

Fair Cake Division Under Monotone Likelihood Ratios

SIDDHARTH BARMAN, Indian Institute of Science

NIDHI RATHI, Indian Institute of Science

This work develops algorithmic results for the classic cake-cutting problem in which a divisible, heterogeneous resource (modeled as a cake) needs to be partitioned among agents with distinct preferences. We focus on a standard formulation of cake cutting wherein each agent must receive a contiguous piece of the cake. While multiple hardness results exist in this setup for finding fair/efficient cake divisions, we show that, if the value densities of the agents satisfy the *monotone likelihood ratio property* (MLRP), then strong algorithmic results hold for various notions of fairness and economic efficiency.

Addressing cake-cutting instances with MLRP, first we develop an algorithm that finds cake divisions (with connected pieces) that are envy-free, up to an arbitrary precision. The time complexity of our algorithm is polynomial in the number of agents and the bit complexity of an underlying Lipschitz constant. We obtain similar positive results for maximizing social (utilitarian) and egalitarian welfare. In addition, we show that, under MLRP, the problem of maximizing Nash social welfare admits a fully polynomial-time approximation scheme (FPTAS).

Many distribution families bear MLRP. In particular, this property holds if all the value densities belong to any one of the following families: Gaussian (with the same variance), linear, binomial, Poisson, and exponential distributions. Furthermore, it is known that linear translations of any log-concave function satisfy MLRP. Therefore, our results also hold when the value densities of the agents are linear translations of the following (log-concave) distributions: Laplace, gamma, beta, Subbotin, chi-square, Dirichlet, and logistic. Hence, through MLRP, the current work obtains novel cake-cutting algorithms for multiple distribution families.

CCS Concepts: • **Theory of computation** → **Algorithmic game theory**;

Additional Key Words and Phrases: Cake Cutting, Monotone Likelihood Ratio Property, Envy-Freeness, Social Welfare, Nash Social Welfare, Egalitarian Welfare

ACM Reference Format:

Siddharth Barman and Nidhi Rathi. 2020. Fair Cake Division Under Monotone Likelihood Ratios. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC '20), July 13–17, 2020, Virtual Event, Hungary*. ACM, New York, NY, USA, 37 pages. <https://doi.org/10.1145/3391403.3399512>

1 INTRODUCTION

Cake division is a quintessential model in the study of fair division. This setup captures the allocation of a divisible resource (metaphorically, the cake) among agents with equal entitlements, but distinct preferences. Over the past several decades, a significant body of work in mathematics, economics, and computer science has been devoted to cake cutting; see [8, 27, 28] for excellent expositions and motivating applications (e.g., border negotiations and divorce settlements) of this framework.

Some of the central solution concepts and axiomatic characterizations in the fair-division literature stem from the cake-cutting context [25]. Indeed, the work of Steinhaus, Banach, and Knaster [31]—which lays the mathematical foundations of fair division—addresses cake division. The notion of *envy-freeness* was also mathematically formalized in this setup [15, 16]. This well-studied notion

ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

EC '20, July 13–17, 2020, Virtual Event, Hungary

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7975-5/20/07...\$15.00

<https://doi.org/10.1145/3391403.3399512>

deems a cake division to be fair if every agent prefers the piece assigned to her over that of any other agent, i.e., if no agent is envious of others.

Formally, the cake is modeled as the unit interval $[0, 1]$ and the cardinal preferences of the agents over pieces of this divisible resource are specified via valuation functions: $v_i(I) \in \mathbb{R}_+$ denotes the value that an agent i has for interval (piece) $I \subset [0, 1]$. These valuations v_i s are typically assumed to be induced by value-density functions f_i s, i.e., $v_i(I) := \int_{x \in I} f_i(x) dx$, for each agent i and interval I .

This work focuses on a standard formulation of cake division in which every agent must receive a contiguous piece of the cake. That is, the goal is to partition the cake $[0, 1]$ into exactly n disjoint intervals (connected pieces) and assign them among the n participating agents. This connectivity requirement is naturally motivated by settings in which a contiguous part of the resource needs to be allocated to every agent [8]; consider, e.g., division of land, transmission spectrum, or processing time on a machine. Specifically, a partition of the cake $[0, 1]$ into intervals I_1, I_2, \dots, I_n —wherein interval I_i is assigned to agent $i \in [n]$ —is said to be envy-free iff $v_i(I_i) \geq v_i(I_j)$ (i.e., iff $\int_{I_i} f_i \geq \int_{I_j} f_i$) for all agents i and j .

The appeal of envy-freeness is substantiated by strong existential results: under mild assumptions, a contiguous envy-free cake division always exists [30, 32, 34]. While these results are built upon interesting mathematical connections,¹ they are, however, nonconstructive. In fact, Stromquist [33] has shown that there does not exist a finite-time algorithm for finding envy-free cake divisions with connected pieces; this result holds in a setup wherein the valuations are provided through an (adversarial) oracle. In addition, the work of Deng et al. [13] establishes PPAD-hardness of finding envy-free cake divisions with contiguous pieces, under ordinal valuations.

Algorithms for envy-free cake division remain elusive even if we relinquish the contiguity requirement. It was not until the work of Brams and Taylor [7] that a bounded-time algorithm was obtained for noncontiguous envy-free cake division. In general, the best-known result for this problem is by Aziz and Mackenzie [5], who develop a hyper-exponential time algorithm for finding envy-free divisions with noncontiguous pieces.²

In light of these algorithmic barriers, identification of computationally-tractable instances in the cake-cutting context stands as a meaningful direction of work. The current paper addresses this consideration and, in particular, identifies an encompassing property—called the *monotone likelihood ratio property*—which enables the development of efficient algorithms for fair cake-cutting (with connected pieces).

Value-density functions f_i and f_j are said to satisfy the monotone likelihood ratio property (MLRP) iff, for every $x \leq y$ in the domain, we have $f_j(x)/f_i(x) \leq f_j(y)/f_i(y)$. In other words, the likelihood ratio $f_j(x)/f_i(x)$ is nondecreasing in the argument $x \in \mathbb{R}$. Intuitively, this property asserts that, in comparison to f_i , the density f_j is higher towards the right end of the domain. We note that MLRP does not require f_i and f_j to be monotonic (or unimodal) by themselves.

In the cake-division context, we will say that a collection $\{f_i\}_{i \in [n]}$ of value densities (of the n agents) satisfies the monotone likelihood ratio property iff for each $i \in [n - 1]$, the likelihood ratio $f_{i+1}(x)/f_i(x)$ is nondecreasing in $x \in [0, 1]$. That is, the agents are indexed with the property that consecutive likelihood ratios bear MLRP. This property is transitive and, hence, in cake-division instances with MLRP, value densities f_i and f_j satisfy MLRP for all $i < j$.

Many distribution families are also known to bear MLRP [9, 22]. In particular, this property holds if all the value densities belong to any one of the following families: Gaussian distributions (with the same variance), Poisson distributions, binomial distributions, and single-parameter exponentials; see

¹For instance, the proof by Su [34] invokes Sperner's lemma.

²The problem of finding an *approximate* envy-free division (not necessarily with connected pieces) admits a fully-polynomial time approximation scheme [23].

Appendix F for details. Furthermore, it is known that linear translations of any log-concave function satisfy MLRP [29]. Specifically, linear translations of the following (log-concave) distributions also satisfy this property: Laplace, uniform, multivariate Gaussian, gamma, beta, Subbotin, chi-square, Dirichlet, and logistic. Hence, the current work obtains novel results for many distribution families in a unified manner.

MLRP is a common assumption on agents' utilities and type distributions in many economic contexts; see [19] for a survey. As a stylized application of MLRP in cake division, consider a setting wherein each agent i has a most preferred point μ_i on the divisible resource (cake) and i 's valuation density decreases as a Gaussian function (with a variance parameter that is common across the agents) of the distance from μ_i . Indeed, the distance here can be geographical (as in case of land division), temporal (i.e., wait time), or it can even be an abstract metric.

Considering similar single-peaked preferences, but with a linear drop in value densities, Wang and Wu [35] developed an efficient algorithm for noncontiguous cake division. Note that while linear densities bear MLRP, this property does not hold for piecewise linear densities. Hence our results do not directly address the setting considered in [35]. However, in absence of the contiguity requirement (as is the case in [35]) one can find a fair cake division by first partitioning the cake into intervals, in each of which the agents' value densities are linear, and then applying the MLRP result separately.³

We focus on cake-division instances with MLRP and develop algorithmic results for almost all the standard notions of fairness and economic efficiency. Our algorithms only require oracle access to the valuations. In particular, the developed algorithms operate under the standard Robertson-Webb model [28], wherein we have access to the agents' valuations through *eval* and *cut* queries; see Section 2 for details. MLRP implies that these cut and eval queries (functions) are λ -Lipschitz (Appendix C.3).

The time complexities of our algorithms depend polynomially on the the bit complexity of this Lipschitz constant $\lambda \geq 1$. Such a runtime dependency on $\log \lambda$ is unavoidable (Appendix G): there exist cake-division instances (with λ -Lipschitz cut and eval queries) wherein for all the agents the value of the cake is almost entirely concentrated in an interval L of length $1/\lambda$. Here, an envy-free cake division can be obtained only by finely partitioning L among the agents. In particular, the cut points that induce an envy-free allocation (and, hence, correspond to the output of a fair-division algorithm) must be $1/\lambda$ close to each other, i.e., the bit complexity of the output has to be $\Omega(\log \lambda)$. In fact, one can construct instances in which a contiguous envy-free division can be obtained only by cutting the cake at irrational points (Appendix G). Hence, in general, (and even under MLRP) one cannot expect an efficient algorithm that outputs an *exact* envy-free division, with contiguous pieces.⁴ Therefore, when considering efficient algorithms for cake division, a precision loss in the output is inevitable. However, our algorithms ensure that this precision loss in value, η , is arbitrarily small; specifically, the developed algorithms run in time $O\left(\text{poly}\left(n, \log \lambda, \log \frac{1}{\eta}\right)\right)$ and, hence, the precision parameter η can be driven exponentially close to zero in polynomial (in the bit complexity of η) time. Note that this bit-precision issue is akin to the one faced in the convex-optimization problems (where again the optimal solutions can be irrational) and our runtime bound, with respect to the precision parameter η , is analogous to the one obtained by the ellipsoid method [17].

Our Results and Techniques: Next, we summarize our results for various notions of fairness and economic efficiency.

Envy-Freeness: We prove that, given a cake-division instance (in the Robertson-Webb query model) with MLRP, an envy-free allocation can be computed, up to an arbitrary precision, in

³Recall that, in contrast to such a result, our focus is on contiguous cake division.

⁴Indeed, the bit complexity of a computationally-bounded algorithm is bounded as well.

time that is polynomial in the number of agents and $\log \lambda$; here λ is the Lipschitz constant of the Robertson-Webb (cut and eval) queries.

To establish this result, we define a class of divisions, referred to as *ripple divisions* (Definition 5.1), and prove that, under MLRP, every ripple division induces a contiguous envy-free cake division (Theorem 5.7). Specifically, a collection of points $x_0 = 0 \leq x_1 \leq \dots \leq x_{n-1} \leq x_n = 1$ (in the cake $[0, 1]$) is said to form a ripple division of the cake if, for each $i \in [n - 1]$, agent i is indifferent between the consecutive intervals $[x_{i-1}, x_i]$ and $[x_i, x_{i+1}]$, i.e., $v_i(x_{i-1}, x_i) = v_i(x_i, x_{i+1})$. Note that a ripple division induces a contiguous cake division—by assigning interval $[x_{i-1}, x_i]$ to agent i —with the property that agent i does not envy agent $i + 1$. That is, in and of itself, a ripple division mandates absence of envy only between consecutive agents, and not between all pairs of agents. We will show that, interestingly, under MLRP this relaxation suffices—the cake division induced by a ripple division is guaranteed to be envy free (Theorem 5.7). Recall that the agents are indexed following the MLRP order: for each $i \in [n - 1]$, the likelihood ratio f_{i+1}/f_i is nondecreasing. Hence, allocating $[x_{i-1}, x_i]$ to agent $i \in [n]$ ensures that the intervals are assigned (left to right on the cake) in accordance with the MLRP order.

We establish the universal existence of ripple divisions through the intermediate value theorem, i.e., a one-dimensional fixed-point argument (Lemma 5.5). Since one can use binary search to find fixed points in the one-dimensional setting, this proof in fact leads to an algorithm for finding ripple divisions and, hence, envy-free divisions. Indeed, the notion of ripple divisions and their connection with envy-freeness, under MLRP, are two key contributions of this work.

Pareto Optimality: We show that in cake-division instances with MLRP, Pareto optimal cake divisions, with connected pieces, conform to the MLRP order (Lemma 6.1). This structural result implies that for maximizing welfare we can restrict attention to allocations wherein the intervals are assigned (left to right on the cake) in accordance with the MLRP order. Intuitively, this leads us to a welfare-maximizing algorithm—specifically, a dynamic program—that recursively finds optimal allocations for intervals placed at the left end of the cake (i.e., for intervals of the form $[0, x]$).

We also establish an extension of Weller’s theorem in the MLRP context. Weller’s theorem [36] asserts that there always exists some cake division—though, not necessarily with connected pieces—which is both envy-free (fair) and Pareto optimal. While this theorem holds in general,⁵ it does not guarantee that envy-freeness and Pareto optimality can be achieved together through contiguous cake divisions. We show that, by contrast, under MLRP *every* contiguous envy-free division is Pareto optimal (Theorem 3.2). Therefore, given a cake-division instance with MLRP, the allocation computed by our algorithm is not only envy-free but also Pareto optimal, up to an arbitrary precision.

Social Welfare: Social (utilitarian/Benthamite) welfare is a standard measure of collective value. For a cake division $\{I_1, I_2, \dots, I_n\}$ it is defined to be the sum of the values that the division generates among the agents, $\sum_i v_i(I_i)$. Maximizing social welfare is a well-studied objective in resource-allocation contexts. In the cake-cutting setup, this maximization problem is known to be APX-hard under general valuations [3]. Complementarily, if the value densities bear MLRP, then we can find (up to an arbitrary precision) a social welfare maximizing division with connected pieces in $O(\text{poly}(n, \log \lambda))$ time (Theorem 3.3). As mentioned previously, our algorithm for this problem is based on a dynamic program.

Egalitarian Welfare: The egalitarian (Rawlsian) welfare of a cake division $\{I_1, \dots, I_n\}$ is defined as the value of the least well-off agent, i.e., $\min_i v_i(I_i)$. From a welfarist perspective, maximizing this minimum value among cake divisions with connected pieces is an important fairness objective. However, no nontrivial approximation guarantees are known for this problem under

⁵Weller’s theorem applies even in the absence of MLRP.

general valuations; the work of Aumann et al. [4] shows that maximizing egalitarian welfare across all contiguous cake divisions is APX-hard. Complementing this hardness result, we develop an algorithm that, under MLRP, maximizes egalitarian welfare (up to an arbitrary precision) and runs in $O(\text{poly}(n, \log \lambda))$ time (Theorem 3.4).

Our algorithm for maximizing egalitarian welfare is based on a “moving-knife” procedure. This procedure, for a given a target value $\tau > 0$, iteratively selects points $x_0 = 0, x_1, x_2, \dots, x_n \leq 1$ such that the each interval $[x_{i-1}, x_i]$ is of value τ to agent $i \in [n]$. Let τ^* denote the optimal egalitarian welfare in the given cake-division instance. The useful observation here is that this moving-knife procedure will succeed for all $\tau \leq \tau^*$. This follows from the fact that here the intervals are assigned (left to right) in the MLRP order⁶ and this ordering is also satisfied by an egalitarian welfare maximizing (in particular, a Pareto optimal) division. Therefore, by performing a binary search with τ , we can find a contiguous division with egalitarian welfare arbitrarily close to the optimal.

Nash Social Welfare: A balance between social and egalitarian welfare is obtained by considering the Nash social welfare [20, 26]. This welfare objective is defined as the geometric mean of the agents’ values. It is known that, in general, it is APX-hard to find a contiguous cake division that maximizes Nash social welfare [3]. Under MLRP, however, the problem of maximizing Nash social welfare admits a fully polynomial-time approximation scheme (Theorem 3.5). We obtain this result via a dynamic program that considers the agents in the MLRP order.

Additional Related Work: Recently, approximation algorithms—with both additive [18] and multiplicative [3] approximation guarantees—have been developed for finding contiguous envy-free cake divisions. In contrast to these results, the current work focuses on cake-division instances with MLRP and shows that in such settings arbitrarily low envy can be achieved among the agents.

The work of Bei et al. [6] also studies contiguous cake division and provides computational results for maximizing social welfare subject to proportional fairness. Under this fairness constraint each agent $i \in [n]$ must receive an interval of value at least $1/n$ times i ’s total value for the cake. Bei et al. [6] show that, if the value densities are linear, then this problem admits a fully polynomial-time approximation scheme (FPTAS). We note that every pair of linear densities bear MLRP and, hence, such value-density functions fall within the purview of the current work. However, our algorithm is incomparable to the FPTAS of Bei et al. [6]—we focus on maximizing social welfare without the fairness constraints. Also, envy-freeness is not addressed in [6].

Another well-studied fairness notion is that of equitability. Specifically, a cake division $\{I_1, I_2, \dots, I_n\}$ is said to be *equitable* iff all the agents derive the same value from the intervals assigned to them, $v_i(I_i) = v_j(I_j)$ for all i and j [2, 14]. Cechlárová and Pillárová [10] showed that—given access to *reverse cutting queries*—contiguous equitable cake divisions can be efficiently computed, up to an arbitrary precision. We note that value densities that satisfy MLRP have, by definition, full support over the cake. In such a case, the reverse cutting queries can be simulated by standard (cut) queries in the Robertson-Webb model. Hence, under MLRP, strong algorithmic results hold for equitability as well.

