Exact Cut Complexity of Equal-Length Proportional Cake Cutting

Yasushi Kawase
Chuo University, Tokyo, Japan

Mohammad Azharuddin Sanpui[1]
Indian Institute of Technology Kharagpur, India

Abstract

We study proportional cake cutting on the interval $[0, 1]$ under an equal-length constraint requiring each of the n agents to receive a bundle of length exactly $1/n$ and to assign value at least $1/n$ to that bundle. We determine the exact worst-case cut complexity of this problem. The exact value is $2n-2$ cuts for every $n \geq 1$. The lower bound follows from a simple identical-valuation instance, and the main contribution is the matching upper bound, since the only all-n upper bound previously available for this problem was quadratic. Our upper-bound proof starts from the constrained necklace-splitting theorem of Jojić, Panina, and Živaljević, which gives the required partition into equal-length bundles when the number of bundles is a prime power. The main difficulty is to convert this prime-power input into an exact all-n cut bound while preserving the equal-length constraint. When r is a prime-power divisor of n and $s = n/r$, our transfer principle constructs r equal-length bundles, builds a balanced fractional assignment of agents to bundles, rounds it by Hall's theorem to an assignment in which each bundle receives exactly s agents, and recurses inside the bundles without additional overhead beyond the recursive cuts. Using the same constrained necklace-splitting theorem, we also show that $2n-2$ cuts suffice for equal-length envy-freeness when n is a prime power. For all n, we give an $O(n^{1.525})$ upper bound via a peeling argument based on the Stromquist–Woodall exact-share theorem. The exact all-n envy-free cut complexity remains open.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases cake cutting, fair division, cut complexity, proportionality, envy-freeness

Digital Object Identifier 10.4230/LIPIcs.MFCS.2026.34

Funding *Yasushi Kawase*: JST ERATO Grant Number JPMJER2301 and JSPS KAKENHI Grant Number JP25K00137.
Mohammad Azharuddin Sanpui: JST Sakura Science Exchange Program, JST LOTUS Programme Grant Number JPMJ25152681.

1 Introduction

Fair division asks how to allocate resources among agents with heterogeneous preferences, and cake cutting is its standard model for divisible resources [2, 9, 27]. Proportional and envy-free cake cutting have been studied extensively, with existential, algorithmic, and complexity results for many variants [3, 7, 13, 14, 23, 29].

In many fair-division settings, the feasible size of each share is fixed independently of the agents' preferences. Appointment slots may have fixed duration, bandwidth blocks may have fixed width, and land plots may have standardized area. These examples differ in domain, but they share a common feature. The amount of resource assigned to an agent is constrained in advance. Fairness must therefore be achieved within this size constraint.

[1] Corresponding author

This paper studies the basic one-dimensional case. The cake is the interval $[0,1]$, and each of the n agents must receive a bundle of length exactly $1/n$. In the terminology of fair division with market values [6], length is the common objective valuation, and the equal-length constraint requires every agent to receive the same market value. More generally, a nonatomic common market valuation can be represented by a measure-preserving parametrization of the cake, under which equal market value becomes equal length. The agents may still have different private valuations, and proportionality is imposed with respect to these private valuations. The cuts must therefore control two quantities at once. For each agent, they must produce a bundle that has private value at least $1/n$ for that agent and common objective value exactly $1/n$.

Even under the equal-length requirement, existence alone does not control how fragmented the allocation is. A fair bundle can be difficult to use if it is scattered across many small fragments. For instance, a schedule split into many short intervals is inconvenient, and a land allotment split into many patches is costly to manage. On a single interval, fragmentation is naturally measured by the number of cuts. This leads to the central question of how many cuts are needed, in the worst case, to give every agent a bundle of length $1/n$ and value at least $1/n$.

We answer this question as an existential cut-complexity problem. Before this work, the only upper bound available for all n was quadratic. It followed from the corresponding envy-free problem with market values [6], which keeps the same common objective-value constraint but strengthens proportionality to envy-freeness for the private valuations.

1.1 Our Results

Let κ_n^{prop} denote the least number of cuts that always suffice for equal-length proportional cake cutting with n agents. Our main result is the exact identity proved in Theorem 15:

$$\kappa_n^{\mathrm{prop}} = 2n - 2 \qquad (n \geq 1).$$

The lower bound is proved by an identical-valuation instance (Proposition 8). The main contribution is the matching upper bound for all n.

The upper-bound proof starts from the constrained necklace-splitting theorem of Jojić, Panina, and Živaljević [19]. For a prime power r, this theorem directly gives the required equal-length proportional allocation with $2r-2$ cuts. Our transfer principle extends this prime-power case to arbitrary n. Choose a prime-power divisor $r \geq 2$ of n and write $n = rs$. We first form r equal-length top-level bundles. We then construct a continuous fractional assignment of agents to these bundles and equalize the total fractional weight assigned to every bundle. Hall's theorem rounds the support of the resulting matrix to an assignment in which each top-level bundle receives exactly s agents. The construction ensures that every rounded assignment sends an agent to a top-level bundle to which the agent assigns value at least $1/r$. Recursing inside the top-level bundles then gives each assigned agent value at least $1/n$, and the pullback from the recursive subproblems introduces no extra cuts beyond the recursive cuts. This yields the recurrence of Lemma 14, which gives the upper bound in Theorem 15. Section 3 states the constrained necklace-splitting theorem in the form used here and derives the prime-power consequences. Section 4 proves the transfer principle and the induction for arbitrary n.

As a further consequence, we study equal-length envy-freeness. Let κ_n^{ef} denote the corresponding worst-case cut complexity. For a prime power r, the same constrained necklace-splitting theorem gives $\kappa_r^{\mathrm{ef}} \leq 2r-2$, and the lower bound for proportionality gives equality

(Proposition 6 and Theorem 15). For all n, a peeling argument based on the Stromquist–Woodall exact-share theorem gives $\kappa_n^{\mathrm{ef}} = O(n^{1.525})$ (Proposition 18). The exact all-n envy-free cut complexity remains open. Section 5 proves these results.

1.2 Related Work

Fair division with market values is the most direct modeling comparison. Barman, Ebadian, Latifian, and Shah [6] study fair division with both private valuations and a common market valuation. By interpreting the common market valuation as length, equal market value becomes exactly the equal-length constraint considered here. In the cake-cutting setting, their result gives the generic upper bound $\kappa_n^{\mathrm{ef}} \le n(n-1)$ for the stronger requirement of envy-freeness with respect to both the private valuations and the common market valuation, and their Open Question 5 asks whether $O(n)$ cuts always suffice. We do not settle this linear bound. Instead, Section 5 gives an $O(n^{1.525})$ upper bound for the length specialization, answering the weaker question of whether a subquadratic bound is possible. Since envy-freeness implies proportionality, their quadratic bound also applied to the proportional problem, but it did not settle the exact proportional cut complexity.

Proportional cake cutting with unequal or weighted entitlements is another comparison class, because it also seeks private-value guarantees while minimizing the number of cuts. In that setting, Segal-Halevi [25] proved a general $O(n \log n)$ upper bound together with a $2n-2$ lower bound, and Crew, Narayanan, and Spirkl [12] improved the upper bound to $3n-4$. These results do not require each allocated bundle to have length exactly $1/n$, or more generally common-measure value exactly $1/n$, and therefore do not imply the present theorem.

Another useful comparison is exact division and continuous necklace splitting, where all agents evaluate all final bundles equally. This stronger requirement yields quadratic cut bounds, for example via Alon [1]. See also Goldberg and Li [15] for the arbitrary-ratio viewpoint. For equal-length proportionality, exact division is stronger than needed. Each agent only has to assign value at least $1/n$ to the assigned bundle. This paper uses the constrained necklace-splitting theorem of Jojić, Panina, and Živaljević [19]. In the special case stated later as Corollary 4, it yields equal-length bundles together with the property that agent i accepts label i. For prime-power numbers of agents, this already gives the optimal $2n-2$ bound.

