Notes for Remote Presentation:

Game Theory/Fairness Modeling

April 27, 2020

Bankruptcy Model:

Fairness axioms for solution methods.

Typical example:

Claims: E is available to distribute

$A = 30 \quad B = 90 \quad C = 120 \; ; \; E = 150$

Fairness axioms for a method:

1. If two claimants have equal claims they should get the same amount.

2. If more money is found to pay off the estate (E grows to a bigger E*) then what a claimant is given should not go down. (Monotonicity in estate size.)

3. Two problems differ in that one claimant Z is asking for more, the other claimants and E (estate size) are unchanged.

Z's share where her claim went up should not now be SMALLER for a fair method.

(Monotonicity in claim sizes.)

Another interesting idea:

Suppose there is a bankruptcy problem solved by method M, and M gives a subset X of claimants D dollars. Now look at the new problem where one uses the amount D to settle the original claims using D instead of E.

M is called *consistent* if the members of a subset of claimants X get the same amount when M is applied to X using D as the "new" estate size as when it used E as the estate size.

The Talmudic method is *consistent*. When applied to 2 claimants it gives the same result as concede and divide.

Other methods are consistent as well.

There are many papers concerned with exactly what axioms a method to solve a bankruptcy problem must hold in order for one to be *forced* to use a particular method.

Axiomatics is important in social choice as well as geometry and algebra.

The "definitive" reference here is:

William Thomson, How to Divide When There Isn't Enough (Subtitle: From Aristotle, the Talmud, and Maimonides to the Axiomatics of Resource Allocation)

Cambridge U. Press, 2019, (481 pages)

Example: (Review)

A = 30 B = 90 C = 120 E = 150

Half claims: A = 15, B =45, C= 60

So these can be met:

A = 15, B = 45, C = 60 (120 used)

Since E was 150 we have 30 units left to settle the other half-claims:

Half claims: A = 15, B =45, C= 60

Give C 15 units to reduce her loss to 45. Now we can give B and C each 7.5 to try to equalize losses as much as possible.

So A: $15 + 0 = 15$

B: $45 + 7.5 = 52.5$

C: $60 + 15 + 7.\ 5 = 82.5$

These add to $E = 150$

A partial test of "consistency."

Check B and C. Their claims were 90 and 120. How would the these claims be settled with the 135 units they were given, using the Talmudic method? Half claims are B = 45 and C = 60; settle these with 105 units, leaving 30 left.

Now settle: B = 45 and C = 60 based on losses. Give C, 15 to reduce C's loss to B's of 45, and split the remaining 15 units between them - 7.5 to each.

B's share is 45 + 7.5 = 52. 5
C's share is 60 + 15 + 7.5 = 82.5

So we settled B and C's claims in a consistent manner!!

(Same as concede and divide.)

We saw many ways to decide an election in an appealing manner; many ways to solve a bankruptcy problem appealingly; many ways to solve an apportionment problem in an appealing manner. However, picking which appealing approach is not so clear, and in may cases there are short lists of fairness properties so that none of these appealing methods obey ALL of the items on a short list of fairness properties! (Arrow's Theorem; Satterthwaire-Gibbard; Balinski-Young; impossibility results due to William Thomson for bankruptcy problems.)

Coalition games:

In bankruptcy problems the players (claimants) are acting on their own for a share of the "pie."

In coalitional games, single or groups of claimants want:

a. To lower their costs

b. To increase their profits

c. Improve outcomes by cooperation

When will individual players cooperate with other players in order to do better?

We can give "normative" advice and also see what actually happens in actual coalition games (European union); lab experiments: behavioral game theory.

Profit sharing:

Corporations get certain amounts separately, but MORE, when they merge.

v({ }) = 0 (the set of no players - the empty set - can't get anything of value!)

v({A}) = 80; v{B}) = 100; v({C}) = 120

v({A,B}) = 184; v({A,C}) = 248; v({B,C}) = 250

v({A,B,C}) = 390

What happens when one shares costs rather than shares profits?

Now coalitions like to see their cost decrease with size!

Thus, a water purification plant system has stand alone costs but costs go down (sometimes) when large units get together to do water purification.

$c(\{\ \}) = 0$ (the set of no players - the empty set - can't get anything of value!)

$c(\{A\}) = 80$; $c(\{B\}) = 100$; $c(\{C\}) = 120$

$c(\{A,B\}) = 150$; $c(\{A,C\}) = 180$;
$c(\{B,C\}) = 210$

$c(\{A,B,C\}) = 270$
(less than summing costs to individual players)

This approach to representing a game, when there are many players, is often called the *characteristic function form*, and goes back to Von Neumann. The goal here is to try to predict and/or offer advice to players of such games as to what coalitions it makes sense to form. Ideally, in many cases, the grand coalition which consists of all of the players will arise - everyone cooperates.

If coalitions form, the major question is how the members of the coalition should share the benefit of cooperating.