Cake-division algorithms for specific classes of valuations have been studied in [12] and [21]. The work of Kurokawa et al. [21] provides a query-efficient algorithm for envy-free, noncontiguous cake division under piecewise linear densities. Cohler et al. [12] also address the noncontiguous version of the problem, and for piecewise constant densities they develop a polynomial-time algorithm that computes an envy-free division with optimal social welfare. In contrast to these results our focus is on contiguous cake division.

⁶Note that the agents are indexed accordingly.

2 NOTATION AND PRELIMINARIES

This work studies the problem of dividing a cake $[0, 1]$ among n agents. Throughout, we will focus on a well-studied formulation of cake cutting which requires that each agent should receive a contiguous piece of the cake, i.e., the goal is to partition the cake $[0, 1]$ into n pairwise disjoint intervals and assign them among the n agents in a fair/efficient manner.

The cardinal preferences of each agent $i \in [n]$ is induced by a value-density function $f_i : [0, 1] \mapsto \mathbb{R}_+$. Following standard conventions, we will assume that each value-density function f_i is (Riemann) integrable. In particular, the (finite) integral of f_i induces agent i 's valuation function over the intervals contained in $[0, 1]$ (i.e., over the pieces of the cake): $v_i(I) := \int_\ell^r f_i(x) dx$ denotes the value that agent $i \in [n]$ has for any interval $I = [\ell, r] \subset [0, 1]$. For notational convenience, we will write $v_i(a, b)$ to denote agent i 's value for interval $[a, b] \subseteq [0, 1]$.

The integrability⁷ and nonnegativity of value-densities f_i imply that the corresponding valuations v_i are (i) nonnegative, (ii) divisible: for every interval $[\ell, r]$ and parameter $\kappa \in [0, 1]$, there exists a $z \in [\ell, r]$ with the property that $v_i(\ell, z) = \kappa v_i(\ell, r)$,⁸ and (iii) sigma additive: $v_i(I \cup J) = v_i(I) + v_i(J)$, for all disjoint intervals $I, J \subset [0, 1]$. The divisibility property ensures that the valuations are non-atomic, i.e., $v_i([x, x]) = 0$ for all $i \in [n]$ and $x \in [0, 1]$. Furthermore, this property allows us, as a convention, to regard two intervals to be disjoint even if they intersect exactly at an endpoint.

We additionally assume that the valuations are normalized such that the value of the entire cake is equal to one for every agent $i \in [n]$, i.e., $\int_0^1 f_i(x) dx = v_i(0, 1) = 1$. Hence, the value-densities f_i s constitute probability density functions over the cake $[0, 1]$. We note that this is a standard assumption in the cake-cutting framework and we conform to it for the purpose of brevity. All of our results hold true, even otherwise.

Problem Instances: A *cake-division instance* C is a tuple $\langle [n], \{f_i\}_{i \in [n]} \rangle$ where $[n] = \{1, 2, \dots, n\}$ denotes the set of $n \in \mathbb{Z}_+$ agents and f_i s denote the value-density functions of the agents.

Robertson-Webb Query Model: While, for exposition, we specify f_i s as part of the problem instance, our algorithms only require oracle access to the valuations. In particular, the developed algorithms operate under the Robertson-Webb model [28], which supports oracle access to agents' valuations in the form of *eval* and *cut* queries:

(i) *Eval queries:* for each agent $i \in [n]$, we have (blackbox) access to function $\text{Eval}_i : [0, 1] \times [0, 1] \mapsto \mathbb{R}_+$, which when queried with any interval $[\ell, r]$ returns (in unit time) the value that agent i has for this interval, i.e., $\text{Eval}_i(\ell, r) = v_i(\ell, r)$.

(ii) *Cut queries:* for each agent $i \in [n]$, we can also query function $\text{Cut}_i : [0, 1] \times \mathbb{R}_+ \rightarrow [0, 1]$, which given an initial point $\ell \in [0, 1]$ and a target value $\tau \in \mathbb{R}_+$, returns $\text{Cut}_i(\ell, \tau) = y$ where $y \in [0, 1]$ is the leftmost point with the property that $v_i(\ell, y) = \tau$. If for a given $\ell \in [0, 1]$ and $\tau \in \mathbb{R}_+$, there does not exist a $y \in [0, 1]$ such that $v_i(\ell, y) = \tau$, then we have, by convention, $\text{Cut}_i(\ell, \tau) = 1$.

Allocations and Cake Divisions: As mentioned previously, the goal is to assign each agent a single interval. Towards this end, for any cake-division instance with n agents, we define an *allocation* to be a collection of n pairwise-disjoint intervals, $\mathcal{I} = \{I_1, I_2, \dots, I_n\}$, where interval I_i is assigned to agent $i \in [n]$ and $\cup_{i \in [n]} I_i = [0, 1]$. Note that here the subscript of each interval identifies unique agent who has been assigned this interval. In addition, we will refer to a collection of pairwise-disjoint intervals $\mathcal{J} = \{J_1, J_2, \dots, J_n\}$ as a *partial allocation* if they do not cover the entire cake, $\cup_{i=1}^n J_i \subsetneq [0, 1]$.

⁷Our results hold for integrable value densities and do not necessarily require f_i s to be continuous. Recall that, by definition, Riemann integrable functions are bounded. Also, every (bounded) continuous function on an (bounded) interval is Riemann integrable, but the converse is not true.

⁸This implication can be obtained by applying the intermediate value theorem to the antiderivative of f_i .

For an allocation $\mathcal{I} = \{I_1, \dots, I_n\}$, the endpoints of the constituent intervals will be referred to as the *cut points* of \mathcal{I} , i.e., if $I_i = [x_{i-1}, x_i]$ for $1 \leq i \leq n$, then the cut-points are $\{x_0 = 0, x_1, \dots, x_n = 1\}$.

We will throughout use the term allocation to specifically refer to partitions of the cake in which each agent receives a connected piece, i.e., receives exactly one interval. More generally, a *cake division* will be used to denote partitions of the cake $\mathcal{D} = \{D_1, D_2, \dots, D_n\}$ in which agent i receives D_i , a finite collection of intervals. Here, the bundles D_i s are pairwise disjoint and their union covers the entire cake $[0, 1]$.

In this work we develop algorithmic results for the following notions of fairness and economic efficiency.

Envy-Freeness: For a cake-division instance C , an allocation $\mathcal{I} = \{I_1, \dots, I_n\}$ is said to be *envy-free* iff each agent prefers its own interval over that of any other agent, $v_i(I_i) \geq v_i(I_j)$ for all agents $i, j \in [n]$.

Pareto Optimality: Given a cake-division instance C , a division $\mathcal{D} = \{D_1, \dots, D_n\}$ is said to *Pareto dominate* another division $C = \{C_1, \dots, C_n\}$ iff $v_i(D_i) \geq v_i(C_i)$ for all agents $i \in [n]$ and, there exists at least one agent $k \in [n]$ such that $v_k(D_k) > v_k(C_k)$. Consequently, a cake division is said to be *Pareto optimal* iff it is not Pareto dominated by any other division.

Recall that a cake division refers to a partition of the cake in which agent receives a finite collection of intervals. By contrast, in an allocation each agent receives a single interval. The algorithms developed in this work compute allocations. Interestingly, though, the Pareto optimality guarantees achieved by our algorithms are stronger in the sense that optimality holds across all cake divisions; specifically, under MLRP, we establish that particular allocations are Pareto optimal not only among the set of all allocations but also among all cake divisions.

Social Welfare: Social welfare is a standard measure of collective value. Specifically, *social welfare* for an allocation $\mathcal{I} = \{I_1, \dots, I_n\}$ is defined to be sum of the agents' valuations, $\text{SW}(\mathcal{I}) := \sum_{i=1}^n v_i(I_i)$.

Egalitarian (Rawlsian) Welfare: For an allocation $\mathcal{I} = \{I_1, \dots, I_n\}$, the *egalitarian welfare* is defined as the minimum value achieved across the agents, $\text{RW}(\mathcal{I}) := \min_{i \in [n]} v_i(I_i)$.

Nash Social Welfare: For an allocation $\mathcal{I} = \{I_1, \dots, I_n\}$, the *Nash social welfare* is defined to be the geometric mean of the agents' valuations, $\text{NSW}(\mathcal{I}) := (\prod_{i=1}^n v_i(I_i))^{1/n}$.

Finding allocations that maximize the above-mentioned welfare notions is known to be APX-hard, in general; see, e.g., [3, 4]. Complementing these negative results, a key contribution of this work is to identify a broad class of cake-division instances that admit strong algorithmic results for these welfare objectives and envy-freeness. Specifically, we focus on value densities (distributions) that satisfy the monotone likelihood ratio property (MLRP). We will next define this property and note that our results hold for multiple distribution families that satisfy MLRP.

Monotone Likelihood Ratio Property: Probability density functions f_i and f_j are said to satisfy the *monotone likelihood ratio property* (MLRP) iff, for every $x \leq y$ in the domain, we have $f_j(x)/f_i(x) \leq f_j(y)/f_i(y)$.

That is, the likelihood ratio $f_j(x)/f_i(x)$ is non-decreasing in the argument $x \in \mathbb{R}$.

Note that MLRP does not require f_i and f_j to be monotonic (or unimodal) by themselves. We also observe that MLRP is transitive: if two pairs of distributions $\{f_i, f_j\}$ and $\{f_j, f_k\}$ satisfy MLRP separately, then the pair $\{f_i, f_k\}$ also conforms to MLRP. Furthermore, this property continues to hold under positive scaling: if f_i and f_j satisfy MLRP, then so do $\gamma_i f_i$ and $\gamma_j f_j$, for any positive scalars $\gamma_i, \gamma_j \in \mathbb{R}_+$. This fact, in particular, allows us to restrict MLRP densities (which are typically defined over the real line) to the cake $[0, 1]$ and, at the same time, assume normalization $\int_0^1 f_i = 1$.

In the cake-division context, we will say that a given collection $\{f_i\}_{i \in [n]}$ of value-density functions satisfies the *monotone likelihood ratio property* iff there exists an order $\pi : [n] \rightarrow [n]$, among the

f_i s, such that, for all $i \in [n - 1]$, the consecutive likelihood ratios $\frac{f_{\pi(i+1)}(x)}{f_{\pi(i)}(x)}$ are non-decreasing in $x \in [0, 1]$. That is, for each $i \in [n - 1]$, the densities $f_{\pi(i)}$ and $f_{\pi(i+1)}$ bear MLRP over $[0, 1]$.

We will refer to this order π as the MLRP order of the value densities. Lemma C.2 (in Appendix C.2) shows that, given a cake-division instance in the Robertson-Webb query model, with the promise that the underlying value densities satisfy MLRP (i.e., given a promise problem), one can efficiently find the MLRP order π . Hence, without loss of generality, we will throughout assume that the n agents are indexed such that π is the identity permutation, i.e., for all $i \in [n - 1]$, the likelihood ratio $\frac{f_{i+1}(x)}{f_i(x)}$ is non-decreasing in $x \in [0, 1]$.

It is relevant to note that, to be well defined, MLRP requires the value densities f_i s to be strictly positive over the cake $[0, 1]$. Hence, for cake-division instances $\langle [n], \{f_i\}_{i \in [n]} \rangle$ with MLRP, we have $f_i(x) > 0$ for all $i \in [n]$ and $x \in [0, 1]$.

Instantiations of MLRP: MLRP induces a total order on linear value densities $f_i(x) = a_i x + b_i$; see Appendix F for details. Hence, our results imply that if, in a cake-division instance, all the value densities are linear, then an envy-free (or welfare-maximizing) allocation can be computed efficiently.

Many other distribution families are also known to bear MLRP, e.g., Gaussian distributions (with the same variance), Poisson distributions, and single-parameter exponentials. Therefore, our algorithmic results address, in particular, cake-division instances wherein all the agents have Gaussian value densities with the same variance, but different means.

In fact, it is known that linear translations of any log-concave function g —i.e., densities of the form $f_\theta(x) := g(x - \theta)$, for $\theta \in \mathbb{R}$ —satisfy MLRP [29]. Hence, our results also hold for linear translations of the following (log-concave) distributions: Laplace, uniform, multivariate Gaussian, gamma, beta, Subbotin, chi-square, Dirichlet, and logistic. These instantiations substantiate the applicability of our algorithmic results which, through MLRP, address a wide range of cake-division instances.

Lipschitz Constant of Cut and Eval Queries: We say that the cut and eval queries in a cake-division instance are λ -Lipschitz iff the following inequalities hold for each agent $i \in [n]$:

$$\begin{aligned} |\text{Eval}_i(\ell', r') - \text{Eval}_i(\ell, r)| &\leq \lambda \|(\ell', r') - (\ell, r)\|_\infty && \text{for all } (\ell', r'), (\ell, r) \in [0, 1] \times [0, 1] \\ |\text{Cut}_i(\ell', \tau') - \text{Cut}_i(\ell, \tau)| &\leq \lambda \|(\ell', \tau') - (\ell, \tau)\|_\infty && \text{for all } (\ell', \tau'), (\ell, \tau) \in [0, 1] \times \mathbb{R}_+ \end{aligned}$$

A useful consequence of MLRP (and the integrability of the value densities) is that the corresponding cut and eval queries are in fact λ -Lipschitz, for a finite $\lambda \geq 1$;⁹ see Proposition C.3. This proposition follows from the fact that (Riemann) integrable value densities f_i s are, by definition, bounded. Furthermore, as mentioned previously, MLRP mandates that the value densities are strictly positive over the cake. Therefore, for each agent $i \in [n]$ and $x \in [0, 1]$ we have $0 < L \leq f_i(x) \leq U$ for some $L, U \in \mathbb{R}_+$. Proposition C.3 asserts that the Lipschitz constant λ can be expressed in terms of these bounding parameters, $\lambda \leq \max\{1/L, U, U/L\}$.

It is worth pointing out that besides MLRP (and, hence, the positivity of the value densities), all the other assumptions made in this work are standard.

The time complexities of our algorithms depend polynomially on the the bit complexity of the Lipschitz constant $\lambda \geq 1$.¹⁰ As mentioned previously, in general, such a runtime dependency on $\log \lambda$ is unavoidable; Appendix G provides an illustrative example. Furthermore, it is possible—even with rational and MLRP value densities—that the exact envy-free/welfare-maximizing allocations

⁹If a function is λ' -Lipschitz then, it is λ -Lipschitz as well, for all $\lambda \geq \lambda'$.

¹⁰In the case of linear value densities, $f_i(x) = a_i x + b_i$, the bit complexity of the Lipschitz constant λ is proportional to the bit complexity of the coefficients a_i s and b_i s (Proposition C.3). Hence, if linear densities are explicitly given as input, then we have a polynomial (in the input size) runtime bound.

are induced by irrational cuts (Appendix G). That is, in general, one cannot expect an efficient algorithm that outputs an exact envy-free (or welfare-maximizing) allocation. Therefore, when considering efficient algorithms for cake division, a precision loss is inevitable. However, our algorithms ensure that this loss is arbitrarily small. Specifically, given a cake-division instance C with MLRP, we can find, in time that is polynomial in $\log(1/\eta)$ (along with n and $\log \lambda$), an allocation $\mathcal{I} = \{I_1, \dots, I_n\}$ such that, $v_i(I_i) \geq v_i(I_j) - \eta$ for all agents $i, j \in [n]$. Since the precision parameter η can be driven exponentially close to zero in polynomial (in the bit complexity of η) time, we will say that an envy-free allocation can be computed up to an arbitrary precision.

Similarly, in the context of maximizing welfare (social or egalitarian), given a cake-division instance wherein the value densities bear MLRP, we can find—in time that is polynomial in $\log(1/\eta)$ —an allocation with the (social or egalitarian) welfare η (additively) close to the optimal. Hence, as in the case of envy-freeness, we assert that a welfare-maximizing allocation can be computed efficiently, up to an arbitrary precision.

3 MAIN RESULTS

This section presents the statements of our key results.

Envy-Freeness: In Section 5 we prove that for cake-division instances, in which the value densities satisfy MLRP, the problem of finding an envy-free allocation essentially admits a polynomial-time algorithm.

THEOREM 3.1. *Let $C = \langle [n], \{f_i\}_{i \in [n]} \rangle$ be a cake-division instance in which the value-density functions satisfy the monotone likelihood ratio property. Then, in the Robertson-Webb query model, an envy-free allocation of C can be computed (up to an arbitrary precision) in $O(\text{poly}(n, \log \lambda))$ time; here $\lambda \in \mathbb{R}_+$ is the Lipschitz constant of the cut and eval queries.*

Recall that, under MLRP, the cut and eval queries are necessarily λ -Lipschitz (Proposition C.3). **Pareto Optimality:** Weller's theorem [36] asserts that there always exists some cake division (though, not necessarily with connected pieces) which is both envy-free and Pareto optimal (among all cake divisions, with or without connected pieces). We show that, in the context of MLRP, every envy-free allocation is in fact Pareto optimal (among all cake divisions). Therefore, for cake-division instances with MLRP, the allocation computed by our algorithm is not only envy-free (fair) but also Pareto optimal, up to an arbitrary precision.

