

Why do we need to make these tedious calculations to conclude that if the time interval between offers is small, the cake is divided equally? The point in here is that we do not get this conclusion if we assume from the very beginning that making offers takes no time. Indeed, in this case the bargaining process would be frictionless. As we demonstrated in the first section, frictionless bargaining process is indeterminate: there must be some force driving the players towards an agreement. So, a frictionless model yields a completely different answer than a model where non-zero friction is "pushed" towards zero.

Next, we extend the model of this section by including some other factors into the model. In the next section, we study how uncertainty about the next round into account affects the equilibrium partition of the cake.

4.4 Alternating offers with a risk of breakdown

In the model of the previous section, the players were certain that there were always a next round when they can make another offer. In this section, we study a variant of the game of Alternating Offers where we relax this assumption. More precisely, we assume that bargaining ends randomly with some probability p, $0 \leq p \leq 1$. In the case of breakdown, we assume that the players obtain some utilities b_A and b_B, respectively. (As in the previous section, these utilities might be zero, but we deal with this more general case.) To simplify our equations, we put $q = 1 - p$, that is, q is the probability that the negotiations continue on the next round. The extensive form of this game is given in Figure 16. (The Player N is *Nature*, who "decides" whether the game continues.) Note that we assume that the bargaining ends immediately after the random breakdown. The material presented here is adapted from Muthoo (1999: 73–97).

In the previous section, we assumed that a perpetual disagreement gives a payoff zero to both of the players. This is not anymore true in the model of this section, and this creates a new twist to analyzing the model. If the players always reject each others offers, the game will end on some round n to the outcome (b_A, b_B). This means that the outcome from perpetual disagreement is not certain, but let us calculate its expected value.

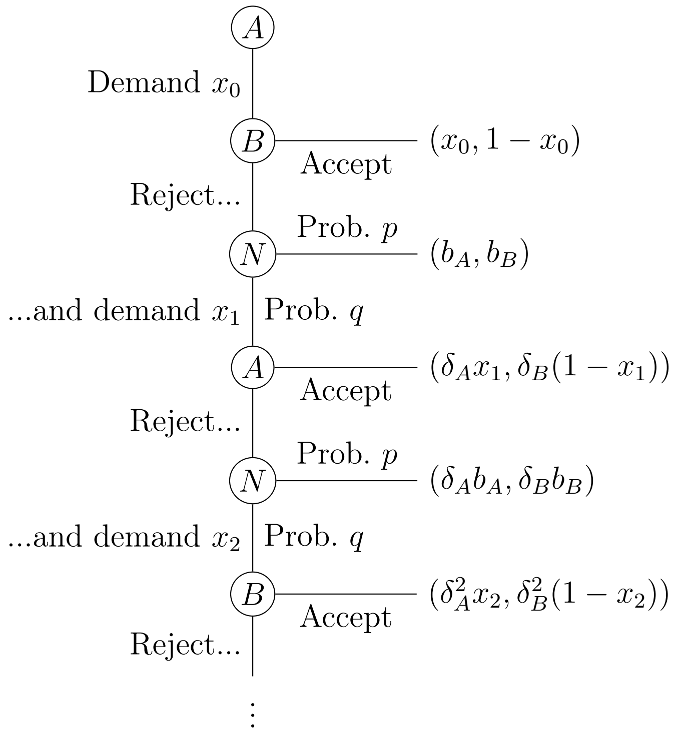

Figure 16: Game $\Gamma(p)$

The probability that the bargaining breaks down at round n is p for $n = 0$ and $q^n p$ for $n \geq 1$. (Note that the first round was labeled zero, so round n for $n \geq 1$ is actually $n + 1^{th}$ round in the game.) If the negotiations break down at round n, the present values of this outcome are $\delta_A^n b_A$ and $\delta_B^n b_B$, respectively. Therefore, the expected value of perpetual disagreement for Player A is

$$pb_A + qp\delta_A b_A + q^2 p\delta_A^2 b_A + \ldots = pb_A \sum_{n=0}^{\infty} (q\delta_A)^n.$$

The value of this geometric series is

$$b_A^* = \frac{pb_A}{1 - q\delta_A}.$$

Similarly, the expected value of perpetual disagreement for Player B is

$$b_B^* = \frac{pb_B}{1 - q\delta_B}.$$

If these values are large enough, we cannot expect an agreement to be reached. Indeed, if $b_A^* + b_B^* > 1$, there is no mutually beneficial division of cake. Therefore, to study the effect of uncertainty in a bargaining outcome, we assume in what follows that $b_A^* + b_B^* \leq 1$.

Theorem 4.7. *Let*

$$x_A^* = b_A^* + \frac{1 - q\delta_B}{1 - q^2\delta_A\delta_B}(1 - b_A^* - b_B^*)$$

and

$$x_B^* = b_B^* + \frac{1 - q\delta_A}{1 - q^2\delta_A\delta_B}(1 - b_A^* - b_B^*).$$

The game $\Gamma(p)$ *has a unique subgame perfect equilibrium where Player* A *always demands* x_A^* *and accepts any offer giving at least the share* $\delta_A x_A^*$*, and Player* B *always demands* x_B^* *and accepts any offer giving at least the share* $\delta_B x_B^*$*. The outcome of this equilibrium is an agreement on the first round, giving shares* x_A^* *and* $\delta_B x_B^*$*, respectively.*

Proof. The difficult part of the proof is in showing that the subgame perfect equilibrium is unique. This follows from similar arguments as given in the proof of Theorem 4.4 (see Muthoo (1999: 76) for required modifications), so we omit the details. Let us only demonstrate how the above given shares are determined.

Again, the players make demands in such a way that the other player is indifferent between accepting and rejecting. If Player B rejects a demand x_A^* made by Player A, there is the probability p that negotiations break down, giving outcome b_B, and probability q that Player B gets to make an offer. In order for Player B to be indifferent between these choices, we need to have $1 - x_A^* = pb_B + q\delta_B x_B^*$. Reversing the roles of the players, we obtain $1 - x_B^* = pb_A + q\delta_A x_A^*$. Solving for x_A be obtain the above result. (Actually, we do not get the above expression immediately, but with a little algebra we see that the answers are the same.) □

The previous theorem generalizes Theorem 4.4. Indeed, if $b_A = b_B = 0$ and $p = 0$, so that there is no uncertainty about the next round, the previous result gives the same partition as in Theorem 4.4.

Note that

$$1 - \frac{1 - q\delta_B}{1 - q^2\delta_A\delta_B} = \frac{1 - q^2\delta_A\delta_B - (1 - q\delta_B)}{1 - q^2\delta_A\delta_B} = q\delta_B \cdot \frac{1 - q\delta_A}{1 - q^2\delta_A\delta_B}.$$

The above given expressions for x_A^* and x_B^* illustrate the following point of the equilibrium partition. First, both of the players obtain the shares b_A^* and b_B^*, respectively, which are their expected values in the case of perpetual disagreement.

Then the rest of the cake $(1 - b_A^* - b_B^*)$ is divided according to Theorem 4.4 with δ_A and δ_B replaced by their probability adjusted discount factors $q\delta_A$ and $q\delta_B$.

Again, we find support for the Nash bargaining solution by taking the limit of the above expressions when the time interval approaches zero. In this case, it seems reasonable to assume that p also tends to zero (and so q tends to one).

Corollary 4.8. *Suppose that $\delta_A = \delta_B$. When the time interval between offers is small, the shares of cake obtained by Player A and Player B in the unique subgame perfect equilibrium of the game $\Gamma(p)$ are approximately*

$$b_A^* + \frac{1}{2}(1 - b_A^* - b_B^*) \quad \textit{and} \quad b_B^* + \frac{1}{2}(1 - b_B^* - b_B^*),$$

respectively.

Note that this is the division obtained in Example 3.9 with $d = (b_A^*, b_B^*)$. So, we may conclude that the Nash bargaining solution gives a good approximation to the equilibrium partition of the model of this section when the time interval between the offers is small. Note, in particular, that to obtain this result we choose the disagreement point d to reflect the players utilities in the case of perpetual disagreement.

4.5 Alternating offers with outside option

So far, we have assumed that the only way the players may benefit from bargaining is by finding an agreement with each others. In this section, we study the effect of an *outside option* on the bargaining outcome by assuming that Player B can quit bargaining with Player A. Of course, Player B will not quit bargaining with Player A unless it benefits him (or her), and the resulting benefit is captured by an outside option. Material in this section is taken from Osborne and Rubinstein (1990: 54–63), see Muthoo (1999: 99–135) for a model where both of the players have an outside option.

To motivate the model of this section, suppose that Player B is willing to sell a house, and Player A is willing to buy the house. As such, this is the game of Alternating Offers studied previously. But now, let us add another factor into the game, namely that Player B has already received an offer b from a third party. This is an *outside option* for Player B, as he (or she) can quit bargaining with

Player A and sell the house to the third party. Note that if Player B uses his (or her) outside option, bargaining with Player A terminates permanently.

Let us first discuss the effect of outside option on the bargaining outcome. Let us assume that Player A is willing to pay the amount \$120 000, and that Player B is willing to sell the house at \$100 000. Now, the question is how the *surplus* \$20 000 is to be divided between the players.