One is concerned with treating the players who join coalitions fairly (in terms of their "benefits," but one also wants the coalitions to be stable - not have cycles where arrangements are broken to get better arrangements, which return to the start arrangement in a cycle.

There usually are MANY ways to share the amount a coalition gets by acting together. Suppose the "grand coalition" forms - how might one divide the cost this coalition is required to pay out to its members?

Let x_i = amount given to player i in a coalitional game; N = set of all players

Rationality or fairness conditions:

a. $x_i \le c(\{i\})$ (player i won't join a coalition unless its costs are less when part of this coalition than acting alone!

Sometimes called individual rationality.

b. The sum of the x_i values for members of the grand coalition is the same as c (N).

Sometimes called group rationality. Similar to Pareto Optimality - want the "pie" that is split to be as large as possible.

This approach is especially important when there are many players in the game.

Intuitively, if there are n-players, each player can be a member of a coalition or not, so there are 2^n different coalitions we have to consider, including the { } the empty set, and set of all players, the grand coalition.

For large corporations with many "subdivisions" one must have each division of the corporation charged for various services - accounting, telephone, utilities, photocopy services, administrative assistants, etc. Some of these coalition games find applications in this kind of "cost allocation."

Consider the cost allocation version:

What fairness conditions will hold?

In games of this kind (cost sharing; profit sharing) if the game has a non-empty CORE, then in principle the grand coalition makes sense because there are allocations to the individuals of the grand coalition that do NOT encourage them to leave the grand coalition to improve their outcomes.

Coalition games can have no core (empty core), many points in the core, or a unique point in the core. The formal definition of CORE is a bit complex but the intuitive idea is that these are ways of sharing the benefits from the grand coalition in a way that does not encourage players to break up into smaller groups, or act alone. However, mathematics does not offer specific advice as to how to do the division - pick out one of many CORE points when the CORE is not empty.

Remaining important fairness models:

a. Two-sided markets

Typical application: pairing medical school graduates with hospital residency positions

b. One-sided markets

Typical application: assigning students rooms in a college dormitory

Two-sided Markets

(Model originally due to David Gale and Lloyd Shapley (Nobel memorial prize in Economics) - now both deceased)

Sometimes described under the title stable matchings and Gale/Shapley

Other major contributors:

Alvin Roth (Nobel memorial prize)
Robert Irving
David Manlove

The paper that started it all!

College admissions and the stability of marriage

D. Gale; L. S. Shapley. The American Mathematical Monthly, Vol. 69, No. 1 (Jan., 1962), 9-15.

Monthly not associated with FIRST PUBLICATION of profound mathematics.

Major applications:

1. College admissions

2. Matching hospitals with residencies to medical school graduates

3. Matching court justices and clerks

4. School choice

5. Kidney exchange

Vanilla (Basic) Rules:

Ladies rank the men without ties
Men rank the ladies without ties

All individuals would rather be paired than be "single."

Rankings are ordinal

Extensions exist where:

a. players on either side of the market might prefer to be unmatched rather than having a specific player paired with them.

b. players can rank the players on the other side of the market with ties allowed

There are different ways of representing the information given by the players in a pairing problem which searches for stable pairings.

Columns can be ranks - entries are players.

Columns can be the people - entries are the ranks.

Problem instance: 4x4 example

4 men rank 4 women without ties

	1 st	2nd	3rd	4th
m1	w1	w2	w3	w4
m2	w2	w1	w4	w3
m3	w1	w4	w2	w3
m4	w2	w3	w4	w1

4 women rank 4 men without ties

	1st	2nd	3rd	4th
w1	m4	m3	m2	m1
w2	m4	m3	m1	m2
w3	m2	m1	m4	m3
w4	m4	m2	m3	m1

Table instance

Men's preferences: (example: w3 is m2's 2nd choice)

	w1	w2	w3	w4
m1	1	4	3	2
m2	1	3	2	4
m3	2	1	4	3
m4	2	3	4	1

Women's preferences: (example m1 is w4's 4th (last) choice)

	m1	m2	m3	m4
w1	4	1	2	3
w2	3	1	4	2
w3	4	3	2	1
w4	4	2	3	1

One sided markets:

Situations where two sets are paired but one of the sets has no "feelings" about the assignment!

Example: Students attending a mathematics research program at a college are provided dorm rooms. The students may have preferences about the rooms but the rooms don't have preferences about the students.

Examples: Students who participated in a mathematics contest are each awarded a different book as a prize for having gotten honorable mention in the contest.

The students may have preferences about the books they get but the books don't have preferences about who owns them.

One approach to such problems that some might view as fair is to assign the rooms or prizes to the individuals at random. However, it might happen that each player gets his/her worst choice. Random may be fair but it is rarely is "efficient" in making the recipients happy.

Goal: Stable matching

What does this mean?