THEOREM 3.2. *Let C be a cake-division instance wherein the value-density functions satisfy the monotone likelihood ratio property. Then, every envy-free allocation in C is also Pareto optimal (over the set of all cake divisions).*

Social Welfare: In Section 7 we show that, up to an arbitrary precision, a social welfare maximizing allocation can be computed efficiently under MLRP.

THEOREM 3.3. *Let $C = \langle [n], \{f_i\}_{i \in [n]} \rangle$ be a cake-division instance in which the value-density functions satisfy the monotone likelihood ratio property. Then, in the Robertson-Webb query model, an allocation that achieves the optimal social welfare in C can be computed (up to an arbitrary precision) in $O(\text{poly}(n, \log \lambda))$ time; here $\lambda \in \mathbb{R}_+$ is the Lipschitz constant of the cut and eval queries.*

Egalitarian Welfare: Appendix A addresses the problem of maximizing egalitarian welfare. Specifically, we prove that, in cake-division instances with MLRP, an allocation with egalitarian welfare arbitrarily close to the optimal can be computed efficiently.

THEOREM 3.4. *Let $C = \langle [n], \{f_i\}_{i \in [n]} \rangle$ be a cake-division instance in which the value-density functions satisfy the monotone likelihood ratio property. Then, in the Robertson-Webb query model,*

an allocation that achieves the optimal egalitarian welfare in C can be computed (up to an arbitrary precision) in $O(\text{poly}(n, \log \lambda))$ time; here $\lambda \in \mathbb{R}_+$ is the Lipschitz constant of the cut and eval queries.

Nash Social Welfare: In Appendix B we show that, under MLRP, the problem of maximizing Nash social welfare admits a fully polynomial-time approximation scheme (FPTAS).

THEOREM 3.5. *Let $C = \langle [n], \{f_i\}_{i \in [n]} \rangle$ be a cake-division instance in which the value-density functions satisfy the monotone likelihood ratio property. Then, in the Robertson-Webb query model, an allocation with Nash social welfare at least $(1 - \varepsilon)$ times the optimal (Nash social welfare) in C can be computed in time that is polynomial in $1/\varepsilon$, $\log \lambda$, and n ; here $\lambda \in \mathbb{R}_+$ is the Lipschitz constant of the cut and eval queries.*

4 CHARACTERIZATIONS OF THE MONOTONE LIKELIHOOD RATIO PROPERTY

This section provides useful characterizations of MLRP (Lemma 4.1). This result will be used in subsequent sections towards the analysis of our algorithms.

LEMMA 4.1. *Let f_i and f_j be two value-density functions. Then, the following properties are equivalent*

- (i) *The likelihood ratios satisfy $\frac{f_j(b)}{f_i(b)} \leq \frac{f_j(c)}{f_i(c)}$ for all $0 \leq b \leq c \leq 1$.*
- (ii) *The values of the intervals satisfy $\frac{\int_a^b f_j}{\int_a^b f_i} \leq \frac{\int_c^d f_j}{\int_c^d f_i}$ for all $[a, b], [c, d] \subseteq [0, 1]$ with $b \leq c$.*
- (iii) *The (normalized) values satisfy $\frac{\int_a^b f_i}{\int_a^b f_i} \leq \frac{\int_a^b f_j}{\int_a^b f_j}$ for all intervals $[a, b] \subseteq [0, 1]$ and all $x \in [a, b]$.*

That is, the value densities f_i and f_j bear MLRP iff they satisfy property (ii) or (iii).

Here, if the likelihood ratio $\frac{f_j(x)}{f_i(x)}$ is strictly increasing, then we have a strict inequality in the corresponding characterizations.

The proof of this lemma is delegated to Appendix C.1. Note that, in terms of the agents' valuations v_i and v_j , property (ii) in Lemma 4.1 can be expressed as $\frac{v_j(a,b)}{v_i(a,b)} \leq \frac{v_j(c,d)}{v_i(c,d)}$ for all $[a, b], [c, d] \subseteq [0, 1]$ with $b \leq c$. Similarly, property (iii) corresponds to $\frac{v_i(x,b)}{v_i(a,b)} \leq \frac{v_j(x,b)}{v_j(a,b)}$ for all $[a, b] \subseteq [0, 1]$ and all $x \in [a, b]$.

It is well-known that MLRP implies first-order stochastic dominance (see Lemma C.1). Interestingly, property (iii) provides a strengthening in terms of the following characterization: over any interval $[a, b]$, the normalized (by $v_j(a, b)$) values of agent j stochastically dominate the normalized (by $v_i(a, b)$) values of agent i .

5 ENVY-FREENESS

In this section we develop an efficient algorithm for finding envy-free allocations in cake division instances with MLRP (Theorem 3.1). Towards this goal, we define a class of cake divisions, referred to as *ripple divisions* (Definition 5.1), and prove that, under MLRP, every ripple division induces an envy-free allocation (Theorem 5.7). Existential and computational guarantees for ripple divisions are established in Section 5.1 (Lemma 5.5) and Section 5.2 (Lemma 5.6), respectively. Section 5.4 builds upon these results to prove our main result (Theorem 3.1) for envy-freeness.

We establish the universal existence of ripple divisions through the intermediate value theorem, i.e., a one-dimensional fixed-point argument (Lemma 5.5). Consequentially, for cake-division instances with MLRP, we develop an alternate proof of existence of envy-free allocations. Since one can use binary search to find fixed points in the one-dimensional setting, this proof in fact leads to an algorithm for finding ripple divisions and, hence, envy-free divisions.

Definition 5.1 (Ripple Division). Given a cake-division instance $C = \langle [n], \{f_i\}_i \rangle$, a set of points $x_0^* = 0 \leq x_1^* \leq \dots \leq x_{n-1}^* \leq x_n^* = 1$ is said to form a ripple division of the cake iff

$$v_i(x_{i-1}^*, x_i^*) = v_i(x_i^*, x_{i+1}^*) > 0 \text{ for all agents } i \in [n-1].$$

For establishing existence of ripple divisions, we first consider a relaxation of Definition 5.1 wherein do not enforce the last cut point (i.e., x_n^*) to be equal to one. Under this relaxation, the intervals $\{[x_{i-1}, x_i]\}_{i=1}^n$ do not cover the entire interval $[0, 1]$ (instead, they span $[0, x_n]$) and, hence, lead to a partial allocation of the cake.

Also, by convention, the agents are indexed following the MLRP order: for each $i \in [n-1]$, the likelihood ratio f_{i+1}/f_i is nondecreasing. Hence, assigning interval $[x_{i-1}, x_i]$ to agent $i \in [n]$ provides an allocation wherein the intervals are assigned (left to right on the cake) in accordance with the MLRP order. We will show that the (partial) allocation obtained by assigning interval $[x_{i-1}, x_i]$ to agent $i \in [n]$ is envy-free (Theorem 5.7).

Definition 5.2 (δ -Ripple Division). Given a cake-division instance $C = \langle [n], \{f_i\}_i \rangle$, a set of points $x_0 = 0 \leq x_1 \leq \dots \leq x_{n-1} \leq x_n \leq 1$ is said to form a δ -ripple division of the cake iff $x_n \geq 1 - \delta$ and

$$v_i(x_{i-1}, x_i) = v_i(x_i, x_{i+1}) > 0 \text{ for all agents } i \in [n-1].$$

Both Definitions 5.1 and 5.2 require that, for all $i \in [n-1]$, agent i 's value for the i th interval ($[x_{i-1}^*, x_i^*]$ and $[x_{i-1}, x_i]$, respectively) is positive. Also, note that a 0-ripple division is an exact ripple division.

To compose the value equalities that define a δ -ripple division (Definition 5.2), we consider $(n-1)$ functions $\text{RD}_i : [0, 1] \mapsto [0, 1]$, for $2 \leq i \leq n$. In particular, focusing on the equalities considered in Definition 5.2 (i.e., $v_i(x_{i-1}, x_i) = v_i(x_i, x_{i+1})$), we note that if we set the first cut point $x_1 = x \in [0, 1]$, then all the subsequent points $x_2, \dots, x_n \in [0, 1]$ are fixed as well. In particular, x_2 is the point that satisfies $v_1(0, x) = v_1(x, x_2)$ and, iteratively, x_{i+1} can be identified from $v_i(x_{i-1}, x_i) = v_i(x_i, x_{i+1})$. The functions RD_i s capture this ‘‘ripple’’ effect and can be expressed as compositions of cut and eval queries. Formally, the functions $\text{RD}_i : [0, 1] \mapsto [0, 1]$, for $i \in \{2, 3, \dots, n\}$, are recursively defined as follows:¹¹ $\text{RD}_2(x) := \text{Cut}_1(x, \text{Eval}_1(0, x))$, and

$$\text{RD}_i(x) := \text{Cut}_{i-1}\left(\text{RD}_{i-1}(x), \text{Eval}_{i-1}\left(\text{RD}_{i-2}(x), \text{RD}_{i-1}(x)\right)\right) \text{ for } i \in \{3, 4, \dots, n\} \quad (1)$$

In particular, $\text{RD}_n(x)$ denotes the value of the last cut point x_n obtained by initializing the ripple effect with $x_1 = x$. Also, note that $\text{RD}_n(0) = 0$ and $\text{RD}_n(1) = 1$; by convention, the response to a cut query $\text{Cut}_i(\ell, \tau)$ is truncated to 1 iff τ is greater than the entire value to the right of ℓ , i.e., if $v_i(\ell, 1) < \tau$.

Since RD_i s can be expressed as a composition of the cut and eval queries, these functions can be efficiently computed in the Robertson-Webb query model. Moreover, using the fact that the cut and eval queries are λ -Lipschitz, in the following proposition we establish that RD_n is also Lipschitz continuous.

PROPOSITION 5.3. *If in a cake-division instance $C = \langle [n], \{f_i\}_i \rangle$ the cut and eval queries are λ -Lipschitz, then the function RD_n is $\lambda^{2(n-1)}$ -Lipschitz.*

The proof of Proposition 5.3 is deferred to Appendix D. The Lipschitz continuity of RD_n turns out to be a key property, both for establishing the existence of ripple divisions and in developing an efficient algorithm for finding them. Specifically, we can apply the intermediate value theorem to RD_n and prove that, for any $\delta > 0$, there exists a δ -ripple division; see Lemma 5.4 below. The

¹¹The first cut point x is specified upfront and, hence, we do not require RD_1 . The functions RD_i s are defined for $i \in \{2, \dots, n\}$.

universal existence of δ -ripple divisions, for all $\delta > 0$, along with a limit argument ($\delta \rightarrow 0$), establishes our existential result (Lemma 5.5) for ripple divisions, i.e., for 0-ripple divisions.

The Lipschitzness of RD_n also ensures that, via binary search, we can find a δ -ripple division (via a point $x_1 \in (0, 1)$ that satisfies $\text{RD}_n(x_1) \geq 1 - \delta$), in time that is polynomial in $\log(\frac{1}{\delta})$. This runtime dependence ensures that the precision parameter δ can be driven exponentially close to zero, in polynomial (in the bit complexity of δ) time; see Lemma 5.6.

5.1 Existence of Ripple Divisions

LEMMA 5.4. *Let $C = \langle [n], \{f_i\}_i \rangle$ be a cake-division instance, with λ -Lipschitz cut and eval queries, and let parameter $\delta \in (0, 1)$. Then, in C , there always exist a point $\widehat{x} \in (0, 1)$ with the property that $\text{RD}_n(\widehat{x}) = 1 - \delta$.*

Furthermore, such a point \widehat{x} initializes a δ -ripple division: the points $x_0 = 0, x_1 = \widehat{x}$, and $x_i = \text{RD}_i(\widehat{x})$, for $2 \leq i \leq n$, form a δ -ripple division in C . Here, the functions $\text{RD}_2, \dots, \text{RD}_n$ are defined with respect to the cut and eval queries of C .

PROOF. The function RD_n is continuous (Proposition 5.3) and it satisfies $\text{RD}_n(0) = 0$ along with $\text{RD}_n(1) = 1$. Hence, the intermediate value theorem ensures that, for any $\delta \in (0, 1)$, there must exist a point $\widehat{x} \in (0, 1)$ which satisfies $\text{RD}_n(\widehat{x}) = 1 - \delta$. This shows that the required point \widehat{x} always exists.

We will complete the proof by establishing that the points $x_0 = 0, x_1 = \widehat{x}$, and $x_i = \text{RD}_i(\widehat{x})$, for $2 \leq i \leq n$, form a δ -ripple division. Note that, since $\delta > 0$, we have $\text{RD}_n(\widehat{x}) < 1$. This strict inequality implies that the intermediate points x_i s were *not* truncated to one. Indeed, in this case, all of the value equalities in Definition 5.2 hold, i.e., $v_i(x_{i-1}, x_i) = v_i(x_i, x_{i+1})$ for all $i \in [n - 1]$. It remains to show that these values are positive.

The bound $\delta < 1$ gives us $\text{RD}_n(\widehat{x}) > 0 = \text{RD}_n(0)$. Therefore, we have $\widehat{x} > 0$, i.e., $x_1 = \widehat{x} > 0 = x_0$. In the current setting, all nonempty intervals have positive value for the agents. This follows from the Lipschitz continuity of the cut queries.¹² Since the interval $[0, x_1]$ is nonempty, agent 1 has a positive value for it, $v_1(0, x_1) > 0$. Furthermore, the value equality $v_1(x_1, x_2) = v_1(0, x_1)$ gives us $v_1(x_1, x_2) > 0$, i.e., $x_2 > x_1$. Extending this argument iteratively shows that $x_0 < x_1 < \dots < x_n < 1$. In other words, each agent receives a positive value under the cut points, i.e., $v_i(x_{i-1}, x_i) = v_i(x_i, x_{i+1}) > 0$ for all $i \in [n - 1]$. Hence, x_i s form a δ -ripple division. \square

Next we use Lemma 5.4 and a limit argument ($\delta \rightarrow 0$) to establish universal existence of ripple divisions. The proof of Lemma 5.5 is delegated to Appendix D.2.

LEMMA 5.5. *Let C be a cake-division instance in which the cut and eval queries are λ -Lipschitz. Then, C necessarily admits a ripple division.*

5.2 Computation of Ripple Divisions

In this section we present an efficient algorithm for computing δ -ripple divisions. Lemma 5.4 implies that the problem of computing a δ -ripple division reduces to finding a point $x \in (0, 1)$ that satisfies $\text{RD}_n(x) \in [1 - \delta, 1)$. The algorithm `BINSEARCH` (Algorithm 1) finds such a point x (and, hence, a δ -ripple division) via binary search. It is well-known that binary search can be used to compute fixed points in the one-dimensional setting. We provide the relevant details here for completeness.

Specifically, in `BINSEARCH`, we initialize $\ell = 0$ along with $r = 1$ and keep bisecting the interval $[\ell, r]$ as long as $\text{RD}_n(\ell) < 1 - \delta$ and $\text{RD}_n(r) = 1$. Recall that $\text{RD}_n(0) = 0$ and $\text{RD}_n(1) = 1$.

Since RD_n is Lipschitz continuous (Proposition 5.3), the intermediate value theorem (applied on $[\ell, r]$ with the bounds $\text{RD}_n(\ell) < 1 - \delta$ and $\text{RD}_n(r) = 1$) guarantees that in each considered

¹²In fact, under MLRP, we explicitly have $f_i(x) > 0$ for all $x \in [0, 1]$.

interval $[\ell, r]$ there exists a point $x \in (\ell, r)$ which satisfies $\text{RD}_n(x) = 1 - \delta$.¹³ In each iteration of this algorithm the length of the considered interval $[\ell, r]$ reduces by a multiplicative factor of two and, hence, after an appropriate number of iterations, the required point x and the midpoint of the interval $(\ell + r)/2$ get sufficiently close. In such a case, one can show (using the Lipschitz continuity of RD_n) that the midpoint $(\ell + r)/2$ itself initializes a 2δ -ripple division. Appendix D.3 formalizes this runtime analysis and provides a proof of the following lemma.

ALGORITHM 1: BINSEARCH

Input: A cake-division instance $C = \langle [n], \{f_i\}_i \rangle$, in the Robertson-Webb query model, and parameter $\delta > 0$.

Output: A δ -ripple division $0 = x_0 \leq x_1 \leq x_2, \dots \leq x_n \leq 1$.

```

1: Initialize  $\ell = 0$  and  $r = 1$ 
2: while  $\ell < r$  do
3:   if  $\text{RD}_n\left(\frac{\ell+r}{2}\right) < 1 - \delta$  then
4:     Update  $\ell \leftarrow (\ell + r)/2$ 
5:   else if  $\text{RD}_n\left(\frac{\ell+r}{2}\right) = 1$  then
6:     Update  $r \leftarrow (\ell + r)/2$ 
7:   else if  $\text{RD}_n\left(\frac{\ell+r}{2}\right) \in [1 - \delta, 1)$  then
8:     Set  $x_0 = 0$ ,  $x_1 = \frac{\ell+r}{2}$ , and  $x_i = \text{RD}_i(x_1)$  for all  $i \in \{2, \dots, n\}$ 
9:     return the collection of points  $x_0, x_1, \dots, x_n$ 
10:  end if
11: end while
    
```

LEMMA 5.6. *Let $C = \langle [n], \{f_i\}_i \rangle$ be a cake-division instance in which the cut and eval queries are λ -Lipschitz. Then, for any $\delta \in (0, 1)$ and in the Robertson-Webb query model, a δ -ripple division of C can be computed in $\mathcal{O}\left(\text{poly}(n, \log \lambda, \log \frac{1}{\delta})\right)$ time.*

5.3 From Ripple Divisions to Envy-Free Allocations

The next theorem establishes the crucial connection between ripple divisions and envy-freeness. In particular, setting $\delta = 0$ in this theorem, we obtain that, under MLRP, every (exact) ripple division induces an envy-free allocation.