Equal-size constraints also appear for indivisible goods. Biswas and Barman [8] study fair division under cardinality constraints for indivisible goods, and Kawase and Mahara [20] study the balanced special case in which each agent must receive exactly the same number of items while seeking fair and efficient allocations. Related models of allocation of indivisible items arise in [10,16,21,22]. Our focus is instead the continuous model, where size is measured by length and cut complexity is the natural structural notion.

Other cake-cutting models impose geometric restrictions, such as non-overlap constraints in multi-layered cake cutting [17] or limits on connected pieces [18,26], but these settings differ from equal-length allocation on a single interval.

2 Preliminaries

We consider allocating the cake $[0,1]$ among n agents, indexed by $\{1,\ldots,n\}$. Each agent i has a valuation v_i, modeled as a nonatomic probability measure over measurable subsets of $[0,1]$ and normalized by $v_i([0,1]) = 1$.

An allocation is a measurable partition $(C_1, \dots, C_n)$ of $[0,1]$. We call it *equal-length* if $\lambda(C_i) = 1/n$ for every $i \in \{1, \dots, n\}$, where λ denotes the standard length measure on $[0,1]$. We call it *proportional* if $v_i(C_i) \geq 1/n$ for every $i \in \{1, \dots, n\}$, and *envy-free* if $v_i(C_j) \leq v_i(C_i)$ for every $i, j \in \{1, \dots, n\}$. In particular, every envy-free allocation is proportional. Indeed, for each agent i, envy-freeness implies that $v_i(C_i)$ is at least the average of $v_i(C_1), \dots, v_i(C_n)$, and hence at least $1/n$.

Given q distinct interior cut points, list them as $0 < t_1 < \cdots < t_q < 1$, put $t_0 = 0$ and $t_{q+1} = 1$, and let $m = q + 1$. For $q = 0$, this list is empty and $m = 1$. For $j \in \{1, \dots, m\}$, define the induced pieces by

$$I_j := \begin{cases} [t_{j-1}, t_j) & \text{if } j \leq m-1, \\ [t_{m-1}, t_m] & \text{if } j = m. \end{cases}$$

We say that an allocation uses at most q cuts if there exist at most q distinct interior cut points such that each bundle is a union of the induced pieces. In configuration spaces below, formal cut coordinates may coincide. Such coincident coordinates create only degenerate intervals and never increase the number of distinct cuts in the original cake. Let κ_n^{prop} denote the least number of cuts that always suffice for equal-length proportional cake cutting with n agents. Similarly, let κ_n^{ef} denote the least number of cuts that always suffice for equal-length envy-free cake cutting with n agents.

For the envy-free results, we use the Stromquist–Woodall exact-share theorem [28]. It provides a region on which several measures all take the same prescribed value. This is the basic input for the exact-share peeling argument in Section 5.

▶ **Theorem 1** (Stromquist–Woodall). *Let $d \geq 2$, let $\eta_1, \dots, \eta_d$ be nonatomic probability measures on $[0,1]$, and let $\alpha \in [0,1]$. Then there exists a measurable set $A \subseteq [0,1]$ such that $\eta_j(A) = \alpha$ for every $j \in \{1, \dots, d\}$. Moreover, A can be chosen as a union of pieces induced by at most $2d-2$ cuts.*[2]

3 The Constrained Necklace-Splitting Theorem

This section states the constrained necklace-splitting theorem used in the proof and explains how it is used for equal-length cake cutting. The section has three parts. Section 3.1 defines the quotient configuration space and states the theorem in interval language. Section 3.2 gives the direct prime-power consequence for equal-length envy-freeness and proportionality. Section 3.3 derives the equalization principle used later for the total fractional weight assigned to each bundle.

3.1 Configuration Space and Theorem Statement

Fix integers $m, r \geq 1$. Let

$$\Delta_m := \left\{ (y_1, \dots, y_m) \in \mathbb{R}^m_{\geq 0} : \sum_{k=1}^{m} y_k = 1 \right\}$$

[2] Stromquist and Woodall [28] originally state the cut-count guarantee for a circle. In our interval notation, this means viewing $[0,1]$ as a circle by identifying 0 and 1. The exact-share set can then be chosen as a union of at most $d-1$ connected pieces of the circle. Opening the circle at the identified endpoint gives a subset of $[0,1]$ with at most $2d-2$ distinct interior boundary points. Since all measures are nonatomic, assigning the opened endpoint to either side does not change any measure.

be the simplex of possible lengths of m consecutive subintervals. For $y \in \Delta_m$, put $x_0(y) = 0$ and $x_k(y) = \sum_{\ell=1}^{k} y_\ell$ for $k \in \{1, \dots, m\}$. These cut positions partition $[0, 1]$ into at most m possibly degenerate subintervals. An (m, r)-*configuration* is a pair (y, f), where $y \in \Delta_m$ and $f : \{1, \dots, m\} \to \{1, \dots, r\}$ assigns a label to each subinterval, with labels on degenerate subintervals identified. We write $\mathcal{C}_{m,r}$ for the resulting quotient space. Thus $y_1, \dots, y_m$ are the subinterval lengths and $f(k)$ is the label of the k-th subinterval. If $y_k = 0$, the corresponding subinterval has no length and its label is ignored. Hence, changing labels on such degenerate subintervals does not change the represented labeled partition. Here, $\mathcal{C}_{m,r}$ is the space of necklace splittings: the m subintervals are the pieces created by $q = m - 1$ cuts, the r labels are the r bundles, and f specifies which agent receives each piece. In cake-cutting terms, $V_i(\xi)$ is the bundle handed to agent i.

Each point of $\mathcal{C}_{m,r}$ is an equivalence class $\xi = [(y, f)]$. The symmetric group S_r acts on $\mathcal{C}_{m,r}$ by relabeling bundles. For $\sigma \in S_r$, define $\sigma \cdot [(y, f)] := [(y, \sigma \circ f)]$.

For a representative (y, f) of ξ and a label i, let $V_i(y, f)$ be the union of the subintervals carrying label i. Changing the representative only changes labels on zero-length subintervals. Therefore, the resulting unions agree up to boundary points. We write $V_i(\xi)$ for this bundle, understood up to boundaries. Under this relabeling rule, $V_{\sigma(i)}(\sigma \cdot \xi)$ and $V_i(\xi)$ agree up to boundaries for every $\sigma \in S_r$ and $i \in \{1, \dots, r\}$. The sets $V_1(\xi), \dots, V_r(\xi)$ form a partition of $[0, 1]$ up to boundaries, and the total number of subintervals across all labeled bundles is at most m. Since $\mathcal{C}_{m,r}$ is a quotient of the compact space $\Delta_m \times \{1, \dots, r\}^{\{1,\dots,m\}}$ and the quotient above is a finite simplicial complex, $\mathcal{C}_{m,r}$ is a compact metric space. In particular, every sequence in $\mathcal{C}_{m,r}$ has a convergent subsequence.

▶ **Definition 2.** *A* preference matrix *on $\mathcal{C}_{m,r}$ is a family $\{A_i^j\}_{i,j\in\{1,\dots,r\}}$ of subsets of $\mathcal{C}_{m,r}$, where A_i^j is the set of configurations at which agent j accepts label i. It is called* closed *if A_i^j is a closed subset of $\mathcal{C}_{m,r}$ for all i and j. It is called* proper *if $\bigcup_{i=1}^{r} A_i^j = \mathcal{C}_{m,r}$ for every agent j. It is called S_r*-equivariant *if $\sigma \cdot A_i^j = A_{\sigma(i)}^j$ for every $\sigma \in S_r$ and every $i, j \in \{1, \dots, r\}$, where $\sigma \cdot A_i^j := \{\sigma \cdot \xi : \xi \in A_i^j\}$.*

The three conditions have standard roles. Closedness ensures that acceptance is preserved under limits. Properness says that, in every configuration, each agent accepts at least one label. Equivariance says that relabeling the bundles only relabels the accepted labels. In the applications below, the sets A_i^j will be defined by valuation comparisons. In the necklace-splitting framework, A_i^j denotes the configurations at which agent j is willing to take bundle i; in cake-cutting terms, agent j finds bundle i acceptable (for instance, of maximum value among the r bundles).

