Fair Division

Uriel Feige

June 4, 2024

1 The cut and choose protocol

Here is modern English translation of a passage from the Bible (Genesis, 13:8–9).

So Abram said to Lot, "Let's not have any quarreling between you and me, or between your herders and mine, for we are close relatives. Is not the whole land before you? Let's part company. If you go to the left, I'll go to the right; if you go to the right, I'll go to the left."

In the setting above there are two players, Abram and Lot, and a divisibe good (land areas). The players wish to divide the good among them. (At the time, this did not mean ownership of the land, but rather where they will take their families and possessions.) Abram offers a partition of the good (into a left part and right part), and lot chooses among the two options. This is probably the earliest recorded application of the *cut and choose* allocation procedure.

In more modern literature, such settings are referred to as *fair division*, and *cake cutting* is often used as a metaphor if the good is divisible.

Is the cut and choose protocol fair? Fairness is a subjective notion, so anyone is entitled to his/her own opinion about it.

Here we discuss a mathematical theory of fairness. It is important to state clearly what assumptions we make, so that it will not be applied in situations in which the assumptions do not hold.

Assumptions:

- There are $n \geq 2$ players. They all have equal entitlement. (This assumption will be relaxed later.)
- There is one good $\mathcal{M}$ (or several goods) that need to be divided among the players. Each player i has a valuation function v_i over parts of the good. Valuations are normalized (not receiving anything has value 0), and monotone non-decreasing – the marginal value of any part is non-negative. (A related branch of fair division deals with allocation of *bads*, also referred to as *chores*, and then valuations are non-increasing.)

- For a divisible good, we assume that valuations are *non-atomic* – every part of the good that has positive marginal value can be further divided. For indivisible goods, $\mathcal{M}$ is a collection of m goods, and each good is atomic – it cannot be divided. (More generally, $\mathcal{M}$ may contain a mixture of divisible and indivisible goods.)

The cut and choose protocol suggests two notions of fairness.

One notion is *comparison based.* Each player compares what she receives with what other players receive. For the case of two players, an allocation $A = (A_1, A_2)$ is called *envy free* (EF) if $v_1(A_1) \geq v_1(A_2)$, and $v_2(A_2) \geq v_2(A_1)$. For the chooser (player 2), there is a simple strategy ensuring (EF) – choose the part (P_1 or P_2) with higher v_2 value. For the cutter (player 1), achieving EF is a more delicate matter. It can be ensured if the cut is into two parts satisfying $v_1(P_1) = v_2(P_2)$. Such a cut exists for a divisible good.

The other notion is *share based.* A share s is a mapping from the entitlement and valuation of an agent to a value, indicating what value (at least) the agent expects to get in a fair allocation. The agent does not care what value other agents receive.

The cut and choose protocol offers a natural notion of a share for the cutter – the maximum value that she can ensure for herself in the cut and choose protocol. This is a max-min notion, referred to as the *maximin share* (MMS). It is the maximum over all partitions, of the value of the less valuable part in the partition.

The cutter in a cut and choose protocol has a strategy that offers her at least her MMS – cut according to the MMS partition. For the chooser, we distinguish between two cases.

If the chooser has an *additive* valuation, then her MMS is equal to her *proportional share*, namely $\frac{1}{2}v_2(\mathcal{M})$, and $\max[v_2(P_1), v_2(P_2)] \geq \frac{1}{2}v_2(\mathcal{M})$. Hence, the chooser can get her MMS. More generally, for additive valuations, in an EF allocation every agent also gets at least her MMS.

If the valuation of the chooser is not additive, she might not be able to get her MMS. In the story of Abram and Lot, suppose that it was the case that all the left (west) part was composed of hills and all the right (east) part was a big valley. If Lot would have liked to have a mixture of hills and valleys (rather than only hills or only a valley), his MMS partition would be into North and South, not East and West. In this case, with the division that Abram offered (into east and west), no part gives Lot his MMS. For agents who care about share based fairness, being the cutter may sometimes offer them better guarantees that being the chooser.

Summarizing, the cutter can guarantee to herself her MMS, the chooser can guarantee for herself no envy.

An allocation $A = (A_1, A_2)$ is *Pareto optimal* (also referred to as *Pareto efficient*) if there is no other allocation $A' = (A'_1, A'_2)$ for which $v_1(A'_1) \geq v_1(A_1)$ and $v_2(A'_2) \geq v_2(A_2)$, with at least one of the inequalities being strict. Pareto optimality is a well accepted critirion for economic efficiency.

Share based fairness notions are consistent with Pareto efficiency. Given an allocation that satisfies them, any allocation that Pareto dominates it also satisfies them. Comparison based fairness notions need not be consistent with Pareto efficiency. In fact, in real life situations, such as two small children arguing over a toy, the final allocation chosen by the parent might be the empty allocation (the parent takes away the toy). That allocation is envy free, but not Pareto optimal.

