A New (Simplified) Derivation of Nash's Bargaining Solution

Hoang Phuong Nguyen[1], Laxman Bokati[2], and Vladik Kreinovich[2,3]
[1]Division Informatics, Math-Informatics Faculty
Thang Long University
Nghiem Xuan Yem Road
Hoang Mai District
Hanoi, Vietnam, nhphuong2008@gmail.com
[2]Computational Science Program
[3]Department of Computer Science
University of Texas at El Paso
500 W. University
El Paso, TX 79968, USA
lbokati@utep.edu,vladik@utep.edu

Abstract

According to the Nobelist John Nash, if a group of people wants to selects one of the alternatives in which all of them get a better deal than in a status quo situations, then they should select the alternative that maximizes the product of their utilities. In this paper, we provide a new (simplified) derivation of this result, a derivation which is not only simpler – it also does not require that the preference relation between different alternatives be linear.

1 Introduction

Practical problem. In many practical situations, a group of people needs to make a joint decision. They can stay where they are – in the "status quo" state. However, they usually have several alternatives in which each of them gets a better deal than in the status quo state. Which of these alternatives should they select?

Nash's bargaining solution to this problem. To solve this problem, in 1950, the Nobelist John Nash formulated several reasonable conditions that the selection must satisfy [4]. He showed that the only way to satisfy all these conditions is to select the alternative that maximizes the product of participants' *utilities* – special functions that describe a person's preferences in decision theory; see, e.g., [1, 2, 3, 5, 6].

Comment. For readers' convenience, in the following section we briefly describe what is utility.

What we do in this paper. In this paper, we provide a new simplified derivation of Nash's bargaining solution, a derivation that is not only simpler – this derivation also uses fewer assumptions: e.g., it does not assume that there is total (linear) pre-order between different alternatives.

2 What Is Utility: A Brief Reminder

How can we describe person's preferences? Let us select one of the participants. For this participants, we have a status quo situation A_0, and we have several other situations A of what this person will get according to different alternatives described by the group. In all these situations, A is better than A_0 – otherwise, this participant will not agree to select this alternative. We will denote this preference by $A_0 < A$.

How can we describe this person's preferences.

For this purpose, let us select a very good ("ideal") alternative A_+, an alternative that for this person will be better than anything that he or she can get in real life. Then, for each real number p from the interval $[0, 1]$, we can form a lottery $L(p)$ in which:

- the participant will get A_+ with probability p and
- the participant will retain his/her status-quo situation in all other cases – i.e., with probability $1 - p$.

To provide a numerical estimate of the quality of an alternative A, we can then ask the participant to compute this alternative A with lotteries $L(p)$ corresponding to different values p.

When $p = 0$, then in the corresponding lottery $L(p)$, we simply get the status quo with probability 1. Thus, this lottery is clearly worse than A: $L(0) < A$. As the probability of the very good outcome increases, the lottery becomes better and better, and for $p = 1$, this lottery will mean getting A_+ with probability 1 – which is, clearly, better than A: $A < L(1)$.

One can show that there exists a threshold

$$u(A) \stackrel{\text{def}}{=} \sup\{p : L(p) < A\} = \inf\{p : A < L(p)\}$$

such that:

- for probabilities $p < u(A)$, we have $L(p) < A$, and
- for probabilities $p > u(A)$, we have $A < L(p)$.

This threshold $u(A)$ is known as the *utility* of the alternative A.

From the above properties of the utility, it follows that for each $\varepsilon > 0$, we have

$$L(u(A) - \varepsilon) < A < L(u(A) + \varepsilon).$$

For very small ε, the probabilities $u(A) - \varepsilon$, $u(A)$, and $u(A) + \varepsilon$ are practically indistinguishable. So, from the practical viewpoint, we can say that the lottery $L(u(A))$ is *equivalent* to the alternative A. We will denote this equivalence by

$$A \equiv L(u(A)).$$

Comment. By definition, the status quo alternative A_0 has utility 0.

Utility is defined modulo scaling. The above definition of utility depends on the selection of a very good alternative A_+. What if we select a different very good alternative A'_+? What is the relation between the utility $u(A)$ corresponding to A_+ and the utility $u'(A)$ corresponding to A'_+?

Without losing generality, let us assume that A_+ is better than A'_+, i.e., that $A_0 < A'_+ < A_+$. In this case, for the alternative A'_+, we can find the A_+-based utility value $u(A'_+)$ for which A'_+ is equivalent to the lottery $L(u(A'_+))$, a lottery in which:

- the participant will get A_+ with the probability $u(A'_+)$ and
- the participant will remain in the status quo state with the remaining probability $1 - u(A'_+)$.

Suppose now that for some alternative A, we know its A'_+-based utility $u'(A)$. By definition of utility, this means that the alternative A is equivalent to a lottery $L'(u'(A))$, in which:

- we get A'_+ with probability $u'(A)$ and
- we get the status quo with the remaining probability $1 - u'(A)$.

