Progress on Envy-Free Cake-Cutting with Few Queries: Towards a Smoothed Analysis

Alejandro Gomez-Leos
UT Austin
alexgomezleos@utexas.edu

Connor Colombe
UT Austin
ccolombe@utexas.edu

December 10, 2024

Abstract

Can we tractably cut cake as to avoid jealousy? In this report* we provide an introductory exposition to cake-cutting in the standard query model. We focus on the open problem of obtaining envy-free allocations in polynomial queries. In particular, we highlight Aziz-Mackenzie (2016) and Chéze (2021). The latter work revisits Webb's super envy-free protocol (1999), of which we provide an extended derivation and analysis. Inspired by the latter works, we present some original progress towards a smoothed query complexity for envy-free cake-cutting.

*This has not been subjected to the usual scrutiny reserved for formal publications.

1 Introduction

In the late 1940's, Hugo Steinhaus had fair division on his mind. Specifically, how should a heterogeneous and divisible resource be fairly partitioned amongst a set of agents? Inspired by a childhood memory, Steinhaus stylized the problem in terms of cutting cake and asked his advisees Banach and Knaster about generalizing the two-person cut-and-choose procedure. In [Ste49], Steinhaus introduced a proportional allocation procedure for n agents. Effectively, the work laid foundations for the last 70 years of intensive research on fair division. In recent decades, the computational questions raised by cake-cutting have sustained a highly active research area in the TCS community—we refer to [BCE+16] for a more comprehensive overview. In this report, we concern ourselves with the following question, which we'll revisit later.

Are envy-free allocations tractably computable?

The problem is commonly formalized as follows. The *cake* is represented by the unit interval $[0, 1]$. There are n agents, collectively associated with measures $\mu_1 \ldots \mu_n$. Each measure is normalized, i.e. $\mu([0, 1]) = 1 \quad \forall i \in [n]$, nonnegative, nonatomic, and dominated by the Lebesgue measure. By this last assumption the associated p.d.f.s exist, so with a slight abuse of notation we denote $\mu_i(x)$ as the density at point $x \in \mathbb{R}$ when clear from context. An *allocation* is a partition $[0, 1] = W_1 \sqcup \cdots \sqcup W_n$ where the i^{th} *piece* belongs to the i^{th} agent. There are standard notions of fairness in cake-cutting, and we review the most necessary for this report. We say an allocation $\bigsqcup_{i=1}^n W_i$ is:

- proportional if $\mu_i(W_i) \geq 1/n \quad \forall i \in [n]$
- equitable if $\mu_i(W_i) = \mu_j(W_j) \quad \forall i, j \in [n]$
- envy-free if $\mu_i(W_i) \geq \mu_j(W_j) \quad \forall i, j \in [n]$

These allocations are also said to satisfy *exact proportionality*, *exact equitability*, and *exact envy-freeness*, respectively. This is in contrast with a standard relaxation of each condition; obtained from the above by fixing an $\epsilon > 0$ and subtracting it from each right-hand side. Doing so yields the conditions for ϵ-proportionality, ϵ-equitability, and ϵ-envy-freeness.

For each notion of fairness, there is an associated algorithmic problem of computing such an allocation. With respect to a query model, the query complexity of a fair cake-cutting problem is the minimum number of queries required from every algorithm to return such an allocation. Often, cake-cutting algorithms operate in the so-called Robertson-Webb (RW) query model [WS07].

Definition 1.1 (Robertson-Webb (RW) Query Model)**.** *A cake-cutting algorithm in the RW query model may only obtain information about the problem instance $\mu_1 \ldots \mu_n$ via oracle access to the following.*

- $\texttt{Eval}(i, x, y)$ *: receive $\mu_i([x, y])$ for any $0 \leq x \leq y \leq 1$ and $i \in [n]$*
- $\texttt{Cut}(i, x, \alpha)$ *: receive "cut-point" y such that $\mu_i([x, y]) = \alpha$ for any $0 \leq x \leq y \leq 1$ and $i \in [n]$*

Remark 1.2. *It is common to describe this model in terms of "asking" the agents for these quantities. However, this brings into question how strategic behavior might affect the resulting allocation. In short, if the agents only "care about" envy-freeness, then the models are equivalent, so we might as well work in the above. We elaborate on this in Sec. 2.2.*

This query model naturally captures the efficiency of so-called *discrete protocols*, which prescribe a turn-based approach in cutting and evaluating pieces of the cake. By contrast, a distinct class of algorithms includes *moving-knife protocols*, which can't be studied in the RW query model.

The problem of envy-free cake-cutting in the RW query model is summarized as follows.

Problem: EF Cake-Cutting
Input: RW query access to measures $\mu_1 \ldots \mu_n$
Output: An envy-free allocation $[0, 1] = \bigsqcup_{i=1}^n W_i$

Note we mandate that *all* of the cake is allocated (the "no-cake" allocation is trivially envy-free). Thus, the initial question is precisely:

Is there an algorithm solving ***EF Cake-Cutting*** *using at most* $\text{poly}(n)$ *queries?*

Interestingly, it was unknown until recently whether there existed a procedure that solved the above in time $O(f(n))$ for some $f : \mathbb{N} \to \mathbb{N}$. Indeed, the celebrated result of Aziz-Mackenzie [AM16] positively affirms this. Prior to their work, the best known procedure was due to Brams-Taylor [BT95], which had the unfortunate drawback that it could take arbitrarily many queries, for any fixed n, by an adversarial choice of input measures.

1.1 Overview

In Sec. 1.2, we briefly survey the state-of-the-art on query complexity bounds for proportional, equitable, and envy-free allocations. An excellent survey of results for other fairness notions is given in [BN22]. Through examples, in Sec. 2 we demonstrate techniques in the classical algorithms for $n = 2, 3, 4$ which are historically significant. In the subsequent sections we focus on ideas presented in [Chè20], which includes a probabilistic analysis of Webb's super envy-free algorithm [Web99]. In Sec. 3 we give a new derivation and analysis of Webb's algorithm. In Sec. 3.3 and Sec. 4, we highlight the main ideas of [Chè20] and discuss a strengthening of their result.

1.2 Related Work

The query complexity of proportional cake-cutting is long-settled to be $\Theta(n \log n)$, with the upper bound established by [EP84] and the matching lower bound thanks to [EP06, WS07]. For equitable allocations, [CP12] proved there is no finite procedure returning a *connected*[1] and exact allocation for general n. Moreover, they showed connected and ϵ-equitable allocations can be found in $O(n \log \frac{n}{\epsilon})$ queries. Extending the aforementioned hardness result, [PW17] established a $\Omega(\log \epsilon^{-1} / \log \log \epsilon^{-1})$ lower bound for algorithms obtaining ϵ-equitable allocations (note the pieces may be disconnected).

For envy-freeness much, less is known. For connected and ϵ-envy-free allocations, [BN22] established an $O(n/\epsilon)$ upper bound and $\Omega(\log \epsilon^{-1})$ lower bound. For exact allocations, [AM16] established an $O(n \uparrow\uparrow 6)$[2] upper bound, whereas [Pro09] gave a $\Omega(n^2)$ lower bound. Recently, [Chè20] revisited the algorithm of [Web99] to establish that the "perturbed query complexity" is $n^{O(1)}$ with high probability, and that the average query complexity satisfies a similar bound. We elaborate more on these last results in Sec. 3.3.

[1] i.e. each piece is a connected subinterval.
[2] Knuth's up-arrow notation, n recursively exponentiated 6 times.

2 EF Cake-Cutting with Small n

Even for small constant n, bounds on the query complexity have been historically difficult to obtain. As a warm-up, we'll begin with the cut-and-choose procedure, known since antiquity.

2.1 $n = 2$: cut-and-choose procedure

Plainly stated, a first agent is arbitrarily chosen. Agent 1 cuts the cake in two pieces of equal μ_1 measure, then agent 2 chooses the piece with higher μ_2 measure. The first step costs 1 `Cut` query (calling $\texttt{Cut}(1, 0, \frac{1}{2})$), and the second step costs 1 `Eval` query (say calling $\texttt{Eval}(2, 0, \frac{1}{2})$)[3] to determine agent 2's preference of the two pieces. Moreover, 2 queries are necessary—the algorithm could err submitting only 1 query since it cannot obtain information about both measures. Envy-freeness follows from the fact that each piece has equal μ_1 measure (so agent 1 doesn't envy agent 2), and agent 2's piece has larger μ_2 measure than agent 1's piece (so agent 2 doesn't envy agent 1).