It seems that Player B has an advantage, since he (or she) has already received an offer. But suppose, for example, that Player A and Player B would agree on the price \$110 000 when bargaining with each others. If the third party has made an offer \$105 000 to Player B, he (or she) cannot use this offer as a credible threat to quit bargaining with Player A. (We are again assuming complete information, which means that Player A knows the offer made by the third party.) On the other hand, if Player B has an offer \$115 000, he (or she) can use this offer as a credible threat to sell the house to the third party. In this case, Player A has to pay at least \$115 000 to get the house, but there is no reason why he (or she) would pay more than this. (Note that, due to the reservation value of Player A, there is still some surplus available for him (or her) also.)

To analyze the game with outside options, we proceed as previously: if we know how to divide a cake of size 1 between the players, we may interpret this division as giving the percentages of how to divide any other amount. We assume for simplicity that the players have the same discount factors, that is, $\delta_A = \delta_B = \delta$. In this case, Theorem 4.4 gives the shares of cake

$$x_A^* = \frac{1-\delta}{1-\delta^2} \qquad \text{and} \qquad x_B^* = \delta \cdot \frac{1-\delta}{1-\delta^2} \tag{4.8}$$

for Player A and Player B, respectively.

The time when Player B may quit bargaining actually plays a role in determining the subgame perfect equilibrium. Therefore, we consider two models: $\Gamma(1)$ where Player B can opt out after an offer from Player A, and $\Gamma(2)$ where Player B can opt out after Player A rejects an offer. Let us first deal with the game $\Gamma(1)$, the extensive form of this game is given in Figure 17 below.

Theorem 4.9. *Let* $0 \le b \le 1$. *The following statements hold:*

(i) If $b < \delta/(1+\delta)$, *then the game* $\Gamma(1)$ *has the same unique subgame perfect*

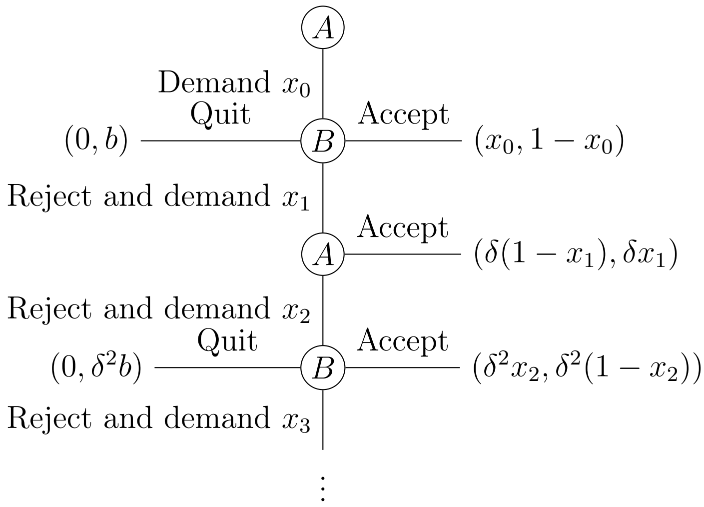

Figure 17: Alternating offers with outside option, $\Gamma(1)$

equilibrium as in Theorem 4.4. In particular, the outcome is an agreement $(1/(1+\delta), \delta/(1+\delta))$ *on the first round.*

(ii) if $b = \delta/(1+\delta)$*, then in every subgame perfect equilibrium the outcome is an agreement* $(1-b, b)$ *on the first round.*

(iii) If $b > \delta/(1+\delta)$*, then the game* $\Gamma(1)$ *has a unique subgame perfect equilibrium where Player* A *always demands* $1-b$ *and accepts any offer which gives at least the share* $\delta(1-b)$*, and Player* B *always demands* $1-\delta(1-b)$ *and accepts any share at least* b *and quits bargaining if a share smaller than* b *is offered. The outcome is an agreement* $(1-b, b)$ *in the first round.*

Proof. See Osborne and Rubinstein (1990: 57–58) for a proper proof. We only discuss briefly some features of this proof. (Again, the difficult part is in showing that the subgame perfect equilibrium in cases (i) and (iii) is unique, but the reasoning is similar as given in the proof of Theorem 4.4.)

Consider statement (i). Note from Equation (4.8) that $\delta/(1+\delta)$ is the outcome that Player B would receive from the Alternating Offers without outside option. Therefore, if the outside option is too small, it does not affect the outcome since Player B cannot use the outside option as a credible threat to quit bargaining.

Next, let us discuss statement (iii). In this case the outside option is large enough and Player B can use this as a credible threat to quit bargaining with Player A. Therefore, Player A needs to make large enough offer so that Player B does not quit bargaining and leave Player A with payoff 0. Since Player B obtains the outcome b by opting out, the maximum amount available to Player A is $1-b$. Since Player A gets to make the first offer, he (or she) demands precisely $1-b$, which leaves Player B indifferent between accepting and rejecting, and so he (or she) accepts this offer. □

Next, let us consider another variant of Alternating Offers where Player B may quit bargaining after Player A has rejected an offer. The extensive form of this game is given in Figure 18 below. Note here that we assume that Player B quits bargaining immediately after Player A rejects an offer. That is, if Player B chooses to opt out, he (or she) obtains the outcome b immediately.

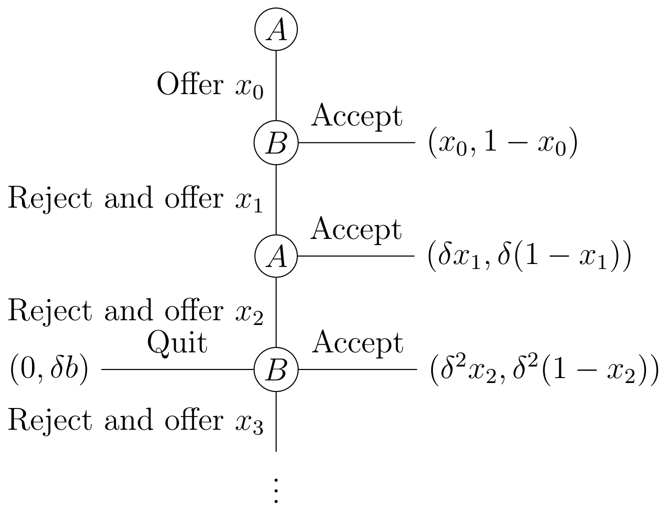

Figure 18: Alternating Offers with outside option, $\Gamma(2)$

Theorem 4.10. *Let* $0 \le b \le 1$. *The following statements hold:*

(i) If $b < \delta^2(1+\delta)$, *then the game* $\Gamma(2)$ *has the same unique subgame perfect equilibrium as in Theorem 4.4. In particular, the outcome is the agreement* $(1/(1+\delta), \delta(1+\delta))$ *on the first round.*

(ii) Suppose that $\delta^2/(1+\delta) \le b \le \delta^2$. *If* $1-\delta \le x \le 1-b/\delta$, *then there exists a subgame perfect equilibrium in the game* $\Gamma(2)$ *which yields the outcome* $(x, 1-x)$ *in the first round. In any subgame perfect equilibrium, Player B obtains at least* $\delta/(1+\delta)$.

(iii) If $b > \delta^2$, *then the game* $\Gamma(2)$ *has a unique subgame perfect equilibrium in which Player A always demands* $1-\delta$ *and accepts any offer, and Player B always demands* 1, *accepts any offer which gives at least the share* δ, *and*

always quits bargaining. In particular, the outcome of this equilibrium is the partition $(1-\delta, \delta)$ *on the first round.*

Proof. See Osborne and Rubinstein (1990: 60–62) for a full proof. Let us only discuss some of these statements shortly.

Consider some round t when Player B may quit bargaining. Opting out gives the outcome b on round t. If Player B continues bargaining, the equilibrium outcome is $\delta(1-\delta)/(1-\delta^2) = \delta/(1+\delta)$ in round $t+1$ in the game without outside options. From the viewpoint of round t, Player B values this as $\delta^2/(1+\delta)$. Therefore, statement (i) follows from the fact that if the outside option is small, it is not credible to quit bargaining.

In case (iii), the outside option is large enough for Player B to quit bargaining. Note that quitting on the first possible moment gives Player B the outcome $\delta b > \delta^3$. Here, δ^3 is the outcome to Player B of obtaining the whole cake in round three, and so Player B will quit bargaining in the first occasion. This gives Player A the outcome 0, and so Player B may credibly demand for the whole cake in his (or her) first demand. □

Note the strict difference between the outcomes given by statements (iii) of the previous two theorems. In the first case, the payoff to Player A depends on the outside option b, but in the second case it depends only on the discount factor δ. Player A has an advantage in the game $\Gamma(1)$ because he (or she) is able to make an offer before Player B considers quitting. Therefore, Player A can make a proposal which Player B finds at least as attractive as the outcome b, and Player A also obtains a share of cake. In the game $\Gamma(2)$, Player B can make a "take it or leave it" offer. Indeed, when we consider the timing of Player B quitting bargaining, we can view his (or her) offer as Player B saying "If You reject this, I will quit." As we saw in the context of the Ultimatum game, this take it or leave it offer gives Player B a huge advantage.

4.6 Alternating offers with inside option

A characteristic property of an outside option studied in the previous section was that exercising this option terminates bargaining permanently. In this section, we continue our study of bargaining situations by examining a model with *inside option*. Unlike an outside option, an inside option gives some utility to players in the case of *temporary* disagreement. This could mean, for example, that in the bargaining situation studied in the previous section the owner of the house attains some utility even if the players do not reach an agreement. Indeed, the owner of the house can continue living in the house and he (or she) does not need to spend time in looking for a new apartment or paying rent. Also, if the players disagree about the price of the house today, they can continue bargaining at some later moment.