THEOREM 5.7. *Let $C = \langle [n], \{f_i\}_i \rangle$ be a cake-division instance in which the value densities satisfy the monotone likelihood ratio property and let parameter $\delta \geq 0$. Then, every δ -ripple division, $x_0 = 0 \leq x_1 \leq \dots \leq x_{n-1} \leq x_n \leq 1$, in C induces an envy-free partial allocation $\{I_i = [x_{i-1}, x_i]\}_{i=1}^n$.*

This theorem asserts that the partial allocation $\mathcal{I} = \{I_1, \dots, I_n\}$ (with $I_i = [x_{i-1}, x_i]$) satisfies $v_i(I_i) \geq v_i(I_j)$, for all $i, j \in [n]$, and at most a δ -length piece of the cake (specifically, $[x_n, 1]$) remains unallocated in \mathcal{I} .

PROOF. We will show that if the points $x_0 = 0 \leq x_1 \leq x_2 \leq \dots \leq x_n \leq 1$ form a δ -ripple division, then the partial allocation $\{I_i = [x_{i-1}, x_i]\}_{i=1}^n$ is envy-free. Here, the definition of a δ -ripple division (Definition 5.2) ensures that, for each agent $i \in [n - 1]$, the values the two consecutive intervals I_i and I_{i+1} are equal and positive

$$v_i(I_i) = v_i(x_{i-1}, x_i) = v_i(x_i, x_{i+1}) = v_i(I_{i+1}) > 0 \quad (2)$$

¹³Note that this argument (in particular, the use of intermediate value theorem) does not require RD_n to be monotonic. Also, recall that RD_i s can be computed efficiently in the Robertson-Webb query model. Hence, evaluating RD_i s in BINSEARCH leads to at most a polynomial overhead in its runtime.

Recall that the agents are indexed following the MLRP order, i.e., for each $i \in [n - 1]$, the likelihood ratio f_{i+1}/f_i is nondecreasing. We fix an agent $i \in [n]$ and establish envy-freeness with respect to i by considering two complementary cases (i) for agents to the left of i , we prove that $v_i(I_1) \leq v_i(I_2) \leq \dots \leq v_i(I_i)$ and (ii) for agents to the right of i , we prove that $v_i(I_i) \geq v_i(I_{i+1}) \geq \dots \geq v_i(I_n)$.

Case (i): Consider any agent $k \in [n]$ such that $k < i$. Given that f_k and f_i bear MLRP (i.e., f_i/f_k is non-decreasing), property (ii) of Lemma 4.1, with $a = x_{k-1}$, $b = c = x_k$, and $d = x_{k+1}$, gives us $\frac{v_i(x_{k-1}, x_k)}{v_k(x_{k-1}, x_k)} \leq \frac{v_i(x_k, x_{k+1})}{v_k(x_k, x_{k+1})}$. That is, for the intervals $I_k = [x_{k-1}, x_k]$ and $I_{k+1} = [x_k, x_{k+1}]$ we have

$$\frac{v_i(I_k)}{v_k(I_k)} \leq \frac{v_i(I_{k+1})}{v_k(I_{k+1})} \quad (3)$$

Instantiating equation (2) for agent k , we can simplify inequality (3) to $v_i(I_k) \leq v_i(I_{k+1})$. Combining this inequality across all $k < i$, we obtain the desired chain of inequalities for agent i , i.e., $v_i(I_1) \leq v_i(I_2) \leq \dots \leq v_i(I_i)$.

Case (ii): Consider any agent $j \in [n]$ such that $j > i$. Given that f_i and f_j bear MLRP (i.e., f_j/f_i is non-decreasing), property (ii) of Lemma 4.1, with $a = x_{j-1}$, $b = c = x_j$, and $d = x_{j+1}$, gives us $\frac{v_j(x_{j-1}, x_j)}{v_i(x_{j-1}, x_j)} \leq \frac{v_j(x_j, x_{j+1})}{v_i(x_j, x_{j+1})}$. That is, for the intervals $I_j = [x_{j-1}, x_j]$ and $I_{j+1} = [x_j, x_{j+1}]$ we have

$$\frac{v_j(I_j)}{v_i(I_j)} \leq \frac{v_j(I_{j+1})}{v_i(I_{j+1})} \quad (4)$$

Instantiating equation (2) for agent j , we can simplify inequality (4) to $v_i(I_j) \geq v_i(I_{j+1})$. Combining this inequality across all $j > i$, we obtain the desired chain of inequalities for agent i , i.e., $v_i(I_i) \geq v_i(I_{i+1}) \geq \dots \geq v_i(I_n)$.

The above two cases establish that agent $i \in [n]$ does not envy any other other agent, i.e., $\mathcal{I} = \{I_1, \dots, I_n\}$ is an envy-free partial allocation. Indeed, if $\delta = 0$, then \mathcal{I} covers the entire cake, i.e., we obtain an envy-free allocation. \square

Notably, Lemma 5.5 and Theorem 5.7 (with $\delta = 0$) provide a stand-alone proof of existence of envy-free allocations in cake-division instances with MLRP. The next section establishes an algorithmic counterpart of this existential result; specifically, we show that using Lemma 5.6 one can directly obtain an efficient algorithm for finding envy-free allocations, under MLRP.

5.4 Proof of Theorem 3.1

This section restates and proves our main result (Theorem 3.1) for envy-freeness.

THEOREM 3.1. *Let $C = \langle [n], \{f_i\}_{i \in [n]} \rangle$ be a cake-division instance in which the value-density functions satisfy the monotone likelihood ratio property. Then, in the Robertson-Webb query model, an envy-free allocation of C can be computed (up to an arbitrary precision) in $O(\text{poly}(n, \log \lambda))$ time; here $\lambda \in \mathbb{R}_+$ is the Lipschitz constant of the cut and eval queries.*

PROOF. Given a cake-division instance C , with MLRP, and precision parameter $\eta > 0$, we invoke Lemma 5.6 to find an $(\frac{\eta}{\lambda})$ -ripple division in $O\left(\text{poly}(n, \log \lambda, \log \frac{1}{\eta})\right)$ time. Write $x_0 = 0 \leq x_1 \leq \dots \leq x_{n-1} \leq x_n \leq 1$ to denote the computed $(\frac{\eta}{\lambda})$ -ripple division and let $\mathcal{I} = \{I_1, \dots, I_n\}$ be the corresponding partial allocation; here $I_i = [x_{i-1}, x_i]$. Theorem 5.7 ensures that \mathcal{I} is envy-free.

We will show that coalescing the unassigned (in \mathcal{I}) piece $[x_n, 1]$ to agent n provides a complete allocation that satisfies envy-freeness, up to η precision. Write $\mathcal{I}^* := \{I_1^*, I_2^*, \dots, I_{n-1}^*, I_n^*\}$ to denote this allocation in which the n th agent receives the interval $I_n^* := [x_{n-1}, 1]$ (equivalently, $I_n^* = I_n \cup [x_n, 1]$) and $I_i^* = I_i$ for the remaining agents $i \in [n - 1]$.

Note that against all agents $j \in [n-1]$, envy-freeness of \mathcal{I}^* directly follows from the fact that the partial allocation \mathcal{I} is envy-free: $v_i(I_i^*) \geq v_i(I_i) \geq v_i(I_j) = v_i(I_j^*)$, for all $i \in [n]$ and all $j \in [n-1]$.

Finally, we address envy against agent n . Recall that x_i s form a $(\frac{\eta}{\lambda})$ -ripple division, hence $x_n \geq 1 - \frac{\eta}{\lambda}$. In addition, the fact that eval queries are λ -Lipschitz gives us $v_i([x_n, 1]) \leq \eta$ for all agents $i \in [n]$. Hence, for all $i \in [n]$ we have $v_i(I_i^*) \geq v_i(I_i) \geq v_i(I_n)$. The last inequality follows from envy-freeness of \mathcal{I} . Given that $I_n^* = I_n \cup [x_n, 1]$, we have $v_i(I_i^*) \geq v_i(I_n^*) - v_i([x_n, 1])$. Finally, since $v_i([x_n, 1]) \leq \eta$, the desired bound holds $v_i(I_i^*) \geq v_i(I_n^*) - \eta$.

Therefore, \mathcal{I}^* is envy-free, up to η precision: $v_i(I_i^*) \geq v_i(I_j^*) - \eta$ for all $i, j \in [n]$. The time complexity obtained via Lemma 5.6 implies that the parameter η can be driven exponentially close to zero, in time that is polynomial in $\log \frac{1}{\eta}$ (i.e., in bit complexity of η). Hence, we can find an envy-free allocation, up to arbitrary precision, in $\mathcal{O}(\text{poly}(n, \log \lambda))$ time. Theorem 3.1 now stands proved. \square

6 PARETO OPTIMALITY

This section shows that, with MLRP in hand, one does not lose out on Pareto optimality by imposing the contiguity requirement. That is, under MLRP, there always exist allocations (i.e., cake divisions with connected pieces) that are Pareto optimal among all cake divisions, with or without connected pieces. Moreover, such allocations conform to the MLRP order.

This structural result implies that for maximizing welfare we can restrict attention to allocations wherein the intervals are assigned (left to right on the cake) in accordance with the MLRP order. Intuitively, this leads us to a welfare-maximizing algorithm—specifically, a dynamic program—that recursively finds optimal allocations for intervals placed at the left end of the cake (i.e., for intervals of the form $[0, x]$); see Section 7 and Appendix B for details.

Subsequently, Section 6.1 (Theorem 3.2) establishes a strong connection between fairness and (Pareto) efficiency in the MLRP context: if the value densities bear MLRP, then every envy-free allocation is necessarily Pareto optimal.

LEMMA 6.1. *Let $C = \langle [n], \{f_i\}_i \rangle$ be a cake-division instance in which the value densities satisfy the monotone likelihood ratio property. Then, for every cake division $\mathcal{D} = \{D_1, \dots, D_n\}$ in C there exists an allocation $\mathcal{J} = \{J_1, \dots, J_n\}$ such that $v_i(J_i) \geq v_i(D_i)$, for $1 \leq i \leq n$.*

Furthermore, every Pareto optimal allocation $\mathcal{I} = \{I_1, \dots, I_n\}$ in C conforms to the MLRP order, i.e., if f_{i+1}/f_i is non-decreasing in $[0, 1]$, then the interval assigned to agent i (i.e., I_i) appears to the left of the interval assigned to the agent $i+1$ (i.e., I_{i+1}).

PROOF. Consider a cake division $\mathcal{D} = \{D_1, \dots, D_n\}$ wherein two consecutive intervals are assigned violating the MLRP order: say, interval $[p, q]$ is assigned to agent j (i.e., this interval is contained in the bundle D_j), the adjacent interval $[q, r]$ is assigned to agent i , and agent i appears before j in the MLRP order (f_j/f_i is nondecreasing over $[0, 1]$).

We will show that in such a case there always exists a point $q' \in [p, r]$ such that $v_i(p, q') \geq v_i(q, r)$ and $v_j(q', r) \geq v_j(p, q)$. That is, one can swap the allocation order between i and j (in the interval $[p, q] \cup [q, r]$) without decreasing the agents' values. Moreover, we note that, if f_j/f_i is strictly increasing in the interval $[p, q] \cup [q, r]$, then this update leads to a strict increase in agent i 's or agent j 's value.

Hence, starting with any cake division $\mathcal{D} = \{D_1, \dots, D_n\}$, we can repeatedly apply the above-mentioned resolution towards the MLRP order and obtain an allocation $\mathcal{J} = \{J_1, \dots, J_n\}$ with the desired property, $v_i(J_i) \geq v_i(D_i)$ for all $i \in [n]$. Note that this resolution process also establishes the second part of the theorem, i.e., it establishes the claim that every Pareto optimal allocation conforms to the MLRP order.

The remainder of the proof addresses the desired point $q' \in [p, r]$. In particular, we will identify q' such that assigning interval $[p, q']$ to agent i (instead of $[q, r]$) and assigning $[q', r]$ to agent j (instead of $[p, q]$) leads to an increment in values.

Write $\beta_i := \frac{v_i(q, r)}{v_i(p, r)}$ to denote the normalized value of agent i under the initial assignment. Define $q' \in [p, r]$ to be the point that satisfies

$$\frac{v_j(p, q')}{v_j(p, r)} = \beta_i \quad (5)$$

Since the value densities satisfy MLRP, property (iii) of Lemma 4.1, applied to the interval $[p, r]$ and $q' \in [p, r]$, gives us $\frac{v_i(q', r)}{v_i(p, r)} \leq \frac{v_j(q', r)}{v_j(p, r)}$. Simplifying further we obtain $1 - \frac{v_j(q', r)}{v_j(p, r)} \leq 1 - \frac{v_i(q', r)}{v_i(p, r)}$,

$$\frac{v_j(p, q')}{v_j(p, r)} \leq \frac{v_i(p, q')}{v_i(p, r)} \quad (6)$$

Therefore, we obtain a value bound for agent i

$$\frac{v_i(p, q')}{v_i(p, r)} \geq \frac{v_j(p, q')}{v_j(p, r)} = \beta_i = \frac{v_i(q, r)}{v_i(p, r)} \quad (\text{via equations (6), (5) and the definition of } \beta_i)$$

That is, agent i 's value is preserved through the reassignment, $v_i(p, q') \geq v_i(q, r)$.

For agent j , via property (iii) of Lemma 4.1, with interval $[p, r]$ and $q \in [p, r]$, we have

$$\frac{v_j(q, r)}{v_j(p, r)} \geq \frac{v_i(q, r)}{v_i(p, r)} = \beta_i = \frac{v_j(p, q')}{v_j(p, r)} \quad (\text{by the definition of } \beta_i \text{ and equation (5)})$$

This inequality reduces to $1 - \frac{v_j(p, q')}{v_j(p, r)} \geq 1 - \frac{v_j(q, r)}{v_j(p, r)}$. Simplifying we obtain $\frac{v_j(p, r) - v_j(p, q')}{v_j(p, r)} \geq \frac{v_j(p, r) - v_j(q, r)}{v_j(p, r)}$. Therefore, we have the desired inequality $v_j(q', r) \geq v_j(p, q)$ and the stated claims follow. \square

6.1 Extending Weller's Theorem

Weller's theorem [36] is a notable result in the cake-cutting literature and it asserts that there always exists some cake division—though, not necessarily with connected pieces—which is both envy-free (fair) and Pareto optimal. While this theorem holds in general, it does not guarantee that envy-freeness and Pareto optimality can be achieved together with allocations. Indeed, there are cake-division instances wherein none of the of envy-free allocations are Pareto optimal.

We show that, by contrast under MLRP, every envy-free allocation is Pareto optimal, among all cake divisions (Theorem 3.2). Therefore, given a cake-division instance with MLRP, the allocation computed by our algorithm (Algorithm 1) is not only envy-free but also Pareto optimal, up to an arbitrary precision.

THEOREM 3.2. *Let C be a cake-division instance wherein the value-density functions satisfy the monotone likelihood ratio property. Then, every envy-free allocation in C is also Pareto optimal (over the set of all cake divisions).*

PROOF. Write $\mathcal{I} = \{I_1, I_2, \dots, I_n\}$ to denote an envy-free allocation in C ; here interval I_i is assigned to agent $i \in [n]$. We assume, towards a contradiction, that there exists a cake division $\mathcal{D} = \{D_1, \dots, D_n\}$ that Pareto dominates \mathcal{I} . Lemma 6.1 implies that in such a case there exists an allocation $\mathcal{J} = \{J_1, \dots, J_n\}$ which also Pareto dominates \mathcal{I} . That is, we have $v_i(J_i) \geq v_i(D_i) \geq v_i(I_i)$, for all agents $i \in [n]$, and there exists some agent $k \in [n]$ such that $v_k(J_k) \geq v_k(D_k) > v_k(I_k)$.

Recall that for an allocation the endpoints of all the constituent intervals are referred to as its cut points. We break our analysis into the following two cases depending on whether \mathcal{I} and \mathcal{J} have the same set of cut points.

Case 1: The cut points of the allocations \mathcal{I} and \mathcal{J} are identical. In this case, there must exist a permutation $\sigma : [n] \mapsto [n]$ such that $I_{\sigma(i)} = J_i$ for all $i \in [n]$. Since \mathcal{I} is envy-free, we have $v_i(I_i) \geq v_i(I_{\sigma(i)}) = v_i(J_i)$ for all agents $i \in [n]$. However, this contradicts the fact that \mathcal{J} Pareto dominates the allocation \mathcal{I} .

Case 2: The cut points of \mathcal{I} and \mathcal{J} are not identical. Since both the allocations form a partition of the same cake $[0, 1]$, there must exist some $s, t \in [n]$ such that the interval J_s is a strict subset of the interval I_t , i.e., $J_s \subset I_t$. Envy-freeness of \mathcal{I} gives us $v_s(I_s) \geq v_s(I_t) > v_s(J_s)$. The last strict inequality follows from the fact that the value density f_s of agent s has full support over $[0, 1]$ and J_s is a strict subset of I_t . This bound $v_s(I_s) > v_s(J_s)$ contradicts the fact that \mathcal{J} Pareto dominates \mathcal{I} and completes the proof. \square

7 SOCIAL WELFARE

This section develops an algorithm for social welfare maximization. Recall that, under MLRP, Pareto optimal allocations conform to the MLRP order (Lemma 6.1). Hence, for maximizing social welfare we can restrict attention to allocations wherein the intervals are assigned (left-to-right on the cake) in accordance with the MLRP order. This observation, in and of itself, leads to a fully-polynomial time approximation scheme for the maximizing social welfare: we can partition the cake into $\text{poly}(n, 1/\epsilon)$ contiguous intervals, each of value at most ϵ , and then solve the problem using a dynamic program. We show that instead of considering a general partition we can identify a set P —of $\mathcal{O}(n^2)$ points—such that the cut points of an optimal allocation are contained in P . This will enable us to execute a dynamic program focusing only on the points in P and establish Theorem 3.3.

