Progress on Envy-Free Cake-Cutting: Beating $O(n^{n^{n^{n^{n^{n}}}}})$ Queries

Alejandro Gomez-Leos, Connor Colombe

UT Austin, Algorithmic Perspective on Microeconomics

November 14, 2024

Objectives

main takeaways

1. cake-cutting : a fundamental model of fair division
2. many open problems, some longstanding
3. "can we efficiently compute envy-free allocations?"

What is cake-cutting?

question: how can we fairly divide a cake amongst people?

1. or heterogenous, divisible good
   - value perceived individually
   - can "chop it up"
2. **cake** : represented by $[0, 1]$
3. **people** : n agents
4. **value** : measures $\mu_1 \ldots \mu_n$
   - $\mu_i([0, 1]) = 1$
   - non-atomic
   - dom. by Lebesgue
5. **allocation** : subsets to players
6. **fairness** : many notions

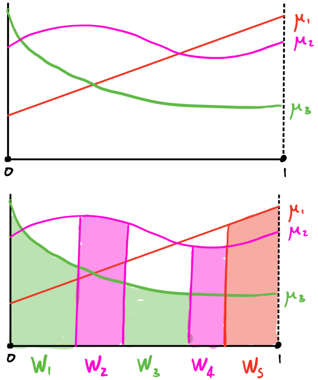

Figure: Example instance and allocation. Yum.

A Notion of Fairness

giving player i piece W_i, the allocation is:

- *proportional* if for all i

$$\mu_i(W_i) \geq 1/n$$

"everyone believes they have proportional slice"

- (exact[1]) *equitable* if for all $i \neq j$

$$\mu_i(W_i) = \mu_j(W_j)$$

"everyone equally satisfied"

- (exact) *envy-free* if for all $i \neq j$

$$\mu_i(W_i) \geq \mu_i(W_j)$$

"every player believes they have the biggest piece"

[1] often these notions admit an ϵ relaxation in the additive sense

A Notion of Efficiency

how to characterize efficiency?

- say alg. gets preferences via questions
- efficient alg. asks few questions as possible

Robertson-Webb Query Model [Woeginger-Sgall '07]

- **Eval**(i, x, y) : get $\mu_i([x, y])$
- **Cut**(i, x, α) : get threshold y such that $\mu_i([x, y]) = \alpha$

query complexity

- $\#$ queries needed to compute fair allocation
- care about bounds in n

Known Results

- proportional :
  - $\Theta(n \log n)$ [Even-Paz '84, Edmonds-Pruhs '06, Woeginger-Sgall '07]
- equitable:
  - connected and exact $\implies \nexists$ algorithm [Cechlarova-Pillarova '12]
  - connected and ϵ-equitable $\implies O(n \log \frac{n}{\epsilon})$ [Cechlarova-Pillarova '12]
  - ϵ-equitable[1] $\implies \Omega(\frac{\log \frac{1}{\epsilon}}{\log \log \frac{1}{\epsilon}})$ [Procaccia-Wang '17]
- envy-free:
  - connected and ϵ-envy-free $\implies O(\frac{n}{\epsilon})$ and $\Omega(\log \frac{1}{\epsilon})$ [Brânzei-Nisan '18]
  - exact $\implies O(n \uparrow\uparrow 6)$ and $\Omega(n^2)$ [Aziz-Mackenzie '16, Procaccia '09]
  - exact and extra assumptions $\implies n^{O(1)}$ [Chèze '21, Webb, '99]

focus of this talk: computing envy-free allocations

[1] strengthens Cechlarova-Pillarova to allow "crumbs"

Overview

1. setting + known results ✓
2. envy-free for $n = 2, 3, 4$.
3. Aziz and Mackenzie ($O(n \uparrow\uparrow 6)$)
4. Webb's Algorithm & Chéze's Result ($n^{O(1)}$)
5. strengthening Chéze's

Examples and Some History of Cake Cutting Algorithms

- Origins of "Cake Cutting"
- $n = 2$: The Cut and Choose Algorithm [The Bible?]
- $n = 3$: Selfridge–Conway Procedure [Selfridge, Conway '60]

Origins of "Cake Cutting"

- Term introduced in the 1940's by Hugo Steinhaus to make the idea of fair division more tangible
- Steinhaus and his colleagues, Knaster and Banach introduced the notion of envy-freeness and worked to develop protocols for **proportional divisions** on n agents
- Were aware of the cut and choose protocol but could not extend it to $n = 3$.