As we have mentioned, the alternative A'_+ is, in its turn, equivalent to the lottery $L(u(A'_+))$. Thus, the original alternative A is equivalent to a two-stage lottery in which:

- first, we select either A'_+ (with probability $u'(A)$) or the status quo – with probability $1 - u'(A)$;
- if we selected A'_+ on the first stage, then on the second stage, we select either A_+ (with probability $u(A'_+)$), or we retain the status quo state with the remaining probability $1 - u(A'_+)$.

In this two-stage lottery, we end up either with A_+ or with the status quo. The probability to get A_+ is equal to the product $u'(A) \cdot u(A'_+)$. By definition of utility, this is exactly the A_+-based utility $u(A)$ of the alternative A. So,

$$u(A) = u'(A) \cdot u(A'_+).$$

Thus, if we change the very good state, all the values of the utility get multiplied by some constant. In this sense, utility is defined modulo such "re-scaling"

$$u(A) \to c \cdot u(A).$$

Need for scale-invariance. For each participant, we can select different very good alternatives A_+ and thus, get different numerical values of his/her utility. This selection is very arbitrary and hence, should not affect what decision the group makes. In other words, the group's decision should not change if we simply re-scale the utilities of one (or more) of the participants.

Comment. Now, we are ready to formulate our main result.

3 A New (Simplified) Explanation of Nash's Bargaining Solution

Let us describe the problem in precise terms. We consider a situation in which n participants need to make a joint decision. In this situation. Each possible alternative can be characterized by a tuple $u = (u_1, \ldots, u_n)$ of the corresponding utility values. We only consider alternatives in which $u_i > 0$ for all i – otherwise why would the i-th person agree to this alternative if he/she does not gain anything – or even lose something?

For some alternatives u and u', the group prefers u' to u. For example, if each person gets more in alternative u' than in u, then u' is clearly better than u. We will denote this preference by $u < u'$.

For some alternatives u and u', the group may consider them equally good; we will denote this by $u \sim u'$. We also allow the possibility that for some pairs of alternatives u and u', the group cannot decide which of them is better. In other words, we do not assume that the preference relation is total (linear).

The relations $<$ and $\sim$ should satisfy natural transitivity conditions: e.g., if u' is better than u and u'' is better than u', then u'' should be better than u. We will call such a pair of relations $(<, \sim)$ a *preference relation*; we will give a precise definition shortly.

We assume that the preference relation is *fair* in the sense that all participants are treated equally. In particular, if we perform any permutation of the utilities, the alternative should remain of the same quality to the group: e.g., (a, b) should be of the same quality as (b, a): $(a, b) \sim (b, a)$.

Finally, since utility u_i of each participants is determined modulo re-scaling $u_i \to c_i \cdot u_i$, relative preference of two different alternatives should not change if we perform such a re-scaling. Thus, we arrive at the following definitions.

Definition 1. *Let A be a set; its elements will be called* alternatives. *By a* preference relation *on the set A, we mean a pair of relations $(<, \sim)$ with the following properties:*

- *for each a, we have $a \sim a$;*
- *for each a and b, if $a \sim b$, then $b \sim a$;*
- *for each a and b, if $a \sim b$, then we cannot have $a < b$;*
- *for each a, b, and c, if $a < b$ and $b < c$, then $a < c$;*
- *for each a, b, and c, if $a < b$ and $b \sim c$, then $a < c$;*
- *for each a, b, and c, if $a \sim b$ and $b < c$, then $a < c$;*
- *for each a, b, and c, if $a \sim b$ and $b \sim c$, then $a \sim c$.*

Definition 2. *Let $A = \mathbb{R}^n_+$ be a set of all n-tuples $u = (u_1, \ldots, u_n)$ on positive numbers, and let $(<, \sim)$ be a preference relation on the set A.*

- *We way that the pre-order is* monotonic *if whenever we have $u_i < u'_i$ for all i, then we should have $u < u'$.*
- *We say that the pre-order is* fair *if for each permutation*

$$\pi : \{1, \ldots, n\} \to \{1, \ldots, n\}$$

  and for each alternative u, we have $(u_1, \ldots, u_n) \sim (u_{\pi(1)}, \ldots, u_{\pi(n)})$.

- *We say that the pre-order is* scale-invariant *if for every two alternatives u and u' and for each tuples $(c_1, \ldots, c_n)$ of positive numbers:*
  - *if $(u_1, \ldots, u_n) < (u'_1, \ldots, u'_n)$ then*

$$(c_1 \cdot u_1, \ldots, c_n \cdot u_n) < (c_1 \cdot u'_1, \ldots, c_n \cdot u'_n);$$

  - *if $(u_1, \ldots, u_n) \sim (u'_1, \ldots, u'_n)$ then*

$$(c_1 \cdot u_1, \ldots, c_n \cdot u_n) \sim (c_1 \cdot u'_1, \ldots, c_n \cdot u'_n).$$

Proposition 1. *There is one and only one monotonic fair scale-invariant preference relation:*

$$(u_1, \ldots, u_n) < (u'_1, \ldots, u'_n) \Leftrightarrow \prod_{i=1}^{n} u_i < \prod_{i=1}^{n} u'_i;$$

$$(u_1, \ldots, u_n) = (u'_1, \ldots, u'_n) \Leftrightarrow \prod_{i=1}^{n} u_i = \prod_{i=1}^{n} u'_i.$$

Comments.