For value densities f_i and f_j , write L_{ij} to denote the set of points at which f_j is at least as large as f_i ; specifically, $L_{ij} := \{x \in [0, 1] : f_j(x) \geq f_i(x)\}$. Since that the densities f_i and f_j are normalized, there must exist a point $x \in [0, 1]$ such that $f_j(x) \geq f_i(x)$; in other words, $L_{ij} \neq \emptyset$. Also, observe that this set is bounded below by 0. Therefore, the greatest lower bound property of \mathbb{R} implies that L_{ij} admits an infimum. We will refer to this infimum as the *switching point* between the two value densities, $p_{ij} := \inf_{x \in L_{ij}} x$.¹⁴ For a cake-division instance $C = \langle [n], \{f_i\}_i \rangle$, we will write P to denote the collection of all switching points p_{ij} s, with $1 \leq i < j \leq n$ (along with 0 and 1), i.e., $P := \{p_{i,j} \in [0, 1] : 1 \leq i < j \leq n\} \cup \{0, 1\}$; we include the endpoints for cake in P for notational convenience. Also, recall the convention that the agents are indexed following the MLRP order.

The next lemma provides a useful property about the switching points in the MLRP context.

LEMMA 7.1. *Let f_i and f_j be two value-density functions that satisfy MLRP, i.e., f_j/f_i is non-decreasing over $[0, 1]$. Then,*

- (a) *For all $y \in [0, p_{i,j}]$, the likelihood ratio satisfies $f_j(y)/f_i(y) < 1$.*
- (b) *For all $z \in (p_{i,j}, 1]$, the likelihood ratio satisfies $f_j(z)/f_i(z) \geq 1$.*

Here, p_{ij} is the switching point between f_i and f_j .

The proof of Lemma 7.1 appears in Appendix E.1. The following corollary asserts that, up to an arbitrary precision, each switching point can be determined efficiently.

COROLLARY 7.2. *Let f_i and f_j be two value densities that bear MLRP and let parameter $\gamma \in (0, 1)$. Then, in the Robertson-Webb model, we can find an interval of length γ that contains the switching point p_{ij} in $\mathcal{O}(\log(1/\gamma))$ time.*

This corollary implies that we can efficiently compute the set of switching points $P = \{p_{i,j} \in [0, 1] : 1 \leq i < j \leq n\}$, up to an arbitrary precision. Next, Lemma 7.3 establishes the usefulness of P . The proofs of Corollary 7.2 and Lemma 7.3 are delegated to Appendix E.2 and E.3, respectively.

¹⁴As noted, p_{ij} exists and satisfies $p_{ij} \in [0, 1]$. However, since that the value densities f_i s are not necessarily continuous (they are only assumed to be Riemann integrable) a point x with the property that $f_i(x) = f_j(x)$ might not exist.

LEMMA 7.3. *Let C be a cake-division instance in which the value densities satisfy the monotone likelihood ratio property. Then, in C , there exists a social welfare maximizing allocation all of whose cut points belong to the set of switching points P .*

We now state the main result of this section and defer its proof to Appendix E.4.

THEOREM 3.3. *Let $C = \langle [n], \{f_i\}_{i \in [n]} \rangle$ be a cake-division instance in which the value-density functions satisfy the monotone likelihood ratio property. Then, in the Robertson-Webb query model, an allocation that achieves the optimal social welfare in C can be computed (up to an arbitrary precision) in $O(\text{poly}(n, \log \lambda))$ time; here $\lambda \in \mathbb{R}_+$ is the Lipschitz constant of the cut and eval queries.*

Remark: For cake-division instances with MLRP, Theorem 3.3 and Theorem 6.1 show that we can efficiently maximize social welfare among all cake divisions, and not necessarily among allocations.

8 CONCLUSION AND FUTURE WORK

The current work studies algorithmic aspects of contiguous cake division under the monotone likelihood ratio property. The scope of this property ensures that the developed algorithms are applicable in various cake-division settings. We also note that while under MLRP the value densities must have full support, the developed framework is somewhat robust to this requirement. For example, our results extend to the class of non-full-support value densities considered in [1]; see Appendix H for details. Generalizing such ideas to address, say, value densities that bear first-order stochastic dominance is an interesting direction of future work.

Another relevant notion of fairness in the context of cake cutting is that of a *perfect division* [2] wherein each of the n agents values every (potentially disconnected) piece at $1/n$. In contrast to the solution concepts considered in the current paper, perfect divisions are not guaranteed to exist under the contiguity requirement—a perfect *allocation* might not exist even with MLRP (Appendix I). However, Alon [2] has shown that a perfect division with $n(n-1)$ cuts always exists. Perfect cake divisions are particularly useful since they lead to truthful mechanisms for cake division [11, 24]. Hence, developing efficient algorithms to find (noncontiguous) perfect divisions under MLRP is a relevant thread for future work. More broadly, it would be interesting to identify tractable classes through MLRP in other computational social choice contexts, such as discrete fair division and voting.

ACKNOWLEDGMENTS

We thank Manjunath Krishnapur for helpful discussions and references. Siddharth Barman gratefully acknowledges the support of a Ramanujan Fellowship (SERB - SB/S2/RJN-128/2015) and a Pratiksha Trust Young Investigator Award. Nidhi Rathi's research is generously supported by an IBM PhD Fellowship.

REFERENCES

- [1] Reza Alijani, Majid Farhadi, Mohammad Ghodsi, Masoud Seddighin, and Ahmad S Tajik. 2017. Envy-free mechanisms with minimum number of cuts. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- [2] Noga Alon. 1987. Splitting necklaces. *Advances in Mathematics* 63, 3 (1987), 247–253.
- [3] Eshwar Ram Arunachaleswaran, Siddharth Barman, Rachitesh Kumar, and Nidhi Rathi. 2019. Fair and Efficient Cake Division with Connected Pieces. In *International Conference on Web and Internet Economics*. Springer, 57–70.
- [4] Yonatan Aumann, Yair Dombb, and Avinatan Hassidim. 2013. Computing socially-efficient cake divisions. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 343–350.
- [5] Haris Aziz and Simon Mackenzie. 2016. 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)*. IEEE, 416–427.
- [6] Xiaohui Bei, Ning Chen, Xia Hua, Biaoshuai Tao, and Endong Yang. 2012. Optimal proportional cake cutting with connected pieces. In *Twenty-sixth AAAI Conference on Artificial Intelligence*.

- [7] Steven J Brams and Alan D Taylor. 1995. An envy-free cake division protocol. *The American Mathematical Monthly* 102, 1 (1995), 9–18.
- [8] Steven J Brams and Alan D Taylor. 1996. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press.
- [9] George Casella and Roger L Berger. 2002. *Statistical inference*. Vol. 2. Duxbury Pacific Grove, CA.
- [10] Katarína Cechlárová and Eva Pillárová. 2012. On the computability of equitable divisions. *Discrete Optimization* 9, 4 (2012), 249–257.
- [11] Yiling Chen, John K Lai, David C Parkes, and Ariel D Procaccia. 2013. Truth, justice, and cake cutting. *Games and Economic Behavior* 77, 1 (2013), 284–297.
- [12] Yuga J Cohler, John K Lai, David C Parkes, and Ariel D Procaccia. 2011. Optimal envy-free cake cutting. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*.
- [13] Xiaotie Deng, Qi Qi, and Amin Saberi. 2012. Algorithmic solutions for envy-free cake cutting. *Operations Research* 60, 6 (2012), 1461–1476.
- [14] Lester E Dubins and Edwin H Spanier. 1961. How to cut a cake fairly. *The American Mathematical Monthly* 68, 1P1 (1961), 1–17.
- [15] Duncan K Foley. 1967. Resource allocation and the public sector. (1967).
- [16] George Gamow and Marvin Stern. 1958. Puzzle-Math.
- [17] Martin Grötschel, László Lovász, and Alexander Schrijver. 2012. *Geometric algorithms and combinatorial optimization*. Vol. 2. Springer Science & Business Media.
- [18] A Hollender, PW Goldberg, and W Suksompong. 2019. Contiguous cake cutting: Hardness results and approximation algorithms. (2019).
- [19] Ian Jewitt. 1991. Applications of likelihood ratio orderings in economics. *Lecture Notes-Monograph Series* (1991), 174–189.
- [20] Mamoru Kaneko and Kenjiro Nakamura. 1979. The Nash social welfare function. *Econometrica: Journal of the Econometric Society* (1979), 423–435.
- [21] David Kurokawa, John K Lai, and Ariel D Procaccia. 2013. How to cut a cake before the party ends. In *Twenty-Seventh AAAI Conference on Artificial Intelligence*.
- [22] Richard J Larsen, Morris L Marx, et al. 2001. *An introduction to mathematical statistics and its applications*. Vol. 5. Prentice Hall Upper Saddle River, NJ.
- [23] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM conference on Electronic commerce*. ACM, 125–131.
- [24] Elchanan Mossel and Omer Tamuz. 2010. Truthful fair division. In *International Symposium on Algorithmic Game Theory*. Springer, 288–299.
- [25] Hervé Moulin. 2004. *Fair division and collective welfare*. MIT press.
- [26] John F Nash Jr. 1950. The bargaining problem. *Econometrica: Journal of the Econometric Society* (1950), 155–162.
- [27] Ariel D Procaccia. 2015. Cake cutting algorithms. In *Handbook of Computational Social Choice, chapter 13*. Citeseer.
- [28] Jack Robertson and William Webb. 1998. *Cake-cutting algorithms: Be fair if you can*. AK Peters/CRC Press.
- [29] Adrien Saumard and Jon A Wellner. 2014. Log-concavity and strong log-concavity: a review. *Statistics surveys* 8 (2014), 45.
- [30] FW Simmons. 1980. Private communication to Michael Starbird. (1980).
- [31] Hugo Steinhaus. 1948. The Problem of Fair Division. *Econometrica* 16 (1948), 101–104.
- [32] Walter Stromquist. 1980. How to cut a cake fairly. *The American Mathematical Monthly* 87, 8 (1980), 640–644.
- [33] Walter Stromquist. 2008. Envy-free cake divisions cannot be found by finite protocols. *the electronic journal of combinatorics* 15, 1 (2008), 11.
- [34] Francis Edward Su. 1999. Rental harmony: Sperner's lemma in fair division. *The American mathematical monthly* 106, 10 (1999), 930–942.
- [35] Chenhao Wang and Xiaoying Wu. 2019. Cake Cutting with Single-Peaked Valuations. In *International Conference on Combinatorial Optimization and Applications*. Springer, 507–516.
- [36] Dietrich Weller. 1985. Fair division of a measurable space. *Journal of Mathematical Economics* 14, 1 (1985), 5–17.

A EGALITARIAN WELFARE

This section presents an algorithm for maximizing egalitarian welfare in cake-division instances with MLRP (Theorem 3.4).

Towards this end, we will define a “moving-knife” procedure that leads to the desired algorithm. Given a target value $\tau > 0$, we consider cut points obtained by iteratively selecting intervals that are of value τ to agents 1 through n , respectively. That is, the first cut point x_1 is chosen to satisfy $v_1(0, x_1) = \tau$. Inductively, for $i \geq 2$, the cut point x_i is selected such that $v_i(x_{i-1}, x_i) = \tau$. Recall that the agents are indexed following the MLRP order. To specify these points, we recursively define n functions, $\text{MK}_i : \mathbb{R}_+ \mapsto [0, 1]$, for $i \in [n]$:

$$\begin{aligned}\text{MK}_1(\tau) &:= \text{Cut}_1(0, \tau) \\ \text{MK}_i(\tau) &:= \text{Cut}_i(\text{MK}_{i-1}(\tau), \tau)\end{aligned}$$

Here, $\text{MK}_n(\tau)$ denotes the last cut point which we obtain by executing the moving-knife procedure with target value of $\tau \in [0, 1]$ for every agent. Given that MK_i s can be expressed as a composition of cut queries, these functions can be efficiently computed in the Robertson-Webb query model.

The following proposition establishes that MK_n is monotonic.

PROPOSITION A.1. *For any cake-division instance $C = \langle [n], \{f_i\}_i \rangle$ and each $i \in [n]$, the function MK_i is monotonically increasing.*

PROOF. Consider two target values $\tau, \tau' \in \mathbb{R}_+$ such that $\tau \leq \tau'$. We will show by inducting over $i \in [n]$ that $\text{MK}_i(\tau) \leq \text{MK}_i(\tau')$. For the base case $i = 1$, observe that $\text{Cut}_1(0, \tau) \leq \text{Cut}_1(0, \tau')$, since $f_1(x) \geq 0$ for all $x \in [0, 1]$. In other words, we have $\text{MK}_1(\tau) \leq \text{MK}_1(\tau')$. Next, with the induction hypothesis $\text{MK}_{i-1}(\tau) \leq \text{MK}_{i-1}(\tau')$ in hand, we will establish $\text{MK}_i(\tau) \leq \text{MK}_i(\tau')$:

$$\text{MK}_i(\tau) = \text{Cut}_i(\text{MK}_{i-1}(\tau), \tau) \leq \text{Cut}_i(\text{MK}_{i-1}(\tau'), \tau) \leq \text{Cut}_i(\text{MK}_{i-1}(\tau'), \tau') = \text{MK}_i(\tau')$$

Here, we use the fact that the value densities are nonnegative. This establishes the stated claim. \square

We now present the main result for egalitarian welfare.

THEOREM 3.4. *Let $C = \langle [n], \{f_i\}_{i \in [n]} \rangle$ be a cake-division instance in which the value-density functions satisfy the monotone likelihood ratio property. Then, in the Robertson-Webb query model, an allocation that achieves the optimal egalitarian welfare in C can be computed (up to an arbitrary precision) in $\mathcal{O}(\text{poly}(n, \log \lambda))$ time; here $\lambda \in \mathbb{R}_+$ is the Lipschitz constant of the cut and eval queries.*

PROOF. For the given cake-division instance C , write $\mathcal{R}^* = \{R_1^*, R_2^*, \dots, R_n^*\}$ to denote an allocation that maximizes egalitarian welfare among all allocations in C . Without loss of generality we can assume that \mathcal{R}^* is Pareto optimal and conforms to the MLRP order (Lemma 6.1). Write $\tau^* := \text{RW}(\mathcal{R}^*) = \min_i v_i(R_i^*)$.

Note that $\text{MK}_n(0) = 0$ and $\text{MK}_n(1) = 1$; by convention, $\text{Cut}_i(\ell, \tau)$ is truncated to 1 iff τ is greater than the entire value to the right of ℓ , i.e., if $v_i(\ell, 1) < \tau$.

We will say that a target value τ is *feasible* iff $v_i(\text{MK}_{i-1}(\tau), \text{MK}_i(\tau)) = \tau$ for all $i \in [n]$. That is, executing the moving-knife procedure with target value τ ensures that each agent receives an interval $[\text{MK}_{i-1}(\tau), \text{MK}_i(\tau)]$ of value τ . Otherwise, we say that τ is *infeasible*. Note that if τ is infeasible, then for some $i \in [n]$ the $\text{MK}_i(\tau)$ gets truncated at one—in such a case, i and subsequent agents $j > i$ do not receive intervals of value τ . Hence, $\tau = 0$ is a feasible target value, whereas $\tau = 1$ is infeasible. Note that we can efficiently determine whether a given τ is feasible or not.¹⁵

Proposition A.1 implies that we have monotonicity with respect to the feasibility of target values. Specifically, (P): for all $\tau \leq \tau'$, if τ is infeasible, then so is τ' . To establish property (P) we note that,

¹⁵As mentioned previously, the functions MK_i s can be computed efficiently in the Robertson-Webb model.

for an infeasible τ , there exists, by definition, an agent $i \in [n]$ such that $v_i(\text{MK}_{i-1}(\tau), \text{MK}_i(\tau)) < \tau$. This happens when $\text{Cut}_i(\text{MK}_{i-1}(\tau), \tau) = \text{MK}_i(\tau)$ gets truncated to 1. Applying Proposition A.1 to agent $i - 1$, with $\tau \leq \tau'$, we obtain $\text{MK}_{i-1}(\tau) \leq \text{MK}_{i-1}(\tau')$. Therefore, for agent i the value of the remaining cake $[\text{MK}_{i-1}(\tau'), 1]$ is less than τ' . Once again $\text{Cut}_i(\text{MK}_{i-1}(\tau'), \tau')$ gets truncated at 1 and we obtain $v_i(\text{MK}_{i-1}(\tau'), \text{MK}_i(\tau')) < \tau'$, i.e., τ' is infeasible.

With a precision parameter $\eta > 0$ in hand, we perform binary search over integer multiples of η , i.e., over the set $\{k\eta\}_{k=0}^{1/\eta}$. Recall that the values of the agents are normalized and, hence, $\tau^* \in [0, 1]$. Also, $\tau = 0$ is feasible, while $\tau = 1$ is infeasible.,

This switch in feasibility (between 0 and 1) along with property (P), imply that there exists a unique index $k_0 \in \{0, 1, \dots, 1/\eta\}$ with the property that $k_0\eta$ is feasible and $(k_0 + 1)\eta$ is infeasible. We can identify k_0 by $\mathcal{O}(\log(1/\eta))$ iterations of binary search.