While simple, the cut-and-choose procedure hints at design principles present in more advanced algorithms. Two broad ideas include the following:

- An agent who can choose their piece before another agent won't be envious of the other agent.
- An agent who partitions the cake into equal pieces w.r.t. their own measure won't be envious given the opportunity to secure one of these.

Of course, the above arguments use the assumption that the agents truthfully report when queried (implicit in Def. 1.1). What if we hadn't made this assumption? Before demonstrating extensions of the above principles, we make a detour to clarify the role of strategic behavior in cake-cutting.

2.2 A Side Remark on Strategic Behavior

We elaborate on the point raised in Remark 1.2, which is particularly straightforward to illustrate using the cut-and-choose procedure. Rather than receiving the responses in Def. 1.1, instead suppose that the algorithm asks the corresponding agent for these quantities. If the utility of each agent is measure of their piece, i.e. $u_i := \mu_i(W_i)$, then the cut-and-choose protocol may be manipulated.

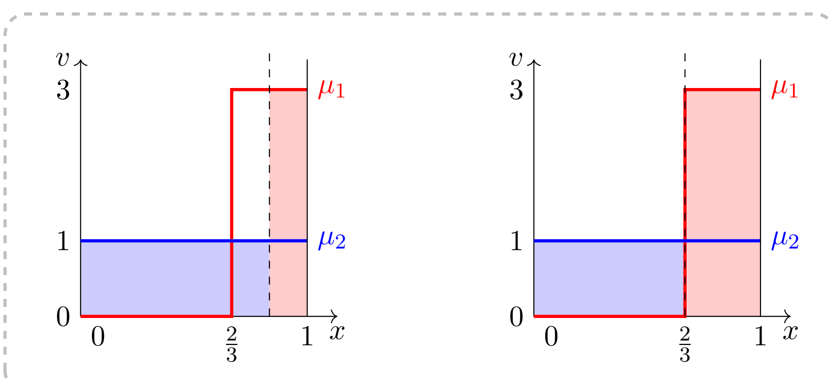

Figure 1: Left: allocations from truthful reporting. Right: allocations from strategic reporting.

[3] Notice a second `Eval` query isn't needed as the cake has unit measure.

The following example is illustrated in Fig. 1. Suppose $\mu_1(x) = \mathbb{I}[x \in [2/3, 1]]$ and $\mu_2(x) = 1$ for all $x \in [0, 1]$. If agents report truthfully, then agent 1 would cut the cake at $x = 5/6$ and receive the allocation $[5/6, 1]$ with utility $1/6$. If agent 1 knows μ_2, however, they may obtain a higher utility by misrepresenting μ_1 when asked for the initial `Cut` query. Specifically, reporting $x = 2/3$ allows them to obtain the allocation $[2/3, 1]$ with utility $1/3$. Note, though, that in either case the resulting allocation is envy-free, which seems to suggest that an algorithm returning envy-free allocations does so regardless of whether the agents are strategic or not. To avoid a nuanced discussion beyond the scope of this report, we simply assume the utilities are given by $u_i := \sum_{j \neq i} \mathbb{I}[i \text{ envies } j]$. Then, truthtelling is a *dominant strategy*, which substantiates our assertion in Remark 1.2.

2.3 $n = 3$: Selfridge-Conway Procedure

An envy-free protocol for $n = 3$ was independently discovered by John Selfridge and John Conway in the 1960's. Building on ideas in the cut-and-choose procedure, their method contributed a design principle to computing envy-free allocations: maintaining a partial allocation, we may allow agents to *trim* already-cut pieces in effort to "balance discrepancy" across agents. To provide a more concrete understanding of these ideas, we detail the full protocol. Visualized below, we indicate the allocations, cuts, and trims of the agents by their colors agent 1, agent 2, and agent 3.

To start, agent 1 cuts the cake into three pieces W_1, W_2, W_3 of equal μ_1 measure (2 `Cut` queries). Then, agents 2 and 3 take turns indicating their favorite of the three pieces (4 `Eval` queries). If they select different pieces, we are done—2 and 3 get their favorite pieces and agent 1 is satisfied with the remaining piece (see Fig. 2).

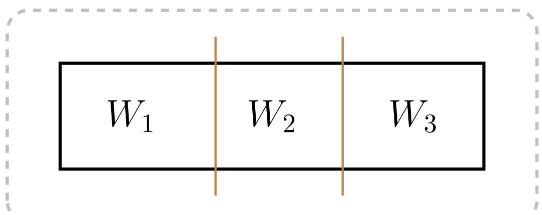

Figure 2: agent 1 cuts the cake such that $\mu_1(W_1) = \mu_1(W_2) = \mu_1(W_3)$

Otherwise, agent 2 and 3 prefer the same piece, WLOG say W_1. Then, agent 2 cuts W_1 into pieces W_1' and W_1'' such that $\mu_2(W_1') = \max(\mu_2(W_2), \mu_2(W_3))$ (1 `Cut` query, we already know the quantities on the right-hand side), i.e. they cut W_1 such that one of the resulting pieces has the same value as their second favorite of the original pieces (see Fig. 3). We set aside W_1'', termed the *residue*. Now the agents choose their favorite of the remaining pieces W_1', W_2, W_3 in the order agent 3 $\to$ agent 2 $\to$ agent 1 (2 total `Eval` queries to learn agent 2 and agent 3's measure of the residue pieces). The only subtlety is that if agent 3 picks a "whole" piece, agent 2 must pick the 'trimmed" piece (which they already claimed was equal in value to their favorite of the remaining pieces, so they will happily take it). This yields a partial allocation that is envy-free. To see why, note that agent 3 chose first, so they mustn't be envious of the others (using the first principle from the previous section). Moreover, since agent 2 determined the trimming, they are guaranteed a piece μ_2 measure equal to their preferred whole piece. Finally, agent 1 is guaranteed a whole piece of the original three they cut (using the second principle from the previous section) (see Fig. 4).

Now it remains to allocate the leftover residue. At first pass, this may resemble the problem we faced initially, so it's natural to doubt that the algorithm would terminate. This is not the case—the

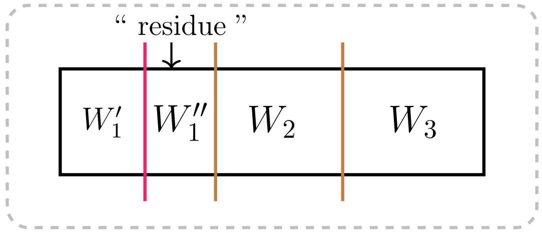

Figure 3: agent 2 *trims* their favorite piece such that $\mu_2(W_1') = \max(\mu_2(W_2), \mu_2(W_3))$.

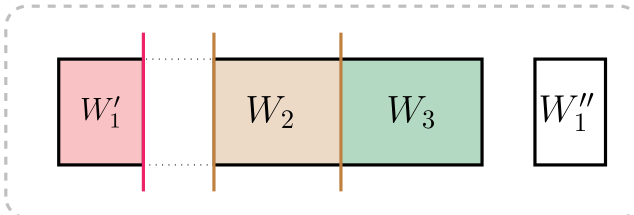

Figure 4: residue placed aside and agents choose their pieces in the order agent 3, agent 2, agent 1. This is an envy-free partial allocation.

remaining problem is subtly distinct from our starting problem! Suppose that in the previous step, agent $i \in \{2, 3\}$ received the trimmed piece and agent $j \in \{2, 3\}$ received the whole piece. Consider agent 1, who holds a whole piece in the partial allocation. Even if the rest of the unallocated cake were to be allocated to agent i, agent 1 would not be envious of i. In this situation, it is said that agent 1 *dominates* agent i. This notion is important in envy-free cake-cutting and was pivotal to the development of Aziz-Mackenzie's breakthrough generalization.

We return to the question of allocating the residue. Agent j cuts the residue into three pieces they view as equal (2 `Cut` queries), then the agents choose their pieces in the order agent $i \to$ agent $1 \to$ agent j (4 `Eval` queries to learn agent 2 and agent 3's measure of the residue partition pieces) (see Fig. 5).

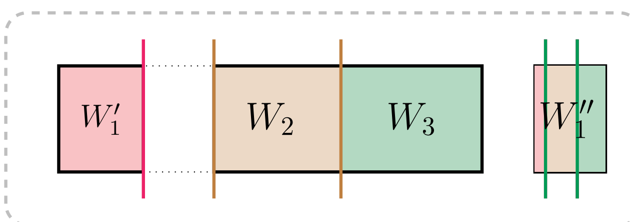

Figure 5: The agent who received the whole piece (agent 3 in this case), cuts W_1'' into three equal pieces and agents choose in the order agent 2, agent 1, agent 3. The final total allocation is envy-free.