- So, we indeed have a new explanation for Nash's bargaining solution.
- Actually, the above preference relation has a stronger property of *strong monotonicity*: that if $u_i \le u'_i$ for all i and $u_i < u'_i$ for some i, then $u < u'$.

Proof. It is easy to check that the preference relation corresponding to Nash's bargaining solution is indeed monotonic, fair, and scale-invariant. So, to complete the proof, it is sufficient to show that every monotonic, fair, and scale-invariant bargaining solution is indeed described by the above formulas.

Indeed, due to symmetry, for all values u_2, u_3, $\ldots$, u_n, we have

$$(1, u_2, u_3, \ldots, u_n) \sim (u_2, 1, u_3, \ldots, u_n).$$

For each value $u_1 > 0$, we can use scale-invariance with $c_1 = u_1$ and $c_2 = \ldots = c_n = 1$ and conclude that

$$(u_1, u_2, u_3, \ldots, u_n) \sim (u_1 \cdot u_2, 1, u_3, \ldots, u_n).$$

So, we replace two values u_1 and u_2 with the product $u_1 \cdot u_2$ and 1 without losing equivalence.

Similarly, we can replace the two values $u_1 \cdot u_2$ and u_3 with the product $(u_1 \cdot u_2) \cdot u_3$ and 1, so

$$(u_1 \cdot u_2, 1, u_3, u_4, \ldots, u_n) \sim (u_1 \cdot u_2 \cdot u_3, 1, 1, u_4, \ldots, u_n)$$

and thus, by transitivity – which is part of the definition of the preference relation – we get

$$(u_1, u_2, u_3, u_4, \ldots, u_n) \sim (u_1 \cdot u_2 \cdot u_3, 1, 1, u_4, \ldots, u_n).$$

We can then similarly absorb u_4, etc., until we get

$$(u_1, \ldots, u_n) = (u_1 \cdot \ldots \cdot u_n, 1, \ldots, 1).$$

So all alternatives with the same value of the product $u_1 \cdot \ldots \cdot u_n$ are equivalent to the same alternative

$$(u_1 \cdot \ldots \cdot u_n, 1 \ldots, 1)$$

and are, thus, equivalent to each other.

Because of this property, each alternative

$$(p, 1, \ldots, 1)$$

is equivalent to

$$(\sqrt[n]{p}, \sqrt[n]{p}, \ldots, \sqrt[n]{p}).$$

When $p < p'$, then $\sqrt[n]{p} < \sqrt[n]{p'}$. So, due to monotonicity,

$$(\sqrt[n]{p}, \sqrt[n]{p}, \ldots, \sqrt[n]{p}) < \left(\sqrt[n]{p'}, \sqrt[n]{p'}, \ldots, \sqrt[n]{p'}\right)$$

and thus,

$$(p, 1, \ldots, 1) < (p', 1, \ldots, 1).$$

So, indeed, alternatives with the larger value of the product are better.

The proposition is proven.

Comment. In the previous text, we dismissed the possibility of alternatives with some of the values 0. It turns out that this dismissal can also be justified on mathematical grounds. Namely, let us show that in this case, no preference relation can satisfy all the above requirements.

Proposition 2. *On the set $A = \mathbb{R}^n_{\ge 0}$ of all tuples with non-negative components, no preference relation is strongly monotonic, fair, and scale-invariant.*

Proof. Due to symmetry, $(1,0) \sim (0,1)$. By using $c_1 = 2$ and $c_2 = 1$, we conclude that $(2,0) \sim (0,1)$, so by transitivity $(1,0) \sim (2,0)$, which contradicts to strong monotonicity. The proposition is proven.

Acknowledgments

This work was supported in part by the US National Science Foundation grants 1623190 (A Model of Change for Preparing a New Generation for Professional Practice in Computer Science) and HRD-1242122 (Cyber-ShARE Center of Excellence).

References

[1] P. C. Fishburn, *Utility Theory for Decision Making*, John Wiley & Sons Inc., New York, 1969.

[2] V. Kreinovich, "Decision making under interval uncertainty (and beyond)", In: P. Guo and W. Pedrycz (eds.), *Human-Centric Decision-Making Models for Social Sciences*, Springer Verlag, 2014, pp. 163–193.

[3] R. D. Luce and R. Raiffa, *Games and Decisions: Introduction and Critical Survey*, Dover, New York, 1989.

[4] J. Nash, "The bargaining problem", *Econometrica*, 1950, Vol. 18, No. 2, pp. 155–162.

[5] H. T. Nguyen, O. Kosheleva, and V. Kreinovich, "Decision making beyond Arrow's 'impossibility theorem', with the analysis of effects of collusion and mutual attraction", *International Journal of Intelligent Systems*, 2009, Vol. 24, No. 1, pp. 27–47.

[6] H. Raiffa, *Decision Analysis*, McGraw-Hill, Columbus, Ohio, 1997.