We present the main ideas of how an inside option affects the bargaining outcome through an example. We could deal with the bargaining situation over the price of a house, but we have chosen an example related to industrial organization: we study an oligopolistic market where a firm is considering to buy another firm and, especially, how the price is determined. In this setting, both of the firms have inside options as they enjoy their profits until agreement on the price is reached. The material in this section is based on Muthoo (1999: 137–185), where the reader may find a continuous version of the model presented. We deal with a discrete version for simplicity.

Let us recall some ideas related to Cornout competition. We assume that Firm A and Firm B produce identical products with unit costs of 0 and $c > 0$, respectively, and that they are the only firms in the market. (Firm A has superior technology in production.) We assume that both of the firms use the same discount factor δ. We also assume that the market demand q and the market price p are related by the equation $p = \alpha - q$, where $\alpha > 2c$ is constant. (This means that the market is large enough; we explain this in a few moments.)

We are interested in determining the price at which one of the firms might buy the other firm. Note that either of the firms could buy the other one: Firm A would do this to eliminate competition, and Firm B would do this to get the more efficient production technology of Firm A. In the case of acquisition, the resulting firm would have a monopoly in the market, so we need to determine how a monopolistic firm would behave. Whether Firm A buys Firm B or *vice versa*, the resulting firm will have the more efficient technology of Firm A. Therefore, after acquisition the (monopolistic) firm will produce an amount q which maximizes its profit $pq = (\alpha - q)q$. Taking the first order condition with respect to q, we find that the optimal output q^* and the profit Π_M with this output are

$$q^* = \frac{\alpha}{2} \quad \text{and} \quad \Pi_M = \frac{\alpha^2}{4}.$$

The present value of profits of a monopolistic firm is, therefore,

$$\sum_{n=0}^{\infty} \delta^n \Pi_M = \frac{\Pi_M}{1-\delta}. \tag{4.9}$$

Next, let us determine what happens before acquisition. We assume that the firms are Cornout competitors, that is, the firms make their choices of production levels simultaneously. When Firm A and Firm B produce some quantities q_A and q_B, respectively, their profits are

$$\Pi_A = (\alpha - q_A - q_B)q_A \quad \text{and} \quad \Pi_B = (\alpha - q_A - q_B)q_B - cq_B,$$

respectively. The equilibrium outputs are determined as follows. We find the best response functions of Firm A and Firm B, respectively, by taking the first order conditions of Π_A and Π_B with respect to q_A and q_B, respectively. The answers are

$$q_A^* = \frac{\alpha - q_B}{2} \quad \text{and} \quad q_B^* = \frac{\alpha - q_A - c}{2},$$

respectively. Substituting q_B^* into the expression for q_A^* and solving for q_A^* gives

$$q_A^* = \frac{\alpha + c}{3} \quad \text{and} \quad q_B^* = \frac{\alpha - 2c}{3}.$$

Note that our assumption $\alpha > 2c$ guarantees that $q_B^* > 0$. These output levels give the profits

$$\Pi_A = \frac{(\alpha + c)^2}{9} \quad \text{and} \quad \Pi_B = \frac{(\alpha - 2c)^2}{9},$$

respectively.

To get to the main point of this section, suppose that one of the firms plans to buy the other firm. We assume that the price is determined by the game of Alternating Offers. If the firms do not reach an agreement at the moment, the firms have inside options as they get their profits during the next time period. Let us see how the price of a firm is determined. Note that the price of a firm is not determined solely by the present value of its profits. Since the resulting firm will enjoy monopoly profits on later periods, and both firms know this, there is a cake, the size of which is given by Equation (4.9), to be divided. The question, then, is how to divide this cake between the firms.

So, the firms bargain over a division of cake of size $\Pi_M/(1-\delta)$. As in the previous models, in a subgame perfect equilibrium[11] Firm A and Firm B adjust their demands x_A and x_B, respectively, such that the other firm is indifferent between accepting and rejecting. Since rejection gives profit Π_B to Firm B in the current period, and after this it is Firm B in turn to make its equilibrium demand x_B, Firm A will make a demand x_A such that

$$\frac{\Pi_M}{1-\delta} - x_A = \Pi_B + \delta x_B.$$

Similarly, in its turn Firm B will make a demand x_B such that

$$\frac{\Pi_M}{1-\delta} - x_B = \Pi_A + \delta x_A.$$

These are practically the same equations as in the proof of Theorem 4.7, so we have (we could write $(1-\delta)/(1-\delta^2) = 1/(1+\delta)$, but we prefer the expression below)

$$x_A^* = \frac{\Pi_A}{1-\delta} + \frac{1-\delta}{1-\delta^2} \cdot (\frac{\Pi_M}{1-\delta} - \frac{\Pi_A}{1-\delta} - \frac{\Pi_B}{1-\delta})$$

and

$$x_B^* = \frac{\Pi_B}{1-\delta} + \frac{1-\delta}{1-\delta^2} \cdot (\frac{\Pi_M}{1-\delta} - \frac{\Pi_A}{1-\delta} - \frac{\Pi_B}{1-\delta}).$$

These demands form a subgame perfect equilibrium where Firm B accepts the first demand x_A^* made by Firm A, and so the firms find an immediate agreement. The outcomes to Firm A and Firm B are

$$x_A^* = \frac{\Pi_A}{1-\delta} + \frac{1-\delta}{1-\delta^2} \cdot (\frac{\Pi_M}{1-\delta} - \frac{\Pi_A}{1-\delta} - \frac{\Pi_B}{1-\delta})$$

and

$$\frac{\Pi_M}{1-\delta} - x_A^* = \frac{\Pi_B}{1-\delta} + \frac{\delta(1-\delta)}{1-\delta^2} \cdot (\frac{\Pi_M}{1-\delta} - \frac{\Pi_A}{1-\delta} - \frac{\Pi_B}{1-\delta}),$$

respectively. Note that neither of the firms exercises its inside option, but these options nevertheless have a significant impact on the outcome. Also, we have a similar interpretation as previously: give both of the firms the present value of their profits they would enjoy without acquisition, and divide the rest of the cake (the term in parenthesis above) according to Theorem 4.4.

[11] We are optimistically conjecturing that an equilibrium exists, see Theorem 4.11.

As previously, an interesting case is the limit of the above expression when the length of the time interval between offers approaches zero. Note that if we divide the time interval into n equally spaced intervals, the firms enjoy profits Π_A/n and Π_B/n in the shorter time intervals. Applying again l'Hospitals rule, we find that

$$\lim_{n\to\infty} \frac{1/n}{1-\delta^{1/n}} = \lim_{n\to\infty} \frac{-1/n^2}{1/n^2 \cdot \ln\delta \cdot \delta^{1/n}} = \lim_{n\to\infty} \frac{-1}{\ln\delta \cdot \delta^{1/n}} = \frac{1}{|\ln\delta|}.$$

Put $r = |\ln\delta|$. In the limit when the time interval approaches zero, the outcomes to Firm A and Firm B are

$$p_A = \frac{\Pi_A}{r} + \frac{1}{2}\cdot(\frac{\Pi_M}{r} - \frac{\Pi_A}{r} - \frac{\Pi_B}{r}) \quad \text{and} \quad p_B = \frac{\Pi_B}{r} + \frac{1}{2}\cdot(\frac{\Pi_M}{r} - \frac{\Pi_A}{r} - \frac{\Pi_B}{r}),$$

respectively. If Firm A buys Firm B, it pays the price p_B, and *vice versa*. Since $\Pi_B < \Pi_A$, we have $p_B < p_A$, and so this simple model illustrates that more efficient firms tend to buy less efficient firms. Note also that Π_A is increasing in c and Π_B is decreasing in c, and so p_A is increasing in c and p_B is decreasing in c. This suggests that as the degree between efficiency of the firms increases, it is more likely that the more efficient firm will buy the less efficient firm. (It is relatively cheap to eliminate competition, which gives a monopoly in the market.) Furthermore, note that both of the prices are increasing in α, which suggests that acquisitions are less likely (more expensive) in large markets.

In the general case, we assume that Player A and Player B have their discount factors δ_A and δ_B, respectively, and that they obtain some utilities g_A and g_B in the case of temporary disagreement. When the size of the cake is one and $g_A/(1-\delta_A) + g_B/(1-\delta_B) < 1$ (otherwise, jointly profitable division of cake does not exist), the subgame perfect equilibrium is given as follows.

Theorem 4.11. *Let*

$$x_A^* = \frac{g_A}{1-\delta_A} + \frac{1-\delta_B}{1-\delta_A\delta_B}\cdot(1 - \frac{g_A}{1-\delta_A} - \frac{g_B}{1-\delta_B})$$

and

$$x_B^* = \frac{g_B}{1-\delta_B} + \frac{1-\delta_A}{1-\delta_A\delta_B}\cdot(1 - \frac{g_A}{1-\delta_A} - \frac{g_B}{1-\delta_B}).$$

Then the unique subgame perfect equilibrium of Alternating Offers with inside options is such that Player A always demands x_A^* and accepts any offer which gives at least the share $\delta_A x_A^*$, and Player B always demands x_B^* and accepts any offer which gives at least the share $\delta_B x_B^*$. The outcome of this equilibrium is an *agreement in the first round, and Player A obtains x_A^* and Player B obtains $\delta_B x_B^*$.*

Proof. Again, the serious part of the proof is in showing that the subgame perfect equilibrium is unique, but this follows from similar arguments as in Theorem 4.4. See Muthoo (1999: 140) for necessary modifications. As above, the given values are determined by the following pair of equations

$$1 - x_A^* = g_B + \delta_B x_B^* \quad \text{and} \quad 1 - x_B^* = g_A + \delta_A x_A^*,$$

which give the above solution. □

Arguing as above, we obtain the following statement in the limit when the time interval between offers approaches zero.