The next theorem is the constrained necklace-splitting theorem used in the rest of the argument. We state Theorem 6.9 of Jojić, Panina, and Živaljević [19] using the notation fixed above. We translate their equivariance convention into the relabeling rule fixed above. Their theorem is stated for continuous probability measures. For Borel probability measures on $[0, 1]$, this is the nonatomic formulation used here, since nonatomicity is equivalent to continuity of the cumulative distribution function. This continuity is what makes the bundle values continuous on the configuration space, and boundary conventions do not affect any measure values. The theorem simultaneously equalizes the measures $\mu_1, \dots, \mu_t$ and guarantees an acceptance pattern after a suitable permutation of the labels. Here, r is the (prime-power) number of equal-length bundles, and t is the number of measures $\mu_1, \dots, \mu_t$ that the splitting equalizes simultaneously across the r bundles. The bound $q = (r - 1)(t + 1)$ on the number of cuts grows linearly in both r and t.

▶ **Theorem 3** (Theorem 6.9 of [19]). *Let $r \geq 2$ be a prime power, let $t \geq 1$, and put $q \coloneqq (r-1)(t+1)$ and $m \coloneqq q+1$. Let $\mu_1, \dots, \mu_t$ be nonatomic probability measures on $[0,1]$, and let $\{A_i^j\}_{i,j\in\{1,\dots,r\}}$ be a closed, S_r-equivariant, proper preference matrix on $\mathcal{C}_{m,r}$. Then, there exist a configuration $\xi \in \mathcal{C}_{m,r}$ and a permutation $\pi \in S_r$ such that $\xi \in A^j_{\pi(j)}$ for every $j \in \{1,\dots,r\}$, and $\mu_\tau(V_i(\xi)) = 1/r$ for every $\tau \in \{1,\dots,t\}$ and every $i \in \{1,\dots,r\}$. In particular, the resulting partition uses at most $q = (r-1)(t+1)$ cuts.*

By equivariance of the preference matrix and symmetry of the measure equalities, the permutation in Theorem 3 can be absorbed by relabeling the bundles.

▶ **Corollary 4.** *Under the assumptions of Theorem 3, there exists a configuration $\xi \in \mathcal{C}_{m,r}$ such that $\xi \in A_j^j$ for every $j \in \{1,\dots,r\}$, and $\mu_\tau(V_i(\xi)) = 1/r$ for every $\tau \in \{1,\dots,t\}$ and every $i \in \{1,\dots,r\}$. The resulting partition uses at most $q = (r-1)(t+1)$ cuts.*

We will use this form in two ways. It ensures that agent i accepts label i, and it simultaneously balances every auxiliary measure across the labeled bundles.

To apply Corollary 4 to valuation-induced preferences, we must verify that the resulting preference matrix is closed. For this, it is enough to know that the value of each labeled bundle varies continuously over the quotient space $\mathcal{C}_{m,r}$. We will also use the same continuity fact later in the limit step of the transfer argument.

▶ **Lemma 5.** *Let ν be a nonatomic finite measure on $[0,1]$, and fix a label $i \in \{1,\dots,r\}$. Then the map $\xi \mapsto \nu(V_i(\xi))$ is continuous on $\mathcal{C}_{m,r}$.*

Proof. Fix a label function $f\colon \{1,\dots,m\} \to \{1,\dots,r\}$. On the fixed-label subset represented by (y,f), let $x_0(y) = 0$ and $x_k(y) = \sum_{\ell=1}^{k} y_\ell$. The bundle $V_i(\xi)$ is the union of those pieces $[x_{k-1}(y), x_k(y))$ with $k < m$ and $[x_{m-1}(y), x_m(y)]$ with $k = m$ for which $f(k) = i$. Hence $\nu(V_i(\xi))$ is a finite sum of terms of the form $\nu([x_{k-1}(y), x_k(y)))$ or $\nu([x_{k-1}(y), x_k(y)])$. These terms are differences of values of the cumulative distribution function of ν at cut positions that depend continuously on y, because the endpoint conventions only change singleton sets. Since ν is nonatomic, this cumulative distribution function is continuous. Hence the map is continuous on each fixed-label subset.

If two representatives differ only by relabeling degenerate pieces, then the corresponding bundle unions differ only by sets of ν-measure zero. Thus the subsetwise formulas agree on overlaps and descend to the quotient $\mathcal{C}_{m,r}$. Since the quotient is obtained by gluing finitely many fixed-label subsets, the pasting lemma gives continuity on $\mathcal{C}_{m,r}$. ◀

3.2 Prime-Power Consequence

For fixed m and valuations $v_1, \dots, v_r$, define the induced preference family on $\mathcal{C}_{m,r}$ by

$$A_i^j \coloneqq \left\{ \xi \in \mathcal{C}_{m,r} : v_j(V_i(\xi)) \geq \max_{i' \in \{1,\dots,r\}} v_j(V_{i'}(\xi)) \right\}.$$

By Lemma 5, each map $\xi \mapsto v_j(V_i(\xi))$ is continuous. Hence each set A_i^j is closed. The family is S_r-equivariant because relabeling the bundles relabels the maximizing indices. It is proper because every finite set of real numbers has a maximizer. Hence Corollary 4 applies.

▶ **Proposition 6.** *If $r \geq 2$ is a prime power, then equal-length envy-free cake cutting with r agents admits a solution with at most $2r-2$ cuts.*

Proof. Apply Corollary 4 with $t = 1$, with measure $\mu_1 = \lambda$, and with the preference family induced by $v_1, \dots, v_r$. The measure equality for λ gives every bundle $V_i(\xi)$ length $1/r$. Moreover, each agent j accepts label j, thus $v_j(V_j(\xi)) \geq v_j(V_i(\xi))$ for every $i \in \{1,\dots,r\}$.

This is exactly envy-freeness: agent j weakly prefers its own bundle $V_j(\xi)$ to every other bundle. Proportionality then follows immediately: since the bundles partition $[0,1]$ up to boundaries and v_j is nonatomic, $\sum_{i=1}^{r} v_j(V_i(\xi)) = 1$, thus the largest of these r values, namely $v_j(V_j(\xi))$, is at least $1/r$. Hence, the allocation is equal-length envy-free, and therefore equal-length proportional. The cut count is $q = (r-1)(1+1) = 2r-2$. ◀

3.3 Equalizing Assignment Weights

Proposition 6 does not directly extend to arbitrary n, because Corollary 4 requires the number of labels to be a prime power. For the all-n result, we use Corollary 4 differently. Instead of assigning final bundles immediately, we use it to equalize a continuous statistic on a prime-power configuration space.

Concretely, we apply Corollary 4 to the maximizing sets of a continuous S_r-equivariant statistic $F\colon \mathcal{C}_{2r-1,r} \to \mathbb{R}^r$. For such an F, we declare label i acceptable at ξ when $F_i(\xi)$ is one of the largest entries of $F(\xi)$. The theorem gives an equal-length configuration at which every label is acceptable. Hence every entry of $F(\xi)$ is largest, and all entries are equal. Later we choose $F_i(\xi)$ to be the total fractional weight assigned to bundle i. This turns a continuous fractional assignment matrix into one in which every bundle receives the same total weight before we convert it into a balanced discrete assignment by Hall's theorem.

The next lemma states this equalization step.

