The Bargaining Problem

9/30/2020

Speaker: Todd Stambaugh

Overview

- Introduction
- Nash Solution
- Kalai-Smorodinski Solution
- Comparison
- Coincidence of the Two
- Utilitarian Solutions
- Egalitarian Solutions
- More Coincidence Results
- Future Work

Introduction
Basics

- Two “highly rational” agents with complete utility functions on a space of possible outcomes
- Equal bargaining skill
- Both potentially stand to benefit from engaging in a bargain

Introduction
Individual Utility Functions

- Each player can decide on what is preferable between any two anticipations (bargains)
- The ordering produced is transitive
- If A is preferable to B which is preferable to C, there is a probability combination of A and C as desirable as B
- Equally desirable states are substitutable
- In short, these are utility functions in the sense of Von Neumann and Morgenstern

Introduction
Individual Utility Functions

- These assumptions guarantee a real valued cardinal utility function for each agent (not unique) with the following property:

For "anticipations" A and B, i=1,2 and $0 \leq p \leq 1$,

$$u_i[pA + (1-p)B] = pu_i(A) + (1-p)u_i(B)$$

Introduction
Two-Person Anticipations

- Nash defines a two person anticipation [A,B] such that the following holds:

$$p[A,B]+(1-p)[C,D]=[pA+(1-p)C\,,pB+(1-p)D]$$

- By convention, we choose utility functions such that the utility associated to the anticipation in which there is no cooperation is 0 for each player (this is actually more significant than mere convention, but for the purposes it will do)

Introduction
Graphical Representation

- Once we have appropriate utility functions for the players, we produce a graph in the real plane by plotting the utilities of each possible anticipation as ordered pairs
- Requirement: The resulting graph S must have the following properties:
  - Convexity
  - Compactness

Introduction
Goal

- The goal of the problem is to find a function which gives a unique point in S which represents a “fair” bargain
- What constitutes fairness in this context?

Nash Solution Basic Notation

- Set of utilities for possible anticipations: $S \subseteq \mathbb{R}^2$
- Utility functions for each player: u_1 , u_2
- Solution point c(S): c is a function from the set of convex, compact subsets of the real plane to the real plane, where c(S) is always in S.

Nash Solution Axioms

- $If\ \alpha\epsilon S\ such\ that\ there\ is\ some\ \beta\epsilon S\ with\ u_i(\alpha) < u_i(\beta)\ for\ i = 1,2\ then\ \alpha \neq c(S)$

  Pareto Optimality

- Each player wants to maximize their own gain in utility

Nash Solution Axioms

- *If S is symmetric about the line* $u_1 = u_2$ *then* $c(S)$ *is on that line*

  Symmetry

- Both agents are equally skilled at negotiating and there is no bias in the negotiation or the arbitrator is not biased

Nash Solution Axioms

- $If\ T\ contains\ S\ as\ a\ subset\ and\ c(T) \epsilon S\ then\ c(T) = c(S)$

  Independence of Irrelevant Alternatives

- If a less restrictive situation yields a solution that remains an option in a more restrictive environment, the same solution should apply in the more restrictive as well
- If a point is optimal in one set, it should remain optimal in a smaller set.

Nash Solution Theorem

- The only function that satisfies these axioms gives the point in S (and in the first quadrant) that maximizes the product $u_1 \bullet u_2$
- Such a point is guaranteed to exist by the compactness of S
- The point is unique by the convexity of S

Nash Solution Proof

- First we transform set S to S* so that the point in question is mapped to (1,1). We can do this because all of the relevant properties are preserved under multiplication of each variable by a constant.
- When we do this, we ensure that for any point (u_1, u_2) in S*, $u_1 + u_2 \leq 2$
- We now construct a square as shown in Figure 1
- By Symmetry and Pareto Optimality, c(S*)=(1,1)
- By IIA, c(S) = c(S*) = (1,1)