Corollary 4.12. *Let $r_A = |\ln \delta_A|$ and $r_B = |\ln \delta_B|$. When the time interval between offers approaches zero, the outcomes to Player A and Player B in the unique subgame perfect equilibrium of the game of Alternating Offers with inside options are*

$$\frac{g_A}{r_A} + \frac{r_B}{r_A + r_B} \cdot \left(1 - \frac{g_A}{r_A} - \frac{g_B}{r_B}\right) \quad \textit{and} \quad \frac{g_B}{r_B} + \frac{r_A}{r_A + r_B} \cdot \left(1 - \frac{g_A}{r_A} - \frac{g_B}{r_B}\right),$$

respectively. In particular, if $\delta_A = \delta_B$, then the term in the parenthesis is split equally between the players.

The previous corollary provides some support for the Nash bargaining solution. Indeed, when we take $\alpha = r_A/(r_A + r_B)$ and $d = (g_A/r_A, g_B/r_B)$, the division given above is the same as obtained in Example 3.19. Note, again, that the disagreement point reflects the players' utilities in the case of perpetual disagreement. Indeed, perpetual disagreement yields the utilities g_A and g_B to Player A and Player B, respectively, on every round. The present values of overall utilities are $g_A/(1-\delta_A)$ and $g_B/(1 - \delta_B)$, respectively, and the numbers g_A/r_A and g_B/r_B are the limits of these values when the time interval approaches zero. Writing

$$\alpha = \frac{r_A + r_B - r_B}{r_A + r_B} = 1 - \frac{r_B}{r_A + r_B},$$

we see that the bargaining power α of Player A increases as r_A increases. Since $r_A = |\ln \delta_A|$, an increase in r_A is equivalent to an increase in δ_A which, roughly, means that Player A becomes more patient. A similar interpretation applies to the bargaining power of Player B.

4.7 Alternating offers with commitment tactics

So far, we have studied how different options during bargaining process affect on the bargaining outcome. In this section, we describe briefly what kind of actions the bargainers might take *before* they start negotiating. We are still dividing a cake of size one, but this time we assume that the players are capable of making some kind of "commitment" before they sit into the bargaining table, and that these commitments partially "tie the hands" of negotiators. We assume that the bargainers are capable of revoking their commitments, but only with some cost. (Otherwise, commitments would not matter in the bargaining situation.) As noted in Schelling (1960: 27–28), this kind of behavior seems to be part of real life bargaining situations, for example, wage or international negotiations. The material in this section is taken from Muthoo (1999).

The agents play a two–stage game. In the first stage, the players choose their commitment levels, and in the second stage they bargain according to the chosen levels. In this section, we are mainly interested in the effect of commitments on the bargaining outcome, so we apply previous results to determine the outcome in the second stage. More precisely, we apply the Nash bargaining solution to solve the second part. (Recall that this gives the same solution as Alternating Offers with small time interval.)

In the first stage, the players choose some numbers $z_A, z_B \in [0, 1]$ which represent their commitments. We may interpret these values as Player i, $i = A, B$, making a public announcement that he (or she) will not settle for a smaller share of cake than z_i. In the second stage, the players start the bargaining process. We assume that if the players reach an agreement (x_A, x_B) on the partition of cake, the cost for Player i, $i = A, B$, is

$$C_i(x_i, z_i) = \begin{cases} 0 & \text{if } x_i \geq z_i, \\ k_i(z_i - x_i) & \text{if } x_i < z_i, \end{cases} \tag{4.10}$$

where $k_i > 0$ is some constant. So, this cost is zero if a player obtains at least the share he (or she) has committed before bargaining. Also, the second equation implies that, if a player obtains less than his (or her) commitment, the cost is increasing in the difference between the commitment and the actual outcome. That is, a large commitment may cause a large cost if the bargaining outcome is small compared to the commitment. (Of course, the obtained share of cake needs to be large enough to give a positive outcome.)

Let us describe the outcomes of this bargaining situation: these outcomes depend on the overall size of commitments. First, assume that $z_A + z_B \leq 1$, that is, the commitments do not exceed the size of the cake. Then neither of the players needs to revoke his (or her) commitment. In this case, we assume that the cake is divided as in Example 3.9. So, the players obtain the shares represented by their commitments and the remaining cake is split equally, that is,

$$U_A = z_A + \frac{1}{2}(1 - z_A - z_B) \quad \text{and} \quad U_B = z_B + \frac{1}{2}(1 - z_A - z_B). \tag{4.11}$$

We need also to describe what happens in the case that players' commitments are incompatible, that is, $z_A + z_B > 1$. This means that, in order to reach an agreement, at least one of the players needs to give up on his (or her) commitment (with a cost). The players' utilities are determined by the share of cake obtained and the cost of revoking a commitment. As above, we assume that the bargaining solution $(x_A, 1 - x_A)$ is given by the Nash solution, that is, as the maximizer of the Nash product

$$[x_A - C_A(x_A, z_A)] \cdot [(1 - x_A) - C_B(1 - x_A, z_B)]. \tag{4.12}$$

The main result of this section is as follows.

Theorem 4.13. *The commitments have a unique Nash equilibrium given by*

$$z_A^* = \frac{1}{1+\gamma} \quad \textit{and} \quad z_B^* = \frac{\gamma}{1+\gamma},$$

where $\gamma = (1 + k_B)/(1 + k_A)$. *The outcome with these commitments is that the players obtain the shares* z_A^* *and* z_B^*, *respectively.*

Proof. For a full proof, see Muthoo (1999: 214–222). We only discuss the proof in some detail.

An easy part of the proof is to show that commitments z_A and z_B with $z_A + z_B < 1$ do not determine a Nash equilibrium. Indeed, in this case one of the players benefits by increasing his (or her) commitment. Put $\varepsilon = 1 - (z_A + z_B)$ and suppose that, Player A for example, changes his (or her) commitment from z_A to $z'_A = z_A + \varepsilon$. Now, $z'_A + z_B = 1$, so the commitments are compatible, and Equation

(4.11) gives Player A the outcome

$$z'_A + \frac{1}{2}(1 - z'_A - z_B) = z_A + \frac{1}{2}(1 - z_A - z_B) + \frac{\varepsilon}{2}.$$

Note here that $z_A + (1 - z_A - z_B)/2$ is the outcome to Player A with original commitments z_A and z_B, respectively. Therefore, Player A benefits by making the commitment z'_A instead of z_A, and the statement follows.

The rest of the proof is rather involved, and we omit the details. The case of incompatible commitments ($z_A + z_B > 1$) requires comparing the Pareto frontiers and the Nash solutions of different bargaining sets. Note that when one of the players changes his (or her) commitment, the parameter z_i in the Nash product changes, and one needs to compare the Nash solutions of these two bargaining problems in order to show that this benefits the player in question. An extra twist in this analysis yields from the fact that the functions C_i are not differentiable at zero, and so the Nash solution might be a corner point of the Pareto frontier. The case $z_A + z_B = 1$ with $z_B/z_A \neq \gamma$ follows from a similar reasoning. Finally, one needs to show that the above given pair is a Nash equilibrium. To explain shortly why this is so, note first that neither of the players has a reason to make a smaller commitment. Also, the cost of revoking guarantees that it does not benefit either of the players to make a greater commitment. (A larger commitment means incompatible commitments. Then the outcome is determined by Equation (4.12) instead of Equation (4.11). It is a property of the Nash solution that the player making a higher commitment also needs to give up on his (or her) commitment in the bargaining solution.) Commitments z^*_A and z^*_B satisfy $z^*_A + z^*_B = 1$, so the outcome is given by Equation (4.11). □

The above given Nash equilibrium reveals a perhaps counter intuitive fact, namely that a high cost of revoking can be a strength in a bargaining situation. Suppose that $k_A > k_B$. Then $\gamma < 1$, so $z^*_A > 1/2$, and so Player A obtains a larger share of cake. We may interpret this in a way that a high cost of revoking ties the hands of Player A in a tighter manner than of Player B. If the players make incompatible commitments, it is easier (less costly) for Player B to give up on his (or her) commitment in order to reach an agreement. Therefore, Player A may credibly commit for a large share of cake. For example, $k_A = 10$ and $k_B = 1$ give Player A approximately 85 % of the cake, but $k_A = 100$ and $k_B = 1$ give Player A approximately 98 % of the cake.[12]

[12]In comparison to the previous games studied, there is no first player advantage in this

Again, we find some support for the Nash bargaining solution from the model of this section. (We applied Nash bargaining solution in the second stage, but we could have used Alternating Offers with a small time interval instead.) Indeed, when we take $\delta = \gamma = 1$ and $\alpha = (1 + k_A)/(2 + k_A + k_B)$ in Example 3.18, we find that the partition of cake is the same as given by Theorem 4.13. Note, once again, that the disagreement point $d = (0, 0)$ reflects the players utilities in the case of perpetual disagreement. We may write

$$\alpha = \frac{2 + k_A + k_B - 1 - k_B}{2 + k_A + k_B} = 1 - \frac{1 + k_B}{2 + k_A + k_B},$$

which supports our finding above that the bargaining power α of Player A is increasing in k_A and decreasing in k_B.