▶ **Lemma 7.** *Let $r \geq 2$ be a prime power. Suppose $F = (F_1, \dots, F_r)\colon \mathcal{C}_{2r-1,r} \to \mathbb{R}^r$ is continuous and satisfies $F_{\sigma(i)}(\sigma \cdot \xi) = F_i(\xi)$ for every $\sigma \in S_r$ and $i \in \{1, \dots, r\}$. Then there exists $\xi \in \mathcal{C}_{2r-1,r}$ such that $\lambda(V_1(\xi)) = \dots = \lambda(V_r(\xi)) = 1/r$, and $F_1(\xi) = \dots = F_r(\xi)$.*

Proof. For each label $i \in \{1, \dots, r\}$, define

$$A_i := \left\{ \xi \in \mathcal{C}_{2r-1,r} : F_i(\xi) \geq \max_{i' \in \{1,\dots,r\}} F_{i'}(\xi) \right\}.$$

Let every agent have the same preference family, i.e., $A_i^j := A_i$ for all $i, j \in \{1, \dots, r\}$. Because F is continuous, each set A_i is closed. Because at least one entry of $F(\xi)$ is maximal at every ξ, the matrix $\{A_i^j\}_{i,j \in \{1,\dots,r\}}$ is proper. Because $F_{\sigma(i)}(\sigma \cdot \xi) = F_i(\xi)$, relabeling labels by any permutation $\sigma \in S_r$ sends maximizing labels at a configuration to maximizing labels at the relabeled configuration. Thus the preference matrix is S_r-equivariant.

Apply Corollary 4 with the single measure λ and with the preference matrix defined above. We obtain a configuration $\xi \in \mathcal{C}_{2r-1,r}$ such that $\lambda(V_1(\xi)) = \dots = \lambda(V_r(\xi)) = 1/r$, and $\xi \in A_j^j = A_j$ for every $j \in \{1, \dots, r\}$. Thus every label corresponds to a largest entry of $(F_1(\xi), \dots, F_r(\xi))$, and therefore $F_1(\xi) = \dots = F_r(\xi)$. ◀

Lemma 7 gives a general equalization principle. Any continuous S_r-equivariant statistic on configurations can be made equal across labels without violating the equal-length constraint.

4 Exact Cut Complexity for Proportionality

In this section we prove that equal-length proportional cake cutting has cut complexity $2n-2$ for every n, that is, $\kappa_n^{\mathrm{prop}} = 2n-2$. Our starting point is Corollary 4, which yields the needed top-level partition when the number of labels is a prime power. The main task is to extend that prime-power case to arbitrary n.

Write $n = rs$, where r is a prime-power divisor of n. We first prove an abstract transfer statement for partitions into r equal-length labeled bundles. It converts a continuous fractional assignment of agents to those bundles into a discrete assignment in which exactly $s = n/r$ agents are assigned to each bundle. We then show that proportionality supplies the required fractional assignment.

We also use the following elementary lower bound, which is analogous to a lower-bound example in the market-value setting [6].

▶ **Proposition 8.** *For every $n \geq 1$, we have $\kappa_n^{\mathrm{prop}} \geq 2n - 2$.*

Proof. Fix $n \geq 1$. The case $n = 1$ is trivial. We may therefore assume $n \geq 2$. Consider the profile in which every agent i has the same valuation $v_i(C') = n \cdot \lambda(C' \cap [0, 1/n])$ for $C' \subseteq [0, 1]$. Let $(C_1, \ldots, C_n)$ be any equal-length proportional allocation. Since every agent values only the prefix $[0, 1/n]$, each bundle C_i must satisfy $\lambda(C_i \cap [0, 1/n]) \geq 1/n^2$. Summing over i shows that equality must hold for every i, namely $\lambda(C_i \cap [0, 1/n]) = 1/n^2$ and $\lambda(C_i \cap (1/n, 1]) = 1/n - 1/n^2$. Therefore every bundle has positive length inside $[0, 1/n]$ and positive length inside $(1/n, 1]$.

Let T be the set of distinct interior cut points used by the allocation. The points of $T \cap (0, 1/n)$ split $[0, 1/n]$ into $|T \cap (0, 1/n)| + 1$ positive-length intervals, each contained in a single induced piece and therefore assigned to a single bundle. Since every one of the n bundles has positive length in $[0, 1/n]$, these intervals must serve all n bundles, and hence $|T \cap (0, 1/n)| + 1 \geq n$. Thus at least $n - 1$ cuts lie in $(0, 1/n)$. This argument does not depend on whether $1/n$ itself is a cut point. If an induced piece crosses $1/n$, its left part still belongs to only the one bundle containing that induced piece.

The same reasoning applied to $(1/n, 1]$ gives $|T \cap (1/n, 1)| \geq n - 1$. The two open intervals are disjoint. Hence every such allocation uses at least $2n - 2$ distinct interior cuts. ◀

4.1 The Transfer Principle

We begin with a discrete rounding step. It converts a fractional assignment matrix with balanced row and column sums into an assignment that sends exactly s agents to each of the r labels. Each agent distributes a fixed total weight across labels, and each label receives total weight 1. We use Hall's marriage theorem in the following standard form. A bipartite graph has a matching covering a given side if every subset of that side has at least as many neighbors as its size.

▶ **Lemma 9.** *Let $M = (M_{ij})$ be an $r \times n$ nonnegative matrix such that every row sums to 1 and every column sums to $1/s$, where $n = rs$ and s is a positive integer. Then there exists a map $\pi \colon \{1, \ldots, n\} \to \{1, \ldots, r\}$ with the following properties, where $\pi^{-1}(i)$ denotes the set of elements mapped to i.*

1. *$|\pi^{-1}(i)| = s$ for every $i \in \{1, \ldots, r\}$, and*
2. *$M_{\pi(j),j} > 0$ for every $j \in \{1, \ldots, n\}$.*

Proof. Consider the bipartite graph with left vertices $\{1, \ldots, r\} \times \{1, \ldots, s\}$ and right vertices $\{1, \ldots, n\}$, where (i, ℓ) is adjacent to j if and only if $M_{ij} > 0$. We first show that this graph has a perfect matching. The two sides both have $n = rs$ vertices. Hence a matching covering the right side is perfect, and by Hall's theorem it is enough to show that every subset $S \subseteq \{1, \ldots, n\}$ of right vertices has at least $|S|$ neighbors.

Let $T \subseteq \{1,\dots,r\}$ be the set of row indices that appear among the neighbors of S. Then all positive entries in columns indexed by S lie in rows indexed by T. Summing the column sums over S gives

$$\frac{|S|}{s} = \sum_{j\in S}\sum_{i=1}^{r} M_{ij} = \sum_{j\in S}\sum_{i\in T} M_{ij} \le \sum_{i\in T}\sum_{j=1}^{n} M_{ij} = |T|.$$

Therefore, the neighborhood of S in $\{1,\dots,r\}\times\{1,\dots,s\}$ has size $s|T| \ge |S|$. Hall's theorem gives a perfect matching.

If column j is matched to (i,ℓ), define $\pi(j) \coloneqq i$. Because there are exactly s copies of each row index i, we have $|\pi^{-1}(i)| = s$. Because edges occur only where $M_{ij} > 0$, we also have $M_{\pi(j),j} > 0$ for every j. Thus the lemma holds. ◀

Before stating the main transfer lemma, Lemma 11, we define the continuous family of fractional assignments that it uses. In our application, $M^{\varepsilon}_{ij}(\xi)$ is the fractional weight that agent j assigns to bundle i at configuration ξ, and positive weight is allowed only for bundles whose value to that agent exceeds $1/r - \varepsilon$. The parameter ε is only a slack parameter. It lets us impose a strict support condition, and Lemma 11 later lets ε tend to zero to obtain the weak proportional bound.

An admissible scheme has three features. Each column sum is $1/s$, the vector of total assignment weights is S_r-equivariant, and positive entries occur only when $v_j(V_i(\xi)) > 1/r-\varepsilon$. This abstract formulation separates the equalization step from the later construction of M^{ε} for proportionality.