$n = 2$: Cut and Choose

- Pretty Old! Appears in the Bible as a way to divide land.
- Players A and B
- Step 1: A cuts the cake into two pieces they think are "equal value"

- Step 2: B chooses the piece they would like and A gets the other.

- This algorithm is envy free!
- A recurring idea in cake cutting: "The person who cuts is fine with getting any of the pieces they cut"
- In the RW model, this takes **3 queries**. One **Cut** query ($\textbf{cut}(1, 0, 1/2)$) and two **eval** queries ($\textbf{eval}(2, 0, x)$, $\textbf{eval}(2, x, 1)$)

$n = 2$: Cut and Choose

- Note that this and the other protocols we discuss are **envy-free protocols**. This means that if an **agent is truthful** about their query responses, then they are **guaranteed to have an envy-free allocation**.
- Envy-free protocol $\not\Rightarrow$ strategy-proof!
- If agents know the valuations of others, they may be able to obtain a "better" allocation.
- However, an agent deviating from truthful reporting can only make themselves envious and **does not affect the envy-free-ness of truthful agents** so we will **assume truthful reporting for the remainder of the presentation**.

$n = 2$: Cut and Choose

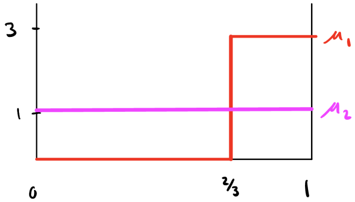

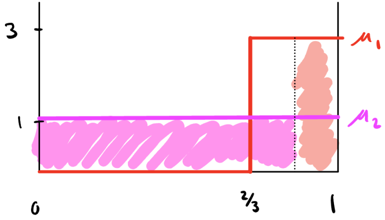

Player 1: True Valuation Cut

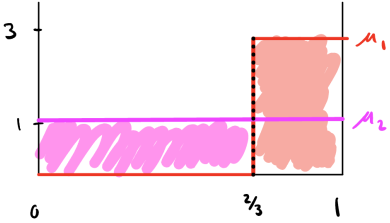

Player 1: False Valuation Cut

$n = 3$: Selfridge–Conway Procedure

- Independently discovered by **John Conway and John Selfridge** in the 60's
- Let the agents be A, B, and C
- Step 1: A cuts the cake into 3 "equal" pieces
  - A would be happy with any **whole piece**, i.e. $\mu_A(1) = \mu_A(2) = \mu_A(3)$

$n = 3$: Selfridge–Conway Procedure

- Step 2: Let B and C pick their favorite pieces of the 3
  - If they choose different pieces, we are done
- Step 3: WLOG they both want piece 1, then B "trims" piece 1 so that **the trimmed piece is equal in value to their second favorite piece**. (WLOG suppose $\mu_B(1') = \mu_B(2)$ here)

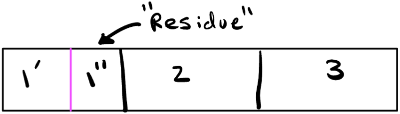

- Step 4: Separate the residue from the cake and have the players choose their pieces from the rest of the cake in the order: $C \to B \to A$.
  - * If C picks 3, then B must pick 1'

$n = 3$: Selfridge–Conway Procedure

- **Claim: The current partial allocation is envy free!**
  - C picked first so they got the best piece in their eyes
  - B is guaranteed a piece they said was most valuable to them ('1 or 2)
  - A cut the pieces originally and is guaranteed a "whole" piece
- It remains to allocate the residue. Suppose WLOG B got 1' in the partial allocation (otherwise swap roles of B and C from here on).
- Note: A **does not care if** B **(player who got the trimmed piece) gets the entire residue**. We say A "dominates" B.