4.8 Alternating offers with incomplete information

All the models studied so far have been perhaps unrealistic in the sense that these models have assumed complete information: both of the players have been completely aware of the characteristics of the other player, such as his (or her) discount factor and outside or inside options. As such, it is not surprising that every outcome has been efficient, that is, agreement has been reached in the first round and no resources (cake) has been wasted. In real life bargaining, however, it might take a considerable amount of time to reach an agreement and resources are wasted during this time, for example, in the form of strikes. In this section, we present a variant of the game of Alternating Offers where one of the players is not certain about all the characteristics of the other player. We deal with a relatively simple model, but we see that this model manages to explain a delay in reaching an agreement. This section is based on Osborne and Rubinstein (1990: 91–120).

A player might hold a large amount of private information concerning bargaining, such as his (or her) preferences or possible options during the bargaining. To keep the model of this section simple, we assume that there is only one uncertain factor affecting the decisions of one of the players. We assume that bargaining entails a fixed cost per round (see Remark 4.5). That is, partition (x_A, x_B) of cake on

game. The players choose their commitments simultaneously, and we applied the Nash solution to determine the bargaining outcome.

round t gives Player i, $i = A, B$, the utility

$$U_i(x, t) = x_i - c_i t.$$

Recall that the constant c_i is the "price" of one round of negotiation to Player i. We assume that Player B knows c_A but Player A is not certain about the bargaining cost c_B of Player B. We assume that there are two possibilities for c_B, namely c_H (high) with probability p and c_L (low) with probability $1-p$, and that Player A knows these values and the corresponding probabilities. Furthermore, we assume that $0 < c_L < c_A < c_H$ and $c_A + c_L + c_H < 1$. The last assumption will be clarified later, but note here that the cases $c_A < c_L, c_H$ and $c_L, c_H < c_A$ require no further analysis since in these cases Player A or Player B, respectively, has all the bargaining power (see Remark 4.5).

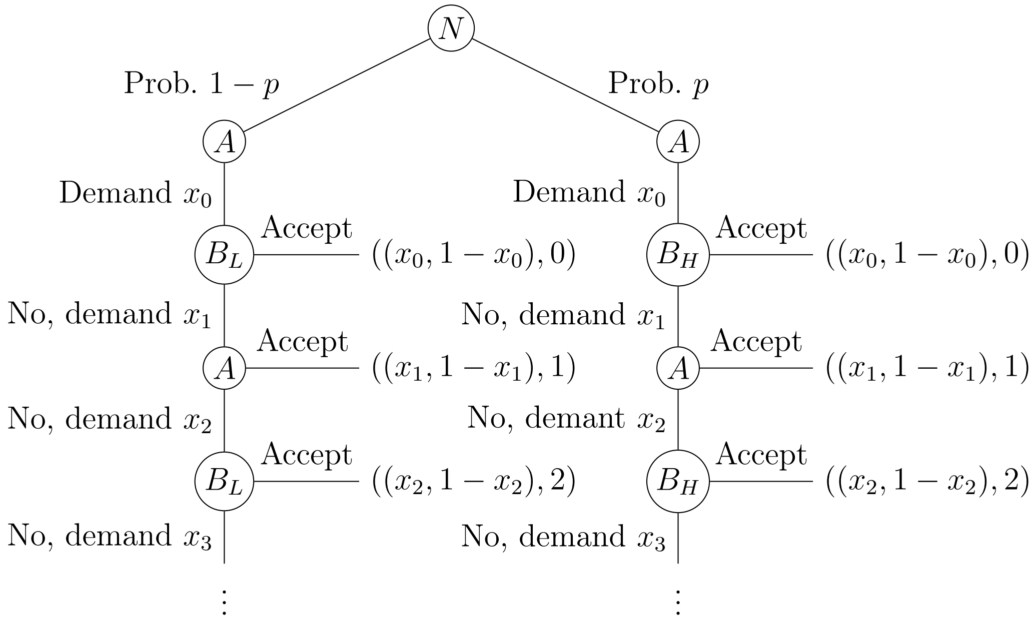

Figure 19: Alternating Offers with incomplete information

To take the different possibilities for the bargaining cost of Player B into account, we introduce two type of players to act in the role of Player B: we denote by B_L and B_H Player B with bargaining cost c_L and c_H, respectively. Our assumption above means that Player A faces Player B_L with probability $1 - p$ and Player B_H with probability p. Recall that Player A is in a weak position against Player B_L and in a strong position against Player B_H. Since Player A is not certain about the bargaining cost of the other player, Player B has an incentive to try to convince Player A that he (or she) is Player B_L. The game tree is given as in Figure 19 above. Here, again, N denotes *Nature* which determines the type of the Player B. (To avoid long expressions in the game tree, we only denote the partition and the round on which this agreement is reached.)

We need to describe what we mean by an equilibrium in this model. Note that subgame perfect equilibrium is of no use in this context, since this game has no proper subgames. (Player A does not know in which branch of the game tree he (or she) is.) In this section, we apply the concept of *sequential equilibrium*, which is an uncertainty version of the concept of subgame perfect equilibrium. However, the definition is rather delicate as we need to describe what kind of conclusions incompletely informed Player A makes during the game. (Player A wishes to figure out whether Player B is of type B_L or B_H.) We present a definition that is suitable for the game of this section; see Fudenberg and Tirole (1995: 321–350) for a more comprehensive discussion about the sequential equilibrium.

Sequential equilibrium consists of a strategy for each of the players A, B_L, and B_H, and a *system of beliefs* for Player A with some (soon to be discussed) assumptions on these strategies and beliefs. The system of beliefs captures the idea that Player A is trying to conclude from the actions of Player B which type of player he (or she) is playing against. Let us first discuss what kind of conclusions Player A is allowed to make.

Player A makes the first decision based on the given probability p, but on later rounds Player A re–evaluates this probability based on his (or her) previous belief on p and actions taken by Player B. If the strategies of both of the players B_L and B_H describe same action, and Player A observes this action, then Player A does not change his (or her) belief on p. (Player A does not obtain any valuable information.) If Player A observes an action which is consistent with the strategy of *only one* of the players B_L or B_H, then Player A concludes which player he (or she) is playing against. (After this, we have either $p = 0$ or $p = 1$.) Furthermore, once Player A makes this conclusion, he (or she) cannot revise this belief in the rest of the game. (See Osborne and Rubinstein (1990: 95–97) for a discussion about this assumption.) Finally, if Player A observes an action that is not part of the strategy of either of the players B_L or B_H, then Player A may form a new belief on p in any way. (Player A may pick any number $0 \leq p \leq 1$ to use henceforth.)

Note that Player A updates his (or her) belief on p only after actions taken by Player B, that is, at the beginning of the turn of Player A. Note also that if the strategy of Player B_H requires to accept some offer and the strategy of Player B_L requires to reject the same offer, and Player B indeed rejects this offer, Player A does not necessarily conclude that he (or she) faces Player B_L. Player A makes this conclusion only after he (or she) receives a counteroffer described by the strategy of Player B_L. (If another offer is observed, Player A is free to form a new belief.) Now, we may formulate the equilibrium concept needed in this section.

Definition 4.14. A *sequential equilibrium* consists of a strategy for each of the three players and a system of beliefs for Player A such that, in every stage of the game, the strategy of both of the players B_L and B_H is optimal given the strategy of Player A, and the strategy of Player A is optimal given the strategies of Players B_L and B_H and the current belief on p.

The previous definition is similar to the concept of subgame perfect equilibrium, as we require that the players' decisions are optimal at every stage of the game. For completely informed players B_L and B_H, this is the same requirement as previously, but for incompletely informed Player A this means that Player A is maximizing the expected payoff from the resulting game based on his (or her) belief on the probability p.

To illustrate the concept of sequential equilibrium, let us discuss some statements that apply to all sequential equilibria.

Lemma 4.15. *As long as $0 < p < 1$, the following statements hold in every sequential equilibrium in every stage of the game:*

(i) If players B_L and B_H reject some offer, they make the same counteroffer.

(ii) If Player B_L accepts some offer, then Player B_H also accepts this offer.

(iii) If Player B_H accepts and Player B_L rejects some demand x_A, then, in the next round, Player A will accept any offer made by Player B_L where the share y_A obtained by Player A satisfies $x_A - c_H \leq y_A \leq x_A - c_L$.

Proof. Shortly, statements (i) and (ii) follow from the fact that the high cost player B_H does not wish to reveal his (or her) identity to Player A. The details are as follows.

(i) Suppose that the strategies of Player B_L and Player B_H call for both of them to reject an offer made by Player A and to demand some shares y and z, respectively, with $y \neq z$. Let us first show that Player A accepts the demand y. When Player A observes either of these demands, Player A knows the type of Player B, and the resulting game is of complete information. Therefore, if Player A rejects y, the equilibrium outcome is the division $(c_L, 1 - c_L)$ in the next round, and if Player A rejects z, the equilibrium outcome is the division $(1, 0)$ in the next round (see Remark 4.5). From the viewpoint of the current round, Player A values these outcomes as $c_L - c_A < 0$ and $1 - c_A$. Therefore, Player A will accept the demand y since this gives the outcome $1 - y \geq 0$ on the current period.

