DS 574 Algorithmic Mechanism Design — **Lecture #18**
Fall 2025 — Prof. Kira Goldner

Cake Cutting

Properties of the Cut and Choose Protocol

Today, we'll focus primarily on fairness properties.

Suppose two people need to split a heterogeneous and divisible good. The usual euphemism in the fair division literature is that of cutting a cake. More practically, the good could be an estate (e.g., in a divorce settlement) or processing time on a computer cluster (perhaps with some times of the day more valuable than others).

Why not just split the good 50/50? This makes sense when the good is homogeneous, but it's not clear what this means with a heterogeneous good. A player may value a part of the good much more than another, and different players can have different opinions about which parts are the most valuable.

We all know a reasonable protocol for two-person cake-cutting—it is mentioned already in the Bible, and is reinvented every year by siblings around the world.

The Cut and Choose Protocol

1. Player 1 splits the good into two pieces A and B, such that the player's value for each is exactly half that of the entire good.
2. Player 2 picks whichever of A, B she likes better.

The description above is the intended behavior of the players in the protocol—we'll talk shortly about whether or not they are incentivized to follow this behavior. We've been led all our lives to believe that this is a "fair" protocol. But is it? How would we formally argue one way or the other?

Here's the formal model. The good, or "cake," is the unit interval $[0, 1]$. (Yes, it's a weird-looking cake.) Each player i has a valuation function v_i, which specifies the value $v_i(S)$ to i of a given subset S of the cake. We'll make the following two assumptions about each valuation function v_i:[1]

1. We normalize each v_i such that $v_i([0, 1]) = 1$, their value for the whole cake. This is more or less without loss of generality, by scaling.
2. v_i is additive on disjoint subsets. That is, if $A, B \subseteq [0, 1]$ are disjoint, then
$$v_i(A) + v_i(B) = v_i(A \cup B).$$

[1] Actually, we also need a "continuity" assumption for everything to make sense—e.g., in the cut and choose protocol, it's important that there exists a cut that makes player indifferent between the two pieces. We omit further discussion of this assumption.

Is the cut and choose protocol strategyproof? For the second player: since they can't affect the split of the cake into A and B, and is supposed to choose the piece they like better, they have no incentive to deviate.

For the first player: if they know something about the second player's valuation function, they might want to deviate from the protocol. For example, suppose the good is a hot fudge sundae, the first player likes all parts of the sundae equally, while the second player likes ice cream but really cares about the cherry. The first player could split the sundae into the cherry and the rest, knowing that the second player would take the cherry, leaving a very valuable piece for the first player. If the first player doesn't know anything about what the second player wants, and assumes that the second player will always leave the piece that is worse for the first player, then the first player is incentivized to follow the protocol (to guarantee herself a piece with value $\frac{1}{2}$). In any case, this is certainly not DSIC.

Is the cut and choose protocol guaranteed to produce a Pareto optimal solution (assuming both players behave as intended)? Again, no. Consider the cake in Figure 1, where the first player only wants the first and third quarters of the cake, while the second player only wants the second and fourth quarters of the cake. One way the first player might split the cake would be into its first and second halves, resulting in both players getting a piece valued at $\frac{1}{2}$. But allocating each piece of cake to the only player who wants it results in both players having value 1.

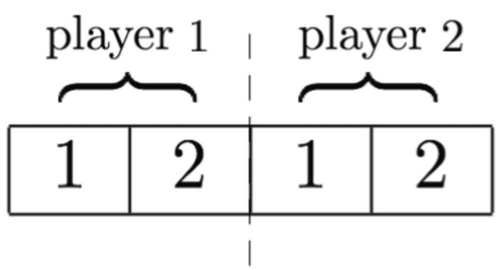

Figure 1: The solution produced by the cut and choose protocol may not be Pareto Optimal.

Can we come up with a protocol that is both Pareto optimal and strategy proof for both players, or is there a fundamental impossibility result? Just for these two properties, if we always give the entire cake to the first player is clearly strategyproof and (under weak assumptions on the v_i's) Pareto optimal.

