Fair Division and Cake Cutting

Subhash Suri

February 27, 2020

1 Fair Division

- Questions about fairness have tantalized thinkers for millennia, going back to the times of ancient Greeks and the Bible. In one of the famous biblical stories, King Solomon attempts to *fairly* resolve a maternity dispute between two women, both claiming to the true mother of a baby, by asking the guards to cut the baby in two and give each woman a half. The baby is finally given to the woman who would rather give the whole baby to the other woman than have it be cut in half.

- One of the classic early versions of fair division occurs in *cutting a cake* to be split between two children. A commonly accepted protocol (solution) for this problem is called the "you cut, I choose:" one child cuts the cake into two "seemingly" equal halves, and the second child gets to pick the piece.

  1. Intuitively, this seems like a fair division, but is there a formal way to prove it?
  2. If so, what assumptions about the cake (goods) are needed in this algorithm?
  3. Can this protocol be extended to n-way division, for any $n \geq 3$?
  4. Are there more than one way to define fairness?
  5. These are some of the questions we plan to study in this lecture.

- Cake-cutting may seem like a problem in recreational mathematics, but ideas behind these algorithms have been applied to a number of important real-world problems, including *international negotiations* and *divorce settlements*, and there are even commercial companies, such as *Fair Outcomes*, based on cake cutting algorithms.

- We tend to think of cakes as *homogeneous* (i.e. uniform) entities, for which fair division seems easy to achieve, our mathematical formulation will deal with *non-homogeneous* cakes; e.g. cakes with highly non-uniform sprinkling of toasted nuts, chocolate icing, or other toppings, and players who may have vastly different *preferences* for those ingredients.

- Formally, our mathematical model for the cake simply assumes that it is a one-dimensional interval $[0, 1]$, and each player i has a valuation function v_i, specifying the value $v_i(S)$ that player i assigns to each subset S of the cake.

- We will make the following two assumptions about each valuation function v_i:

  – v_i is normalized with $v_i([0, 1]) = 1$, which is easily achieved by scaling.

  – v_i is additive on disjoint subsets. That is, if $A, B \subset [0, 1]$ are disjoint, then

$$v_i(A \cup B) \quad = \quad v_i(A) \ + \ v_i(B)$$

- There are at least two different but natural ways to define fairness.

  1. **Proportionality.** An allocation $A_1, A_2, \ldots, A_n$ of cake to n players is *proportional* if, for every player i, we have

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

     That is, each player feels that he received at least $1/n$ fraction of the cake (under his valuation).

  2. **Envy-freeness.** An allocation $A_1, A_2, \ldots, A_n$ of cake to n players is *envy-free*, for every pair of players i, j, we have

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

     That is, no player is envious of any other player's piece.

- Observe that an allocation may be proportional but still not envy-free: a player may have $v_i(A_i) \geq 1/n$ yet may still be envious of some other player's piece. (Thus, proportionality does not imply envy-freeness.)

- One can show, however, that envy-freeness implies proportionality, as follows. For each player, we have $\sum_{j=1}^{n} v_i(A_j) = v_i([0, 1]) = 1$. Thus, if player i likes A_i better than any other piece (by envy-freeness), then it must satisfy $v_i(A_i) \geq 1/n$.

- In the case of $n = 2$ players, both definitions are satisfied: if you have a piece of value $\geq 1/2$, swapping cannot get you a piece with value $> 1/2$.

- **Homework exercise.** Show an example with $n \geq 3$ where a proportional allocation is not envy-free.

2 Proportional Cake Cutting

2.1 2-Way Division

- It is easy to see that the "cut-and-choose" algorithm finds a proportional cake division for $n = 2$ players.

  1. The first player cuts the cake into two equal pieces A_1, A_2, and so his valuations must satisfy $v_1(A_1), v_1(A_2) \geq 1/2$.
  2. The player 2 clearly chooses one for which $v_2(A_i) \geq 1/2$, and so both get a proportional division.

- With $n = 3$ players, an obvious extension of the cake cutting algorithm is the following:

  1. Alice cuts the cake into 3 equal pieces (by her measure)
  2. Bob then chooses the piece he consider largest (by his measure)
  3. Carole then chooses the larger (by her measure) of the two remaining pieces
  4. Alice gets the last piece.

- Is this protocol fair?

  – The protocol is certainly fair for Alice, because she can ensure each piece is worth 1/3.

  – The protocol is also fair for Bob because at least one of the three pieces is worth $\geq 1/3$ to him.

  – Carole, on the other hand, is not guaranteed a fair share. In fact, Alice and Bob can easily collude to cheat Carole, by letting Alice cutting one large piece and two tiny pieces.