Next, Player A rejects the demand z. Indeed, accepting this demand cannot be part of equilibrium strategies: if Player A accepts both demands y and z, one of the players B_L or B_H could benefit by changing his (or her) demand to larger of the demands y and z. In conclusion, Player B_H receives the outcome 0 in the next period, and Player B_H values this as $0 - c_H < 0$. This is worse than obtaining y in the current period, which is possible if Player B_H mimics Player B_L and makes the demand y in the first place. Therefore, in equilibrium Player B_H makes the same demand as Player B_L.

(ii) Suppose that the equilibrium strategy of Player B_L calls to accept some offer made by Player A on some round t. If Player B_H also accepts this demand, the outcome to Player B_H from the viewpoint of round t is positive. If the strategy of Player B_H requires to reject the offer, then Player A knows that Player B has high bargaining cost. The resulting game is again of complete information, and the equilibrium outcome is the division $(1 - c_A, c_A)$ on round $t + 1$. From the viewpoint of round t, Player B_H values this outcome as $c_H - c_A < 0$. In conclusion, Player B_H benefits by acting like Player B_L.

(iii) Suppose that Player B rejects the demand x_A and offers the share y_A to Player A in the next round. By the described strategies, Player A concludes that he (or she) is playing against a low cost player, and so Player A accepts the share y_A. (As above, rejection yields a negative outcome to Player A since Player A is now in a weak position.) Since Player B_H does not benefit by acting like Player B_L, we must have $1 - x_A \geq 1 - y_A - c_H$, or $y_A \geq x_A - c_H$. Since Player B_L does not accept the demand made by Player A, we must also have $1 - x_A \leq 1 - y_A - c_L$, or $y_A \leq x_A - c_L$. □

The main result of this section is as follows. Note that, this time, we have a large number of equilibria.

Theorem 4.16. *(i) If $p > 2c_A/(c_A + c_H)$, then the expected outcome to Player A is at least $p + (1-p) \cdot (1 - c_A - c_H)$ in every sequential equilibria.*

(ii) If $p \leq 2c_A/(c_A + c_H)$, then for every $x^* \in [c_A, 1 - c_A + c_L]$ there exists a sequential equilibrium in which Player A makes the demand x^* in the first round and both of the player B_L and B_H accept this demand.

(iii) If $(c_A + c_L)/(c_A + c_H) \leq p \leq 2c_A/(c_A + c_H)$, then for every $x^* \geq c_H$ there exists a sequential equilibrium in which Player A demands the share x^* in the first round, Player B_H accepts this offer, but Player B_L rejects this offer and demands the share $1 - x^* + c_H$ in the second round, which Player A accepts.

Proof. For a full proof, see Osborne and Rubinstein (1990: 99–104). Again, we only discuss some features of this statement.

The first statement follows from similar arguments as Theorem 4.4. Statements (ii) and (iii) are perhaps surprising, since these imply that there are sequential equilibria in which Player A obtains a major share of cake even if the probability that Player A is in a weak position is high. Actually, this follows from the fact that we have not restricted how Player A updates the probability p in the case of inconsistent actions.

For statement (ii), equilibrium strategies are given as follows. Player A demands the share x^* and accepts any offer which gives at least the share $x^* - c_A$ (which is the value of obtaining x^* in the next round), and both types of Player B demand the share $1 - x^* + c_A$ and B_L accepts an offer which gives at least the share $1 - x^*$ and B_H accepts an offer which gives at least the share $1 - x^* - c_H + c_A$. If Player B is revealed to be B_L, the play proceeds as in the case of complete information. (By the previous lemma, Player B is revealed to be B_H only in the case that Player B_H accepts some offer, but the game ends after this.) If Player A observes an action which is not described in the above strategy for either of the players B_L or B_H, Player A makes an optimistic guess that $p = 1$. After this, the game proceeds as the game of complete information between Player A and Player B_H, and Player A has all the bargaining power in the resulting game. The outcome of these strategies is the agreement $(x^*, 1 - x^*)$ in the first round.

Statement (iii) follows from a similar set of strategies, and the crucial point of these strategies is the optimistic guess made by Player A. □

When all the bargaining costs are small (that is, c_A, c_L, and c_B are close to zero) and p is large, statement (i) implies that Player A obtains a large share of cake in all sequential equilibria. This is not surprising, since in this case Player A knows that the opponent is weak bargainer with high probability. Otherwise, the concept of sequential equilibrium is not very informative with the game of this section. Even when the probability $1 - p$ that Player A faces a strong opponent is close to one, statements (ii) and (iii) of the previous theorem imply that practically any outcome can be supported by sequential equilibrium. In the above described strategies, these outcomes rely strongly on the assumption that Player A is free to form new beliefs in the case of inconsistent actions. Even if the probability of facing a strong opponent is large, in the case of inconsistent action Player A makes an optimistic guess that this probability is zero. Once Player A makes this guess, he (or she) sticks with it and the resulting game is practically between Player A and Player B_H. This optimistic guess gives credibility to a tough bargaining strategy of Player A and allows for a wide range of equilibria to be generated.

There are still other sequential equilibria where agreement is reached later than in round two, see Osborne and Rubinstein (1990: 104–107). As noted, the above described strategies are based on the optimistic guess made by Player A in the case of any deviation made by Player B. Some deviations of Player B_L, however, can be rationalized. Suppose that Player A offers some share x_B to Player B and that Player B rejects and demands y_B with $x_B + c_L < y_B < x_B + c_H$. In this case, it is actually unreasonable for Player A to conclude that he (or she) is playing against Player B_H. Had Player B accepted the share x_B, the outcome to Player B would have been x_B on the given round. But if Player A accepts the demand y_B, the outcome to Player B is y_B with one period of delay. Player B_L values this as $y_B - c_L > x_B$ and Player B_H values this as $y_B - c_H < x_B$. Therefore, it is reasonable that Player A observes rejection and the above demand only from Player B_L. With the above described strategy, Player A makes immediately the conclusion that he (or she) faces Player B_H without worrying whether such an action is reasonable for Player B_H, and this gives him (or her) bargaining power. Restricting the concept of sequential equilibrium by assuming that the beliefs of Player A are *rationalizing* in a similar manner as just discussed, the set of equilibria can be narrowed dramatically. A problem with this approach, however, is that Player A is trying to rationalize *any* deviation of Player B. See Osborne and Rubinstein (1990: 107–112) for further details.

4.9 Some additional comments

In this chapter, we have modeled the bargaining procedure using Alternating Offers, that is, the players make offers and counteroffers each in turn after equal times. There are also other approaches modeling the bargaining procedure. For example, Player A may make all the offers and Player B may only accept or reject these offers, or the player to make an offer is chosen randomly on every round, or the players may use different amounts of time to respond to an offer. See Osborne and Rubinstein (1990: 173–188) or Muthoo (1999: 187–210) for a discussion about different trading procedures and their effect on the equilibrium outcomes.

In Section 3, we have used fixed utility functions $U_A(x) = U_B(x) = x$ and discount factors to evaluate the outcome to the players in the case of agreement. In the original paper Rubinstein (1982), Rubinstein does not use any particular utility functions. Instead, he begins with some assumptions on players' time preferences (that is, how the players preferences depend on the time when some amount of cake is obtained) and concludes that preferences satisfying these assumptions can be represented by suitable utility functions. In particular, this approach includes the utility functions used in this thesis, but also allows for more general functions. This approach is also discussed in Osborne and Rubinstein (1990: 29–67). The proof given in Section 3 is not the original proof presented by Rubinstein, but follows Shaked and Sutton (1984) where a simplified version of the proof was presented. Muthoo (1999) also presents results for more general utility functions for the models of this chapter.

In Section 5, we have concentrated on models of alternating offers where only one of the players has an outside option. See Muthoo (1999: 99–135) for models where both of the players have outside options. For simplicity, we have studied the effect of an inside option separately in Section 6. However, inside and outside options can be combined into same model. This means that, in every stage of the game, the players have three options, namely accept the offer, exercise the inside option by rejecting and make a counteroffer in the next round, or exercise the outside option and quit permanently. See Muthoo (1999: 146–153) for a model which includes both of these options.

In Section 7, we studied commitment tactics. The choice of commitment levels is similar to Nash demand game presented in Nash (1953). In this game, the players announce simultaneously demands x and y. If $x + y \leq 1$, then the players obtain the requested shares of cake, but neither obtains anything if the demands are incompatible, that is, $x+y > 1$. Nash presented this game to support his axiomatic solution treated in Chapter 3. See Binmore (2007b: 496–500), Montet and Serra (2003: 240–241), or Osborne and Rubinstein (1990: 76–81) for a discussion about this game.

Games of incomplete information have drawn a lot of attention and a number of variations (also with both sided incomplete information) have been studied. See Fudenberg and Tirole (1995: 397–434) for a treatment of some of these models, or Osborne and Rubinstein (1990: 119–120) for a short discussion about different models, or the collection Roth (1985) for a number of articles in this field.

5 FOLK THEOREM AND COASE THEOREM

In the previous chapter, we studied some game theoretic models of bargaining and how different options during negotiations affect on bargaining outcome. Due to sequential structure of bargaining situations studied, we were mostly interested in the subgame perfect (or sequential) equilibria of the resulting models. In this short chapter, we do not consider any particular models of bargaining, but we discuss two related issues, namely Folk Theorem and Coase Theorem.