What about "fairness?" What are possible notions? Equally happy? This property (both get 1/2) is also not satisfied by the cut and choose protocol: the first player is guaranteed to get a piece that they value at $\frac{1}{2}$, while the second player might well end up with a piece that they value at greater than $\frac{1}{2}$.

One definition of "fairness" is that each player receives at least her fair share, at least from her perspective.

Definition 1. An allocation $A_1, A_2, \ldots, A_n$ of cake to n players is *proportional* if

$$v_i(A_i) \geq \frac{1}{n}$$

for every player i.

The cut and choose protocol satisfies proportionality—the first player gets a piece that they value at $\frac{1}{2}$ and the second player does at least as well. Obviously, the dictator protocol does not satisfy proportionality.

A second definition is that no player wants to trade places with any other player.

Definition 2. An allocation $A_1, A_2, \ldots, A_n$ of cake to n players is *envy-free* if

$$v_i(A_i) \geq v_i(A_j)$$

for every pair i, j of players.

This means that while player j might like her piece more than i likes her own, to player i's own tastes, her piece is better than that of player j's.

Do you notice a relationship between these? Protocols that satisfy one but not the other? The second definition is more stringent than the first. To see this, note that for every i, $\sum_{j=1}^{n} v_i(A_j) = v_i([0,1]) = 1$ (by our assumptions on the v_i's). So if player i likes A_i better than every other piece (as dictated by envy-freeness), it must be that $v_i(A_i) \geq \frac{1}{n}$.

The converse also holds for the special case of $n = 2$ (if you get a piece of cake that you value at least $\frac{1}{2}$, then swapping would net you a piece that you value at most $\frac{1}{2}$). In particular, the cut and choose protocol is envy-free. (This is also easy to see directly: the first player is indifferent between the two pieces, while the second player gets her favorite one.) For $n \geq 3$, there can be proportional allocations that are not envy-free (exercise).

Beyond Two Players

The obvious next question is to ask about analogs of the cut and choose protocol that work with 3 or more players. We focus on envy-freeness and ignore strategyproofness issues.

$n = 3$ **players:** Selfridge and Conway (∼1960) independently designed the same envy-free protocol for this case, see exposition in [5].

- The third player cuts the cake into three pieces such that she has value $\frac{1}{3}$ for each of them.
- The first and second players then select their favorite pieces.
- If they select different pieces, then we are done (why?). Otherwise, if the first and second players both prefer the same piece, say the first (wlog):
  - Then, the second player is asked to trim off a subpiece of the first piece such that she is indifferent between the trimmed piece and her second-favorite of the original three pieces.
  - The protocol then recurses on the trimmings with the roles of the players swapped (with the second player cutting the trimmings into three pieces of equal value to her).

The protocol does not need to recurse further, because of the special structure of the recursively defined subproblem. In the worst case, the Selfridge-Conway protocol makes 5 cuts (the first two cuts by the third player, the trim by the second player, and two more cuts in the recursive call).

$n \geq 4$ **players:** In 1995, Brams and Taylor [3], gave a finite-but-unbounded protocol for computing an envy-free allocation with any number of players: for every $n \geq 4$ and T, there is a choice of $v_1, \ldots, v_n$ such that the protocol requires more than T steps to terminate.

In 2016, breakthrough results by Aziz and Mackenzie [2, 1] finally gave bounded protocols for envy-free protocols for $n \geq 4$ players. For $n = 2$, there are at most 203 cuts. For general n, it's a tower of 6 n's, meaning $n^{n^{\cdot^{\cdot^{n}}}}$ with 6 layers. As for lower bounds, the best known is $\Omega(n^2)$ [6]. Now there's a gap that's in need of narrowing!

Rent Division: Fair Division in Practice

One place where fair division protocols are used in practice is on `spliddit.org`, which has been used by tens of thousands of people. One of the problems that spliddit solves is the rent division problem, where there are n people, n rooms, and a rent of R. The goal is to assign people and rents to rooms, with one person per room and with the sum of rents equal to R, in the "best" way possible.