The resulting allocation is envy-free. Indeed, see that agent i chooses first, so will not be envious of the others pieces. Agent 1 dominates agent i and chooses before agent j, so agent 1 will not envy either. Finally, agent j is indifferent to whichever piece they receive since they cut the pieces. Moreover, this was obtained in 15 queries, although it is not clear whether this is tight. Evidently, the principle of domination serves as a nice complement to the aforementioned principles.

2.4 $n = 4$: Brams-Taylor and Aziz-Mackenzie

In [BT95], Brams and Taylor invented the first envy-free discrete protocol handling $n = 4$ and more agents. Their procedure relied on the trimming technique developed by Selfridge and Conway—searching for partial envy-free allocations whilst recursing on the residue. However, as mentioned in Sec. 1, their algorithm could not be bounded based on the number of agents alone. Open for more than two decades, this problem was resolved by [AM15] for $n = 4$ and generalized in [AM16] to $n \geq 4$ agents. Extending the aforementioned techniques, their additional innovation was allowing agents to permute their pieces as to ensure envy-freeness, but also increase the number of dominating agents. In turn, this enables one to remove a dominating agent and call a Selfridge-Conway-like procedure on the remaining cake and agents. Currently, the best bound on this procedure's queries is $O(n \uparrow\uparrow 6)$. For the remainder of this section, we present an example that highlights their procedure on the $n = 3$ case, and omit the query tally for clarity.

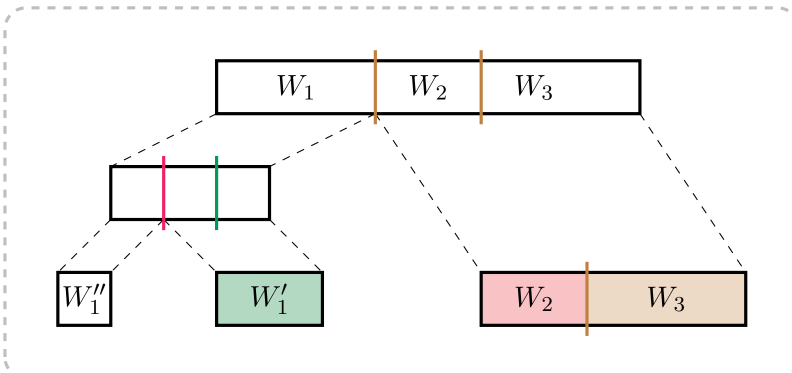

Figure 6: The initial steps of the procedure. agent 1 cuts the cake into three equal pieces w.r.t. their measure. If agent 2 and agent 3 prefer the same piece W_1, they each mark it such that the right-hand portion of their mark is equal in measure to their respective second favorite of W_2 and W_3. W_1 is cut into W_1' and W_1'' at the mark farthest left, in this case agent 2's mark. agent 3 receives W_1', agent 2 picks their favorite of the remaining pieces, and agent 1 recieves the remaining piece.

The procedure begins in a manner similar to the Selfridge-Conway procedure (see Fig. 6). Like Selfridge-Conway, the resulting partial allocation is envy-free, and leaves a residue W_1''. Notably, agent 1 dominates agent 3 since agent 1 receieved a whole piece. However, agent 1 doesn't dominate 2 since agent 2 could potentially get the residue. If we could augment the partial allocation such that agent 1 dominates agents 2 as well, then agent 1 could be removed from the procedure entirely! Following this, we may recurse on the smaller subproblem to deal with the residue.

For instance, consider a case in which we have achieved two successive partial allocations (Fig. 7). Suppose that in each case, agent 3 receives the trimmed piece (i.e. that they marked a line to the right of agent 2's line). Each time, they recieve a "bonus," i.e. the piece between the red and green lines in Fig. 6, say with measure b_1 and b_2. The next step advocates to swap the pieces of agents 2 and 3 in the allocation where b_i is the smallest, say the first partial allocation. This breaks envy-freeness in the first allocation but is offset by the fact that agent 3 received a larger bonus in the second partial allocation—the overall partial allocation remains envy-free. Notably, both agents 2 and 3 have received trims from pieces originally cut by agent 1 and are therefore dominated by

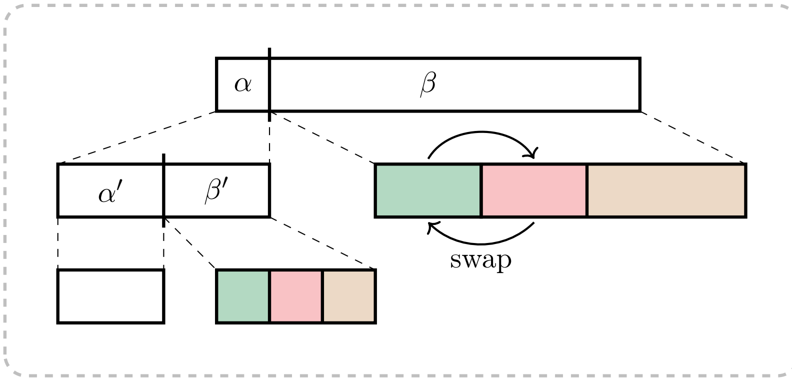

Figure 7: Assume agent 3 receives the trimmed piece both times, and recieves "bonus" b_1 and b_2 in each allocation, respectively. Suppose that agent 3 more additional measure in the second(lower) partial allocation. Then by swapping the pieces of agent 2 and agent 3 in the first(upper) partial allocation, we retain an overall envy-free partial allocation, but now agent 1 dominates the remaining agents and can be removed from future allocations.

agent 1. We can now remove agent 1 and run the same protocol on the remaining residue with agents 2 and 3, which amounts to cut-and-choose for this instance. For a detailed depiction of the procedure on $n = 4$ agents, see [AM20].

In the next section, we introduce Webb's algorithm—an important pre-requisite for understanding Cheze's result and our follow-up work.

3 Webb's Algorithm

3.1 Notation

Let a, b be real scalars. For $\epsilon > 0$, we say $a \approx_\epsilon b$ if $|a - b| \leq \epsilon$. We say $a \lesssim b$ if $a \leq Cb$ for some absolute constant $C > 0$ and $a \asymp b$ if $a \lesssim b$ and $a \gtrsim b$.

3.2 Derivation

Remarkably, an envy-free allocation can be found in $n^{O(1)}$ additional queries given a "good enough" starting partition. Somewhat orthogonal to the previously mentioned techniques, this observation is captured in Webb's algorithm [Web99], which we'll derive in a way we imagine the author might've. Suppose we begin with an initial "guess" partition $[0,1] = \bigsqcup_{k=1}^n W_k$. No matter which partition we begin with, we may attempt to subdivide these pieces into a finer partition $[0,1] = \bigsqcup_{k=1}^n (\bigsqcup_{i=1}^n W_k^i)$, allocating the piece $\bigcup_{k=1} W_k^i$ to the i^{th} agent (see Fig 8). Of course, we want this final allocation to be envy-free. It suffices for these pieces to satisfy the following system of inequalities.

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

Such an allocation satisfies the more stringent condition of being *super* envy-free. Why should we aim for a stronger property? Conveniently, we already have a simple and complete existential characterization, due to Barbanel.

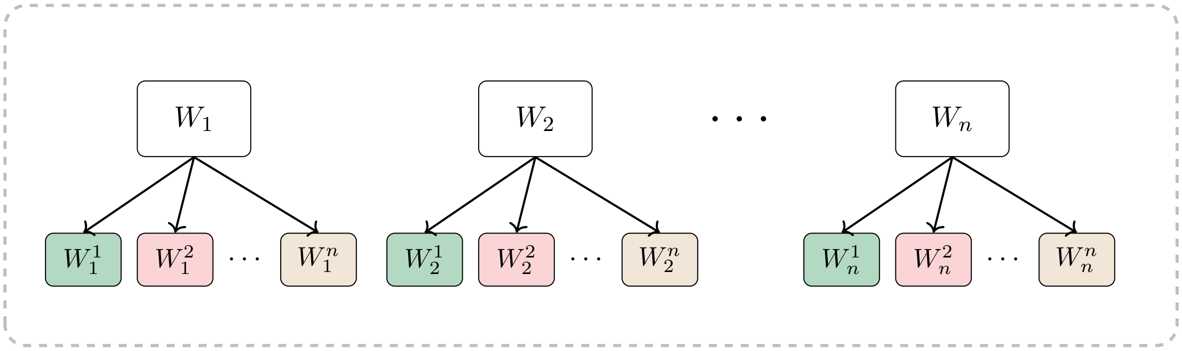

Figure 8: Subdivision of initial "guess" partition.