▶ **Definition 10.** *Let $n = rs$, where s is a positive integer, and fix $\varepsilon > 0$. An ε-*admissible fractional assignment scheme *for an n-agent profile is a continuous map $M^{\varepsilon} \colon \mathcal{C}_{2r-1,r} \to \mathbb{R}^{r\times n}_{\ge 0}$ with the following properties.*

1. *For every configuration ξ and every agent $j \in \{1,\dots,n\}$, the column sum satisfies $\sum_{i=1}^{r} M^{\varepsilon}_{ij}(\xi) = 1/s$.*
2. *The vector of total assignment weights $F^{\varepsilon}(\xi) \coloneqq \left(\sum_{j=1}^{n} M^{\varepsilon}_{1j}(\xi),\dots,\sum_{j=1}^{n} M^{\varepsilon}_{rj}(\xi)\right)$ satisfies $F^{\varepsilon}_{\sigma(i)}(\sigma\cdot\xi) = F^{\varepsilon}_{i}(\xi)$ for every $\sigma \in S_r$ and $i \in \{1,\dots,r\}$.*
3. *Whenever $M^{\varepsilon}_{ij}(\xi) > 0$, we have $v_j(V_i(\xi)) > 1/r - \varepsilon$.*

Lemma 11 is the abstract transfer step. It converts approximate fractional solutions into an exact top-level partition and a balanced discrete assignment. The equalization step gives equal-length labeled bundles, but the assignment of agents must also satisfy the cardinality constraint that exactly s agents are assigned to each top-level bundle. For each fixed ε, we enforce balance at the level of aggregate fractional weights and use Hall's theorem to choose a balanced assignment from the positive entries of M^{ε}. We then let ε tend to zero. The finite number of balanced assignments lets us pass to a subsequence with one fixed assignment, and only the resulting value inequalities are passed to the limit.

▶ **Lemma 11.** *Let $n = rs$, where $r \ge 2$ is a prime power and s is a positive integer, and let $v_1,\dots,v_n$ be an n-agent profile. Suppose that for every $\varepsilon > 0$ this profile admits an ε-admissible fractional assignment scheme. Then this profile admits a partition of $[0,1]$ into r bundles $C_1,\dots,C_r$ with at most $2r-2$ cuts together with a map $\pi \colon \{1,\dots,n\} \to \{1,\dots,r\}$ such that*

1. *$\lambda(C_1) = \cdots = \lambda(C_r) = 1/r$,*
2. *$|\pi^{-1}(i)| = s$ for all $i \in \{1,\dots,r\}$, and*
3. *$v_j(C_{\pi(j)}) \ge 1/r$ for every agent $j \in \{1,\dots,n\}$.*

Proof. Fix $\varepsilon > 0$, and choose an ε-admissible fractional assignment scheme M^ε. Let F^ε be its vector of total assignment weights. Since M^ε is continuous, F^ε is continuous as well. By Lemma 7, there exists $\xi_\varepsilon \in \mathcal{C}_{2r-1,r}$ such that $\lambda(V_1(\xi_\varepsilon)) = \cdots = \lambda(V_r(\xi_\varepsilon)) = 1/r$ and $F_1^\varepsilon(\xi_\varepsilon) = \cdots = F_r^\varepsilon(\xi_\varepsilon)$.

The total sum of all entries of $M^\varepsilon(\xi_\varepsilon)$ is $\sum_{j=1}^n \sum_{i=1}^r M_{ij}^\varepsilon(\xi_\varepsilon) = n \cdot 1/s = n/s = r$, while all row sums are equal. Hence every row sum must equal 1.

Applying Lemma 9 to $M^\varepsilon(\xi_\varepsilon)$, we obtain a map $\pi_\varepsilon \colon \{1,\dots,n\} \to \{1,\dots,r\}$ such that each label has exactly s preimages and $M^\varepsilon_{\pi_\varepsilon(j),j}(\xi_\varepsilon) > 0$ ($j \in \{1,\dots,n\}$). By admissibility, $v_j(V_{\pi_\varepsilon(j)}(\xi_\varepsilon)) > 1/r - \varepsilon$.

Let Π be the set of maps $\pi \colon \{1,\dots,n\} \to \{1,\dots,r\}$ that assign exactly s agents to each label. Choose a sequence $\varepsilon_k > 0$ with $\varepsilon_k \to 0$. For each k, apply the argument above with $\varepsilon = \varepsilon_k$, and write (ξ_k, π_k) for the resulting pair. Since Π is finite, pass to a subsequence with $\pi_k = \pi$ for all k. Since $\mathcal{C}_{2r-1,r}$ is a compact metric space, pass to a further subsequence with $\xi_k \to \xi$. The subsequence fixes the final assignment π.

Set $C_i \coloneqq V_i(\xi)$. The sets $C_1,\dots,C_r$ form a partition of $[0,1]$ up to boundaries, because ξ has a representative with finitely many cut positions and labels. Assign the finitely many boundary points according to the half-open convention of Section 2. Since all relevant measures are nonatomic, no length or value equality or inequality changes. By Lemma 5, for each fixed agent j the map $\xi \mapsto v_j(V_{\pi(j)}(\xi))$ is continuous. Since $v_j(V_{\pi(j)}(\xi_k)) > 1/r - \varepsilon_k$ and $\varepsilon_k \to 0$, continuity implies $v_j(C_{\pi(j)}) \geq 1/r$ for every $j \in \{1,\dots,n\}$. Indeed, for every $\delta > 0$, all sufficiently large k satisfy $v_j(V_{\pi(j)}(\xi_k)) > 1/r - \delta$, and hence the limit is at least $1/r - \delta$. The same continuity lemma applied to λ gives $\lambda(C_i) = 1/r$ ($i \in \{1,\dots,r\}$). Since $\xi \in \mathcal{C}_{2r-1,r}$, it has a representative with at most $2r-2$ formal interior cut positions. Repeated positions only create degenerate intervals. Hence, the resulting partition uses at most $2r-2$ distinct cuts. ◀

Thus approximate fractional feasibility is enough to obtain the exact top-level partition and balanced assignment needed for the recurrence.

To apply the transfer lemma to proportionality, it remains to construct the required fractional assignment schemes. We assign weight only to bundles that are sufficiently valuable for each agent. For each agent j, the quantity $g_{ij}(\xi)$ measures how far bundle i lies above the threshold $1/r - \varepsilon$. We normalize these excesses to make agent j distribute total assignment weight $1/s$ only among bundles above that threshold. The equalization lemma then makes all row sums equal, and Lemma 9 rounds the support to a discrete assignment.

▶ **Lemma 12.** *Let $n = rs$, where $r \geq 2$ is a prime power and s is a positive integer, and let $v_1,\dots,v_n$ be an n-agent profile. Then for every $\varepsilon > 0$, this profile admits an ε-admissible fractional assignment scheme.*

Proof. Fix $\varepsilon > 0$. For each configuration $\xi \in \mathcal{C}_{2r-1,r}$, bundle index $i \in \{1,\dots,r\}$, and agent $j \in \{1,\dots,n\}$, define $g_{ij}(\xi) \coloneqq \max\{0,\ v_j(V_i(\xi)) - (1/r - \varepsilon)\}$. The maps $\xi \mapsto v_j(V_i(\xi))$ are continuous on $\mathcal{C}_{2r-1,r}$ by Lemma 5. Hence each g_{ij} is continuous. For every fixed j, the bundle values sum to 1. Hence at least one bundle has value at least $1/r$, and therefore $d_j(\xi) \coloneqq \sum_{k=1}^r g_{kj}(\xi)$ is continuous and strictly positive for every ξ.

Now define $M_{ij}^\varepsilon(\xi) \coloneqq (r/n)\, g_{ij}(\xi)/d_j(\xi)$.

Since g_{ij} and d_j depend only on the configuration, M^ε is a well-defined continuous nonnegative map. For each $j \in \{1,\dots,n\}$ we have $\sum_{i=1}^r M_{ij}^\varepsilon(\xi) = r/n = 1/s$. If $M_{ij}^\varepsilon(\xi) > 0$, then $g_{ij}(\xi) > 0$, hence $v_j(V_i(\xi)) > 1/r - \varepsilon$. Finally, relabeling sends row i of $M^\varepsilon(\xi)$ to row $\sigma(i)$ of $M^\varepsilon(\sigma \cdot \xi)$. Indeed, $g_{\sigma(i),j}(\sigma \cdot \xi) = g_{ij}(\xi)$, and therefore $d_j(\sigma \cdot \xi) = d_j(\xi)$ and $M^\varepsilon_{\sigma(i),j}(\sigma \cdot \xi) = M_{ij}^\varepsilon(\xi)$. Summing over j gives $F^\varepsilon_{\sigma(i)}(\sigma \cdot \xi) = F_i^\varepsilon(\xi)$. ◀

Thus, for proportionality, Lemma 12 and Lemma 11 give a partition into r equal-length top-level bundles using at most $2r-2$ cuts. They also give a balanced assignment of exactly s agents to each bundle, and every assigned agent assigns value at least $1/r$ to that bundle. This top-level partition and assignment will be used in the recurrence below.