- Assume that each person i has a value v_{ij} for each room j.
- Normalize i's values so that their sum for all rooms is $\sum_j v_{ij} = R$.
- People have quasi-linear utility for their room minus rent paid for their room.[2]

A solution to a rent division problem is envy-free if

$$v_{iM(i)} - p_{M(i)} \geq v_{iM(j)} - p_{M(j)} \tag{1}$$

for every pair i, j of players, where $M(i)$ denotes the room to which i is assigned and p_j denotes the rent assigned to the room j. That is, no one wants to trade places with anyone else (where trading places means swapping both rooms *and* rents).

This is essentially looking for a Walrasian Equilibrium, except that we have the added constraint that the prices must sum to the total rent. The good news is that an envy-free solution is guaranteed to exist, and that one can be computed efficiently.

The bad news is that there can be many envy-free solutions, and not all of them are reasonable. For example, suppose there are two players and two rooms, that the total rent R is 1000, and that the first player only wants the first room ($v_{11} = 1000$ and $v_{12} = 0$) while the second player only wants the second room ($v_{21} = 0$ and $v_{22} = 1000$). **What are envy free allocations and prices?**

The only reasonable room assignment is to give each person the room that they want. Intuitively, by symmetry, each person should pay 500 in rent. But every division of the rent is envy-free! Even if you make the first person pay almost 1000 for her room, she still doesn't want to swap with the other person.

The upshot is that we need a method for selecting one out of the many envy-free solutions. One can imagine several ways of doing this; here's what happens on spliddit (given v_{ij}'s and R):

[2] Rent division isn't really a special case of cake cutting, since the rooms are indivisible.

1. Choose the room assignment M to maximize $\sum_i v_{iM(i)}$.
2. Set the room rents so that envy-freeness (1) holds, and subject to this, maximize the minimum utility:

$$\max_p \left(\min_{i=1}^{n} (v_{iM(i)} - p_{M(i)}) \right).$$

Algorithmically, they just solve this by computing a maximum-weight bipartite matching via linear programming. See [4] for further details.

Fair Division of Indivisible Goods

We now move on to indivisible goods, which is actually a much harder setting to achieve the same guarantees for. To see this, consider the example of two players and a single indivisible good.

We reiterate the model for this setting. Again, we will normalize valuations for each player such that if they receive *all of the goods*, their value is 1.

Definition 3. We have a set $N = [n]$ of n players and a set $\mathcal{M}$ of m goods. Each agent $i \in N$ has a monotone valuation function $v_i : 2^M \to \mathbb{R}_+$. We assume the valuations are normalized, i.e. $v_i(\emptyset) = 0,\ v_i(\mathcal{M}) = 1$ for all $i \in N$.

In many cases, we will also restrict our attention to the case of additive valuations as they are much easier to work with. A function v_i is said to be additive if it can be written as:

$$v_i(S) = \sum_{j \in S} v_i(j).$$

Now, that we have defined an instance of the fair division problem, let us discuss what we wish to achieve: a fair division. We begin by defining the notion of an allocation.

Definition 4. Given an instance of the fair division problem, we define an allocation $A = (A_1, \ldots, A_n)$ as a tuple of disjoint sets $A_i \subseteq \mathcal{M}$, i.e. $A_i \cap A_j = \emptyset$ for all $i \neq j$. We call the allocation a *partition* if the allocation allocates all the goods, i.e.

$$\bigcup_{i=1}^{n} A_i = \mathcal{M}.$$

Notions of Envy Freeness

Our definition of envy-freeness (EF) is the same as before. Since we have already seen that it is not always possible to achieve an EF allocation, we relax this condition and define two lesser notions of envy and thus fairness.

Definition 5. For an instance of the fair division problem $(N, \mathcal{M}, (v_1, \ldots, v_n))$ and an allocation $A = (A_1, \ldots, A_n)$:

1. We say that the allocation A is *envy-free up to one good* (EF1) if for all $i, j \in N$, we have that:

   $$v_i(A_i) \geq v_i(A_j \setminus \{\alpha\}) \quad \text{for some } \alpha \in A_j.$$

   This means that by removing some good from the bundle of agent j, agent i no longer envies agent j.