Theorem 3.1 ([Bar96])**.** *A super envy-free subdivision of $W \subseteq [0,1]$ exists if and only if $\mu_1 \dots \mu_n$ are linearly independent measures, i.e. $\sum c_i \mu_i = 0$ only for the trivial $\vec{c}$.*

Hence, we'll assume throughout this section that the input measures are linearly independent. Our overall goal is to re-cast the problem of obtaining a finer partition satisfying (1) as solving a particular linear system. To this end, suppose we pick some $\delta > 0$ (to be decided later) such that

$$\mu_i\left(\sum_{k=1}^{n} W_k^j\right) = \begin{cases} 1/n + \delta & j = i \\ 1/n - \delta/(n-1) & j \neq i \end{cases} \tag{2}$$

We momentarily ignore the $1/(n-1)$ factor in the case $j \neq i$. Recalling that W_k denotes the k^{th} piece of the original partition, we can re-write the left-hand side.

$$\mu_i\left(\sum_{k=1}^{n} W_k^j\right) = \sum_{k=1}^{n} \mu_i(W_k^j) = \sum_{k=1}^{n} \mu_i(W_k) \cdot \frac{\mu_i(W_k^j)}{\mu_i(W_k)} := \sum_{k=1}^{n} \mu_i(W_k) \cdot R_{k,j,i} \tag{3}$$

Here, $R_{k,j,i}$ has the natural interpretation as the "fractional worth of the j^{th} subdivision of the k^{th} piece as seen by agent i". Notably, $\sum_{j=1}^{n} R_{k,j,i} = 1 \quad \forall k, i$. But (2) does not quite reflect a linear system, so we'll assume that all agents "see the same fractional worth" of each subdivision. More precisely, we assume that for all $i \neq i'$ we have $R_{k,j,i} = R_{k,j,i'}$. In this case, we can instead deal with the following refinement of (2).

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

Now we genuinely have a linear system. Explicitly,

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

We'll later see there is a natural choice for δ based on $\mathbf{M}$. With this in mind, once $\mathbf{M}$ is determined and if it is nonsingular, then one can recover $\mathbf{R}$ and attempt to subdivide as in Fig. 8

such that the entries of $\mathbf{R}$ have the endowed interpretation. If this subdivision is successful, then by the above we satisfy (2) and are done. In summary, Webb's algorithm is "trying" to do the following.

1. Obtain $\mathbf{M}$, stop if $\mathbf{M}$ is singular
2. Pick δ depending on $\mathbf{M}$, compute $\mathbf{R} = \mathbf{M}^{-1}\mathbf{N}_\delta$
3. Subdivide each of $W_1 \dots W_n$ such that for each k, j we have $\mu_i(W_k^j)/\mu_i(W_k) = \mathbf{R}_{k,j} \quad \forall i \in [n]$
4. Allocate $\bigcup_{k=1} W_k^i$ to the i^{th} agent, i.e. as in Fig. 8

It remains to specify the choice of δ and the details of step 3. Revisiting the choice of $\delta > 0$, recall that $\mathbf{R}$ must be row stochastic and nonnegative. Indeed, the first desiderata is satisfied since

$$\mathbf{R}\mathbf{1} = \mathbf{M}^{-1}(\mathbf{N}_\delta \mathbf{1}) = \mathbf{M}^{-1}(\mathbf{1}) = \mathbf{M}^{-1}(\mathbf{M}\mathbf{1}) = \mathbf{1}.$$

Notably, this elucidates the $1/(n-1)$ renormalization in the off-diagonals of $\mathbf{N}_\delta$. For the nonnegativity, first note the above shows $\mathbf{M}^{-1}$ is row stochastic. Denoting t as the minimum entry of $\mathbf{M}^{-1}$,

$$\begin{aligned}
\mathbf{R}_{kj} &= \sum_{i=1}^{n} (\mathbf{M}^{-1})_{ki} (\mathbf{N}_\delta)_{ij} \\
&= \sum_{i=1}^{n} (\mathbf{M}^{-1})_{ki} \left(\left(\frac{1}{n} + \delta\right) \mathbb{I}[i = j] + \left(\frac{1}{n} - \frac{\delta}{n-1}\right) \mathbb{I}[i \neq j] \right) \\
&= (\mathbf{M}^{-1})_{kj} \left(\frac{1}{n} + \delta\right) + \left(1 - (\mathbf{M}^{-1})_{kj}\right) \left(\frac{1}{n} - \frac{\delta}{n-1}\right) \\
&= (\mathbf{M}^{-1})_{kj} \left(\frac{1}{n} + \delta - \frac{1}{n} + \frac{\delta}{n-1}\right) + \frac{1}{n} - \frac{\delta}{n-1} \\
&= (\mathbf{M}^{-1})_{kj} \left(\delta + \frac{\delta}{n-1}\right) + \frac{1}{n} - \frac{\delta}{n-1} \\
&\geq t \left(\delta + \frac{\delta}{n-1}\right) + \frac{1}{n} - \frac{\delta}{n-1}.
\end{aligned}$$

Thus, to have $\mathbf{R}_{kj} \geq 0$ it suffices to impose

$$\begin{aligned}
t \left(\delta + \frac{\delta}{n-1}\right) + \frac{1}{n} - \frac{\delta}{n-1} \geq 0 &\iff t\,(\delta n) + \frac{n-1}{n} - \delta \geq 0 \\
&\iff \frac{n-1}{n} \geq \delta(1 - tn) \\
&\iff \delta \leq \frac{n-1}{n(1-tn)}.
\end{aligned}$$

Hence, any choice of $\delta \in (0, \frac{n-1}{n(1-tn)}]$ will satisfy our purposes.

We now handle step 3 using a blackbox reduction to obtaining so-called ϵ-exact partitions.

Definition 3.2 ("ϵ-exact partitions"). *Let $W \subseteq [0,1]$ and fix some measures $\mu_1 \dots \mu_n$. Letting $\epsilon > 0$ and $\vec{\alpha} = (\alpha_1, \dots, \alpha_n)$ be a point on the simplex in $\mathbb{R}^n$, we say a partition $W = W^1 \sqcup \dots \sqcup W^n$ is ϵ-exact for fractions $\vec{\alpha}$ if $\forall i, j \quad \frac{\mu_i(W^j)}{\mu_i(W)} \approx_\epsilon \alpha_j$.*

Fortunately, we can efficiently obtain these partitions efficiently, as the next result shows.

Theorem 3.3 ([RW97], [RW98]). *There's an algorithm* $\textit{\texttt{NearExact}}(W, \vec{\alpha}, \epsilon)$ *which outputs an* ϵ*-exact partition* $W = W_1 \sqcup \cdots \sqcup W_n$ *in* $\lesssim n^{2.5}/\epsilon$ *RW queries on any instance* $\mu_1 \ldots \mu_n$.

The procedure asks the first agent to provide k partition points so that the μ_1 measure of each slice is at most $1/k$ (inducing $\lesssim k$ `Cut` queries), then has the second agent "add" $\leq k-1$ more partition points so that the μ_2 measure of each slice is at most $1/k$ (inducing $\lesssim k$ `Cut` queries), and so on. For k sufficiently large, one can show that the final subintervals can be glued together to form a coarser partition which satisfies the desired properties[4].

Thus, by taking small enough ϵ, we can simply call $\texttt{NearExact}(W_k, \vec{\alpha}, \epsilon)$ on each W_k such that the obtained allocations still satisfy (1). For instance, as the off-diagonal-diagonal gap in (2) is $\geq \delta$, then we can take[5] $\epsilon \leq \frac{\delta}{3}$ so that we ultimately satisfy (1). Webb's algorithm is summarized in algorithm 1.