Nash Solution Proof

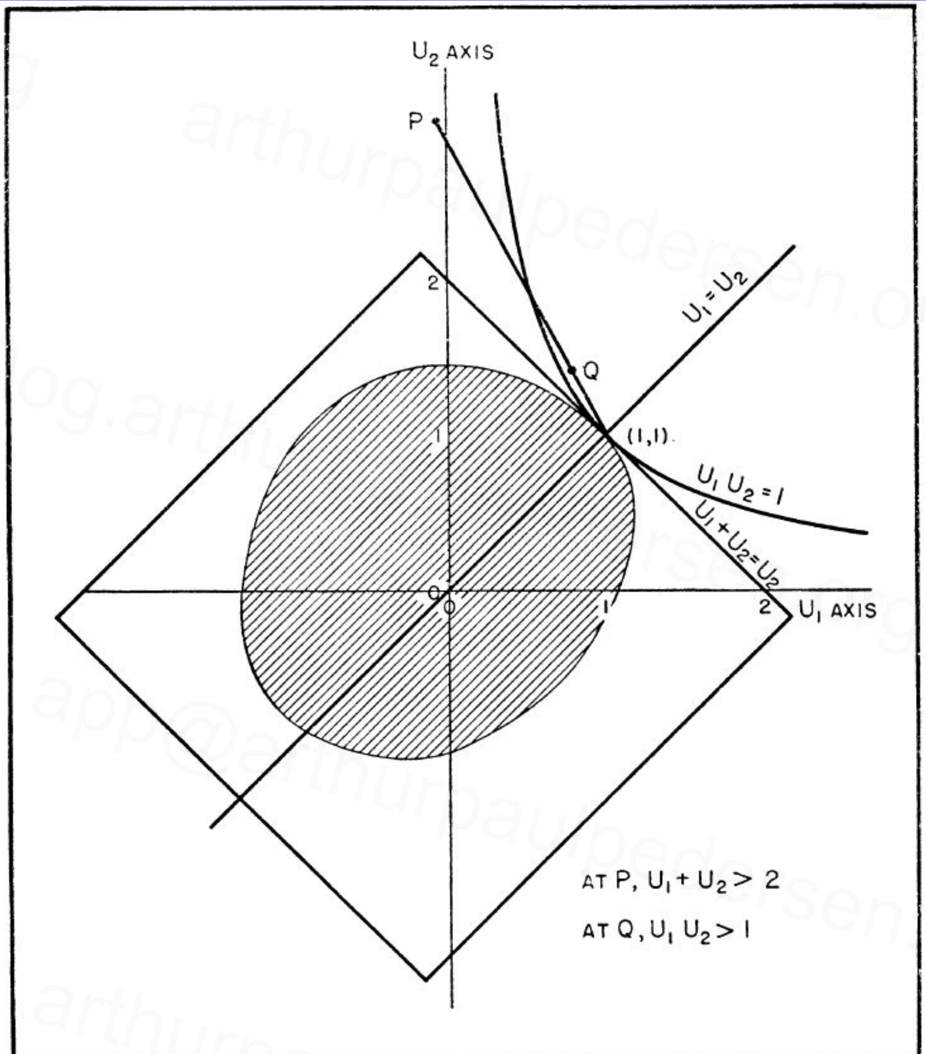

FIGURE 1

Nash Solution Example

Bill's goods	*Utility to Bill*	*Utility to Jack*
book	2	4
whip	2	2
ball	2	1
bat	2	2
box	4	1
Jack's goods		
pen	10	1
toy	4	1
knife	6	2
hat	2	2

Bill gives Jack: book, whip, ball, and bat,
Jack gives Bill: pen, toy, and knife.