2.2 n-Way Cake Cutting: Dubins-Spanier (Moving Knife) Algorithm

- The general case of n-way division was solved by **Dubins and Spanier**, using what's often called a "moving knife" algorithm:

  1. A neutral party, called "referee," slowly moves a knife over the cake, from left to right.
  2. When the knife reaches a point at which the cake to the left is worth $1/n$ to one of the players, this playe shouts "stop." (Ties can be broken arbitrarily.)

3. The referee makes a cut at that point, and gives the piece to this player, who then drops out of the competition, and the game continues with the remaining players.
4. The last player receives the unclaimed piece.

- That this algorithm produces proportional division can be argued as follows.

  1. Clearly, each player other than the last receives a piece that he values at $1/n$.
  2. The value of the last player for each of the unclaimed pieces is $\leq 1/n$, and so by additivity, the value of the last piece is at least $1 - (n-1)/n = 1/n$.

- One can implement the Dubins-Spanier algorithm through a "discrete" procedure, and without the impartial referee, as follows.

  1. At each stage, each remaining player makes a mark so that the piece of the cake to the left of the mark is worth $1/n$ to him.
  2. We then cut the cake at the *leftmost* of these marks, and give the piece to the player who made this mark.
  3. We then continue with the remaining players.

- It is easy to see that this is a discrete simulation of the continuous moving knife algorithm.

- Dubins-Spanier algorithm does not necessarily produce envy-free allocations: a player certainly does envy other players who come earlier in the allocation, they he can easily envy recipients of later allocations.

- We will return to envy-free division later, but for now let us turn to *algorithmic complexity* of proportional division.

2.3 Computational Complexity of Cake Cutting

- The discrete Dubins-Spanier algorithm makes $\Theta(n^2)$ (tentative) cuts—$O(n)$ cuts in each of the n rounds.

  – Can one divide the cake with fewer cuts?
  – What is the minimum number of cuts necessary?

- In order to answer these questions, we should begin with a formal model of what constitutes an "elementary operation" in cake-cutting algorithms.

- A standard model is one proposed by Robertson and Webb in their 1998 book "Cake Cutting Algorithms: Be Fair If You Can."

- As before, the cake is represented by the unit interval $[0, 1]$, and for any piece S, which is a subset of $[0, 1]$, we use $v_i(S)$ to denote the value of player i for piece S.
- In the Robertson-Webb model, a cake-cutting algorithm uses only the following two operations:

  1. **Evaluation.** Given two points $x, y \in [0, 1]$, with $x \leq y$, ask a player i for his value of the cake interval $[x, y]$. That is,

  $$eval_i(x, y) = v_i([x, y])$$

  2. **Cut.** Ask a player i to make a subinterval worth a given value α, starting at a given point x. That is,

  $$cut_i(x, \alpha) \;=\; y \quad \text{such that} \quad v_i([x, y]) = \alpha$$

- This may seem like a restricted model, but all known cake cutting algorithms can be simulated using these two operations.
- For instance, Cut-and-Choose makes only two queries: $cut_1(0, 0.5)$, which returns a point y such that $v_1([0, y]) = 1/2$, and $eval_2(0, y)$. If the answer to the second query is $\geq 1/2$, then player 2 gets the piece $[0, y]$; otherwise, he gets the remaining piece.
- In this model, the Dubins-Spanier algorithm clearly needs $O(n^2)$ queries.
- We now describe an algorithm by Even-Paz that produces n-way proportional division with only $O(n \log n)$ queries.

2.4 Even and Paz Algorithm

- The Even-Paz algorithm uses divide-and-conquer.

  1. If $n = 2$, solve the problem with 1 cut.
  2. Otherwise, ask each of the n players to make a mark x_i such that the cake to the left of x_i has $v_i([0, x_i]) = 1/2$. That is, instead of asking the player to propose a piece of size $1/n$, we ask them to propose the piece of half the size.
  3. Choose the *median* mark x^* among $\{x_1, x_2, \ldots, x_n\}$, and make a cut at x^*.
  4. Divide the players into two groups L and R, where L contains all players whose mark is left of (or coincident with) x^*, and R contains the remaining players.
  5. By the median property, $|L| = |R| \leq \lceil n/2 \rceil$.

6. Recursively solve the problem for the left half-cake with players L and for the right half-cake with players R.

- For ease of analysis, assume that n is a power of 2; otherwise, in each round we ask the players to cut the cake of size $\lceil n/2 \rceil$.
- The correctness of the algorithm is easy to prove: all players in L value their piece to have values $\geq 1/2$, and by additivity, all players in R have the same valuation for the right half-cake. This property holds recursively, and so at termination each player get a piece worth $1/n$ to him.
- **Query Complexity.** How many cut and eval queries are made by this algorithm?