Algorithm 1 Webb's Super Envy-Free Protocol

Input: RW query access to $\mu_1 \ldots \mu_n$, starting partition $[0,1] = \sqcup_{i=1}^n W_i$
Output: A (super) envy-free allocation $W_1^* \ldots W_n^*$

1: $\mathbf{M}_{ij} = \mu_i(W_j)$, stop if not invertible

2: $t :=$ min entry of $\mathbf{M}^{-1}$, $\delta = \frac{n-1}{n(1-tn)}$

3: Compute $\mathbf{M}^{-1}\mathbf{N}_\delta = \mathbf{R} := \begin{pmatrix} \vec{R}_1 \\ \vdots \\ \vec{R}_n \end{pmatrix}$

4: For each $i \in [n]$, call $\texttt{NearExact}(W_k, \vec{R}_k, \frac{\delta}{3})$ to get $W_k^1 \ldots W_k^n$

5: **Return** $W_i^* = W_1^i \sqcup W_2^i \sqcup \cdots \sqcup W_n^i \quad \forall i \in [n]$

The runtime analysis is particularly simple. Step 1 takes n^2 `Eval` queries. Step 4 takes $\lesssim \frac{n^{2.5}}{\delta} \cdot n \lesssim n^{3.5}(1-tn)$ `Cut` queries by Theorem 3.3. Notably $t \leq 0$, for otherwise the dot product of the i^{th} row of $\mathbf{M}^{-1}$ and the j^{th} column of $\mathbf{M}$ is positive, which contradicts the off-diagonals of $\mathbf{M}^{-1}\mathbf{M} = \mathbf{I}$. Moreover, $|t|^2 \leq \|\mathbf{M}^{-1}\|_F^2 \leq \|\mathbf{M}^{-1}\|_2^2$. Observing $\|\mathbf{M}\|_2 \geq 1$, the above implies step 4 takes $\lesssim n^{4.5}\|\mathbf{M}^{-1}\|_2 \leq n^{4.5}\|\mathbf{M}\|_2\|\mathbf{M}^{-1}\|_2 = n^{4.5}\kappa(\mathbf{M})$ queries in the worst case.

Combined with our derivation, we've shown the following result.

Theorem 3.4 ([Web99]). *If (i)* $\mu_1 \ldots \mu_n$ *are linearly independent and (ii) the matrix* $\mathbf{M}_{ij} = \mu_i(W_j)$ *is nonsingular for a starting partition* $[0,1] = \bigsqcup_{i=1}^n W_i$*, then algorithm 1 returns a (super) envy-free allocation in* $\lesssim n^{4.5}\kappa(\mathbf{M})$ *RW queries.*

We've discussed that condition (i) is necessary for a super envy-free allocation. On the other hand, (ii) may not hold for a arbitrary starting partition. For instance, consider the $n = 2$ instance in Fig. 9 where $\theta := 1/2$. Taking $W_1 = [0, 1/2 - \epsilon]$ and $W_2 = [1/2 - \epsilon, 1]$ yields a singular $\mathbf{M} = [1/2 - \epsilon, 1/2 + \epsilon; 1/2 - \epsilon, 1/2 + \epsilon]$ due to the "cancellation across the peaks", even though the measures are linearly independent.

[4]For more details, we refer the reader to Theorem 10.2 [RW98].

[5]In the original work $\epsilon := \delta/n^2$, which seems unnecessarily small.

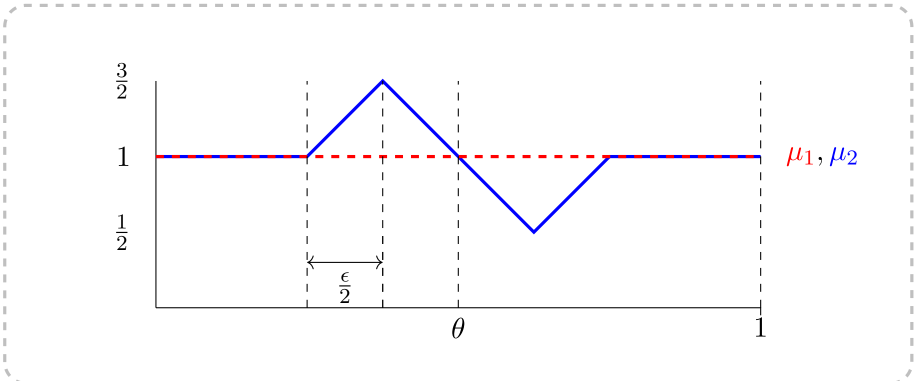

Figure 9: (i) doesn't imply (ii) in Theorem 3.4.

One may hope that choosing a random partition may circumvent these hard instances, i.e. sampling $n-1$ partition points from Unif$[0,1]$. However, the same example above can be used to construct a hard instance family. Consider an adversary who draws $\theta \sim$ Unif$[0,1]$ and takes ϵ to approach 0. Indeed, for $\mathbf{M}$ to be nonsingular at least one point must land in the region $[\theta-\epsilon, \theta+\epsilon]$. However, the Lebesgue density of this region is $\lesssim \epsilon$ so the probability that the partition witnesses the linear independence[6] of the measures is vanishingly small, i.e. $\mathbf{M}$ will be singular with overwhelming probability.

Although [Web99] didn't acknowledge these hard instances, we imagine the author had them in mind as they advocated for a brute-force approach to find a nonsingular $\mathbf{M}$. Specifically, they advocate for a sequential procedure: at iteration $i = 1, 2, 3, \ldots$, partition $[0,1]$ into 2^i intervals and test the singularity of $\mathbf{M}$ under all possible partitions to n players comprised of these intervals. Imagining an instance of the above hard example where $\epsilon \lesssim 2^{-n}$, it's clear that $\gtrsim 2^n$ intervals will be needed to obtain a nonsingular $\mathbf{M}$. Hence, this iteration of Webb's proposed solution would need to try at least $\gtrsim n^{2^n}$ partitions, costing at least as many queries.

3.3 Chèze's Result

Disembarking from recent algorithmic ideas for envy-free cake-cutting, in [Chè20] the author revisits the algorithm of [Web99]. In the spirit of smoothed analysis, the nice insight of this work is that the hard instance families for Webb's algorithm are seemingly "brittle". Their main result is the following.

Theorem 3.5 ([Chè20])**.** *Consider an instance of algorithm 1 inducing a matrix $\mathbf{M}$. Fix a $\sigma > 0$. If $\mu_i(x) > \sigma$ everywhere and $\mathbf{E}$ is a random matrix with iid entries in $(-\sigma, \sigma)$, then replacing the matrix $\mathbf{M}$ in algorithm 1, step 1 with the matrix $\tilde{\mathbf{M}}_{ij} = (\mathbf{M}_{ij} + \mathbf{E}_{ij})/(\sum_{j=1}^n (\mathbf{M}_{ij} + \mathbf{E}_{ij}))$ results in a procedure using more than $C_\sigma n^{O(1)}$ queries with probability $o(n^{-1})$. Here, C_σ is a quantity depending on σ but not on n.*

Although not quite a smoothed analysis result, the above morally suggests that Webb's is efficient for "realistic" inputs. By a runtime characterization similar to the bound in Sec. 3, the proof follows by an anti-concentration of the minimum singular value of $\tilde{\mathbf{M}}$, thanks to a result from Tao-Vu [VT07]. Unfortunately, the analysis in [VT07] does not track the dependence on σ. As a result,

[6] One should keep in mind that linear independence isn't necessary for efficient envy-free allocation. As an extreme example, consider the number of queries sufficing for $\mu_1 = \cdots = \mu_n$.

Theorem 3.5 does not explain the dependence of σ on the query complexity (e.g. it doesn't rule out exponential dependence on σ^{-1}). Moreover, it is not known whether the allocation returned by this algorithm is envy-free with respect to the original measures — it seems plausible that for a sufficiently small perturbation this could be the case. Finally, it is natural to ask whether the restriction to positive densities can be relaxed.

In the next section, we present work resolving some of these concerns for certain "natural" perturbation models. Namely, that the dependence on σ^{-1} is polynomial and the restriction to positive densities can be removed.

4 Towards a Smoothed Analysis

Inspired by the ideas in [Chè20], in this section we make progress towards obtaining a smoothed query complexity for **EF Cake-Cutting**. To define our space of perturbations, recall the input to Webb's algorithm is fully specified by a partition, say $\mathcal{W}$, and input measures $\mu := (\mu_1, \dots \mu_n)$.

Definition 4.1 (Perturbation Model for Webb's)**.** *Fix $\sigma > 0$, and let φ be a p.d.f. with nonnegative support. For Webb's instance $\mathrm{I}_{\mu,\mathcal{W}}$, let $\mathrm{I}^{\sigma}_{\mu',\mathcal{W}}$ be the random instance given as follows. Let $\mathbf{P} \in \mathbb{R}^{n\times n}$ be a random matrix with iid entries "drawn from" φ, and set for each $i \in [n]$*

$$\mu_i'(x) = \mu_i'([0,1])^{-1}(\mu_i(x) + \sigma \mathbf{P}_{ij}) \quad \forall x \in \mathcal{W}_j$$

(see Fig. 10 for a visualization). Notably,

- *under the original input $\mathrm{I}_{\mu,\mathcal{W}}$, algorithm 1 works with $\mathbf{M}_{ij} = \mu_i(\mathcal{W}_j)$*
- *under the random perturbed input $\mathrm{I}^{\sigma}_{\mu',\mathcal{W}}$, algorithm 1 works with $\tilde{\mathbf{M}} := \mathbf{D}(\mathbf{M} + \sigma\mathbf{P})$ where $\mathbf{D} = diag\left[(1+\Sigma_{j=1}^n \sigma \mathbf{P}_{1j})^{-1}, (1+\Sigma_{j=1}^n \sigma \mathbf{P}_{2j})^{-1}, \dots, (1+\Sigma_{j=1}^n \sigma \mathbf{P}_{nj})^{-1})\right]$*

Two natural perturbation models are given by taking φ to be either Unif$[0,1]$, or the half-normal distribution $|\mathcal{N}(0,1)|$. We conjecture that for these models the smoothed query complexity is $\mathrm{poly}(n, \sigma^{-1})$.