5.1 Folk Theorem

Folk Theorem[13] is a statement about repeated interaction between players. It applies to a wide variety of games, but to make the ideas transparent let us concentrate on Prisoner's Dilemma. In this game, both of the player have two strategies, D (defect) and C (cooperate), and the payoff matrix is given as follows.

	D	C
D	-10,-10	0,-12
C	-12,0	-1,-1

This is a famous example where individual rationality yields suboptimal outcome. The game has unique Nash equilibrium (D, D) (D is a dominant strategy for both of the players), but (C, C) would give a better outcome to both of the players. The problem is that the players make their decisions to maximize their personal outcome and they are not able to make a binding contract to play C. (This would require some enforcement mechanism, for example, established by government). This kind of reasoning applies also to a number of other situations, for example, to nuclear weapons: it would be better to have a world without nuclear weapons, but having such weapons is dominant strategy. See McCain (2010: 349–372) for an elementary discussion and other applications of Prisoner's Dilemma.

Prisoner's Dilemma yields the non–cooperative outcome when it is played only once. Let us think what happens in the case that the players play this game more than once. One might hope that a repeated situation would give an incentive for the players to cooperate, but this is not always the case: the outcome depends on

[13]Folk comes from folklore. This statement seemed to be folklore of game theorists for a long time.

how many times the game is played, and in finitely repeated games cooperating is not equilibrium behavior. Suppose that the game is played N times and that both of the players know this number. We apply backward induction in our reasoning. In the last round, the players have no reason to cooperate, and so both players defect. This means that the players have no reason to cooperate in the second to last round (building reputation as a cooperative player is made vain by the next round), and so the players defect also on this round. Continuing the reasoning, we see that the players defect on every round.

The point in here is that the players know the precise number N of rounds. All promises to cooperate or attempts to build reputation as a cooperative player are made empty by the last round where the players have a strong incentive to switch to dominant strategy D. But note that this argument relies strongly on the assumption that both of the players know the precise number of rounds.

Suppose now that the players play Prisoner's Dilemma repeatedly but the number of rounds is not fixed, that is, the play might continue for infinite number of rounds. In this setting, we see that the players have an incentive to cooperate. This game allows for strategies which *punish* non–cooperative behavior. One such strategy is *tit–for–tat*, where player begins by cooperating, and on later rounds makes the same decision as the other player in the previous round. With this strategy, a player begins by playing cooperatively and continues cooperating as long as the other player does. If the other player defects, the player in question continues cooperating only after the other player has done so. Other punishment strategies are different *trigger strategies*: a player plays cooperatively as long as the other player cooperates, and after observing a defection plays defect for some fixed number (which might be infinite) of rounds. The point of these strategies is that they support cooperation by punishing non–cooperative behavior.

Folk Theorem deals with infinitely repeated games, and the conclusion is very different compared to the case of finitely repeated games. For this statement, we need the following definitions. First, the *payoff region* is the convex hull of player's outcomes in the one–shot game. Second, the *security level* of a player is the smallest outcome the player can guarantee from the one–shot game. Folk Theorem is given as follows.

Theorem 5.1. *If the players are not too impatient, then any outcome in the payoff region where the players get at least their security levels can be supported by Nash equilibrium in infinitely repeated game.*

Note that the condition that the players are not too impatient guarantees that they do not value first rounds at the expense of later rounds. Otherwise, the players would seek for large outcome in the first rounds and they would have a strong incentive to defect on early rounds. When the players are patient, they value long–time cooperation. See Fudenberg and Tirole (1995: 150–160) for a more precise discussion about this theorem.

Previously, we have been able to single out some outcomes as equilibria, but Folk Theorem says that one can expect practically any outcome to occur, so one might ask what is the point of Folk Theorem. The important point of Folk Theorem is not on what it says about possible outcomes, but the fact that cooperation can be *self–sustaining* (equilibrium behavior). Repeated interactions give players incentives to find cooperative outcomes *without need for any outside parties*. This is in sharp contrast compared to finitely repeated games, where cooperation is not equilibrium behavior. Also, to escape the logic of finitely repeated games, it is not necessary that there exists an infinite number of rounds: it is enough that the game continues with high enough probability. Non–cooperative outcome of finitely repeated game is based on the existence of last round, which makes cooperation unattainable on every round of the game. However, if the interaction continues with high enough probability, the players have an incentive to cooperate to support cooperation also in the future. See Hargreaves Heap and Varoufakis (1995: 170–189) for a further discussion on this subject.

Few human interactions are of the type of finitely repeated games, where the number of interactions is precisely known. This is sometimes used to explain people's behavior in such games, like in Ultimatum game discussed previously. It seems that social norms, built by a large number of interactions, affect on people's behavior in such games, and they subconsciously apply cooperative behavior in games where game theory tells that such behavior should not be observed. See Binmore (2007a: 3–22) and references therein for a further discussion on this subject.

5.2 Coase Theorem

In the previous two chapters, we have assumed that the players are capable of negotiating over partition of a cake, that is, the players are somehow entitled to obtain a share of a cake. If the players do not possess such rights, they simply cannot bargain with each others over how to divide the cake. In this section, we discuss some ideas originally presented in Coase (1960) which deal with property rights and economic situations involving externalities. Last term is defined slightly differently depending on the literature. We apply the following description given in Perman, Ma, Common, and Maddison (2011): *Externality* occurs when the production or consumption decision of an agent affects utility (or profit) of an another agent in an unintended way, and when no compensation is paid by the generator of the impact.

A standard example of externality is pollution. Let us consider a firm producing some good and releasing some amount of pollution as a side effect of production. This affects the well–being of nearby residents and might affect negatively on production possibilities of other firms. For example, if pollution is released in water, a fishery company suffers from production decisions of polluting firm. Note in the above given description that we require that no compensation is paid. Of course, even if some compensation is paid, a fishery company still suffers from pollution, but lost production possibilities are now taken into account in the form of compensation. Coase Theorem is a statement about property rights and externalities and is given as follows.

Theorem 5.2. *If the transaction costs are low enough, the agents may achieve economically efficient equilibrium regardless to whom the property rights are given. If the transaction costs are high, laws and property rights affect the equilibrium, and it may not be efficient.*

The term "Coase Theorem" is perhaps misleading since the previous statement is not given explicitly in Coase (1960). Rather, Coase discusses a number of externality related examples, many of which were settled in court, to support his view that the agents may find efficient outcomes once property rights are established. It is not uncommon to see that only the first statement given above is labeled as Coase Theorem, but a large part of the original paper is devoted to analyzing cases where the transaction costs are high.

Let us discuss the above statement more carefully through an example. We deal with an example of a polluting firm and a victim of pollution. If the firm may pollute freely, it produces an amount which releases the amount E of pollution in Figure 20. (At this point, the marginal benefit of the firm is zero, so this is the profit maximizing amount of pollution.)

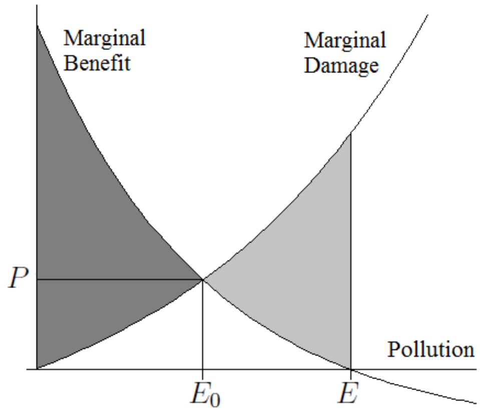

Figure 20: Coase Theorem

Efficient level of pollution is given at the point E_0, where marginal benefit and marginal damage are equal. That is, the total benefit is maximized at this level of pollution. (On the left side of E_0, marginal benefit is larger than marginal damage. A small increase in pollution would give the polluter larger benefit than the victim suffers, and so this small increase in pollution increases total benefit. Similarly, on the right side of E_0 a small reduction increases total benefit.) However, if the firm is free to make its production decision, it will produce at the level E, and there is a loss of efficiency.

The loss in efficiency yields from the fact that the firm is interested only in the profit it makes and it need not take the caused damage into account. Suppose now that property rights have been established, that is, either the firm has the right to pollute or the victim has the right for pollution–free environment. If the polluter has the right to pollute, moving from E to E_0 creates a cake the size of which is represented by the light gray area in Figure 20. Indeed, the area below the marginal benefit curve represents the forgone profit of polluter, and the total area below the marginal damage curve represents the damage lessened by the victim. Therefore, the net change is the difference of these areas, which is the light gray area. Similarly, if the victim has the property rights, the size of a cake is represented by the dark gray area.

Now, if the parties find an agreement over how to divide either of these cakes, they have reached the efficient level of pollution via bargaining. Note that although the cakes are of different sizes, agreement over either of these cakes gives the efficient level of pollution, and so it does not matter which party has the property rights. This is the content of the first statement in Coase Theorem. Of course, the resulting welfare of the agents depends to whom the property rights are given: the cakes to be divided are of different sizes, and one of the parties ends up paying compensation to the other one.

The first statement of Coase Theorem is often taken to mean that property rights create markets, and market forces drive the agents towards efficient outcome. This would mean that the market price P in Figure 20 is achieved. This interpretation, however, is sometimes problematic. In the above case this would create a market of two agents. In a competitive market, we need a large number of buyers and sellers for market forces to exist, and in a monopolistic market we need a large number of buyers whose willingness to pay determines the market demand curve. See Hahnel and Sheeran (2009) for a discussion about problems of this interpretation, and Butler and Garnett (2003) for a discussion about misinterpretations of Coase Theorem.