2. We say that the allocation A is *envy-free up to any good* (EFX) if for all $i, j \in N$, we have that:
$$v_i(A_i) \geq v_i(A_j \setminus \{\alpha\}) \quad \text{for all } \alpha \in A_j.$$
   This means that by removing any good from the bundle of agent j, agent i no longer envies agent j.

When can we achieve an EF1 allocation? Let's try for $n = 2$ bidders, any m items, and additive valuations. [*Hint:* Try something like cut-and-choose.]

Agent 1 will cut the items into two sets $(A, M \setminus A)$ such that the allocation is EF1 for him and agent 2 will choose the set he prefers.

To divide the items agent 1 uses the following algorithm: He creates two empty bundles A_1 and A_2. Then, he arbitrarily chooses an item and adds it to the lesser valued bundle. If the value of the two bundles is equal, he chooses one of them arbitrarily. We claim that during this process, the allocation is EF1 (for agent 1), and thus the final partition is EF1 as well.

As such, the bundle agent 1 receives is EF1. Moreover, agent 2 chooses the bundle he prefers and thus necessarily doesn't envy agent 1. All in all, the allocation is EF1.

Example 1. Suppose we have two players, three goods, and additive valuations given by the following table:

Player	A	B	C
1	10	9	2
2	1	8	4

In this case, notice that agent 1 might provide the cut $P_1 = (\{A\}, \{B, C\})$ and agent 2 will choose the set $\{B, C\}$. Thus, the final allocation is P_1 which is indeed EF1. Similarly, if agent 2 cuts the items, the bundles he would provide could be $P_2 = (\{A, C\}, \{B\})$; in this case, agent 1 would choose the set $\{A, C\}$ and the final allocation would be P_2 which is EF1.

We now consider the general case of n players, m items and additive valuations. It turns out that an EF1 allocation is always achievable.

The Round Robin Mechanism First, set an arbitrary order of the players; let us denote with 1 the first agent, with 2 the second agent, etc. Then, it proceeds as follows: until all the items are allocated, it goes through the players in the order $1, 2, \dots, n$ in a cyclic manner. When it is the turn of agent i, he chooses the most valuable item (for him) which hasn't been allocated.

Theorem 1. *For any instance of the fair division instance $(N, \mathcal{M}, v_1, \dots, v_n)$ with additive valuations, there exists an EF1 allocation. Moreover, such an allocation is produced by the Round Robin Mechanism.*

Proof. Consider two players i, j in the order of the mechanism. If $i < j$, then in each round, i chose before agent j, and thus clearly doesn't envy j. This gives an EF allocation.

Now, consider the case where $i > j$. In this case, remove the first item chosen by agent j. Thus, we can compare the item picked in round k by agent i with the item picked in round $k+1$ by agent j. Thus, agent i doesn't envy agent j when removing the first item picked by agent j, giving an EF1 allocation. $\square$

But, we claim that the Round Robin Mechanism can produce allocations which are intuitively very unfair. Consider the following example:

Example 2. Suppose there are n players with identical values for the items. There are $n-1$ copies of items with a value of n, and n copies of items with a value of 1. The round-robin algorithm will produce an allocation where the first $n-1$ players receive $\{n, 1\}$ and the last agent receives only a single item valued at 1. This allocation is EF1, but it is clearly unfair. Instead, compare to the allocation that gives the first $n-1$ agents $\{n\}$ and gives the last agent all n copies of the 1's. Thus, all bundles would be equally valued at n and in particular agent n would receive a bundle of value n.

This notion of splitting the items into bundles and then choosing the bundle last relates to a notion of fairness called MMS and leads us to share-based fairness.

Share Based Fairness

We begin by defining two shares of central importance in the study of share-based fairness.

Definition 6. Given a fair division instance $(N, \mathcal{M}, v_1, \ldots, v_n)$,

- We define the max min share MMS_i of agent i as the maximum value which can be achieved by dividing the items into n bundles and receiving the least valuable bundle. Formally:

$$\mathrm{MMS}_i = \max_{(A_1,\ldots,A_n)\in\Pi(\mathcal{M})} \min_{1\leq j\leq n} v_i(A_j),$$

  where $\Pi(\mathcal{M})$ is the set of all partitions of $\mathcal{M}$.