4.2 Recurrence and Induction

The previous subsection gives a top-level bundle C_i and an assigned group of s agents for each label i. To finish the construction, we must divide each C_i among its assigned agents. The definition of κ_s^{prop} applies to one interval, whereas C_i may be a union of several intervals. We handle this by turning the part of the instance inside C_i into a standard instance on $[0,1]$. We concatenate the components of C_i, rescale by the total length of C_i, and normalize the restricted measures. This standard instance on $[0,1]$ is the normalized subproblem. After solving it, we return to the original cake by taking inverse images under the concatenation map. The next lemma shows that this operation preserves the relevant lengths and values and adds no extra cuts beyond the recursive cuts.

Let $C \subseteq [0,1]$ be a finite union of intervals. Suppose that every interior endpoint of these intervals has already been counted as a cut in the original cake, and suppose $\lambda(C) > 0$. Concatenate the components of C in left-to-right order and rescale by total length to obtain a map $\phi\colon C \to [0,1]$ that preserves relative length. For each nonatomic finite measure ν on C with $\nu(C) > 0$, define the normalized pushforward $\widehat{\nu}$ on $[0,1]$ by $\widehat{\nu}(Y) = \nu(\phi^{-1}(Y))/\nu(C)$. Measures with $\nu(C) = 0$ give only zero values on C and impose no nontrivial constraint there. The normalized subproblem is the problem on $[0,1]$ obtained from these normalized pushforward measures.

If an allocation of the normalized subproblem gives a set $B \subseteq [0,1]$, its pulled-back set in the original cake is $\phi^{-1}(B) \subseteq C$. The pulled-back cuts are the inverse images under ϕ of the cuts used in $[0,1]$.

▶ **Lemma 13.** *If the normalized subproblem has a solution using at most k cuts, then the pulled-back solution uses at most k new cuts in the original cake inside C. Lengths are multiplied by $\lambda(C)$, and values for each measure ν with $\nu(C) > 0$ are multiplied by $\nu(C)$. Thus equal-length conditions, value inequalities, envy comparisons, and equality constraints in the normalized solution become the corresponding conditions inside C.*

Proof. An interior cut point used in the normalized solution pulls back either to an interior point of one component of C or to a boundary between two consecutive components in the concatenation. In the first case it contributes one new cut in the original cake. In the second case the corresponding boundary in $(0,1)$ was already cut. Hence at most k new cuts are added. If such a point lies on a component boundary, its ownership can be fixed by the same half-open convention used in Section 2. Since ϕ is length-preserving up to the common factor $\lambda(C)$ on each concatenated component, normalized equal-length constraints pull back to the corresponding absolute lengths inside C. For every nonatomic finite measure ν with $\nu(C) > 0$, the definition of $\widehat{\nu}$ gives $\widehat{\nu}(Y) = \nu(\phi^{-1}(Y))/\nu(C)$. Thus passing from the normalized problem back to C only multiplies all values for ν by the positive constant $\nu(C)$. This preserves inequalities, envy comparisons, and equality constraints. Boundary conventions do not matter because all measures are nonatomic. ◀

Applying the top-level partition and assignment together with this lemma gives the recurrence.

▶ **Lemma 14.** *Let $n = rs$, where $r \geq 2$ is a prime power and s is a positive integer. Then $\kappa_n^{\mathrm{prop}} \leq (2r-2) + r\kappa_s^{\mathrm{prop}}$.*

Proof. Consider an arbitrary n-agent profile $v_1, \ldots, v_n$. Apply Lemma 12 and then Lemma 11. We obtain bundles $C_1, \ldots, C_r$ of length $1/r$ and agent groups $S_i := \pi^{-1}(i)$ ($i \in \{1, \ldots, r\}$), each of size s, such that every $j \in S_i$ satisfies $v_j(C_i) \geq 1/r$.

Fix $i \in \{1, \ldots, r\}$. Each component boundary of C_i in $(0,1)$ is already present as a cut in the original cake. Concatenate the interval components of C_i in left-to-right order and rescale to $[0,1]$. For each $j \in S_i$, we have $v_j(C_i) \geq 1/r > 0$, so we may restrict v_j to C_i and normalize by dividing by $v_j(C_i)$. This gives an s-agent equal-length proportional instance on $[0,1]$. By the definition of κ_s^{prop}, this normalized instance has a solution with at most κ_s^{prop} cuts. By Lemma 13, pulling it back uses at most κ_s^{prop} additional cuts and yields s bundles, each of absolute length $1/n$. Each assigned agent $j \in S_i$ receives value at least $v_j(C_i)/s \geq (1/r)/s = 1/n$.

Applying this independently inside each bundle gives an equal-length proportional allocation of the whole cake. The top-level partition contributes at most $2r-2$ cuts, and the recursive cuts add across the r disjoint bundles. Hence the total number of cuts is at most $(2r-2) + r\kappa_s^{\mathrm{prop}}$. ◀

The transfer principle yields the recurrence in Lemma 14. Strong induction now gives the exact formula for κ_n^{prop}.

▶ **Theorem 15.** *For every $n \geq 1$, we have $\kappa_n^{\mathrm{prop}} = 2n - 2$.*

Proof. We prove the upper bound $\kappa_n^{\mathrm{prop}} \leq 2n-2$ by strong induction on n.

The case $n = 1$ is trivial. Now let $n \geq 2$, and assume for every $t < n$ that $\kappa_t^{\mathrm{prop}} \leq 2t - 2$.

Choose any prime-power divisor $r > 1$ of n, and write $n = rs$. By Lemma 14 and the induction hypothesis applied to $s < n$, we obtain $\kappa_n^{\mathrm{prop}} \leq (2r-2) + r\kappa_s^{\mathrm{prop}} \leq (2r-2) + r(2s-2) = 2rs - 2 = 2n - 2$. This proves $\kappa_n^{\mathrm{prop}} \leq 2n-2$ for every n. Together with Proposition 8, this gives $\kappa_n^{\mathrm{prop}} = 2n - 2$. ◀

5 Cut Complexity for Envy-Freeness

This section studies equal-length envy-freeness. Since every equal-length envy-free allocation is also proportional, Proposition 8 gives $\kappa_n^{\mathrm{ef}} \geq \kappa_n^{\mathrm{prop}} \geq 2n-2$ for every $n \geq 1$. We show that $\kappa_r^{\mathrm{ef}} = 2r-2$ for prime powers r and that $\kappa_n^{\mathrm{ef}} = O(n^{1.525})$ for arbitrary n, while the exact all-n cut complexity remains open.

In the terminology of Barman, Ebadian, Latifian, and Shah [6], this is the length specialization of the cake-cutting problem behind their Open Question 5, which asks whether $O(n)$ cuts always suffice. We do not settle that linear bound. The results below improve the previously known quadratic upper bound $\kappa_n^{\mathrm{ef}} \leq n(n-1)$ to $O(n^{1.525})$, answering the weaker question of whether a subquadratic bound is possible in this specialization.

The proof strategy differs from the proportional transfer in Section 4. We repeatedly peel off regions obtained from the Stromquist–Woodall exact-share theorem (Theorem 1). These regions equalize all agent valuations and auxiliary measures. This ensures that each agent outside a subproblem assigns value exactly $1/n$ to every bundle produced inside that subproblem.