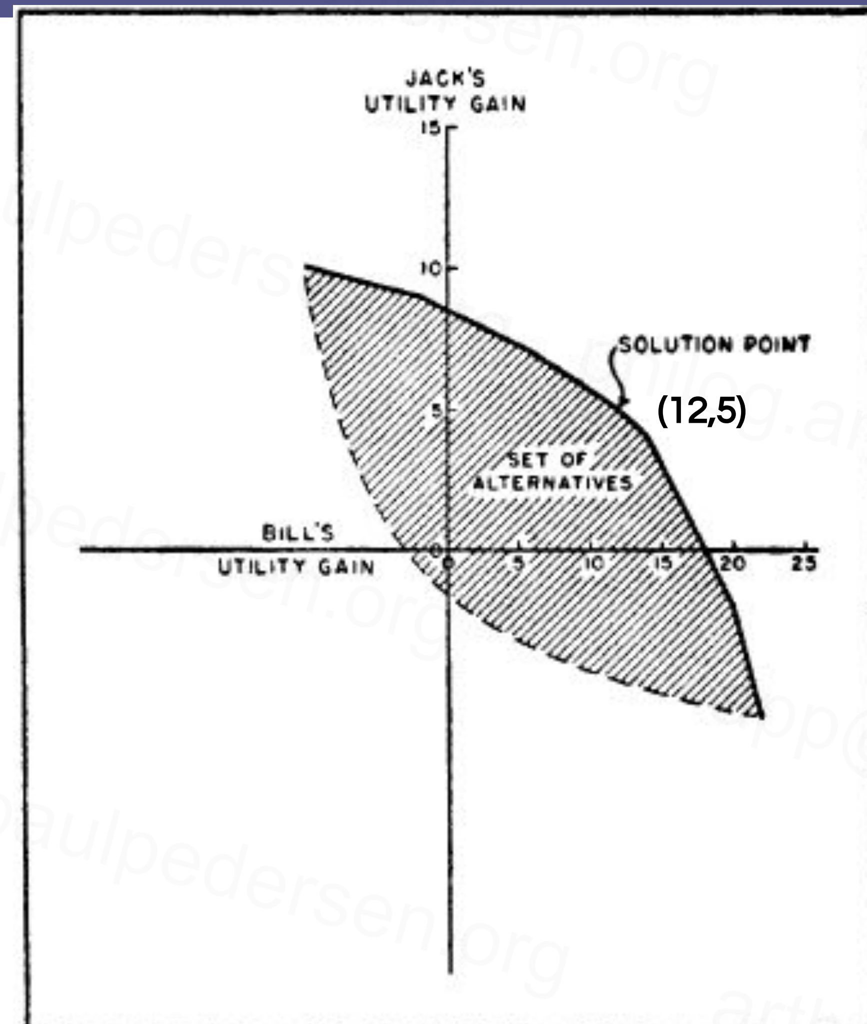

FIGURE 2

Kalai-Smorodinsky Solution Modified Formulation

- We define a bargaining pair (a,S) to be a set S as before, with $a = (a_1, a_2)$ the point in S which represents the utilities associated with no cooperation (the *disagreement point*). With this modification, the Nash solution (with an extra axiom), η(S), will be the point in S that maximizes the product $(x-a_1)(y-a_2)$.
- S is convex, compact
- New assumption: $a \leq x$ for all x in S (for Nash, this means S is contained in the first quadrant)

Kalai-Smorodinsky Solution Axioms

Let f(a,S) be the solution point for the pair

- Pareto Optimality: Same as Nash
- Symmetry: Let T(x, y)=(y, x). Then for any (a,S), f(T(a),T(S)) = T(f(a,S))
- Invariance w.r.t. Affine Transformation (also added to make Nash's solution work in this more general context): If A is an affine transformation of the plane of the form

  $A(x,y) = (c_1x + d_1, c_2y + d_2)$, then

  f(A(a),A(S)) = A(f(a,S))
- One additional axiom to come

Kalai-Smorodinsky Solution Final Axiom Motivation

- IIA is controversial in a variety of contexts, especially considering that it is not supported by empirical evidence.
- Consider:

  S_1 = convex hull {(0,1), (1,0), (.75, .75), (0,0)}

  S_2 = convex hull {(0,1), (1,0), (1, .7), (0,0)}

Kalai-Smorodinsky Solution Final Axiom Motivation

- Nash's solutions:

For S_1: (.75, .75)

For S_2: (1, .7)

Nash's approach gets agent 2 a worse payoff in the second scenario. This is problematic.