The relevant observation here is that $\tau^* \in [k_0\eta, (k_0 + 1)\eta]$. Indeed, the optimal value τ^* is feasible: $\mathcal{R}^* = \{R_1^*, \dots, R_n^*\}$ conforms to the MLRP order and $v_i(R_i^*) \geq \tau^*$ for all $i \in [n]$. Hence, property (P) ensures that τ^* cannot be greater than the infeasible target $(k_0 + 1)\eta$. Furthermore, the optimality of τ^* gives us $\tau^* \geq k_0\eta$.

For any feasible τ , the moving-knife procedure provides an allocation with egalitarian welfare at τ . Specifically, for the computed index k_0 , consider allocation $\widehat{\mathcal{R}} = \{\widehat{R}_1, \widehat{R}_2, \dots, \widehat{R}_n\}$ in which $\widehat{R}_i = [\text{MK}_{i-1}(k_0\eta), \text{MK}_i(k_0\eta)]$, for $i \in [n - 1]$ and $\widehat{R}_n = [\text{MK}_{n-1}(k_0\eta), \text{MK}_n(k_0\eta)] \cup [\text{MK}_n(k_0\eta), 1]$. For notational convince here we set $\text{MK}_0(k_0\eta) = 0$. Feasibility of $k_0\eta$ ensures that $v_i(\widehat{R}_i) = k_0\eta \geq \tau^* - \eta$ for all $i \in [n]$. Therefore, we have $\text{RW}(\widehat{\mathcal{R}}) \geq \tau^* - \eta$.

For the runtime analysis, note that the binary search finds the desired index k_0 in $\mathcal{O}(\text{poly}(n, \log \frac{1}{\eta}))$ time. The dependency on λ stems from the fact that the bit-complexity of the output (i.e., of the computed cut points) can be $\mathcal{O}(\log \frac{\eta}{\lambda})$.

Overall, these arguments show that we can find an allocation $\widehat{\mathcal{R}}$ with egalitarian welfare η close to the optimal in time $\mathcal{O}(\text{poly}(n, \log \lambda, \log \frac{1}{\eta}))$. The precision parameter η can be driven exponentially close to zero in time that is polynomial in $\log \frac{1}{\eta}$ and, hence, the stated claim follows. \square

B NASH SOCIAL WELFARE

This section presents an FPTAS for maximizing Nash social welfare in cake-division instances with MLRP.

THEOREM 3.5. *Let $C = \langle [n], \{f_i\}_{i \in [n]} \rangle$ be a cake-division instance in which the value-density functions satisfy the monotone likelihood ratio property. Then, in the Robertson-Webb query model, an allocation with Nash social welfare at least $(1 - \varepsilon)$ times the optimal (Nash social welfare) in C can be computed in time that is polynomial in $1/\varepsilon$, $\log \lambda$, and n ; here $\lambda \in \mathbb{R}_+$ is the Lipschitz constant of the cut and eval queries.*

PROOF. Given a cake-division instance C with MLRP, write $\mathcal{A}^* = \{A_1^*, A_2^*, \dots, A_n^*\}$ to denote an allocation that maximizes the Nash social welfare in C . Arunachaleswaran et al. [3] have shown that the bundles in any Nash optimal allocation $\mathcal{A}^* = \{A_1^*, A_2^*, \dots, A_n^*\}$ satisfy $v_i(A_i^*) \geq \frac{1}{4}v_i(A_j^*)$ for all $i, j \in [n]$; this result holds even in the absence of MLRP. For a fixed agent $i \in [n]$, we sum the inequalities $v_i(A_i^*) \geq \frac{1}{4}v_i(A_j^*)$ over all $j \in [n]$ to obtain $v_i(A_i^*) \geq \frac{1}{4n}$. Recall that $v_i(0, 1) = 1$ for all $i \in [n]$.

For an approximation parameter $\varepsilon > 0$, we consider points $c_0 = 0 < c_1 < c_2 < \dots < c_N = 1$ such that $v_i(c_{t-1}, c_t) \leq \frac{\varepsilon}{8n}$ for all $t \in [N]$ and for all agents $i \in [n]$. In particular, we start with $c_0 = 0$ and iteratively select points $c_t := \min_i \text{Cut}_i(c_{t-1}, \frac{\varepsilon}{8n})$, for all $t \geq 1$. Note that the following upper bound holds for total number of selected points $N \leq \frac{8n^2}{\varepsilon}$.

We round each cut point in the optimal allocation \mathcal{A}^* to its closest point in the collection $\{c_i\}_{i=0}^N$. This leads us to another allocation, $\widehat{\mathcal{A}} = \{\widehat{A}_1, \widehat{A}_2, \dots, \widehat{A}_n\}$, wherein interval \widehat{A}_i is assigned to agent $i \in [n]$. By construction, the cut points of $\widehat{\mathcal{A}}$ are contained in the set $\{c_t\}_{t=0}^N$. Also, since \mathcal{A}^* conforms to the MLRP order (Lemma 6.1), so does $\widehat{\mathcal{A}}$. Next we show that the Nash social welfare of $\widehat{\mathcal{A}}$ is comparable to that of \mathcal{A}^* . For all $i \in [n]$, we have

$$\begin{aligned} v_i(\widehat{A}_i) &\geq v_i(A_i^*) - \frac{\varepsilon}{4n} && \text{(since } v_i(c_{t-1}, c_t) \leq \frac{\varepsilon}{8n} \text{ for all } t \in [N]) \\ &\geq v_i(A_i^*) - \varepsilon v_i(A_i^*) && \text{(since } v_i(A_i^*) \geq \frac{1}{4n}) \\ &\geq (1 - \varepsilon) v_i(A_i^*) \end{aligned}$$

Multiplying the above inequality over $i \in [n]$, we obtain $\text{NSW}(\widehat{\mathcal{A}}) = \left(\prod_{i=1}^n v_i(\widehat{A}_i)\right)^{1/n} \geq (1 - \varepsilon) \left(\prod_{i=1}^n v_i(A_i^*)\right)^{1/n} = (1 - \varepsilon) \text{NSW}(\mathcal{A}^*)$.

Therefore, there exists an allocation $\widehat{\mathcal{A}}$ with the properties that (i) $\widehat{\mathcal{A}}$ has Nash social welfare at least $(1 - \varepsilon)$ times the optimal (Nash social welfare), (ii) the cut points of $\widehat{\mathcal{A}}$ are contained in the set $\{c_t\}_t$, and (iii) $\widehat{\mathcal{A}}$ conforms to the MLRP order. To complete the proof of the theorem we will show that, among all the allocations that satisfy (ii) and (iii), we can find (via a dynamic program) one that maximizes the Nash social welfare.

For $t \in [N]$ and $k \in [n]$, we write $H(k, t)$ to denote the optimal Nash product (i.e., the product of valuations) that one can achieve by allocating the interval $[0, c_t]$ among the first k agents (in order).

$$H(k, t) := \max_{1 \leq t' \leq t} \{H(k-1, c_{t'}) \cdot v_k(c_{t'}, c_t)\}$$

Here, we initialize the dynamic program by setting $H(1, t) := v_1(0, c_t)$ for all $1 \leq t \leq N$. One can show, via induction, that $H(n, N)$ is the optimal Nash product among allocations that satisfy properties (ii) and (iii). That is, $(H(n, N))^{1/n} \geq (1 - \varepsilon) \text{NSW}(\mathcal{A}^*)$. Hence, the dynamic program computes the desired allocation using $\mathcal{O}(n^2 N)$ eval queries.

Therefore, we can find an allocation with Nash social welfare at least $(1 - \varepsilon)$ times the optimal in $\mathcal{O}(\text{poly}(n, 1/\varepsilon, \log \lambda))$ time; the dependency on $\log \lambda$ stems from the fact that the bit-complexity of the output (i.e., of the computed cut points) can be $\mathcal{O}(\log \frac{\varepsilon}{\lambda})$. Overall, we get that maximizing Nash social welfare admits an FPTAS under MLRP. \square

C CHARACTERIZATIONS AND IMPLICATIONS OF MLRP

C.1 Proof of Lemma 4.1

In this section we restate and prove Lemma 4.1.

LEMMA 4.1. *Let f_i and f_j be two value-density functions. Then, the following properties are equivalent*

- (i) *The likelihood ratios satisfy $\frac{f_j(b)}{f_i(b)} \leq \frac{f_j(c)}{f_i(c)}$ for all $0 \leq b \leq c \leq 1$.*
- (ii) *The values of the intervals satisfy $\frac{\int_a^b f_j}{\int_a^b f_i} \leq \frac{\int_c^d f_j}{\int_c^d f_i}$ for all $[a, b], [c, d] \subseteq [0, 1]$ with $b \leq c$.*
- (iii) *The (normalized) values satisfy $\frac{\int_x^b f_i}{\int_a^b f_i} \leq \frac{\int_x^b f_j}{\int_a^b f_j}$ for all intervals $[a, b] \subseteq [0, 1]$ and all $x \in [a, b]$.*

That is, the value densities f_i and f_j bear MLRP iff they satisfy property (ii) or (iii).

PROOF. First, we will prove that (i) implies (ii). Given MLRP and with $b \leq c$, we have

$$\frac{f_j(x)}{f_i(x)} \leq \frac{f_j(b)}{f_i(b)} \leq \frac{f_j(y)}{f_i(y)} \quad \text{for all } x \in [a, b] \text{ and for all } y \in [c, d] \quad (7)$$

Recall that MLRP value densities are, by definition, positively valued, $f_i(x) > 0$ for all $x \in [0, 1]$. Therefore, equation (7) gives us $f_j(x) \leq \frac{f_j(b)}{f_i(b)} f_i(x)$ for all $x \in [a, b]$. Integrating we obtain

$\int_a^b f_j(x)dx \leq \int_a^b \frac{f_j(b)}{f_i(b)} f_i(x)dx$ and, hence,¹⁶

$$\frac{\int_a^b f_j(x)dx}{\int_a^b f_i(x)dx} \leq \frac{f_j(b)}{f_i(b)} \tag{8}$$

Starting with equation (7) and integrating over the interval $[c, d]$, we can also establish the following equality

$$\frac{f_j(b)}{f_i(b)} \leq \frac{\int_c^d f_j(x)dx}{\int_c^d f_i(x)dx} \tag{9}$$

Equations (8) and (9) lead to property (ii):

$$\frac{\int_a^b f_j(x)dx}{\int_a^b f_i(x)dx} \leq \frac{\int_c^d f_j(x)dx}{\int_c^d f_i(x)dx}.$$

For the reverse direction, i.e., (ii) implies (i), consider any two points $0 \leq s < t \leq 1$ and an arbitrarily small $\varepsilon > 0$ (with $\varepsilon < |t - s|/2$). Instantiating property (ii) over the disjoint intervals $[s, s + \varepsilon]$ and $[t, t + \varepsilon]$ gives us

$$\frac{\int_s^{s+\varepsilon} f_j(x)dx}{\int_s^{s+\varepsilon} f_i(x)dx} \leq \frac{\int_t^{t+\varepsilon} f_j(x)dx}{\int_t^{t+\varepsilon} f_i(x)dx} \tag{10}$$

Dividing all the terms in equation (10) by $\varepsilon > 0$ and taking limits, we obtain

$$\lim_{\varepsilon \rightarrow 0} \frac{\frac{1}{\varepsilon} \int_s^{s+\varepsilon} f_j(x)dx}{\frac{1}{\varepsilon} \int_s^{s+\varepsilon} f_i(x)dx} \leq \lim_{\varepsilon \rightarrow 0} \frac{\frac{1}{\varepsilon} \int_t^{t+\varepsilon} f_j(x)dx}{\frac{1}{\varepsilon} \int_t^{t+\varepsilon} f_i(x)dx}.$$

Since the value densities f_i s are positive, via the algebra of limits, the previous inequality reduces to

$$\frac{\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \int_s^{s+\varepsilon} f_j(x)dx}{\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \int_s^{s+\varepsilon} f_i(x)dx} \leq \frac{\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \int_t^{t+\varepsilon} f_j(x)dx}{\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \int_t^{t+\varepsilon} f_i(x)dx}$$

Finally, invoking the fundamental theorem of calculus, we obtain MLRP (i.e., property (i)) for $s \leq t$

$$\frac{f_j(s)}{f_i(s)} \leq \frac{f_j(t)}{f_i(t)}.$$

¹⁶Recall that the integral of a positive function is positive.

To establish that property (ii) implies property (iii), we instantiate (ii) over the intervals $[a, x]$ and $[x, b]$: $\frac{\int_a^x f_j}{\int_a^x f_i} \leq \frac{\int_x^b f_j}{\int_x^b f_i}$. Cross multiplying the terms¹⁷ and adding one to both sides of the inequality, gives us $\frac{\int_a^x f_j}{\int_x^b f_j} + 1 \leq \frac{\int_a^x f_i}{\int_x^b f_i} + 1$. Simplifying further we obtain $\frac{\int_a^x f_j + \int_x^b f_j}{\int_x^b f_j} \leq \frac{\int_a^x f_i + \int_x^b f_i}{\int_x^b f_i}$. This gives us the desired bound

$$\frac{\int_x^b f_i}{\int_a^b f_i} \leq \frac{\int_x^b f_j}{\int_a^b f_j} \quad (11)$$

Finally, we complete the chain of implications by observing that (iii) implies (ii). Cross multiplying the terms in inequality (11) and subtracting one from both sides yield $\frac{\int_a^b f_j}{\int_x^b f_j} - 1 \leq \frac{\int_a^b f_i}{\int_x^b f_i} - 1$. This inequality simplifies to $\frac{\int_a^x f_j}{\int_x^b f_j} \leq \frac{\int_a^x f_i}{\int_x^b f_i}$. That is, we obtain property (ii) for intervals $[a, x]$ and $[x, b]$. Reapplying this bound (with $a, x,$ and b set appropriately) shows that (iii) implies (ii). This completes the proof. \square

C.2 Efficiently Finding the MLRP Order

This section shows that, if in a cake-division instance the underlying value densities bear MLRP, then we can efficiently find the MLRP order in the Robertson-Webb query model. Specifically, through eval queries $\{\text{Eval}_i(1/2, 1)\}_{i \in [n]}$ we find the value that each agent $i \in [n]$ has for the piece $[1/2, 1]$ of the cake, and sort the agents according to the increasing order (with ties broken arbitrarily) of these values. Lemma C.2 shows that this sorting order is in fact the MLRP order of the underlying value densities.

To prove this claim, we first state a well-known result (in Lemma C.1) that MLRP implies first-order stochastic dominance; we provide a proof here for completeness. Recall that, given two probability density functions f_i and f_j over $[0, 1]$, density f_j is said to have *first-order stochastic dominance* over f_i iff $\int_t^1 f_j(x)dx \geq \int_t^1 f_i(x)dx$, for all $t \in [0, 1]$, and there exists at least one $t' \in [0, 1]$ such that $\int_{t'}^1 f_j(x)dx > \int_{t'}^1 f_i(x)dx$.

LEMMA C.1. *Let f_i and f_j be two value-density functions that satisfy the monotone likelihood ratio property: for every $0 \leq x \leq y \leq 1$ we have $\frac{f_j(x)}{f_i(x)} \leq \frac{f_j(y)}{f_i(y)}$. Then, the density f_j has first-order stochastic dominance over f_i .*

PROOF. Given that f_i and f_j bear MLRP, we consider property (iii) of Lemma 4.1, with $a = 0$, $b = 1$, and $x = t$, for any $0 \leq t \leq 1$, to obtain $\int_t^1 f_i \leq \int_t^1 f_j$. Here, we use the fact that the valuations are normalized, $\int_0^1 f_i = \int_0^1 f_j = 1$.

Furthermore, since the two densities are distinct, there exists a point $t' \in [0, 1]$ such that $\int_{t'}^1 f_j$ is not equal to $\int_{t'}^1 f_i$. For such a point t' , a strict inequality must hold, $\int_{t'}^1 f_j > \int_{t'}^1 f_i$. \square

Note that, since both MLRP and first-order stochastic dominance are transitive properties, the above lemma directly extends to a collection of probability density functions.

Next we prove that, given a cake-division instance, with the promise that the underlying value densities satisfy MLRP, one can find the MLRP order efficiently.

¹⁷Recall that the value densities are strictly positive.

LEMMA C.2. *Let C be a cake-division instance in which the value-density functions satisfy the monotone likelihood ratio property. Then, in the Robertson-Webb query model, we can find the MLRP order of C in polynomial time.*

PROOF. Through eval queries $\{\text{Eval}_i(1/2, 1)\}_{i \in [n]}$, we find the value that each agent $i \in [n]$ has for the piece $[1/2, 1]$ of the cake, and sort the agents according to the increasing order (with ties broken arbitrarily) of these values. We will show that that this sorting order, say $\pi : [n] \mapsto [n]$, is in fact the MLRP order of the underlying value densities. This directly provides an efficient algorithm for finding the MLRP order.

Consider two agents $i, j \in [n]$ such that i appears before j in the MLRP order, i.e., the likelihood ratio $f_j(x)/f_i(x)$ is non-decreasing in $x \in [0, 1]$. Below, we will address the settings in which the value densities of i and j are distinct. Otherwise, if the value densities are identical, then $\text{Eval}_i(1/2, 1) = \text{Eval}_j(1/2, 1)$ and they will be placed appropriately in the sorting order π .