$n = 3$: Selfridge–Conway Procedure

- Step 5: To allocate the residue: C cuts the cake into three "equal" pieces, B picks, A, picks and then C picks.

  - B thinks they got the best piece
  - A doesn't care what B got but thinks they got a better piece than C
  - C is equally happy with any piece
  - The total allocation is envy-free!

  - In the worse case, we need **15 RW queries to achieve an envy-free allocation**

$n \geq 4$?

- "We figured out $n = 3$, surely $n = 4$ can't be that bad right?"
- $n \geq 4$ was considered a major open problem in mathematics
- A major breakthrough in 1995 with The first algorithm for any n [Brams and Taylor '95].
  - While guaranteed to terminate, the number of cuts was dependent on the valuation functions of the agents
  - For any constant c, you can find valuation functions even with $n = 4$ to make the number of cuts in the protocol exceed c
- Does there exist a bounded protocol for $n \geq 4$?
- The first algorithm (with bounded complexity) for $n = 4$ [Aziz and Mackenzie '16]
- Later generalized to any n [Aziz and Mackenzie '17]
- Recently $n = 4$ envy-freeness was shown to be achievable in fewer than 200 queries. [Amanatidis et al. '18]

First Envy Free With Bounded Query Complexity,[Aziz-Mackenzie '17]

- First bounded protocol for general number of agents n
- But it might take a while $\mathcal{O}(n^{n^{n^{n^{n^{n}}}}})$ ...
- At a high-level the protocol works to **find envy-free partial allocations** in which a **subset of players dominates the rest**.
  - We can then remove these players and work on a smaller subproblem
- A key novel idea is to **allow players to swap portions of their allocated cake** to achieve a domination. **Kick out the dominating players and solve smaller instance**.
- The entire protocol is very complex. The runtime is due to many iterations looking over permutations of players, over permutations of allocations, etc.

Overview

1. setting + known results ✓
2. envy-free for $n = 2, 3, 4$, Aziz and Mackenzie ($O(n \uparrow\uparrow 6)$) ✓
3. Webb's Algorithm & Chéze's Result ($n^{O(1)}$)
4. strengthening Chéze's

Webb's Envy-Free Algorithm (1/8)

computes envy-free in $n^{O(1)}$ under assumptions

strategy:

- guess a partition
- subdivide each into n pieces

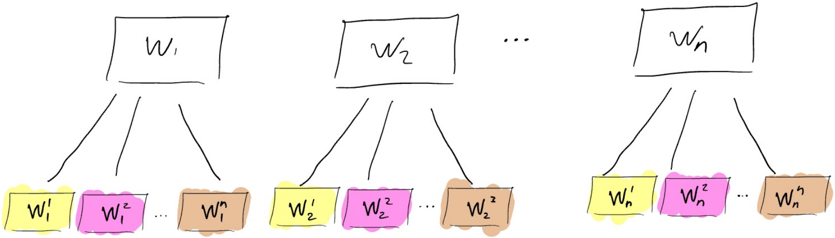

- give each player a union of pieces

Webb's Envy-Free Algorithm (2/8)

- specifically, allocate

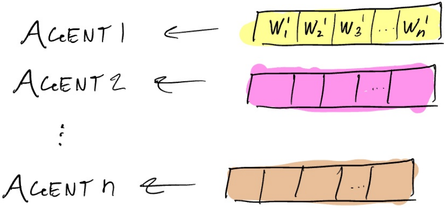

- ... such that division is envy-free

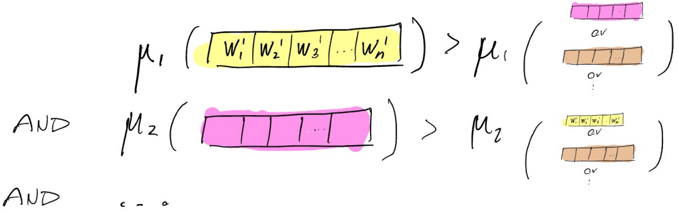

suffices to get "super envy-free" allocation, i.e.

$$\mu_i(W_1^j + W_2^j + \cdots + W_n^j) = \begin{cases} > 1/n & j = i \\ < 1/n & j \neq i \end{cases}$$

why doable?