5.1 The Exact-Share Recurrence

For the recurrence, we use the auxiliary quantity $\kappa^{\mathrm{ef}}_{n,h}$. Consider a subproblem with n agents and additional nonatomic probability measures $\nu_1, \ldots, \nu_h$ on $[0,1]$. The n agents receive the n final bundles. The measures $\nu_1, \ldots, \nu_h$ impose only equality constraints and do not receive bundles. We call them auxiliary measures. Then $\kappa^{\mathrm{ef}}_{n,h}$ denotes the least number of cuts that always suffice to produce an equal-length envy-free allocation for the agents such that each auxiliary measure ν_ℓ assigns the same value to all n final bundles. Thus $\kappa^{\mathrm{ef}}_n := \kappa^{\mathrm{ef}}_{n,0}$. We also use the base case $\kappa^{\mathrm{ef}}_{1,h} = 0$ for $h \geq 0$, since with one agent the unique final bundle is the whole cake and already equalizes every auxiliary measure. In the recurrence below, valuations of agents assigned to the complementary side of a split are carried into the subproblem as auxiliary measures. This controls cross-bundle envy. Consequently, in the recurrence for a fixed top-level n, every agent outside a recursive side assigns value exactly $1/n$ in the original scale to each bundle produced on that side. Thus local envy-freeness on each side is enough to rule out envy across the split.

For prime powers, Proposition 6 gives the upper bound $\kappa^{\mathrm{ef}}_r \leq 2r-2$. Combined with the lower bound $\kappa^{\mathrm{ef}}_r \geq \kappa^{\mathrm{prop}}_r = 2r-2$ from Theorem 15, this yields the exact prime-power identity $\kappa^{\mathrm{ef}}_r = 2r-2$.

We next derive a recurrence from exact-share splitting. Applying Theorem 1 to length, the auxiliary measures, and the agent valuations produces a peeled region on which all of these measures take the same prescribed value.

▶ **Theorem 16.** *For every $n \geq 2$, every $h \geq 0$, and every $a \in \{1, \ldots, n-1\}$,*

$$\kappa^{\mathrm{ef}}_{n,h} \leq 2(n+h) + \kappa^{\mathrm{ef}}_{a,h+n-a} + \kappa^{\mathrm{ef}}_{n-a,h+a}.$$

Proof. Let the agent valuations be $\mu_1, \ldots, \mu_n$, let the auxiliary measures be $\nu_1, \ldots, \nu_h$, and let λ denote length. Apply the Stromquist–Woodall theorem, Theorem 1, to the $m = n+h+1$ probability measures $\lambda, \nu_1, \ldots, \nu_h, \mu_1, \ldots, \mu_n$ with target ratio $\alpha = a/n$. Here $m \geq 3$. We obtain a set $A \subseteq [0,1]$ using at most $2m-2 = 2(n+h)$ cuts such that $\lambda(A) = a/n$, $\nu_j(A) = a/n$ for every j, and $\mu_i(A) = a/n$ for every i. Let $B := [0,1] \setminus A$. Then $\lambda(B) = (n-a)/n$, $\nu_j(B) = (n-a)/n$ for every j, and $\mu_i(B) = (n-a)/n$ for every i.

The component boundaries of A and B are already cut by the initial exact-share split. Thus Lemma 13 applies to both sides, and only the recursive cuts themselves add new cuts in the original cake.

Choose an arbitrary subset $S \subseteq \{1, \ldots, n\}$ with $|S| = a$. On A, consider the a agents $\{\mu_i : i \in S\}$ together with the $h+n-a$ auxiliary measures consisting of $\nu_1, \ldots, \nu_h$ and the measures μ_i for $i \notin S$. After restricting to A, normalizing by the common total mass a/n, and using A's component boundaries from the exact-share split, we solve the normalized subproblem on $[0,1]$ and pull the solution back to A using at most $\kappa^{\mathrm{ef}}_{a,h+n-a}$ additional cuts. Every resulting bundle $X \subseteq A$ then has original length $\lambda(A)/a = 1/n$, every original auxiliary measure assigns it value $1/n$, and every valuation μ_i with $i \notin S$ assigns it value $\mu_i(A)/a = 1/n$.

Similarly, on B we solve the $(n-a)$-agent subproblem with agents $\{\mu_i : i \notin S\}$ and auxiliary measures consisting of $\nu_1, \ldots, \nu_h$ and the measures μ_i for $i \in S$, which again has all component boundaries from the exact-share split and therefore costs at most $\kappa^{\mathrm{ef}}_{n-a,h+a}$ additional cuts after pulling the normalized solution back to B. Every resulting bundle contained in B also has original length $1/n$, every original auxiliary measure assigns it value $1/n$, and every valuation μ_i with $i \in S$ assigns it value $1/n$.

It remains to check envy-freeness. Let $i \in S$, and let X_i denote the bundle assigned to agent i inside A. By local envy-freeness in the subproblem on A, X_i is weakly preferred by agent i to every other bundle $X \subseteq A$. Thus $\mu_i(X_i) \geq \mu_i(X)$ for every such bundle. Summing over the a bundles in A gives $\mu_i(X_i) \geq \mu_i(A)/a = 1/n$. On the other hand, every final bundle $Y \subseteq B$ satisfies $\mu_i(Y) = \mu_i(B)/(n-a) = 1/n$ because agent i appears there as an auxiliary measure. Hence $\mu_i(X_i) \geq \mu_i(Y)$ for every bundle $Y \subseteq B$. Therefore agent i envies no bundle outside A. Local envy-freeness inside A handles the bundles within A. The case $i \notin S$ is symmetric, and the recurrence follows. ◀

5.2 Peeling and a Subquadratic Upper Bound

We iterate the exact-share recurrence with $a = 1$, peeling off one bundle at a time until only a prime-power core remains.

▶ **Theorem 17.** *Let $n \geq 2$, let $2 \leq r \leq n$ be a prime power, and put $t \coloneqq n - r$. Then*

$$\kappa_n^{\mathrm{ef}} \leq 2nt + (r-1)(t+2) = 3nt - t^2 - 3t + 2n - 2 \leq 3nt + 2n.$$

Proof. We iterate the case $a = 1$ of Theorem 16 until a prime-power core remains, and then apply the exact prime-power bound to that core.

Using $\kappa_{1,m}^{\mathrm{ef}} = 0$ for all $m \geq 0$, taking $a = 1$ in Theorem 16 gives

$$\kappa_{n,h}^{\mathrm{ef}} \leq 2(n+h) + \kappa_{n-1,h+1}^{\mathrm{ef}}.$$

Starting from $\kappa_{r+t,0}^{\mathrm{ef}}$, we apply this recurrence repeatedly while decreasing the number of agents by one and increasing the number of auxiliary measures by one. Since $r + t = n$, for each $k = 0, \dots, t-1$ we obtain

$$\kappa_{n-k,k}^{\mathrm{ef}} = \kappa_{r+t-k,k}^{\mathrm{ef}} \leq 2(r+t) + \kappa_{r+t-k-1,k+1}^{\mathrm{ef}} = 2n + \kappa_{n-k-1,k+1}^{\mathrm{ef}},$$

since the sum of the two indices stays equal to n throughout. Summing these t inequalities gives

$$\kappa_n^{\mathrm{ef}} \leq 2nt + \kappa_{r,t}^{\mathrm{ef}}.$$

To bound the remaining term, fix any prime-power core instance with r agent valuations and t auxiliary measures $\nu_1, \dots, \nu_t$. Apply Corollary 4 with the $t+1$ measures $\lambda, \nu_1, \dots, \nu_t$, and the preference family induced by the r agent valuations. The length measure makes every core bundle have length $1/r$, each auxiliary measure assigns the same value to all core bundles, and the maximizing preferences make the diagonal assignment envy-free. Since the number of measures is $t+1$, the cut bound from Corollary 4 is $(r-1)(t+2)$. Hence

$$\kappa_{r,t}^{\mathrm{ef}} \leq (r-1)(t+2).$$

This proves the first displayed bound. Since $r = n - t$, we also obtain

$$\kappa_n^{\mathrm{ef}} \leq 2nt + (n-t-1)(t+2) = 3nt - t^2 - 3t + 2n - 2,$$

and the simpler bound $\kappa_n^{\mathrm{ef}} \leq 3nt + 2n$ is immediate. ◀

Therefore, to obtain a subquadratic upper bound for κ_n^{ef}, it suffices to choose a prime power $r \leq n$ for which the gap $t = n - r$ is small. The next proposition uses a theorem of Baker, Harman, and Pintz [4] stating that, for all sufficiently large x, the interval $[x, x + x^{0.525}]$ contains a prime.