Given that f_i and f_j are distinct we get, via Lemma C.1, that f_j has first-order stochastic dominance over f_i . Therefore, agent j 's value for the piece $[1/2, 1]$ is at least as large as agent i 's value for this piece, $\int_{1/2}^1 f_j \geq \int_{1/2}^1 f_i$, i.e., $\text{Eval}_j(1/2, 1) \geq \text{Eval}_i(1/2, 1)$. We will show that, under MLRP, this inequality is in fact strict and, hence, π correctly identifies the MLRP order among each pair of agents j and i .

The first-order stochastic dominance between f_j and f_i also ensures that there exists a point $t' \in [0, 1]$ such that

$$\int_{t'}^1 f_j > \int_{t'}^1 f_i \quad (12)$$

We consider two complementary cases (i) $t' \leq 1/2$ and (ii) $t' > 1/2$. In both of these cases we assume, towards a contradiction, that $\int_{1/2}^1 f_j = \int_{1/2}^1 f_i$.

Case (i): $t' \leq 1/2$. Note that in this case equation (12) expands to $\int_{t'}^{1/2} f_j + \int_{1/2}^1 f_j > \int_{t'}^{1/2} f_i + \int_{1/2}^1 f_i$. Since, $\int_{1/2}^1 f_j = \int_{1/2}^1 f_i$, we have $\int_{t'}^{1/2} f_j > \int_{t'}^{1/2} f_i$.

On the other hand, applying Lemma 4.1, property (ii), to intervals $[t', 1/2]$ and $[1/2, 1]$ gives us $\frac{\int_{t'}^{1/2} f_j}{\int_{t'}^{1/2} f_i} \leq \frac{\int_{1/2}^1 f_j}{\int_{1/2}^1 f_i} = 1$. This inequality leads to the desired contradiction, $\int_{t'}^{1/2} f_j \leq \int_{t'}^{1/2} f_i$.

Case (ii): $t' > 1/2$. Here, equation (12) and the assumed equality $\int_{1/2}^1 f_j = \int_{1/2}^1 f_i$ give us $\int_{1/2}^{t'} f_j < \int_{1/2}^{t'} f_i$.

Note that (due to normalization) we have $\int_0^{1/2} f_j = \int_0^{1/2} f_i$. This equality contradicts an application of Lemma 4.1, property (ii), to the intervals $[0, 1/2]$ and $[1/2, t']$

$$\frac{\int_0^{1/2} f_j}{\int_0^{1/2} f_i} \leq \frac{\int_{1/2}^{t'} f_j}{\int_{1/2}^{t'} f_i} < 1$$

Therefore, by sorting the values of the eval queries $\{\text{Eval}_i(1/2, 1)\}_{i \in [n]}$, we can efficiently find the MLRP order. \square

C.3 Lipschitzness of Cut and Eval Queries under MLRP

Here, we prove that if the value densities satisfy MLRP, then the cut and eval queries of the corresponding cake-division instance are Lipschitz continuous. Recall, that the value densities f_i s are assumed to be Riemann integrable and hence, by definition, f_i s are bounded. Furthermore, as mentioned previously, MLRP mandates that the value densities are strictly positive over the

cake. Therefore, for each agent $i \in [n]$ and $x \in [0, 1]$ we have $0 < L \leq f_i(x) \leq U$ for some $L, U \in \mathbb{R}_+$. Proposition C.3 below shows that we can express the Lipschitz constant λ in terms of these bounding parameters, $\lambda \leq \max\{1/L, U, U/L\}$. This establishes the Lipschitz continuity of the cut and eval queries under MLRP.

PROPOSITION C.3. *Let $C = \langle [n], \{f_i\}_{i \in [n]} \rangle$ be a cake-division instance in which the value densities are positively-bounded, i.e., there exists $L, U \in \mathbb{R}_+$ such that $0 < L \leq f_i(x) \leq U$ for each agent $i \in [n]$ and all $x \in [0, 1]$. Then the cut and eval queries, $\{\text{Cut}_i\}_i$ and $\{\text{Eval}_i\}_i$, are λ -Lipschitz with $\lambda \leq \max\{1/L, U, U/L\}$.*

PROOF. Fix an agent $i \in [n]$. First, we will prove that the eval query (function) Eval_i is U -Lipschitz. Then, we will prove that Cut_i is $\max\{1/L, U/L\}$ -Lipschitz. These two results imply the stated claim.

Recall that $\text{Eval}_i : [0, 1] \times [0, 1] \mapsto \mathbb{R}_+$ is said to be U -Lipschitz iff $|\text{Eval}_i(\ell, r) - \text{Eval}_i(\ell', r')| \leq U \max\{|\ell - \ell'|, |r - r'|\}$. We will establish this inequality component wise.

For a fixed first component $\ell_0 \in [0, 1]$, consider the function $g(r) := \text{Eval}_i(\ell_0, r) = \int_{\ell_0}^r f_i(x) dx$ for $r \in [0, 1]$. Since f_i is Riemann integrable, via the fundamental theorem of calculus, we have that the derivative of $g(\cdot)$ is equal $f_i(\cdot)$, i.e., $\frac{d}{dr}g(r) = f_i(r)$ for $r \in [0, 1]$. Recall that, if a real-valued function has a bounded derivative, then it is Lipschitz. Hence, the inequality $|f_i(r)| \leq U$ implies that $g(r)$ is U -Lipschitz. Therefore, $\text{Eval}_i(\ell, r)$ is U -Lipschitz in the second argument r .

Similarly, considering the function $\tilde{g}(\ell) := \text{Eval}_i(\ell, r_0)$, with a fixed $r_0 \in [0, 1]$, we get that derivative of $\tilde{g}(\ell)$ is equal to $-f_i(\ell)$. Hence, $\text{Eval}_i(\ell, r)$ is U -Lipschitz in the first argument ℓ as well. The two parts imply that Eval_i is U -Lipschitz.

For $\text{Cut}_i(\ell, \tau)$ we again establish Lipschitz continuity by considering the two arguments ℓ and τ separately. Fix $\ell_0 \in [0, 1]$ and for target values $\tau \leq \tau'$, let $y := \text{Cut}_i(\ell_0, \tau)$ and $y' := \text{Cut}_i(\ell_0, \tau')$. Note that $y \leq y'$ and $\tau' - \tau$ is equal to agent i 's value for the interval $[y, y']$. This value can be lower bounded by observing that $f_i(x) \geq L$, for all $x \in [0, 1]$; specifically, agent i 's value for the interval $[y, y']$ is at least $L(y' - y)$. Therefore, we get that $\text{Cut}_i(\ell_0, \tau') - \text{Cut}_i(\ell_0, \tau) = y' - y \leq \frac{1}{L}(\tau' - \tau)$. That is, $\text{Cut}_i(\ell, \tau)$ is $\frac{1}{L}$ -Lipschitz in $\tau \in \mathbb{R}_+$.

For the remaining analysis, fix target value $\tau_0 \in \mathbb{R}_+$, and for $\ell \leq \ell'$, let $y := \text{Cut}_i(\ell, \tau_0)$ and $y' := \text{Cut}_i(\ell', \tau_0)$. A useful observation here is that agent i 's value for the interval $[y, y']$ is equal to i 's value for the interval $[\ell, \ell']$, i.e., $v_i(y, y') = v_i(\ell, \ell')$.

Specifically, if $y \leq \ell'$, then we can write δ to denote agent i 's value for the interval $[y, \ell']$ (i.e., $\delta := v_i(y, \ell')$) and note that $v_i(\ell, \ell') = v_i(\ell, y) + v_i(y, \ell') = \tau_0 + \delta$. The equality then follows from observing that $v_i(y, y') = v_i(y, \ell') + v_i(\ell', y') = \delta + \tau_0$. Complementarily, if $y > \ell'$, then we write $\delta := v_i(\ell', y)$ and use the following equations to obtain the desired equality: $v_i(\ell, \ell') = v_i(\ell, y) - v_i(\ell', y) = \tau_0 - \delta$ and $v_i(y, y') = v_i(\ell', y') - v_i(\ell', y) = \tau_0 - \delta$.

The bounds on the value densities, $0 < L \leq f_i(x) \leq U$, and the equality $v_i(y, y') = v_i(\ell, \ell')$ imply $L(y' - y) \leq v_i(y, y') = v_i(\ell, \ell') \leq U(\ell' - \ell)$. Hence, we have $|\text{Cut}_i(\ell', \tau_0) - \text{Cut}_i(\ell, \tau_0)| = |y' - y| \leq \frac{U}{L}|\ell' - \ell|$. That is, $\text{Cut}_i(\ell, \tau)$ is $\frac{U}{L}$ -Lipschitz in $\ell \in \mathbb{R}_+$.

Combing the above-mentioned bounds we get that the Cut_i and Eval_i queries are λ -Lipschitz with $\lambda \leq \max\{U, 1/L, U/L\}$. \square

D MISSING PROOFS FROM SECTION 5

D.1 Proof of Proposition 5.3

Here we restate and prove Proposition 5.3

PROPOSITION 5.3. *If in a cake-division instance $C = \langle [n], \{f_i\}_i \rangle$ the cut and eval queries are λ -Lipschitz, then the function RD_n is $\lambda^{2(n-1)}$ -Lipschitz.*

PROOF. Applying strong induction over $i \in \{2, 3, \dots, n\}$, we will prove that the the function RD_i is $\lambda^{2(i-1)}$ -Lipschitz. For the base case of RD_2 consider the following bound, with points $x, x' \in [0, 1]$:

$$\begin{aligned} |\text{RD}_2(x') - \text{RD}_2(x)| &= |\text{Cut}_1(x', \text{Eval}_1(0, x')) - \text{Cut}_1(x, \text{Eval}_1(0, x))| && \text{(by definition of } \text{RD}_2\text{)} \\ &\leq \lambda \max\{|x' - x|, |\text{Eval}_1(0, x') - \text{Eval}_1(0, x)|\} && \text{(Cut}_1 \text{ is } \lambda\text{-Lipstchiz)} \\ &\leq \lambda \max\{|x' - x|, \lambda|x' - x|\} && \text{(Eval}_1 \text{ is } \lambda\text{-Lipstchiz)} \\ &= \lambda^2|x' - x| && \text{(since } \lambda \geq 1\text{)} \end{aligned}$$

Hence, RD_2 is λ^2 -Lipschitz. Next, assuming that, for all $k \leq i - 1$, RD_k is $\lambda^{2(k-1)}$ -Lipschitz, we will prove that RD_i is $\lambda^{2(i-1)}$ -Lipschitz. Note that the following bound holds for all $x', x \in [0, 1]$

$$\begin{aligned} &\text{Eval}_{i-1}(\text{RD}_{i-2}(x'), \text{RD}_{i-1}(x')) - \text{Eval}_{i-1}(\text{RD}_{i-2}(x), \text{RD}_{i-1}(x)) \\ &\leq \lambda \max\{|\text{RD}_{i-2}(x') - \text{RD}_{i-2}(x)|, |\text{RD}_{i-1}(x') - \text{RD}_{i-1}(x)|\} && \text{(since } \text{Eval}_{i-1} \text{ is } \lambda\text{-Lipschitz)} \\ &\leq \lambda \max\{\lambda^{2(i-3)}|x' - x|, \lambda^{2(i-2)}|x' - x|\} && \text{(using the induction hypothesis for } i - 2 \text{ and } i - 1\text{)} \\ &= \lambda \lambda^{2(i-2)}|x' - x| && \text{(since } \lambda \geq 1\text{)} \\ &= \lambda^{2i-3}|x' - x| && (13) \end{aligned}$$

We can now bound $|\text{RD}_i(x') - \text{RD}_i(x)|$ for all $x', x \in [0, 1]$:

$$\begin{aligned} &|\text{RD}_i(x') - \text{RD}_i(x)| \\ &= |\text{Cut}_{i-1}(\text{RD}_{i-1}(x'), \text{Eval}_{i-1}(\text{RD}_{i-2}(x'), \text{RD}_{i-1}(x'))) - \text{Cut}_{i-1}(\text{RD}_{i-1}(x), \text{Eval}_{i-1}(\text{RD}_{i-2}(x), \text{RD}_{i-1}(x)))| \\ &\leq \lambda \max\{|\text{RD}_{i-1}(x') - \text{RD}_{i-1}(x)|, \lambda^{2i-3}|x' - x|\} && \text{(since } \text{Cut}_{i-1} \text{ is } \lambda\text{-Lipstchiz and by equation (13))} \\ &\leq \lambda \max\{\lambda^{2(i-2)}|x' - x|, \lambda^{2i-3}|x' - x|\} && \text{(using the induction hypothesis for } i - 1\text{)} \\ &= \lambda^{2(i-1)}|x' - x| && \text{(since } \lambda \geq 1\text{)} \end{aligned}$$

Therefore, RD_i is $\lambda^{2(i-1)}$ -Lipschitz for all $2 \leq i \leq n$. Setting $i = n$, gives us the desired claim. \square

D.2 Proof of Lemma 5.5

This section restates Lemma 5.5 and proves universal existence of ripple divisions.

LEMMA 5.5. *Let C be a cake-division instance in which the cut and eval queries are λ -Lipschitz. Then, C necessarily admits a ripple division.*

PROOF. Lemma 5.4 asserts that, for any $\delta \in (0, 1)$, there exists a collection of points $x_0^\delta = 0 < x_1^\delta < \dots < x_{n-1}^\delta < x_n^\delta = 1 - \delta$ that form a δ -ripple division. Note that here $x_n^\delta = 1 - \delta$. We consider the sequence of these δ -ripple divisions, $\mathcal{S} := \langle (x_0^\delta, x_1^\delta, \dots, x_n^\delta) \rangle_{\delta \in (0,1)}$.

Since \mathcal{S} is a nonempty and bounded sequence in $[0, 1]^n$, the Bolzano-Weierstrass theorem ensures that \mathcal{S} contains a convergent subsequence, say $\langle (x_0^{\delta_j}, x_1^{\delta_j}, \dots, x_n^{\delta_j}) \rangle_{\delta_j > 0}$. Write $(x_0^*, x_1^*, \dots, x_n^*) \in [0, 1]^n$ to denote the limit of this subsequence as δ_j tends to zero. We will show that the points $x_0^*, x_1^*, \dots, x_n^*$ form a ripple division in C (see Definition 5.1), i.e., establish that $x_n^* = 1$ and $v_i(x_{i-1}^*, x_i^*) = v_i(x_i^*, x_{i+1}^*) > 0$ for all agents $i \in [n - 1]$.

First, note that $x_n^* = 1$, since the sequence $\langle 1 - \delta_j \rangle \rightarrow 1$ as $\delta_j \rightarrow 0$. Also, $x_0^* = 0$, since the constant sequence $\langle 0 \rangle$ tends to 0.

We will next prove that $v_i(x_{i-1}^*, x_i^*) = v_i(x_i^*, x_{i+1}^*)$ for all agents $i \in [n-1]$. Given that the collection $(x_0^\delta, x_1^\delta, \dots, x_n^\delta)$ forms a δ_j -ripple division, we have $v_i(x_{i-1}^{\delta_j}, x_i^{\delta_j}) = v_i(x_i^{\delta_j}, x_{i+1}^{\delta_j})$, for all $i \in [n-1]$.

Recall that the value-density function v_i (equivalently, Eval_i) is Lipschitz continuous over $[0, 1]^2$. Therefore, the sequential criterion of continuity¹⁸ gives us $\langle v_i(x_{i-1}^{\delta_j}, x_i^{\delta_j}) \rangle \rightarrow v_i(x_{i-1}^*, x_i^*)$ and $\langle v_i(x_i^{\delta_j}, x_{i+1}^{\delta_j}) \rangle \rightarrow v_i(x_i^*, x_{i+1}^*)$, as δ_j tends to 0. Applying the algebra of limits¹⁹ on the two sequences $\langle v_i(x_{i-1}^{\delta_j}, x_i^{\delta_j}) \rangle$ and $\langle v_i(x_i^{\delta_j}, x_{i+1}^{\delta_j}) \rangle$, we obtain that their limits must be equal as well, i.e., $v_i(x_{i-1}^*, x_i^*) = v_i(x_i^*, x_{i+1}^*)$.

We now complete the proof by showing that these two equal values must be positive, $v_i(x_{i-1}^*, x_i^*) = v_i(x_i^*, x_{i+1}^*) > 0$. Since $x_n^* = 1$, we have $x_1^* > 0$; otherwise, if $x_1^* = 0$, then following the value equalities we would have $x_n^* = \text{RD}_n(0) = 0$. In the current setting, all nonempty intervals have positive value for the agents. This follows from the Lipschitz continuity of the cut queries. Hence, $x_1^* > x_0^* = 0$ gives us $0 < v_1(x_0^*, x_1^*) = v_1(x_1^*, x_2^*)$. This bound also implies $x_2^* > x_1^*$. Extending this argument iteratively shows that $0 = x_0^* < x_1^* < \dots < x_n^* = 1$. Hence, the agents' values are positive, $v_i(x_{i-1}^*, x_i^*) > 0$. Overall, we get that the set of points $0 = x_0^* < x_1^* < \dots < x_n^* = 1$ form a ripple division in C . \square

D.3 Proof of Lemma 5.6

This section restates Lemma 5.6 and shows that `BINSEARCH` (Algorithm 1) satisfies this claim.