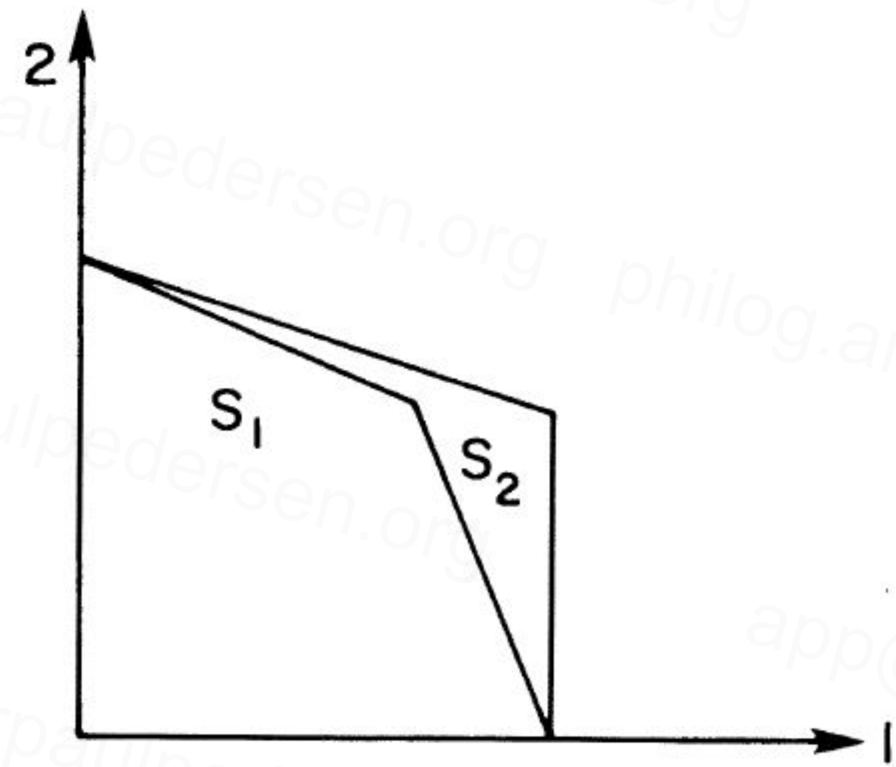

FIGURE 1

Kalai-Smorodinsky Solution Some Notation

- Let $b(S) = (b_1(S), b_2(S))$, where
  - $b_1(S) = \sup\{x: (x,y) \text{ is in } S \text{ for some } y\}$
  - $b_2(S) = \sup\{y: (x,y) \text{ is in } S \text{ for some } x\}$
- $b(S)$ is called the *utopia point*
- $g_S(x) = y$ if (x,y) is in S and there is no point in S strictly greater than (x,y)

  $= b_2(S)$ if no such y exists
- $g_S(x)$ is meant to represent the best agent 2 can get if 1 gets at least x
- We call (a,S) *normalized* if $a = (0,0)$ and $b(S) = (1,1)$

Kalai-Smorodinsky Solution Final Axiom

- Monotonicity: If (a,S_1), (a,S_2) are bargaining pairs such that $b_1(S_1) = b_1(S_2)$ and $g_1(x) \leq g_2(x)$ for all x, then $f_2(a,S_1) \leq f_2(a,S_2)$ (agent 2's payoff)
- If for every level of utility agent 1 could demand, agent 2 has a higher maximum potential utility, then 2's ultimate payoff should be higher as well.

Kalai-Smorodinsky Solution Theorem

- There is a unique function μ satisfying the axioms given where μ(a,S) is the maximal element in S on the line L(a,b(S)) which joins a and b(S).
- It is easy to show that this μ is well defined, by showing that L(a,b(S)) intersects S and the partial order on the plane induces a total order on L(a,b(S)).

Kalai-Smorodinsky Solution Proof

- μ is clearly symmetric
- Pareto optimality follows from the compactness and convexity of S
- Affine transformations preserve the ordering in question, map lines to lines, and send b(S) to b(A(S))
- Monotonicity & Uniqueness...

Comparison of η and μ

- As before, let

  S_1 = convex hull {(0,1), (1,0), (.75, .75), (0,0)}

  S_2 = convex hull {(0,1), (1,0), (1, .7), (0,0)}
- $\mu(0, S_1) = (.75, .75)$ (same as Nash)
- $\mu(0, S_2) \approx (.77, .77)$ (where Nash gave (1, .7))
- We see that the two are not in general the same (they are also different for Nash's example, for which Nash gave (12,5) and the ratio of utilities for μ must be 11:5)

Coincidence of η and μ

- It is clear that the two solutions will be the same in the case of a symmetric set, but it is easy to construct an example of a set which is not symmetric, for which the two are the same.

- Example: The convex hull of (0,0), (1,0), (0,1), (.9,.9) and (.5, 1).
- What property do these sets have in common?