Theorem (Barbanel, '96)

A super envy-free subdivision of $W \subseteq [0,1]$ exists iff $\mu_1 \ldots \mu_n$ are linearly independent measures, i.e. $\sum c_i \mu_i = 0$ only for the trivial $\vec{c}$.

(assuming linearly independence throughout)

suffices for $\delta > 0$ that

$$\sum_k \mu_i(W_k^j) = \begin{cases} 1/n + \delta & j = i \\ 1/n - \delta/(n-1) & j \neq i \end{cases}$$

Webb's Envy-Free Algorithm (4/8)

$$\sum_k \mu_i(W_k^j) = \sum_k \mu_i(W_k) \cdot \frac{\mu_i(W_k^j)}{\mu_i(W_k)}$$

$$:= \sum_k \mu_i(W_k) \cdot R_{k,j,i} \quad (\text{note } \sum_j R_{k,j,i} = 1)$$

suppose for all $i \neq i'$, $R_{k,j,i} = R_{k,j,i'}$

"all agents believe W_k^j / W_k the same"

then have

$$\sum_k \mu_i(W_k) \cdot R_{k,j} = \begin{cases} 1/n + \delta & j = i \\ 1/n - \delta/(n-1) & j \neq i \end{cases}$$

Webb's Envy-Free Algorithm (5/8)

$$\sum_k \mu_i(W_k) \cdot R_{k,j} = \begin{cases} 1/n + \delta & j = i \\ 1/n - \delta/(n-1) & j \neq i \end{cases}$$

i.e.

$$\underbrace{\begin{pmatrix} \mu_1(W_1) & \dots & \mu_1(W_n) \\ \mu_2(W_1) & \dots & \mu_2(W_n) \\ \vdots & \ddots & \vdots \\ \mu_n(W_1) & \dots & \mu_n(W_n) \end{pmatrix}}_{\mathbf{M}} \mathbf{R} = \underbrace{\begin{pmatrix} \frac{1}{n}+\delta & \frac{1}{n}-\frac{\delta}{n-1} & \dots & \frac{1}{n}-\frac{\delta}{n-1} \\ \frac{1}{n}-\frac{\delta}{n-1} & \frac{1}{n}+\delta & \dots & \frac{1}{n}-\frac{\delta}{n-1} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{1}{n}-\frac{\delta}{n-1} & \frac{1}{n}-\frac{\delta}{n-1} & \dots & \frac{1}{n}+\delta \end{pmatrix}}_{\mathbf{N}_\delta}$$

(pick δ so $\mathbf{R}$ row stochastic and nonnegative)

Webb's Envy-Free Algorithm (6/8)

if **M** invertible, we can:

- compute $\mathbf{R} = \mathbf{M}^{-1}\mathbf{N}_\delta$
- subproblem: chop $W_1 \dots W_n$ so that $\mathbf{R}_{kj} = \frac{\mu_i(W_k^j)}{\mu_i(W_k)}$

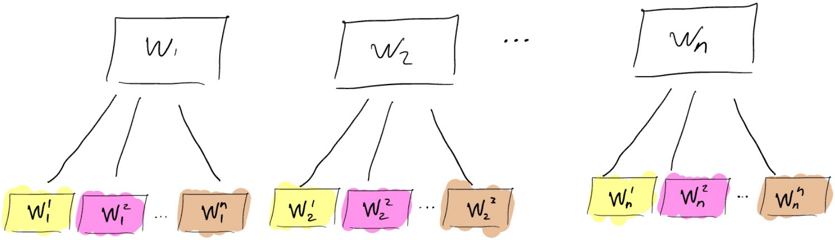

- these pieces will be satisfy what we want

Webb's Envy-Free Algorithm (7/8)

solving "subproblem: chop $W_1 \dots W_n$ so that $\mathbf{R}_{kj} = \frac{\mu_i(W_k^j)}{\mu_i(W_k)}$"?