Conjecture 1. Let $\sigma > 0$, and suppose φ is either Unif$[0,1]$ or $|\mathcal{N}(0,1)|$ in the above perturbation model. Denoting $Q(I)$ as the number of queries used by Webb's algorithm on input I, then

$$\max_{\mathrm{I}_{\mu,\mathcal{W}}} \mathbb{E}_{\mathrm{I}\sim \mathrm{I}^{\sigma}_{\mu',\mathcal{W}}}[Q(\mathrm{I})] \lesssim \frac{\mathrm{poly}(n)}{\sigma^2}.$$

Recalling the runtime bound from the last section, such a conjecture is morally equivalent to the following.

Conjecture 2. Let $\mathbf{M}$ be a nonnegative and row stochastic matrix. Let $\mathbf{P} \in \mathbb{R}^{n\times n}$ be a random matrix with iid entries "drawn from" φ, where φ is either Unif$[0,1]$ or $|\mathcal{N}(0,1)|$. Let $\sigma > 0$, and $\mathbf{D} = \mathrm{diag}\left[(1+\Sigma_{j=1}^n \sigma \mathbf{P}_{1j})^{-1}, (1+\Sigma_{j=1}^n \sigma \mathbf{P}_{2j})^{-1}, \dots, (1+\Sigma_{j=1}^n \sigma \mathbf{P}_{nj})^{-1})\right]$. Then,

$$\mathbb{E}_{\mathbf{P}_{ij}\sim\varphi}[\kappa(\mathbf{D}(\mathbf{M} + \sigma\mathbf{P}))] \lesssim \frac{\mathrm{poly}(n)}{\sigma^2}.$$

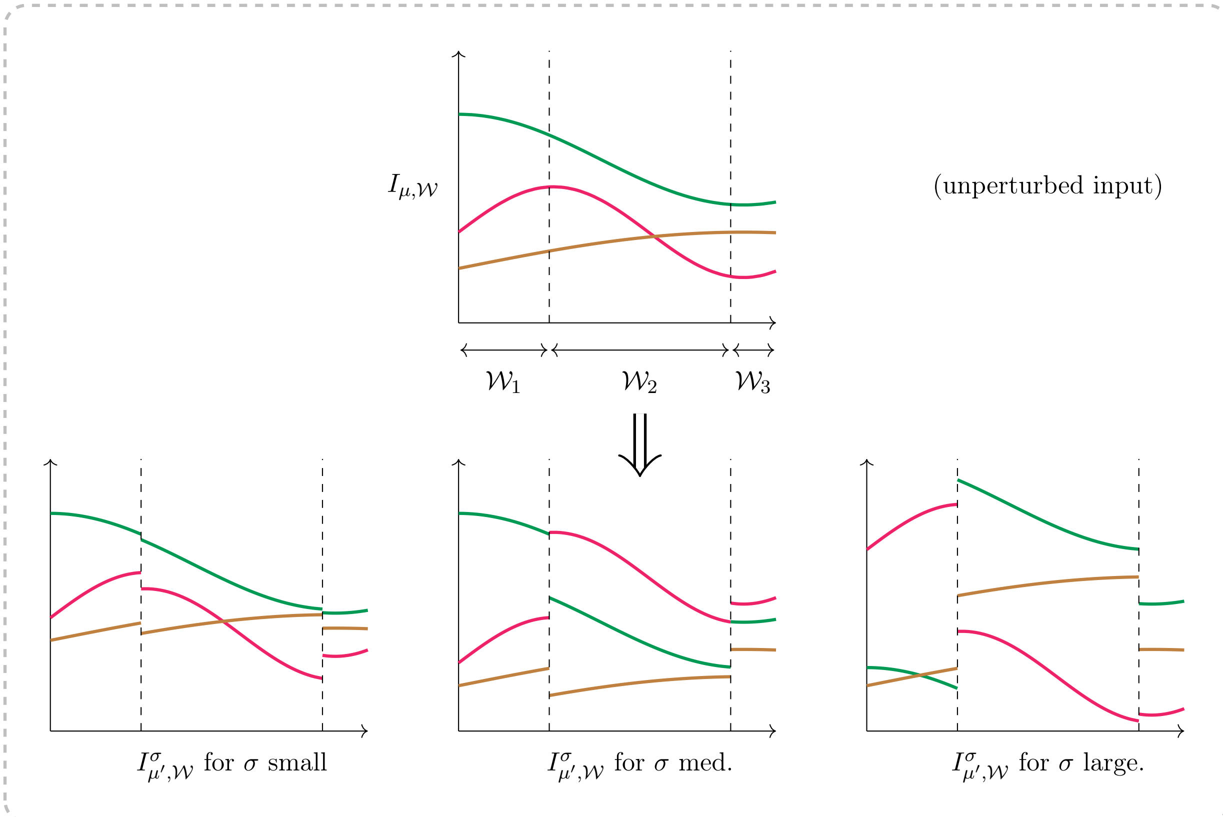

Figure 10: Examples of realized instances under Def. 4.1.

How do we hope to prove the above? A natural approach is to control the integral-over-tails representation of the expectation via bounds on the extreme singular values. By Cauchy-Schwarz and standard arguments, one can show $\|\mathbf{D}(\mathbf{M}+\sigma\mathbf{P})\|_2 \lesssim (1+\sigma)\sqrt{n}$ with exponentially small failure probability, so the key challenge lies in obtaining a anti-concentration for $\sigma_n(\mathbf{D}(\mathbf{M}+\sigma\mathbf{P}))$. There is a rich literature on anti-concentrations for randomly perturbed matrices (e.g. [RV09, VT07, SST06]), with an excellent exposition of the main strategies in [Tao12]. A key technique uses the characterization of the smallest singular value as the largest distance from any fixed row of $\mathbf{M}+\sigma\mathbf{P}$ to the hyperplane spanned by the other rows. Concretely, this amounts to obtaining $\mathbb{P}[|\langle \mathbf{t}, \mathbf{v}\rangle| \leq \epsilon] \lesssim \epsilon/\sigma^2$ where $\mathbf{v}$ is a row of $\mathbf{M}+\sigma\mathbf{P}$ and $\mathbf{t}$ is the normal unit vector to the induced hyperplane. When $\mathbf{P}$ has Gaussian entries, then $\langle \mathbf{t}, \mathbf{v}\rangle \sim \mathcal{N}(0,\sigma^2)$ so one can use a "small rectangle approximation" of the p.d.f. to achieve this. Closer to our setting, the p.d.f. of a linear combination Unif$[0,1]$ (or $|\mathcal{N}(0,1)|$) does not have a easy-to-bound closed form (e.g. [Fle71, RP49]). However, it seems plausible that such a condition could be proved, for instance, using ideas from anti-concentrations of polynomials [MNV16].

On a side note, it is a standard technique to approximate the p.d.f. of $\langle \mathbf{t}, \mathbf{v}\rangle$ by a Gaussian under the correct normalization via Berry-Esseen type estimates. Recall that such estimates give a quantitative form the of the CLT.

$$\mathbb{P}\left[x \leq \frac{\langle \mathbf{t}, \mathbf{v}-\mathbb{E}[\mathbf{v}]\rangle}{\sqrt{\mathbf{Var}(\langle \mathbf{t}, \mathbf{v}\rangle}}) \leq y\right] \approx_\lambda \mathbb{P}_{G\sim\mathcal{N}(0,1)}\left[x \leq G \leq y\right], \quad \lambda \lesssim \frac{\sum_{i=1}^n \mathbf{t}_i^3 \mathbb{E}[|\mathbf{v}_i - \mathbb{E}[\mathbf{v}_i]|^3]}{\mathbf{Var}(\langle \mathbf{t}, \mathbf{v}\rangle)^{3/2}}$$

However, since $\mathbf{t}$ is a random vector with a complicated distribution, bounding $\mathbf{t}$ by a worst-case choice diminishes hope for controlling such a term, which would be aided by a lower bound on the entries of $\mathbf{t}$ (it could be very sparse).