It is important to take a careful look about some implicit assumptions lying behind the above argument. When we claim that the agents may reach the efficient level of pollution by bargaining, we need to assume that the agents have complete information. Indeed, the agents must know the size of the cake they are bargaining over, and this means that the agents must know the marginal functions of each others. In addition, we need to assume that bargaining costs are low and that costs of establishing the property rights are low (that is, transaction costs are low). For otherwise these costs might exceed possible gains of negotiating, and the efficient level cannot be attained via bargaining. Furthermore, nothing is said about the actual bargaining process: it is enough that the agents find an agreement.

Real life bargaining is rarely costless, and possible gains might very well be too small relative to the transaction costs. Under incomplete information, which seems to be more realistic assumption than complete information, achieving the efficient level via bargaining seems to be rather difficult. For one thing, both of the parties have an incentive to try to convince the other one that their marginal functions are of different shape, as this affects the size of a cake. Samuelson (1985) studies the content of Coase theorem under incomplete information and the main results are as follows.

(i) the parties affected by externality are not always capable of negotiating an efficient agreement, or if such an agreement is reached it may come only after costly delay.

(ii) the efficiency of the negotiated agreement may depend on which party the property rights are given and also the bargaining process used.

(iii) Efficiency can be increased by allocating property rights via competitive bid.

The most important part of Coase's work is not on the argument that problems involving externalities can be solved by bargaining between the agents involved. Instead, the true impact of Coase's ideas seems to be that these provided a new way to think about externality problems. Prior to Coase, a common view was that problems involving pollution can be solved by taxing the polluter. Introducing a tax to the polluter transforms the marginal benefit curve of the polluter, and the efficient level of pollution can be obtained. (In Figure 20, this would require a tax which transforms the marginal benefit curve downwards such that it intersects the x–axis at E_0.) Coase does not claim that taxes are not an appropriate way to handle externalities, and admits that sometimes governmental intervention is needed, but he argues that this is not *always* the most efficient way. Also, the above discussed bargaining solution requires heavy informational assumption, but so does the taxing approach: tax authorities should know the true shape of the marginal benefit function of the firm to determine the appropriate tax level in a seek for efficient level of pollution. See Medema (2013) for a more detailed discussion.

6 CONCLUSIONS

The purpose of this thesis was to study some theoretical models of two–person bargaining situations. The thesis surveys research on this subject and we have not presented new results. It was our intention to proceed in such a way that each of the chosen models illustrates what kind of effect some factor present in bargaining situation has on the outcome. We have tried to favor simplicity and clarity over generality, and we have not presented very general models: all the models treated in the thesis can ge generalized. We have also tried to avoid a large number of technicalities. Many of the statements would have required proper reasoning, but we have chosen to present verbal explanations rather than technical expositions. Each of the results presented serves some chosen purpose, but the most important results of the thesis are without doubt Theorems 3.7 and 4.4, as discussed previously.

It is certainly reasonable to ask how accurately the studied models describe real life bargaining situations. A drawback of any model is that one needs to make a number of simplifying assumptions to keep the model in reasonable boundaries: all the factors present in real life bargaining situations simply cannot be taken into account. We have used equilibrium concepts of game theory to analyze the models, and so the theory presented suffers from same weaknesses as game theory in general. In particular, in comparison to real life bargaining situations we might assume too much on players' abilities when we assume that each player is *homo economicus*. This means that the results presented cannot be taken to accurately describe peoples' behavior in bargaining situations, but this is neither the point of studying bargaining models. Instead, these results provide approximations about how reasonable people behave *on average*. Also, a key element in a bargaining situation is the position, bargaining power, of a bargainer, and the models have illustrated how bargaining power can be improved.

Bargaining is a central element of peoples' everyday life, ranging from marriage to acquisitions of multinational companies, and thereby certainly an interesting object of research. In real life, it is such a complex process that it seems overly optimistic to claim that it can be described by some set of equations. A central aim of research on bargaining, however, is not to give precise predictions, but to understand different driving forces in bargaining situations. It seems that there is still a lot of research to be done in this area.

BIBLIOGRAPHY

Binmore, K. (1994). *Playing Fair: Game Theory and the Social Contract I.* Cambridge: MIT Press.

Binmore, K. (1998). *Just Playing: Game Theory and the Social Contract II.* Cambridge: MIT Press.

Binmore, K. (2007a). *Does game theory work? The bargaining challenge.* Cambridge: The MIT Press.

Binmore, K. (2007b). *Playing for Real: A Text on Game Theory.* Oxford: Oxford University Press.

Butler, M. R. & Garnett, R. F. (2003). Teaching the Coase Theorem: Are We Getting It Right? *Atlantic Economic Journal* 31, 133–145.

Chatterjee, K. (2014). Game Theory and the Practice of Bargaining. In: Chatterjee, K. & Samuelson, W. (eds.). *Game Theory and Business Applications.* New York: Springer Science+Business Media, 189–206.

Chiang, A. F. & Wainwright, K. (2005). *Fundamental Methods of Mathematical Economics* (4th edition). Singapore: McGraw-Hill.

Coase, R. H. (1960). The Problem of Social Cost. *Journal of Law and Economics* 3, 1–44.

Cowell, F. (2006). *Microeconomics: Principles and Analysis.* Oxford: Oxford University Press.

Cross, J. G. (1965). A theory of the bargaining process. *American Economic Review* 55(1), 67–94.

Fudenberg, D. & Tirole, J. (1995). *Game Theory.* Cambridge: The MIT Press.

Gravelle, H. & Rees, R. (2004). *Microeconomics* (3rd edition). Harlow: Pearson Education Limited.

Hahnel, R. & Sheeran, K. A. (2009). Misinterpreting the Coase Theorem. *Journal of Economic Issues* 43(1), 215–237.

Hargreaves Heap, S. P. & Varoufakis, Y. (1995) *Game Theory: A Critical Introduction.* London: Routledge.

Harsanyi, J. C. & Selten, R. (1972). A Generalized Nash Solution for Two-Person Bargaining Games with Incomplete Information. *Management Science* 18(5), 80–106.

Kalai, E. & Smorodinsky, M. (1975). Other Solutions to Nash's Bargaining Problem. *Econometrica*, 43(3), 513–518.

Korth, C. (2009) *Fairness in Bargaining and Markets*. Berlin: Springer-Verlag.

Luce, R. D. & Raiffa, H. (1957). *Games and Decisions*. New York: John Wiley & Sons, Inc.

Maschler, M., Solan, E. & Zamir, S. (2013). *Game Theory*. Cambridge: Cambridge University Press.

McCain, R. A. (2010). *Game theory: A Nontechnical Introduction to the Analysis of Strategy*. Singapore: World Scientific Publishing Co. Pte. Ltd.

Medema, S. G. (2013). The Importance of Being Misunderstood: The Coase Theorem and the Legacy of 'The Problem of Social Cost'. *Journal of Natural Resources Policy Research*, 5(4), 249–253.

Montet, C. & Serra, D. (2003). *Game Theory and Economics*. New York: Palgrave MacMillian.

Muthoo, A. (1999). *Bargaining Theory with Applications*. Cambridge: Cambridge University Press.

Napel, S. (2002). *Bilateral Bargaining: Theory and Applications*. Berlin: Springer-Verlag.

Nash, J. F. (1950). The Bargaining Problem. *Econometrica*, 18(2), 155–162.

Nash, J. (1953). Two-Person Cooperative Games. *Econometrica*, 21(1), 128–140.

Osborne, M. J. (2009). *An Introduction to Game Theory*. Oxford: Oxford University Press, Inc.

Osborne, M. J. & Rubinstein, A. (1990). *Bargaining and Markets*. San Diego: Academic Press, Inc.

Owen, G. (2008). *Game Theory* (3rd edition). Bingley: Emerald Group Publishing.

Perman, R., Ma, Y., Common, M., Maddison, D., & McGilvray, J. (2011). *Natural resource and environmental economics* (4th edition). Harlow: Pearson Education Limited.

Peters, H. J. M. (2010). *Axiomatic Bargaining Game Theory*. Dordrecht: Kluwer Academic Publishers.

Rasmusen, E. (2007). *Games and Information: An Introduction to Game Theory* (4th edition). Oxford: Blackwell Publishing.

Roth, A. E. (1977). Individual Rationality and Nash's Solution to the Bargaining Problem. *Mathematics of Operations Research*, 1, 64–65.

Roth, A. E. (1979). *Axiomatic Models of Bargaining.* Berlin: Springer-Verlag.

Roth, A. E. (1985). *Game-Theoretic Models of Bargaining.* Cambridge: Cambridge University Press.

Rubinstein. A. (1982). Perfect Equilibrium in a Bargaining Model. *Econometrica*, 50(1), 97–109.

Samuelson, W. (1985). A Comment on the Coase Theorem. In: Roth, A. E. (ed.). *Game-Theoretic Models of Bargaining.* Cambridge: Cambridge University Press, 321–339.

Schelling, T. (1960). *The Strategy of Conflict.* Cambridge: Harvard University Press.

Shaked, A., & Sutton, J. (1984). Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model, *Econometrica*, 52, 1351–1364.

Ståhl, I. (1972). *Bargaining Theory.* Stockholm: Economics Research Institute Stockholm School of Economics.

Tadelis, S. (2013). *Game Theory: An Introduction.* New Jersey: Princeton University Press.