  1. In the first stage, each player makes one cut, for a total of n. In the next round, we have two groups of $n/2$ players, and each makes one cut again, and so on.
  2. In each round, the size of the cake drops by half, so the number of rounds before each group has at most one player is $O(\log n)$.
  3. Thus, the total number of cut queries is $O(n \log n)$.
  4. The number of eval queries is only $O(n)$, needed at the end to decide which of the two pieces to allocate.

- **Theorem.** A cake can be divided proportionately among n players using $O(n \log n)$ cut-and-eval queries.
- A natural question now is whether the complexity can be improved further, perhaps to $O(n)$ queries.
- Answering this questions partially Woeginger and Sgall (" On the complexity of cake cutting," Discrete Optimization, 2007) proved the following:

  any cake-cutting algorithm that allocates only connected pieces of the cake to players requires at least $\Omega(n \log n)$ queries.

- Can allocating disconnected pieces of cake lead to better algorithms?
- No! Edmonds and Pruhs ("Cake cutting is not a piece of cake," Proc. of Symp. on Discrete Algorithms, 2006) extended the lower bound of Woeginger-Sgall and proved the $\Omega(n \log n)$ lower bound for the general case.

3 Envy-free Cake Cutting

- While proportionality in cake cutting is well understood, envy-freeness is a far more elusive property.
- The first advance was made around 1960 when Selfridge and Conway independently discovered the same protocol for envy-free division among *three players*! (Neither saw it worthy of publishing, so it came to light only years later.)
- **Selfridge-Conway Algorithm.**
  1. **Stage 0:** Player 1 divides the cake into 3 equal parts, according to his valuation. Player 2 "trims" the largest piece (cuts off a slice) such that there is a tie between the two largest pieces, according to his valuation. *Set aside this trimming as bonus slice, to be allocated later.*
  2. **Stage 1:** Player 3 chooses one of the three cake pieces. If player 3 *does not choose* the trimmed piece, then *give the trimmed piece to player 2.* Otherwise, player 2 chooses one of the other two pieces. Thus, either player 2 or player 3 receives the trimmed piece; we denote that player by T, and other player by T'. Player 1 is allocated the remaining (untrimmed) piece.
  3. **Stage 2:** T' divides the bonus slice into three equal parts according to his valuation. Player T, 1, and T' choose the pieces of this bonus in this order.
- **Proof of envy-freeness.**
  1. We first argue that the division of the cake (without the bonus slice) is envy-free. Player 3 chooses first; player 2 likes the two tied pieces equally (the trimmed one and second one), and so is guaranteed to receive one of those pieces. Finally, since player 1 made the original cuts, he is indifferent between the two untrimmed pieces, and he receives one of those.
  2. Division of the bonus slice is more subtle. Player T goes first, and so does not envy the others; and player T' is indifferent because he cut it into three pieces.
  3. Player 1 does not envy T' but may prefer the bonus cake slide allocated to T to his own. However, at the end of Stage 1, player 1 has a "irrevocable advantage" over T: even if we allocated all of the bonus slice to T, we would reconstruct just one of the three original pieces cut by player 1 (each worth 1/3 to him); but player 1 already received a piece worth 1/3 at the end of stage 1.
- **Extending Envy-free Division to n Players.** Unfortunately, extending envy-free division to more than 3 players has proved to be far more challenging that expected. In fact, the problem defied any progress for *three decades* until Brams and Taylor made a

stunning breakthrough in 1992—*the first ever envy-free cake cutting algorithm for any number of players!* ("An envy-free cake division protocol," American Mathematical Monthly, 1995.)

- The algorithm is maddeningly complex—in fact, solving the 4-player game is already extremely complicated, and just the special case of 4-players algorithm requires 20 steps!

- The Brams-Taylor result, as celebrated a result as it is, also suffers from a major computational flaw: the number of steps is unbounded. While for any fixed valuations $v_1, v_2, \ldots, v_n$ the algorithm halts in finite number of steps, for every $n \geq 4$ and T, there is a choice of v_i's for which the protocol requires more than T steps to terminate.

- A major advance was made recently by Aziz-Mackenzie whoose envy-free division algorithm needs *only*

$$n^{n^{n^{n^{n^{n}}}}}$$

steps (a tower of 6 n's). ("A discrete and bounded envy-free cake cutting protocol for any number of agents," FOCS 2016.)

4 Fair Rent Division

- The website *spliddit.org* is used by tens of thousands for a fair rent division. (The website provides fair division protocols, not just for rent.)

- In the rent division problem, a group of n people want to rent a house with n room, and want a fair allocation of rooms and rent. In other words, the goal is to assign people and rents to rooms, with one person per room and with the sum of rents equal to the total rent R in the "best" possible way.

- Assume that each person i has a value v_{ij} for each room j. We normalize these values so that $\sum_j v_{ij} = R$. (This is simply acknowledging that rent must get paid.)