Before we move on, we remark that at this point we only consider fairness for deterministic allocations. The use of randomized allocation mechanisms can offer options for achieving fairness guarantees in situations in which deterministic allocations do not appear to be fair enough. A well known example is *random serial dictatorship*, used in programs such as *Mechir Lamishtaken* in Israel (in which the players are eligible buyers, and the items to be allocated are apartments at discounted prices).

1.1 Cut and choose for multiple agents

What would be the analog of the cut and choose protocol for $n > 2$ players? This is not so clear. However, let us consider the roles of the cutter and the chooser separately.

For the cutter, it makes sense that she is required to cut into n parts, one for each agent. Among these parts, she should be willing to accept the worst one, as otherwise she will make the worst part of negligible value (or even empty), effectively cutting to fewer than n pieces. This gives the share based fairness notion of the MMS.

Definition 1.1 *Given a good $\mathcal{M}$, the* maximin share *(MMS) $MMS(\mathcal{M}, v_i, \frac{1}{n})$ of an agent i with valuation v_i and entitlement $\frac{1}{n}$ is $\max_{\{P_1,\dots P_n\}} \min_j v_i(P_j)$, where $\{P_1, \dots P_n\}$ ranges over all partitions of $\mathcal{M}$ into n parts. An allocation $A_1, \dots, A_n$ is an MMS allocation if each agent i gets a part of value at least her MMS, namely, $v_i(A_i) \geq MMS(\mathcal{M}, v_i, \frac{1}{n})$.*

For the chooser, faced with a partition of $\mathcal{M}$ into n parts, it makes sense that she will be allowed to choose her most preferred part, as it might be that the cutter put all value in one part, leaving almost no value for the other parts. This gives the comparison based notion of EF.

Definition 1.2 *Given a good $\mathcal{M}$, an allocation $A_1, \dots, A_n$ is* envy-free *(EF) if for every agent i and for every j it holds that $v_i(A_i) \geq v_i(A_j)$.*

2 Divisible good, MMS allocation

We consider a single divisible good $\mathcal{M}$. With non-additive valuations, an MMS division need not exist. Hence, we assume that all valuations are additive. In this case, we seek a proportional allocation, in which each agent i gets value at least

$\frac{1}{n}v_i(\mathcal{M})$. By scaling valuations by constant multiplicative factors, we may assume that for every agent $v_i(\mathcal{M}) = 1$.

We view the divisible good at the interval from 0 to 1. The valuation function of each agent is a nonnegative bounded density function over the interval. The value of a subinterval is the integral of this density function. In a *contiguous* allocation, each agent gets a single subinterval. We show that there is an MMS allocation that is contiguous.

Each agent cuts the interval into n consecutive pieces, each of value $\frac{1}{n}$ for the agent. Then we allocate subintervals in n rounds, in a greedy fashion. In round j, the agent to get the next subinterval is the one (among those remaining) whose jth cut point is earliest.

The above protocol requires each agent to report $n-1$ cut points. Reporting a cut point might require infinite precision, an issue that we ignore here. Still, it is desirable to reduce the number of reported cuts.

A more efficient protocol of Even and Paz [2] reports only $\lceil \log n \rceil$ cut points per agent. Each agent reports its $\lfloor \frac{n}{2} \rfloor$ the cut point. Finding the median of these cut points, the $\lfloor \frac{n}{2} \rfloor$ agents with lower cut points continue recursively with the prefix of the interval, and the $\lceil \frac{n}{2} \rceil$ agents with higher cut points continue recursively with the suffix of the interval. Even and Paz also present a randomized allocation protocol whose expected total number of cut points is $O(n)$. It is based on randomized algorithms for finding the median. The expected number of comparisons to find the median is $O(n)$, but the expected number of pivot elements, against which comparisons are made, is only $O(\log n)$. Agents holding pivot elements (the cut points are the elements) need to report their cut points. Other agents need only report the outcome of the comparison of their cut point with the pivot cut point.)

3 Divisible good, envy-free division

Here, we do not assume that valuations are additive (though recall that they are normalized, monotone and non-atomic).

It seems reasonable that EF allocations exists, or at least allocations in which the envy tends to 0. At an intuitive level, one may imagine cutting the unit interval (cake) into infinitely small crumbs, and allocating them in periodic fashion to the players.

A more interesting question is whether there is an EF allocation (exactly EF, not nearly EF) in which the good is cut into finitely many pieces. Perhaps the best we can hope for is a contiguous allocation that is EF. The following theorem is due to Stromquist [4].

Theorem 3.1 *For every instance of allocating a divisible good among equally entitled agents with valuation functions that are normalized, monotone and non-atomic, there is a contiguous envy free allocation.*

Proof: We sketch a proof which is presented in a very readable form in Su [5]. It is based on Sperner's lemma, that we encountered in the context of proving existence of a Nash equilibrium.