Definition

A partition $W = W^1 \sqcup \cdots \sqcup W^n$ is ϵ-*exact* for fractions $(\alpha_1, \dots, \alpha_n)$ if

$$\forall i, j \quad \frac{\mu_i(W^j)}{\mu_i(W)} \approx_\epsilon \alpha_j$$ "all agents believe W^j/W roughly α_j"

Theorem (Robertson-Webb, '04)

There's an algorithm $\mathtt{NearExact}(W, \vec{\alpha}, \epsilon)$ *which outputs an* ϵ*-exact partition* $W = W_1 \sqcup \cdots \sqcup W_n$ *in* $O(n^{2.5}/\epsilon)$ *queries.*

- call for each W_k

Webb's Envy-Free Algorithm (8/8)

Webb's Envy-Free Algorithm (1999)

1. starting partition $[0,1] = W_1 \sqcup \cdots \sqcup W_n$
2. $\mathbf{M}_{ij} = \mu_i(W_j)$, stop if not invertible $\qquad \triangleright\, O(n^2)$ queries
3. $t :=$ min entry of $\mathbf{M}^{-1}$, $\delta = \frac{n-1}{n(1-tn)}$
4. compute $\mathbf{M}^{-1}\mathbf{N}_\delta = \mathbf{R} = \begin{pmatrix} \vec{R}_1 \\ \vdots \\ \vec{R}_n \end{pmatrix}$
5. $\forall$ pieces, get `NearExact`$(W_k, \vec{R}_k, \frac{\delta}{3n})$ $\qquad \triangleright\, O(n^{4.5}|t|)$ queries
6. allocate: agent i gets $W_1^i \sqcup W_2^i \sqcup \cdots \sqcup W_n^i$

Recap of Webb's

what did we show?

Theorem (Webb, '99)

If μ_i linearly independent and $\mathbf{M}$ nonsingular for starting partition, then Webb's returns a (super) envy-free allocation in $O(n^{4.5} \cdot \kappa(\mathbf{M}))$ queries.

1. this for existence of super envy-free division
2. this for correctness

note: (1) $\not\Rightarrow$ (2)

- efficiently find starting partition?
- checking candidate $O(n^2)$, but exponentially many
- we don't know satisfying answer

Chéze's Result

- notice Webb's is "brittle"
- we understand random matrices now
- hit **M** with one?

Theorem (Chéze '21)

Suppose $\mu_i(x) > \epsilon$ everywhere. Let $\mathbf{E}$ be random matrix with iid entries in $(-\epsilon, \epsilon)$. Then, Webb's ran on the matrix $\tilde{\mathbf{M}}_{ij} = \frac{\mathbf{M}_{ij}+\mathbf{E}_{ij}}{\sum_j (\mathbf{M}_{ij}+\mathbf{E}_{ij})}$ uses more than $C_\epsilon n^{O(1)}$ queries with probability $o(\frac{1}{n})$.

- satisfying?
  - relationship with final allocation?
  - ϵ dependence?
  - instances with 0 densities ?
- smoothed query complexity?

Overview

1. setting + known results ✓
2. envy-free for $n = 2, 3, 4$, Aziz and Mackenzie ($O(n \uparrow\uparrow 6)$) ✓
3. Webb's Algorithm & Chéze's Result ($n^{O(1)}$) ✓
4. strengthening Chéze's

Towards a Smoothed-Analysis

Recall input to Webb's is μ and starting partition, say $\mathcal{P}$

Definition

Fix $\sigma > 0$. For Webb's instance $\mathrm{I}_{\mu,\mathcal{P}}$ let $\mathrm{I}^{\sigma}_{\mu,\mathcal{P}}$ be a random instance:

- $\hat{\mathbf{G}}$: random matrix iid entries $|\mathcal{N}(0, \sigma^2)|$
- $\tilde{\mu}_i(x) = \frac{1}{\tilde{\mu}_i([0,1])}(\mu_i(x) + \hat{\mathbf{G}}_{ij})$ for $x \in \mathcal{P}_j$