Nonetheless, we obtain a preliminary result towards Conjecture 1.

Proposition 4.2. *Let $\sigma \in (0, 1/n)$, and suppose $\varphi \equiv |\mathcal{N}(0,1)|$ in the perturbation model. Denoting $Q(I)$ as the number of queries used by Webb's algorithm on input I, then*

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

The guarantee is much weaker than conjectured, but substantially smaller than the iterated tower estimate in [AM16]. Moreover, our result suggests the dependence on the perturbation is indeed polynomial. Note that we average over a "two-draw" neighborhood. We should emphasize that the guarantee holding for two draws (as opposed to, say, n draws) morally suggests that the effect is not due to concentrations. At a high-level, our proof in Sec. 4.1 bounds $\mathbb{E}_{\mathbf{G}_{ij}\sim\mathcal{N}(0,\sigma^2)}[\kappa(\mathbf{D}(\mathbf{M}+\mathbf{G}))]$ using tools from [SST06], which then translates to a bound on the quantity $\mathbb{E}_{\mathbf{P}_{ij}\sim|\mathcal{N}(0,1)|}[\kappa(\mathbf{D}(\mathbf{M}+\sigma\mathbf{P}))]$ via a change-of-measure argument.

4.1 Proof of Prop. 4.2

In this section we use the following facts.

Fact 4.3 (Cauchy-Schwarz)**.** *If $\mathbf{A}, \mathbf{B}$ are square & nonsingular, then $\kappa(\mathbf{AB}) \leq \kappa(\mathbf{A})\kappa(\mathbf{B})$*

Fact 4.4 (Union bound+normality)**.** *Let $X_1 \dots X_n$ be iid draws from $\mathcal{N}(\mu, \sigma^2)$. If $\mu = O(1)$, then for any $s > 0$ we have $\mathbb{P}[\max_i |X_i| \geq \sqrt{2\sigma^2 \log 2n} + s] \lesssim e^{-s/2\sigma^2}$.*

Fact 4.5 (Paley-Zygmund Inequality)**.** *Let Z be a nonnegative random variable with finite variance. For any $\theta \in [0,1]$, we have $\mathbb{P}[Z > \theta\mathbb{E}[Z]] \geq (1-\theta)^2 \frac{\mathbb{E}[Z]^2}{\mathbb{E}[Z^2]}$.*

Fact 4.6 ([Wsc04, SST06])**.** *If $\mathbf{A} \in \mathbb{R}^{n\times n}$ with $\|\mathbf{A}\|_2 \leq \sqrt{n}$, $\mathbf{G}_{ij} \sim \mathcal{N}(0, \sigma^2)$, $\sigma^2 \leq 1$, then for all $x \geq 1$, $\mathbb{P}[\kappa(\mathbf{A}+\mathbf{G}) \geq x] \lesssim \frac{\sqrt{n \log x}}{x\sigma}$.*

Fix an instance μ and $\mathcal{W}$. To prove Prop. 4.2, by Theorem 3.4 and Def. 4.1 it suffices to show

$$\mathbb{E}_{\mathbf{P}_1,\mathbf{P}_2\sim|\mathcal{N}(0,1)|}\left[\min\{\kappa(\tilde{\mathbf{M}}_1), \kappa(\tilde{\mathbf{M}}_2)\}\right] \lesssim 2^{n^2}/\sigma^2. \tag{4}$$

To accomplish this, we'll instead show[7]

$$\mathbb{E}_{\mathbf{G}_1,\mathbf{G}_2\sim\mathcal{N}(0,\sigma^2)}\left[\min\{\kappa(\mathbf{D}_1(\mathbf{M}+\mathbf{G}_1)), \kappa(\mathbf{D}_2(\mathbf{M}+\mathbf{G}_2))\}\right] \lesssim \mathrm{poly}(n)/\sigma^2, \tag{5}$$

which we can use to establish (4) as follows. Denote h and h' as the p.d.f.s of $\mathcal{N}(0,\sigma^2)$ and $|\mathcal{N}(0,\sigma^2)|$, respectively. Viewing $f(\mathbf{G}_1, \mathbf{G}_2) := \min\{\kappa(\mathbf{D}_1(\mathbf{M}+\mathbf{G}_1)), \kappa(\mathbf{D}_2(\mathbf{M}+\mathbf{G}_2))\}$ as a function of $2n^2$ variables, we have

$$\mathbb{E}_{\mathbf{P}_1,\mathbf{P}_2\sim|\mathcal{N}(0,1)|}\left[\min\{\kappa(\tilde{\mathbf{M}}_1), \kappa(\tilde{\mathbf{M}}_2)\}\right] = \mathbb{E}_{\mathbf{G}_1,\mathbf{G}_2\sim|N(0,\sigma^2|}[f(\mathbf{G}_1,\mathbf{G}_2)]$$

[7] Actually, we'll show this is bounded by $\lesssim n/\sigma^2$.

$$
\begin{aligned}
&= \int \cdots \int_{\mathbb{R}^{2n^2}} f(\vec{x}) \prod_{i=1}^{2n^2} h'(x_i) dx_i \\
&= \int \cdots \int_{\mathbb{R}^{2n^2}} f(\vec{x}) \prod_{i=1}^{2n^2} \frac{h'(x_i)}{h(x_i)} h(x_i) dx_i \\
&\le \int \cdots \int_{\mathbb{R}^{2n^2}} f(\vec{x}) \prod_{i=1}^{2n^2} 2 \cdot h(x_i) dx_i \\
&= 2^{2n^2}\, \mathbb{E}_{\mathbf{G}_1, \mathbf{G}_2 \sim \mathcal{N}(0,\sigma^2)} \left[f(\mathbf{G}_1, \mathbf{G}_2) \right] \\
&\lesssim 2^{n^2}/\sigma^2.
\end{aligned}
$$

Therefore, we just need to prove (5). To this end, suppose that the following claim is true.

Proposition 4.7. *Let $\sigma \in (0, 1/n)$. Letting $X_{ij} \sim \mathcal{N}(0, \sigma^2)$, denote $\mathbf{D}$ as the random diagonal matrix with $\mathbf{D}_{ii} = \frac{1}{1+\sum_j X_{ij}}$. Then, $\mathbb{P}[\kappa(\mathbf{D}) \ge t] \lesssim n/z^{1.5}$*

Proof. Let $Y_i := 1 + \sum_j X_{ij}$, noting these may be treated as iid samples from $\mathcal{N}(1, n\sigma^2)$. We have

$$
\kappa(\mathbf{D}) = \frac{\max_i |(1 + \sum_j X_{ij})^{-1}|}{\min_i |(1 + \sum_j X_{ij}))^{-1}|} = \frac{\max_i |Y_i|}{\min_i |Y_i|}.
$$

We may write

$$
\begin{aligned}
&\mathbb{P}\left[\frac{\max_i |Y_i|}{\min_i |Y_i|} \ge z \right] \\
&\le \mathbb{P}\left[\frac{\max_i |Y_i|}{\min_i |Y_i|} \ge z \mid \max_i |Y_i| < \sqrt{2\mathbf{Var}(Y) \log 2n} + s \right] + \mathbb{P}\left[\max_i |Y_i| > \sqrt{2\mathbf{Var}(Y) \log 2n} + s \right] \\
&\le \mathbb{P}\left[\frac{\sqrt{2n\sigma^2 \log 2n} + s}{\min_i |Y_i|} \ge z \right] + e^{-s/2n\sigma^2} \quad \text{(by Fact 4.4)} \\
&\le \mathbb{P}\left[\frac{\sqrt{2(1/n) \log 2n} + s}{\min_i |Y_i|} \ge z \right] + e^{-sn/2} \quad (\sigma \le 1/n) \\
&\lesssim \mathbb{P}\left[\min_i |Y_i| \le \frac{s}{z} \right] + e^{-sn/2} \quad (n \text{ large})
\end{aligned}
$$

Moreover,

$$
\mathbb{P}\left[\min_i |Y_i| \le \frac{s}{z} \right] = 1 - \mathbb{P}\left[\min_i |Y_i| > \frac{s}{z} \right] = 1 - \mathbb{P}_{Y \sim N(1, n\sigma^2)}\left[Y^2 > \frac{s^2}{z^2} \right]^n \tag{6}
$$