- We use a quasi-linear utility function, a standard assumption in economics, by which each player i want to maximize its *utility*, namely, value v_{ij} minus the rent she pays for room j.

- A solution to a rent division problem is envy-free if, for every pair i, j of players,

$$v_{i\alpha(i)} - r_{\alpha(i)} \quad \geq \quad v_{i\alpha(j)} - r_{\alpha(j)}$$

where $\alpha(i)$ is the room assigned to i, and r_j is the rent assigned to room j.

- That is, under envy-free rent division, no player wants to trade places with anyone else, where trading places means swapping *both* the rooms and the rents.
- The good news is that an envy-free solution is guaranteed to exist, and that that it can be computed efficiently.

5 Fair Allocations and Social Welfare

- Envy-freeness and proportionality are fairness ideas aimed at *individuals*. What implications do they hold if we wish to promote interests of the society as a whole?
- The social welfare is a quantification of the "happiness" of the society as a whole. Typically, the welfare is considered in flavors" *utilitarian*, which is the *sum* of individual player's values for their allocations, and *egalitarian*, which is determined by the *lowest* value of any player. In other words, the former uses a *max-sum* measure while the latter uses *max-min*.
- Intuitively, there is often an inevitable tension between interests of individuals and of society. Along the lines of *price of anarchy*, the computer scientists have tried to study this tension using the ratio *price of fairness*: this is the worst-case ratio between social welfare of an optimal allocation and the social welfare of the best *fair allocation*.
- For instance, a price of fairness of 2 means that the social welfare of the best fair allocation is at most 50% of what it could be if the fairness restriction was removed.
- First, when considering the interests and happiness of different individuals, we need to make sure a common yardstick exists. We will assume that the *all players have the same value for the whole cake*, which we can assume is $1.
- In a fair allocation, we wish to ensure that *all players* have about the same value of their allocations. We show how requirement may cause an *economically inefficient* division.
- Suppose there are $\sqrt{n}$ "large" players, where the ith player only values the ith $1/\sqrt{n}$-size slice of the cake. That is, if we partition the cake in $\sqrt{n}$ pieces each of size $1/\sqrt{n}$, then each of these large players wants a unique piece.
- More precisely, the ith player valuation is $v_i([\frac{i-1}{\sqrt{n}}, \frac{i}{\sqrt{n}}]) = 1$, 0 for the rest of the cake. Within this desirable slice, his value is *uniform*, namely, 1/10 for a tenth of that piece.
- The remaining $n - \sqrt{n}$ "small" players value the entire cake uniformly, meaning, their value is x for any piece of size x.

- Under a fair allocation, each of the players must get $1/n$ size slice. The total value (social welfare) of this fair allocation is

$$(n - \sqrt{n}) \times \frac{1}{n} \ + \ \sqrt{n} \times \frac{1/n}{1/\sqrt{n}} \quad \approx \quad 2$$

  (After allocating to the small players, the leftover cake has size $1/\sqrt{n}$, whose total value among all the large players is 1.)

- In contrast, if we dropped the fairness requirement, we can allocate the cake to just the $\sqrt{n}$ large players, achieving a total value (social welfare) of $\sqrt{n}$.

- Thus, the price of fairness is at least $\sqrt{n}/2$.

- The price of fairness for envy-free division is also at least this high because envy-freeness implies proportionality.

- Social fairness has other counter-intuitive properties as well, including "dumping paradox," where by discarding a piece of the cake one can increase the social welfare of the best proportional or best envy-free allocation.

- A high price of fairness means that there are examples where fair allocations are severely sub-optimal from the society's point of view. Nevertheless, these examples may be rare and do not preclude the possibility of usually obtaining high social welfare even under fairness constraints.

- This is an interesting and active area of research in algorithmic fairness.

6 Notions of Fairness

- Fairness has been extensively studied, especially in social science, welfare economics, and engineering.

- There are many subjective interpretations of fairness, and no single universally accepted single principle. Nevertheless, there are general theories of justice and equity that figure prominently in the literature, on which most fairness schemes are based.

- Among the most prominent, the oldest theory of justice is Aristotle's equity principle, according to which resources should be allocated in proportion to some preexisting claims or rights to the resources that each player has. (Such a principle makes sense, for instance, for allocating profit among shareholders.)

- A second theory, widely considered in economics in the 19th century, is classical utilitarianism, which dictates an allocation of resources to maximize the sum of player utilities. (This principle has been criticized on ethical grounds: in maximizing the sum of utilities, the utility of some players may be greatly reduced to confer benefit to the system.)

- A third approach is based on Rawls theory of justice, which gives priority to those that are the least well off, so as to maximize the smallest utility of any player.

- A final theory uses Nash equilibria, in which resource is transferred from one player to another if former's utility loss is smaller than latter's gain.