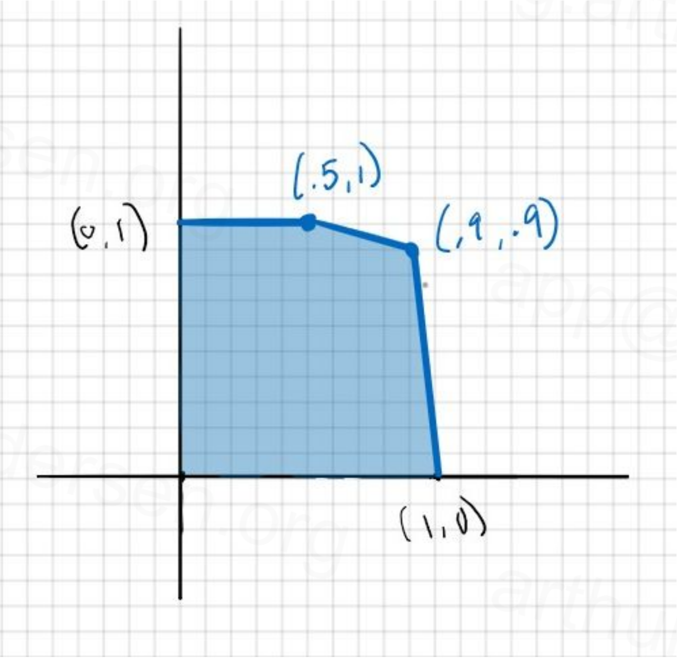

Coincidence of η and μ

- Theorem: Let (0,S) be a normalized bargaining pair.

$\eta(0,S) = \mu(0,S)$ iff for every (x,y) in S, $x+y \leq 2\mu_1$

$= \mu_1 + \mu_2$

- Corollary: Let (a,S) be a bargaining pair.

$\eta(a,S) = \mu(a,S)$ iff for every (x,y) in S,

$$(\mu_2 - a_2)x + (\mu_1 - a_1)y \leq 2\mu_1\mu_2 + a_1\mu_2 + a_2\mu_1$$

Utilitarian Solutions
General Definition

- One of the most natural classes of solutions, a bargaining solution f is called utilitarian if there are weights $\lambda_1, \lambda_2 > 0$ such that f maximizes the sum of the weighted utility gains for the players
- The point specified by such a function may not be unique except under certain conditions
- These solutions are scale-dependent, so we will not address them as generally as the previous cases

Utilitarian Solutions
The Pure Utilitarian Solution

- We will restrict our attention to one such solution, the pure utilitarian solution, υ
- Let $\upsilon^*(a, S) = \underset{x \epsilon S}{\operatorname{argmax}}[(x_1 - a_1) + (x_2 - a_2)]$

$$= \underset{x \epsilon S}{\operatorname{argmax}}[x_1 + x_2]$$

$$\upsilon(a, S) = \underset{x \epsilon \upsilon^*(a,S)}{\operatorname{argmin}}[|(x_1 - a_1) - (x_2 - a_2)|]$$

$$= \underset{x \epsilon \upsilon^*(a,S)}{\operatorname{argmin}}[|x_1 - x_2|]$$

Egalitarian Solutions General Definition

- Similar to utilitarian solutions, the family of egalitarian solutions is quite natural, but is also scale-dependent
- A solution f is egalitarian if there are weights $\lambda_1, \lambda_2 > 0$ such that f(a,S) is Pareto optimal in S and the weighted utility gains are equal
- Egalitarian solutions are invariant under translations and homogeneous, that is they are invariant under equal scaling when the base point is 0

Egalitarian Solutions
The Pure Egalitarian Solution

- Let the pure egalitarian solution γ(a,S) be defined as the maximal point in S for which the players' gains are equal, i.e. $x_1 - a_1 = x_2 - a_2$
- Since like the K-S solution, this solution is characterized by finding the maximal point on a line, the two should be easy to compare

Coincidence of η, μ, and υ

- From the previous coincidence theorem, it is plain to see that for normalized bargaining pairs (and a slightly larger class of pairs) the Nash and K-S solutions coincided precisely when they were purely utilitarian
- More precisely, for any pair $(0, S)$ with $b_1 = b_2$, if $\eta(0, S) = \mu(0, S)$, then $\eta(0, S) = \mu(0, S) = \upsilon(0, S)$

Coincidence of μ and γ

- Clearly the two solutions coincide for normalized bargaining pairs, and the pure egalitarian solution has some invariance properties, so if (a,S) can be obtained by equal scaling in both directions, then a translation from a normalized pair, the two will still coincide
- That is, $\mu(a, S) = \gamma(a, S)$ if and only if

$$b_1 - a_1 = b_2 - a_2$$

Other Observations and Questions

- Future Work:
  - Other solutions and their coincidence properties
  - Further refinement of coincidence theorems
  - New solutions

THANK YOU