Consider the special case of three agents. A contiguous allocation of the unit interval can be specified by a vector of dimension 3, (x_1, x_2, x_3), for the sizes of the pieces. The vector satisfies the constraints $x_1, x_2, x_3 \geq 0$ and $x_1 + x_2 + x_3 = 1$. Hence, geometrically, all vectors form a triangle. (With $n > 3$ agents, we have a simplex in a higher dimension.)

Subdivide the simplex into arbitrarily small simplices. For the triangle, this can be done by drawing lines parallel to the sides of the triangle. Name the vertices of the subdivided simplex by names of agents in a "nameful" way, meaning that the vertices of each small simplex all have different names. This can be done by naming the vertices on the left side in the cyclic order A, B, C (going from top to bottom), and naming vertices of each row (from left to right) in the same cyclic order. (For higher dimensional simplicies, a bit more care is needed in subdividing the simplex so as to have a namerful naming. See details in [5]).

Color each vertex of the subdivided simplex by the part that the player naming the vertex likes best, among the three parts implied by the coordinates (x_1, x_2, x_3) of the vertex (breaking ties arbitrarily, though never preferring an empty part). This coloring satisfies the conditions of Sperner's lemma (due to the fact that empty parts are never chosen). Hence we must have a colorful simplex. In the vertices of the simplex, each player is satisfied with a different part (has no envy).

By refining the subdivision more and more, we converge to an EF contiguous allocation. $\square$

The proof of Theorem 3.1 is not algorithmic. An algorithm (in a model in which players can announce cuts with infinite precision) of Aziz and Mackenzie [1] finds a non-contiguous EF allocation. and the number of queries that it makes is is bounded by a function of n (a tower of exponents, of height 6).

4 Allocation of indivisible goods

We now consider allocation of indivisible goods. Clearly, there need not be any EF allocation. For example, this is the case when there are fewer goods than agents. Also, for general monotone valuations, MMS allocations need not exist. For example, this is the case when there are two agents and four items arranged in a two by two matrix, where one agent likes any row of the matrix, and the other agent likes any column of the matrix. However, for additive valuations, it appears that MMS allocations may exist (the case of fewer items than agents is not a negative example, because in this case the MMS of every agent is 0). Perhaps surprisingly, with $n \geq 3$ agents, there are allocation instances with additive valuations in which no MMS allocation exists [3].

As neither EF nor MMS are feasible, one is led to consider relaxations of these notions.

4.1 Comparison based fairness for indivisible goods

For EF, well studied relaxations are EF1 and EFX. The terminology of *minimal*, *mild* and *strong* used in the following definition is introduced here, and is not commonly used in related literature.

Definition 4.1 *In an allocation $A_1, \ldots, A_n$ we say that agent i envies agent j if $v_i(A_i) < v_i(A_j)$. The envy is* minimal *if for every item $e \in A_j$, it holds that $v_i(A_i) \geq v_i(A_j \setminus \{e\})$. The envy is* mild *if for some item $e \in A_j$, it holds that $v_i(A_i) \geq v_i(A_j \setminus \{e\})$. The envy is strong if for every item $e \in A_j$, it holds that $v_i(A_i) < v_i(A_j \setminus \{e\})$.*

*An EFX allocation (*envy free up to any item*) is an allocation in which for every two agents i and j, either i does not envy j, or the envy is minimal. An EF1 allocation (*envy free up to one item*) is an allocation in which for every two agents i and j, either i does not envy j, or the envy is mild.*

The proof that contiguous EF allocations exist for a divisible good implies somewhat weaker fairness for indivisible goods. Arranging the goods on a line and allowing to cut within a good (extending the valuation function in some consistent way to bundles that may contain fractions of items), a contiguous envy free allocation gives every agent a bundle in which at most two items are cut. Allocate every cut item to the agent who got the rightmost part. The envy between an agent i and j can be eliminated by removing one item from A_j (the cut item that ended up in A_j), and addiing one item to A_i (the cut item that agent i gave up). We refer to such allocations as EF_1^1.

We will see that EF1 allocations also exist. The question of whether EFX allocations always exist is a central open question in fair division.

To be continued in next lecture

References

[1] Haris Aziz, Simon Mackenzie: A Discrete and Bounded Envy-Free Cake Cutting Protocol for Any Number of Agents. FOCS 2016: 416–427.

[2] Shimon Even, Azaria Paz: A note on cake cutting. Discret. Appl. Math. 7(3): 285–296 (1984).

[3] David Kurokawa, Ariel D. Procaccia, Junxing Wang: Fair Enough: Guaranteeing Approximate Maximin Shares. J. ACM 65(2): 8:1–8:27 (2018).

[4] Walter Stromquist. How to cut a cake fairly. The American Mathematical Monthly, 87(8):640-–644, 1980.

[5] Francis Edward Su. Rental harmony: Sperner's lemma in fair division. The American mathematical monthly, 106(10):930—942, 1999.