Suppose (m,w) are paired in some matching. w can't find another man m* who she prefers to m and where m* prefers w to the woman m* is paired with in the matching!

Similarly for m in this pair.

Such a pair is called a blocking pair.

When a matching M has no blocking pair it is called stable.

When a matching M has a blocking pair (m, w) one or both members of this pair have incentive to work outside the system and pair up with their "better" choice.

Gale and Shapley showed there is a simple algorithm to guarantee that there is at least one stable marriage. Usually there are MANY stable pairs.

Unintuitively, among all the stable marriages there is one which is BEST from the ladies point of view and one which is BEST from the men's point of view - though sometimes these two coincide in which case there is ONLY one stable marriage for the preferences involved.

These stable marriages are known as the MALE OPTIMAL stable marriage and the FEMALE optimal stable marriage.

Assume the FEMALE and MALE optimal stable marriages are different.

The FEMALE optimal stable marriage is the worst from the male point of view among all stable marriages.

The MALE optimal stable marriage is the worst from the female point of view among all stable marriages.

If n women and n men rate each other there are n! possible matchings, for a given problem instance. The number of stable marriages can be as few as 1 for an instance but there is a family of marriage problem instances where the number of stable marriages grows exponentially in n. Exactly how large this number can be is still an *unsolved* problem.

Men's preferences:

	w1	w2	w3	w4
m1	1	4	3	2
m2	1	3	2	4
m3	2	1	4	3
m4	2	3	4	1

Women's preferences:

	m1	m2	m3	m4
w1	4	1	2	3
w2	3	1	4	2
w3	4	3	2	1
w4	4	2	3	1

The Gale-Shapley deferred acceptance algorithm proceeds in rounds.

The FEMALE optimal stable marriage is produced when the girls PROPOSE to the boys; The MALE optimal marriage occurs when the boys propose to the girls.

Female optimal:

Set up: Each boy is at a table in front of the "gym."

In each round any un-paired girl proposes to the next boy on her preference list who has not already turned her down!

To propose go to the table of the next person on your list.

When one or more girls propose to a boy he temporarily picks the one who is ranked highest on this list. If in a future round a better choice arrives he (perhaps) temporarily changes to this better choice.

(For the male optimal matching reverse the roles of girls and boys - boys propose.)

Again:

When girls propose they do best.

When boys propose they do best.

Best, in the sense of getting their highest ranked person of the opposite sex they can in ANY stable pairing.

Men's preferences: (example: w2 is m3's first choice)

	w1	w2	w3	w4
m1	1	4	3	2
m2	1	3	2	4
m3	2	1	4	3
m4	2	3	4	1

Women's preferences: (m1 is w4's 4th choice

	m1	m2	m3	m4
w1	4	1	2	3
w2	3	1	4	2
w3	4	3	2	1
w4	4	2	3	1

Round 1: m1 m2 m3 m4
(all) w1,w2 w3,w4

m2 accepts w1; m4 accepts w4; w2 and w3 moves to Round 2

Men's preferences:

	w1	w2	w3	w4
m1	1	4	3	2
m2	1	3	2	4
m3	2	1	4	3
m4	2	3	4	1

Women's preferences:

	m1	m2	m3	m4
w1	4	1	2	3
w2	3	1	4	2
w3	4	3	2	1
w4	4	2	3	1

Round 2: m1 m2 m3 m4
(w2,w3); w1 w3 w2,w4

m2 accepts w1; m3 accepts w3; m4 accepts w4; w2 to Round 3

Men's preferences:

	w1	w2	w3	w4
m1	1	4	3	2
m2	1	3	2	4
m3	2	1	4	3
m4	2	3	4	1

Women's preferences:

	m1	m2	m3	m4
w1	4	1	2	3
w2	3	1	4	2
w3	4	3	2	1
w4	4	2	3	1

Round 3: m1 m2 m3 m4
(w2); w2 w1 w3 w4

m1 accepts w2; m2 accepts w1; m3 accepts 3, m4 accepts w4

Men's preferences:

	w1	w2	w3	w4
m1	1	4	3	2
m2	1	3	2	4
m3	2	1	4	3
m4	2	3	4	1

Women's preferences:

	m1	m2	m3	m4
w1	4	1	2	3
w2	3	1	4	2
w3	4	3	2	1
w4	4	2	3	1

Stable: m1 m2 m3 m4

w2 w1 w3 w4

To find the female or male optimal marriage is not HARD but requires careful coordination between the information generated as the algorithm is carried out and the use of the two preferences tables to carry out the algorithm! When done by hand it requires careful practice.

In the medical residency version the algorithm that is used is students propose but in the past the hospital's optimal version of the algorithm was used!

In Great Britain, the same problem arises and they did not use a stable matching algorithm and the market unraveled - that is, people worked outside the centralized system to find pairings.

This suggests that using an algorithm that finds stable solutions is valuable.