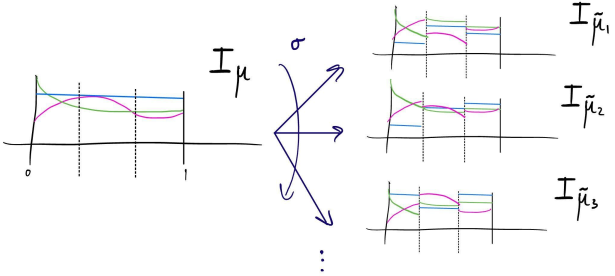

Figure: Some realizations of instance perturbation

Towards a Smoothed-Analysis

Conjecture (smoothed query complexity of Webb's)

Fix $\sigma > 0$. Denoting $Q(\cdot) = \#$ queries by Webb's,

$$\max_{\mathrm{I}_{\mu,\mathcal{P}}} \mathbb{E}_{\mathrm{I} \sim \mathrm{I}^{\sigma}_{\mu,\mathcal{P}}} [Q(\mathrm{I})] = O\left(\frac{poly(n)}{\sigma^2}\right)$$

- i.e. the σ-smoothed query complexity is not bad
- linear independence doesn't matter anymore (satisfied w.p. 1).
- what can we prove today?

Towards a Smoothed-Analysis

we can "almost" prove:

Conjecture ("two-draw" smoothed query complexity of Webb's)

Fix $\sigma \in (0, 1/n)$. Denoting $Q(\cdot) = \#$ queries by Webb's,

$$\max_{\mathrm{I}_{\mu,\mathcal{P}}} \mathbb{E}_{\mathrm{I},\mathrm{I}' \sim \mathrm{I}^{\sigma}_{\mu,\mathcal{P}}} \left[\min\{Q(\mathrm{I}), Q(\mathrm{I}')\} \right] = O\left(\frac{n}{\sigma^2}\right)$$

next: what we have, and what we need

What we have

Recall Webb's runtime $O(n^{4.5} \cdot \kappa(\mathbf{M}))$

In perturbation, deal with $\tilde{\mathbf{M}}_{\mathbf{G}} := \underbrace{\mathbf{D}}_{\text{renormalize}} \Big(\underbrace{\mathbf{M}}_{\substack{\text{orig. matrix} \\ \mu,\ \mathcal{P}}} + \underbrace{\mathbf{G}}_{\text{shifts}} \Big)$

Claim (this work)

For any instance $\mu, \mathcal{P}$ *giving rise to* $\mathbf{M}$, *and*

- $\mathbf{G}$ *: random matrix with iid* $\mathcal{N}(0, \sigma^2)$ *entries*

Then,

$$\mathbb{E}[\min\{\kappa(\tilde{\mathbf{M}}_{\mathbf{G}}), \kappa(\tilde{\mathbf{M}}_{\mathbf{G}}{}')\}] = O\left(\frac{n}{\sigma^2}\right)$$

- not our perturbation model (these might "sign" measures)

What we need

if this true then our "two-draw" conjecture follows

Conjecture

- $\mathbf{G}$ *: random matrix iid entries* $\mathcal{N}(0, \sigma^2)$
- $\hat{\mathbf{G}}$ *: random matrix iid entries* $|\mathcal{N}(0, \sigma^2)|$

Then, for any square (nonnegative row stochastic) $\mathbf{M}$*, we have*

$$\mathbf{E}[\kappa(\mathbf{M}_{\hat{\mathbf{G}}})] \leq O(poly(n, \sigma^{-1})) \cdot \mathbf{E}[\kappa(\mathbf{M}_{\mathbf{G}})]$$

- stability of our perturbation model isn't "much worse" than gaussian
- intuitive - models pretty similar
- numerically supported

Overview

1. setting + known results ✓
2. envy-free for $n = 2, 3, 4$, Aziz and Mackenzie ($O(n \uparrow\uparrow 6)$) ✓
3. Webb's Algorithm & Chéze's Result ($n^{O(1)}$) ✓
4. strengthening Chéze's ✓

Takeaways

- cake-cutting is **NOT** a piece of cake!
- envy-freeness
  - important and interesting fairness criterion
  - historically restricted to small n
  - evidently tractable for "real" inputs
- very active field, trying to bring the complexity down so that we can get to enjoying our cake!