▶ **Proposition 18.** *For all sufficiently large n, we have $\kappa_n^{\mathrm{ef}} \leq 3n^{1.525} + 2n$. In particular, $\kappa_n^{\mathrm{ef}} = O(n^{1.525}) = o(n^2)$.*

Proof. Set $x := n - n^{0.525}$. By Baker, Harman, and Pintz [4], for all sufficiently large n there exists a prime $r \in [x, x + x^{0.525}]$. Then

$$x + x^{0.525} = n - n^{0.525} + \left(n - n^{0.525}\right)^{0.525} \le n,$$

and hence $r \in [n - n^{0.525}, n]$. Thus r is a prime power and $t = n - r \le n^{0.525}$. Substituting into Theorem 17 yields

$$\kappa_n^{\mathrm{ef}} \le 3n \cdot n^{0.525} + 2n = 3n^{1.525} + 2n. \blacktriangleleft$$

▶ Remark 19. The only number-theoretic input in this argument is the size of the gap $t = n - r$. Accordingly, any stronger bound on the distance from n to a prime power improves the all-n bound from Theorem 17. Concretely, if every sufficiently large n admits a prime power $r \in [n - \Delta(n), n]$, then $t = n - r \le \Delta(n)$ and Theorem 17 yields $\kappa_n^{\mathrm{ef}} \le 3n\Delta(n) + 2n$. Under the Riemann hypothesis, standard prime-gap estimates imply $\Delta(n) = O(\sqrt{n} \log n)$ and therefore $\kappa_n^{\mathrm{ef}} = O(n^{3/2} \log n)$ [24]. The scale $(\log n)^2$ is a classical conjectural benchmark for prime gaps. It originates in Cramér's probabilistic model and has been discussed recently in the context of probabilistic models for large prime gaps [5, 11]. Under such a Cramér-type hypothesis, one would have $\Delta(n) = O((\log n)^2)$, and hence $\kappa_n^{\mathrm{ef}} = O\big(n(\log n)^2\big)$.

References

1. Noga Alon. Splitting necklaces. *Advances in Mathematics*, 63(3):247–253, 1987. doi:10.1016/0001-8708(87)90055-7.
2. Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. Fair division of indivisible goods: Recent progress and open questions. *Artif. Intell.*, 322:103965, 2022. doi:10.1016/J.ARTINT.2023.103965.
3. 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, 2016. doi:10.1109/FOCS.2016.52.
4. R. C. Baker, G. Harman, and J. Pintz. The difference between consecutive primes, II. *Proceedings of the London Mathematical Society*, 83(3):532–562, 2001. doi:10.1112/plms/83.3.532.
5. William D. Banks, Kevin Ford, and Terence Tao. Large prime gaps and probabilistic models. *Inventiones Mathematicae*, 233(3):1471–1518, 2023. doi:10.1007/s00222-023-01199-0.
6. Siddharth Barman, Soroush Ebadian, Mohamad Latifian, and Nisarg Shah. Fair division with market values. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(13):13589–13596, 2025. doi:10.1609/aaai.v39i13.33484.
7. Xiaohui Bei, Ning Chen, Xia Hua, Biaoshuai Tao, and Endong Yang. Optimal proportional cake cutting with connected pieces. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2012)*, pages 1263–1269, 2012. doi:10.1609/aaai.v26i1.8243.
8. Arpita Biswas and Siddharth Barman. Fair division under cardinality constraints. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI 2018)*, pages 91–97, 2018. doi:10.24963/ijcai.2018/13.
9. Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. Fair allocation of indivisible goods. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, pages 284–310. Cambridge University Press, 2016. doi:10.1017/CBO9781107446984.013.
10. Xiaolin Bu, Zihao Li, Shengxin Liu, Jiaxin Song, and Biaoshuai Tao. Fair division with allocator's preference. In *Proceedings of the 19th Conference on Web and Internet Economics (WINE 2023)*, pages 77–94, 2023. doi:10.1007/978-3-031-48974-7_5.
11. Harald Cramér. On the order of magnitude of the difference between consecutive prime numbers. *Acta Arithmetica*, 2:23–46, 1936.

12 Tom Crew, Bhargav Narayanan, and Sophie Spirkl. Disproportionate division. *Bulletin of the London Mathematical Society*, 52(5):885–890, 2020. doi:10.1112/blms.12368.

13 Ágnes Cseh and Tamás Fleiner. The complexity of cake cutting with unequal shares. *ACM Transactions on Algorithms (TALG)*, 16(3):1–21, 2020. doi:10.1145/3380742.

14 Jeff Edmonds and Kirk Pruhs. Cake cutting really is not a piece of cake. *ACM Transactions on Algorithms*, 7(4):51:1–51:12, 2011. doi:10.1145/2000807.2000819.

15 Paul Goldberg and Jiawei Li. Consensus Division in an Arbitrary Ratio. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, pages 57:1–57:18, 2023. doi:10.4230/LIPIcs.ITCS.2023.57.

16 Paul Gölz and Hannane Yaghoubizade. Fair division among couples and small groups. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(20):16963–16970, 2026. doi:10.1609/AAAI.V40I20.38743.

17 Hadi Hosseini, Ayumi Igarashi, and Andrew Searns. Fair division of time: Multi-layered cake cutting. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI 2020)*, pages 214–220, 2020. doi:10.24963/ijcai.2020/30.

18 Ayumi Igarashi and Frédéric Meunier. Envy-free division of multi-layered cakes, 2021. arXiv:2106.02262.

19 Duško Jojić, Gaiane Panina, and Rade T. Živaljević. Splitting necklaces, with constraints. *SIAM Journal on Discrete Mathematics*, 35(2):1268–1286, 2021. doi:10.1137/20M1331949.

20 Yasushi Kawase and Ryoga Mahara. Fair and efficient balanced allocation for indivisible goods. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(20):17067–17075, 2026. doi:10.1609/AAAI.V40I20.38755.

21 Yasushi Kawase, Bodhayan Roy, and Mohammad Azharuddin Sanpui. Simultaneously Fair Allocation of Indivisible Items Across Multiple Dimensions. In *45th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2025)*, pages 41:1–41:19, 2025. doi:10.4230/LIPIcs.FSTTCS.2025.41.

22 Max Dupré la Tour. Bad news for couples: Tight lower bounds for fair division of indivisible items. *arXiv preprint arXiv:2601.01012*, 2026. doi:10.48550/arXiv.2601.01012.

23 Oleg Pikhurko. On envy-free cake division. *The American Mathematical Monthly*, 107:736–738, 2000. URL: http://www.jstor.org/stable/2695471.

24 Lowell Schoenfeld. Sharper bounds for the Chebyshev functions $\theta(x)$ and $\psi(x)$. II. *Mathematics of Computation*, 30(134):337–360, 1976. doi:10.2307/2005406.

25 Erel Segal-Halevi. Cake-cutting with different entitlements: How many cuts are needed? *Journal of Mathematical Analysis and Applications*, 480(1):123382, 2019.

26 Erel Segal-Halevi. Fair multi-cake cutting. *Discrete Applied Mathematics*, 291:15–35, 2021. doi:10.1016/j.dam.2020.10.011.

27 Hugo Steinhaus. Sur la division pragmatique. *Econometrica*, 17:315, 1949.

28 Walter Stromquist and D. R. Woodall. Sets on which several measures agree. *Journal of Mathematical Analysis and Applications*, 108(1):241–248, 1985. doi:10.1016/0022-247X(85)90021-6.

29 Francis Edward Su. Rental harmony: Sperner's lemma in fair division. *American Mathematical Monthly*, 106:930–942, 1999.