- On the other hand, the proportional share of agent i is the value of the whole bundle divided by the number of players. Thus: $\mathrm{PS}_i = v_i(\mathcal{M})/n$. In the case of normalized valuations, this is simply $1/n$.

In the case of additive valuations, the MMS of agent i is at most his proportional share. These notions of shares lead to natural notions of fairness, in particular:

Definition 7. For a fair division instance $(N, \mathcal{M}, v_1, \ldots, v_n)$, and a given partition $A = (A_1, \ldots, A_n)$,

1. We say that the allocation A is *MMS-fair* if for all $i \in N$, we have that

$$v_i(A_i) \geq \mathrm{MMS}_i.$$

2. Moreover, we say that the allocation A is *proportionally fair* if for all $i \in N$, we have that $v_i(A_i) \geq \mathrm{PS}_i$.

Firstly, proportional fairness cannot be achieved in general. The case of two agents and a single item shows this.

Thus, naturally, we ask when can we achieve an MMS-fair allocation and begin by considering the case of two players with additive valuations. Indeed, in this case, an MMS allocation can be ensured by the Cut-and-Choose mechanism.

Lemma 1. *For the case of $n = 2$ with additive valuations, the Cut-and-Choose mechanism produces an MMS-fair allocation.*

Proof. Firstly, notice that agent 1 can choose the cut $(A, \mathcal{M} \setminus A)$ which guarantees his MMS. In other words, whether agent 1 receives A or $\mathcal{M} \setminus A$, they both have a value of at least MMS_1. Moreover, notice that $v_2(A) + v_2(\mathcal{M} \setminus A) = v_2(\mathcal{M}) = 1$ and thus at least one of the bundles must be valued at least $1/2 = \text{PS}_i \geq \text{MMS}_i$. $\square$

But, it turns out that the Cut-and-Choose mechanism is not sufficient outside of the additive case.

Example 3. Indeed, consider a case with two players and four items, $\mathcal{M} = \{1, 2, 3, 4\}$. Suppose that the valuations of the players are as follows:

$$v_1(\{1,2\}) = v_1(\{3,4\}) = v_1(\mathcal{M}) = 1,$$
$$v_2(\{1,3\}) = v_2(\{2,4\}) = v_2(\mathcal{M}) = 1,$$

where we define the values of all other subsets to be 0, except the subsets which must have a value of 1 due to monotonicity. In this case, A will choose the cut $(\{1, 2\}, \{3, 4\})$ and no matter what B chooses, he will receive a bundle of value 0, while his MMS is 1.

Thus, we shift our focus to the case of n players while restricting our attention to additive valuations. A main negative result in this case is the following:

Theorem 2. *Even for the case of $n = 3$ with additive valuations, an MMS-fair allocation is not feasible.*

Directions of research here are often approximate MMS, approximate envy-freeness, and EFX for small specific cases, as this is a very very hard problem.

Acknowledgements

This lecture was developed in primarily using lecture notes by Moshe Babaioff and materials by Tim Roughgarden, and in particular, lectures from "CS269I: Incentives in Computer Science" [7].

References

[1] Haris Aziz and Simon Mackenzie. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 416–427. IEEE, 2016.

[2] Haris Aziz and Simon Mackenzie. A discrete and bounded envy-free cake cutting protocol for four agents. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 454–464, 2016.

[3] Steven J Brams and Alan D Taylor. An envy-free cake division protocol. *The American Mathematical Monthly*, 102(1):9–18, 1995.

[4] Ya'akov Gal, Moshe Mash, Ariel D Procaccia, and Yair Zick. Which is the fairest (rent division) of them all? In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 67–84, 2016.

[5] Erica Klarreich. How to cut cake fairly and finally eat it too. *Quanta Magazine*.

[6] Ariel D Procaccia. Thou shalt covet thy neighbor's cake. In *IJCAI*, pages 239–244, 2009.

[7] Tim Roughgarden. CS269I: Incentives in Computer Science, 2016.