Here we invoke Fact 4.5 for the random variable Y^2, noting that

- $\mathbb{E}[Y^2] = \mathbf{Var}(Y) + \mathbb{E}[Y]^2 = n\sigma^2 + 1$
- $\mathbb{E}[Y^4] = 1 + 6n\sigma^2 + 3n^2\sigma^4$ (fourth moment of $\mathcal{N}(\mu, \sigma^2)$ is $\mu^4 + 6\mu^2\sigma^2 + 3\sigma^4$)

Taking $\theta := \frac{s^2}{z^2} \frac{1}{n\sigma^2+1}$ (and noting that $s \le z$ is required), we have

$$
\mathbb{P}\left[Y^2 > \theta \mathbb{E}[Y^2] \right] = \mathbb{P}\left[Y^2 > \frac{s^2}{z^2} \right]
$$

$$\geq \left(1 - \frac{s^2}{z^2}\frac{1}{n\sigma^2+1}\right)^2 \frac{(n\sigma^2+1)^2}{1+6n\sigma^2+3n^2\sigma^4}$$
$$\gtrsim \left(1 - \frac{s^2}{z^2}\right)^2 \quad (n \text{ large}, \sigma \leq 1/n)$$

Thus, $\mathbb{P}\left[\min_i |Y_i| \leq \frac{s}{z}\right] \lesssim 1 - \left(1 - \frac{s^2}{z^2}\right)^{2n}$. Overall, for any $s > 0$ and $s \leq z$, we have

$$\mathbb{P}\left[\kappa(\mathbf{D}) \geq z\right] \lesssim 1 - \left(1 - \frac{s^2}{z^2}\right)^{2n} + e^{-sn/2}$$

Taking $s = z^{1/4}$, and using Taylor expansion of $1-(1-x)^n$ for $x \in [0,1]$, and $n \geq 2$

$$\mathbb{P}(\kappa(\mathbf{D}) \geq z) \lesssim 1 - \left(1 - \frac{1}{z^{1.5}}\right)^{2n} + e^{-z^{1/4}n/2} \lesssim n/z^{1.5} + e^{-z^{1/4}} \lesssim n/z^{1.5}.$$

$\square$

We are now ready to prove (5). Fix some $\epsilon > 0$ to be decided later. We have that

$$\begin{aligned}
&\mathbb{E}_{\mathbf{G}_1,\mathbf{G}_2\sim N(0,\sigma^2)}\left[\min\{\kappa(\mathbf{D}_1(\mathbf{M}+\mathbf{G}_1)), \kappa(\mathbf{D}_2(\mathbf{M}+\mathbf{G}_2))\}\right] \\
&\lesssim \int_1^\infty \mathbb{P}(\min\{\kappa(\mathbf{D}_1(\mathbf{M}+\mathbf{G}_1)), \kappa(\mathbf{D}_2(\mathbf{M}+\mathbf{G}_2))\} \geq x)dx \\
&= \int_1^\infty \mathbb{P}(\kappa(\mathbf{D}(\mathbf{M}+\mathbf{G})) \geq x)^2 dx \\
&\lesssim \int_1^\infty \mathbb{P}(\kappa(\mathbf{D})\kappa(\mathbf{M}+\mathbf{G}) \geq x)^2 dx \quad \text{(by Fact 4.3)} \\
&\lesssim \int_1^\infty \left(\mathbb{P}(\kappa(\mathbf{D})\kappa(\mathbf{M}+\mathbf{G}) \geq x \text{ and } \kappa(\mathbf{D}) < x^\epsilon) + \mathbb{P}(\kappa(\mathbf{D}) \geq x^\epsilon\right)^2 dx \\
&\lesssim \int_1^\infty \left(\mathbb{P}(\kappa(\mathbf{M}+\mathbf{G}) \geq x^{1-\epsilon} \text{ and } \kappa(\mathbf{D}) < x^\epsilon) + \mathbb{P}(\kappa(\mathbf{D}) \geq x^\epsilon\right)^2 dx \\
&\lesssim \int_1^\infty \left(\mathbb{P}(\kappa(\mathbf{M}+\mathbf{G}) \geq x^{1-\epsilon}) + \mathbb{P}(\kappa(\mathbf{D}) \geq x^\epsilon\right)^2 dx \\
&\lesssim \int_1^\infty \left(\frac{\sqrt{(1-\epsilon)n\log x}}{x^{2(1-\epsilon)}\sigma} + \frac{n}{x^{1.5\epsilon}}\right)^2 dx \quad \text{(by Prop. 4.7 and Fact 4.6)}
\end{aligned}$$

Taking $\epsilon = 1/3 + \Delta$ for $\Delta > 0$ yields that this integral is $\lesssim \frac{n}{\sigma^2} + n^2$. And $\sigma^2 \leq 1/n^2 \implies 1/\sigma^2 \geq n^2$, yielding overall this integral is $\lesssim n/\sigma^2$. Hence, by this point we've shown (5), so we are done.

References

[AM15] Haris Aziz and Simon Mackenzie. A discrete and bounded envy-free cake cutting protocol for four agents. *CoRR*, abs/1508.05143, 2015.

[AM16] 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.

[AM20] Haris Aziz and Simon Mackenzie. A bounded and envy-free cake cutting algorithm. *Commun. ACM*, 63(4):119–126, March 2020.

[Bar96] Julius B Barbanel. Super envy-free cake division and independence of measures. *Journal of Mathematical Analysis and Applications*, 197(1):54–60, 1996.

[BCE+16] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. *Handbook of computational social choice*. Cambridge University Press, 2016.

[BN22] Simina Brânzei and Noam Nisan. The query complexity of cake cutting. *Advances in Neural Information Processing Systems*, 35:37905–37919, 2022.

[BT95] Steven J Brams and Alan D Taylor. An envy-free cake division protocol. *The American Mathematical Monthly*, 102(1):9–18, 1995.

[Chè20] Guillaume Chèze. Envy-free cake cutting: A polynomial number of queries with high probability. *arXiv preprint arXiv:2005.01982*, 2020.

[CP12] Katarína Cechlárová and Eva Pillárová. On the computability of equitable divisions. *Discrete Optimization*, 9(4):249–257, 2012.

[EP84] Shimon Even and Azaria Paz. A note on cake cutting. *Discrete Applied Mathematics*, 7(3):285–296, 1984.

[EP06] Jeff Edmonds and Kirk Pruhs. Cake cutting really is not a piece of cake. In *SODA*, volume 6, pages 271–278, 2006.

[Fle71] Joseph L Fleiss. On the distribution of a linear combination of independent chi squares. *Journal of the American Statistical Association*, 66(333):142–144, 1971.

[MNV16] Raghu Meka, Oanh Nguyen, and Van Vu. Anti-concentration for polynomials of independent random variables. *Theory of Computing*, 12(1):1–17, 2016.

[Pro09] Ariel D Procaccia. Thou shalt covet thy neighbor's cake. In *Twenty-First International Joint Conference on Artificial Intelligence*, 2009.

[PW17] Ariel D Procaccia and Junxing Wang. A lower bound for equitable cake cutting. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 479–495, 2017.

[RP49] Herbert Robbins and EJG Pitman. Application of the method of mixtures to quadratic forms in normal variates. *The annals of mathematical statistics*, pages 552–560, 1949.

[RV09] Mark Rudelson and Roman Vershynin. Smallest singular value of a random rectangular matrix. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 62(12):1707–1739, 2009.

[RW97] Jack M Robertson and William A Webb. Near exact and envy free cake division. *Ars Combinatoria*, 45:97–108, 1997.

[RW98] Jack Robertson and William Webb. *Cake-cutting algorithms: Be fair if you can.* AK Peters/CRC Press, 1998.

[SST06] Arvind Sankar, Daniel A Spielman, and Shang-Hua Teng. Smoothed analysis of the condition numbers and growth factors of matrices. *SIAM Journal on Matrix Analysis and Applications*, 28(2):446–476, 2006.

[Ste49] Hugo Steinhaus. Sur la division pragmatique. *Econometrica: Journal of the Econometric Society*, pages 315–319, 1949.

[Tao12] Terence Tao. *Topics in random matrix theory*, volume 132. American Mathematical Soc., 2012.

[VT07] Van H Vu and Terence Tao. The condition number of a randomly perturbed matrix. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 248–255, 2007.

[Web99] William A Webb. An algorithm for super envy-free cake division. *Journal of mathematical analysis and applications*, 239(1):175–179, 1999.

[WS07] Gerhard J Woeginger and Jiří Sgall. On the complexity of cake cutting. *Discrete Optimization*, 4(2):213–220, 2007.

[Wsc04] Mario Wschebor. Smoothed analysis of κ (a). *Journal of Complexity*, 20(1):97–107, 2004.