LEMMA 5.6. *Let $C = \langle [n], \{f_i\}_i \rangle$ be a cake-division instance in which the cut and eval queries are λ -Lipschitz. Then, for any $\delta \in (0, 1)$ and in the Robertson-Webb query model, a δ -ripple division of C can be computed in $O(\text{poly}(n, \log \lambda, \log \frac{1}{\delta}))$ time.*

PROOF. Lemma 5.4 implies that the problem of computing a δ -ripple division reduces to finding a point $x \in (0, 1)$ that satisfies $\text{RD}_n(x) \in [1 - \delta, 1)$. Here, we will show that `BINSEARCH` (Algorithm 1) finds such a point x (and, hence, a δ -ripple division) in $O(\text{poly}(n, \log \lambda, \log \frac{1}{\delta}))$ time.

By design, `BINSEARCH` maintains two points $\ell \leq r$ and keeps iterating till it finds a point $x = (\ell + r)/2$ that satisfies the required property $\text{RD}_n(x) \in [1 - \delta, 1)$. Hence, at termination the algorithm indeed finds a δ -ripple division. Below we complete the proof by showing that the time complexity of the algorithm is as stated (in particular, the algorithm necessarily terminates).

Note that, throughout the algorithm's execution, the maintained points ℓ and r satisfy $\text{RD}_n(\ell) < 1 - \delta$ and $\text{RD}_n(r) = 1$. This invariant holds at the beginning of the algorithm, where we initialize $\ell = 0$ and $r = 1$; recall that $\text{RD}_n(0) = 0$ and $\text{RD}_n(1) = 1$. Furthermore, in each iteration of the while-loop in `BINSEARCH`, the left endpoint ℓ gets updated $\ell \leftarrow (\ell + r)/2$ iff the midpoint $(\ell + r)/2$ satisfies $\text{RD}_n\left(\frac{\ell+r}{2}\right) < 1 - \delta$. That is, even after the update we have $\text{RD}_n(\ell) < 1 - \delta$. Similarly, for the right endpoint r , during the algorithm's execution we have $\text{RD}_n(r) = 1$.

This invariant implies that between ℓ and r we always have a point $x \in (\ell, r)$ which satisfies $\text{RD}_n(x) = 1 - \delta/2$. This observation is obtained by applying the intermediate value theorem to the (continuous) function RD_n on the interval $[\ell, r]$. Hence, the interval under consideration, $[\ell, r]$, continues to contain a required point.

Furthermore, in each iteration of the while-loop of `BINSEARCH`, the difference between ℓ and r reduces by a multiplicative factor of two. Hence, after T iterations we must have $r - \ell \leq \frac{1}{2^T}$. In

¹⁸For any continuous function $g : [0, 1]^2 \mapsto \mathbb{R}$, if a sequence $\langle a_t \rangle_t \in [0, 1]^2$ converges to some $a \in [0, 1]^2$, then the sequence $\langle g(a_t) \rangle_t$ must converge to $g(a)$.

¹⁹For any two convergent sequences $\langle a_t \rangle_t \rightarrow a$ and $\langle b_t \rangle_t \rightarrow b$, if we have $a_t = b_t$ for all $t \geq 1$, then their limits must be equal as well, i.e., $a = b$.

particular, the difference between $(\ell + r)/2$ and the desired point $x \in (\ell, r)$ is at most $\frac{1}{2^T}$, after T iterations. Therefore, the following bound holds after T iterations

$$\begin{aligned} \left| \text{RD}_n(x) - \text{RD}_n\left(\frac{\ell + r}{2}\right) \right| &\leq \lambda^{2(n-1)} \left| x - \left(\frac{\ell + r}{2}\right) \right| && \text{(Proposition 5.3: } \text{RD}_n \text{ is } \lambda^{2(n-1})\text{-Lipschitz)} \\ &\leq \frac{\lambda^{2(n-1)}}{2^T} \end{aligned}$$

Note that, if the iteration count T is greater than $2(n-1) \log\left(\frac{2\lambda}{\delta}\right)$, then $\left| \text{RD}_n(x) - \text{RD}_n\left(\frac{\ell+r}{2}\right) \right| < \delta/2$. Since $\text{RD}_n(x) = 1 - \delta/2$, we get that $\text{RD}_n\left(\frac{\ell+r}{2}\right) \in (1 - \delta, 1)$. That is, in at most $2(n-1) \log\left(\frac{2\lambda}{\delta}\right)$ iteration `BINSEARCH` will certainly find a point $\left(\frac{\ell+r}{2}\right)$ that satisfies the termination criterion $\text{RD}_n\left(\frac{\ell+r}{2}\right) \in [1 - \delta, 1)$. This shows the time complexity of `BINSEARCH` is $\mathcal{O}(\text{poly}(n, \log \lambda, \log \frac{1}{\delta}))$ and completes the proof. \square

E MISSING PROOFS FROM SECTION 7

E.1 Proof of Lemma 7.1

This section restates and proves Lemma 7.1.

LEMMA 7.1. *Let f_i and f_j be two value-density functions that satisfy MLRP, i.e., f_j/f_i is non-decreasing over $[0, 1]$. Then,*

- (a) *For all $y \in [0, p_{i,j})$, the likelihood ratio satisfies $f_j(y)/f_i(y) < 1$.*
- (b) *For all $z \in (p_{i,j}, 1]$, the likelihood ratio satisfies $f_j(z)/f_i(z) \geq 1$.*

Here, $p_{i,j}$ is the switching point between f_i and f_j .

PROOF. As observed previously, $p_{i,j} \in [0, 1]$ exists and is unique. We begin by proving part (a) of the stated claim. Consider any point $y \in [0, p_{i,j})$ and assume, towards a contradiction, that the likelihood ratio satisfies $\frac{f_j(y)}{f_i(y)} \geq 1$. This implies that the point y belongs to the set $L_{ij} = \{x \in [0, 1] : f_j(x) \geq f_i(x)\}$. Since $y < p_{i,j}$, we get a contradiction to the fact that $p_{i,j}$ is the infimum of the set L_{ij} .

For proving part (b), consider any point $z \in (p_{i,j}, 1]$. Assume, towards a contradiction, that at z we have $\frac{f_j(z)}{f_i(z)} < 1$. Since the likelihood ratio $\frac{f_j(x)}{f_i(x)}$ is non-decreasing over $[0, 1]$ (by definition of MLRP), we have $\frac{f_j(t)}{f_i(t)} < 1$ for all $0 \leq t \leq z$. That is, there does not exist a point $t \in [0, z]$ with the property that $f_j(t) \geq f_i(t)$. Hence, z constitutes a lower bound for the set $L_{ij} = \{x \in [0, 1] : f_j(x) \geq f_i(x)\}$. Since $p_{i,j} < z$, we get a contradiction to the fact that $p_{i,j}$ is the infimum (greatest lower bound) of the set L_{ij} . This completes the proof. \square

E.2 Proof of Corollary 7.2

Here we prove Corollary 7.2.

COROLLARY 7.2. *Let f_i and f_j be two value densities that bear MLRP and let parameter $\gamma \in (0, 1)$. Then, in the Robertson-Webb model, we can find an interval of length γ that contains the switching point $p_{i,j}$ in $\mathcal{O}(\log(1/\gamma))$ time.*

PROOF. Consider intervals of the form $B_k = \left[(k-1)\frac{\gamma}{2}, k\frac{\gamma}{2}\right]$ for $k \in \{1, 2, \dots, \frac{2}{\gamma}\}$, i.e., for analysis, we discretize the cake $[0, 1]$ evenly into intervals each of length $\gamma/2$. Write k^* to denote the index with the property that $p_{i,j} \in B_{k^*} = \left[(k^*-1)\frac{\gamma}{2}, k^*\frac{\gamma}{2}\right]$. Part (a) of Lemma 7.1 implies that $v_j(B_k) < v_i(B_k)$ for all $k < k^*$. Similarly, part (b) of Lemma 7.1 gives us $v_j(B_k) \geq v_i(B_k)$ for all $k > k^*$.

Therefore, applying binary search, we can, in $O(\log(1/\gamma))$ iterations, find the smallest index $k' \in \{1, 2, \dots, \frac{2}{\gamma}\}$ such that $v_j(B_{k'}) < v_i(B_{k'})$ and $v_j(B_{k'+1}) \geq v_i(B_{k'+1})$. Note that the computed index k' satisfies $k' \in \{k^* - 1, k^*\}$: if, for contradiction, we have $k' \leq k^* - 2$, then it must be the case that $v_j(B_{k'+1}) < v_i(B_{k'+1})$. This inequality contradicts the selection criterion of k' . Also, the inequality $k' \geq k^* + 1$ would lead to the contradiction $v_j(B_{k'}) \geq v_i(B_{k'})$.

The value comparisons required to execute the binary search can be performed using eval queries. Hence, in $O(\log(1/\gamma))$ iterations we can find an interval $B_{k'} \cup B_{k'+1}$ of length γ that contains p_{ij} . \square

E.3 Proof of Lemma 7.3

Next we prove Lemma 7.3

LEMMA 7.3. *Let C be a cake-division instance in which the value densities satisfy the monotone likelihood ratio property. Then, in C , there exists a social welfare maximizing allocation all of whose cut points belong to the set of switching points P .*

PROOF. Among all allocations that maximize social welfare, consider the ones that conform to the MLRP order; Lemma 6.1 ensures that this collection is nonempty. Furthermore, among these optimal allocations select one $\mathcal{S} = \{S_1, S_2, \dots, S_n\}$ that minimizes $|\{s_0 = 0, s_1, s_2, \dots, s_n = 1\} \setminus P|$; here s_i s denote the cut points of \mathcal{S} . That is, \mathcal{S} is a social welfare maximizing allocation that uses as many points from P as possible. We will show that $\{s_i\}_i \setminus P = \emptyset$ and, hence, the claim follows.

Towards a contradiction, assume that there exists a cut point s_t of the allocation \mathcal{S} that does not belong to P . Let $S_i = [s, s_t]$ and $S_j = [s_t, s']$ be the two *nonempty* intervals in \mathcal{S} that are separated by s_t . Interval S_i is to the (immediate) left of S_j and, since the allocations conform to the MLRP order, we that $i < j$.

Given that $s_t \notin P$, we know that $s_t \neq p_{ij}$; here p_{ij} is the switching point between f_i and f_j . We will show that in this case we can always move s_t towards p_{ij} and obtain another social welfare maximizing allocation that uses more cut points from P than \mathcal{S} . This contradicts the choice of \mathcal{S} and establishes the stated claim. Towards this goal, consider two complementary cases

Case (i): $s_t < p_{ij}$. In this case we can move s_t to the right without decreasing the social welfare. In particular, if $p_{ij} \in S_i \cup S_j = [s, s']$, then, instead of S_i and S_j , we can assign intervals $[s, p_{ij}]$ and $[p_{ij}, s']$ to agents i and j , respectively. Since $f_j(x) < f_i(x)$ for all $x \in [s_t, p_{ij}]$ (Lemma 7.1, part (a)), such an update increases the social welfare. This contradicts the optimality (with respect to social welfare) of \mathcal{S} . A similar argument holds if $p_{ij} > s'$. Here, we can assign $[s, s']$ entirely to agent i (and an empty set to agent j). For all $x \leq s' < p_{ij}$, we have $f_j(x) < f_i(x)$ (Lemma 7.1, part (a)). Therefore, the reassignment increase the social welfare and leads to a contradiction.

Case (ii): $s_t > p_{ij}$. In this case we can move s_t to the left (towards p_{ij}). If we have $p_{ij} \in S_i \cup S_j$, then assigning intervals $[s, p_{ij}]$ and $[p_{ij}, s']$ to agents agents i and j , respectively, does not decrease the social welfare (Lemma 7.1, part (b)). Though, at the same time, it does provide an allocation that uses more cut points from P and, hence, contradicts the choice of \mathcal{S} . On the other hand, if $p_{ij} < s$, then we can assign the entire interval $[s, s']$ to agent j . As before, the reassignment does not decrease the social welfare. However, it does decrease the cardinality of the set difference between the cut points and P . This contradicts the selection criterion of \mathcal{S} .

Hence, the cut points of \mathcal{S} satisfy $\{s_i\}_i \subseteq P$ and the stated claim follows. \square

E.4 Proof of Theorem 3.3

This section restates Theorem 3.3 and presents its proof.

THEOREM 3.3. *Let $C = \langle [n], \{f_i\}_{i \in [n]} \rangle$ be a cake-division instance in which the value-density functions satisfy the monotone likelihood ratio property. Then, in the Robertson-Webb query model, an allocation that achieves the optimal social welfare in C can be computed (up to an arbitrary precision) in $O(\text{poly}(n, \log \lambda))$ time; here $\lambda \in \mathbb{R}_+$ is the Lipschitz constant of the cut and eval queries.*

PROOF. Given a cake-division instance C with MLRP, write $\mathcal{S}^* = \{S_1^*, S_2^*, \dots, S_n^*\}$ to denote the allocation identified in Lemma 7.3; in particular, \mathcal{S}^* is a social welfare maximizing allocation whose cut points $\{0 = s_0^*, s_1^*, \dots, s_n^* = 1\}$ belong to the set of switching points P .

For a precision parameter $\eta > 0$ and for each switching point $p_{ij} \in P$, we invoke Corollary 7.2 to $\widehat{p}_{ij} \in [0, 1]$ with the property that $|\widehat{p}_{ij} - p_{ij}| \leq \frac{\eta}{n\lambda}$. Write \widehat{P} to denote the set of these estimates, $\widehat{P} := \{\widehat{p}_{ij} : 1 \leq i < j \leq n\} \cup \{0, 1\}$.²⁰ Applying Corollary 7.2 to each \widehat{p}_{ij} , we get that the set \widehat{P} can be computed in $O(\text{poly}(n, \log \lambda, \log \frac{1}{\eta}))$ time.

Next we will show that there exists an allocation $\widehat{\mathcal{S}}$ whose cut points are contained in \widehat{P} and this allocation has near-optimal social welfare, $\text{SW}(\widehat{\mathcal{S}}) \geq \text{SW}(\mathcal{S}^*) - \eta$. Specifically, for each cut point s_i^* of \mathcal{S}^* , let \widehat{s}_i denote its closest point in \widehat{P} , i.e., $\widehat{s}_i := \arg \min_{\widehat{p} \in \widehat{P}} |s_i^* - \widehat{p}|$; here we break ties, say, lexicographically. By construction of \widehat{P} we have $|s_i^* - \widehat{s}_i| \leq \frac{\eta}{n\lambda}$, for each $0 \leq i \leq n$.

For the optimal allocation $\mathcal{S}^* = \{S_1^*, \dots, S_n^*\}$ we have $S_i^* = [s_{i-1}^*, s_i^*]$ for all $i \in [n]$. Write allocation $\widehat{\mathcal{S}} := \{\widehat{S}_1, \widehat{S}_2, \dots, \widehat{S}_n\}$, where interval $\widehat{S}_i = [\widehat{s}_{i-1}, \widehat{s}_i]$ is assigned to agent $i \in [n]$. The cut points of allocations $\widehat{\mathcal{S}}$ are contained in \widehat{P} . Also, given that \mathcal{S}^* conforms to the MLRP order, so does $\widehat{\mathcal{S}}$.

Using the bounds $|\widehat{s}_i - s_i^*| \leq \frac{\eta}{n\lambda}$, for $0 \leq i \leq n$, and the fact that eval queries are λ -Lipschitz, we obtain $|v_i(\widehat{S}_i) - v_i(S_i^*)| \leq \frac{\eta}{n}$, for all agents $i \in [n]$. Summing over all $i \in [n]$, we obtain the desired social-welfare bound: $|\sum_{i \in [n]} v_i(\widehat{S}_i) - \sum_{i \in [n]} v_i(S_i^*)| = |\text{SW}(\widehat{\mathcal{S}}) - \text{SW}(\mathcal{S}^*)| \leq \eta$.

Therefore, there exists an allocation $\widehat{\mathcal{S}}$ with the properties that (i) $\widehat{\mathcal{S}}$ has near-optimal social welfare, (ii) cut points of $\widehat{\mathcal{S}}$ are contained in \widehat{P} , and (iii) $\widehat{\mathcal{S}}$ conforms to the MLRP order. To complete the proof of the theorem we will show that, among all allocations that satisfy properties (ii) and (iii), we can find one that maximizes social welfare. We accomplish this algorithmic result by a simple dynamic program.

Recall that cardinality of the set \widehat{P} is $O(n^2)$ and this set can be computed in $O(\text{poly}(n, \log \lambda, \log \frac{1}{\eta}))$ time using eval queries. We index the elements of the computed set $\widehat{P} = \{\widehat{p}_t\}_t$ such that $0 = \widehat{p}_0 < \widehat{p}_1 < \dots < \widehat{p}_{|\widehat{P}|} = 1$.

For each $k \in [n]$ and $1 \leq t \leq |\widehat{P}|$, we write $M(k, t)$ to denote the maximum social welfare that one can achieve by allocating the interval $[0, \widehat{p}_t]$ among the first k agents (in order).²¹

The following recursive equation for $M(k, t)$ gives us the desired dynamic program: $M(k, t) := \max_{1 \leq t' \leq t} \{M(k-1, t') + v_k(\widehat{p}_{t'}, \widehat{p}_t)\}$. Here, we initialize $M(1, t) := v_1(0, \widehat{p}_t)$ for all $1 \leq t \leq |\widehat{P}|$. One can directly show, via induction, that $M(n, |\widehat{P}|)$ is equal to the optimal social welfare among allocations that satisfy the above-mentioned properties (ii) and (iii). That is, $M(n, |\widehat{P}|) \geq \text{SW}(\mathcal{S}^*) - \eta$. Hence, the dynamic program gives us the desired allocation. For the runtime analysis, note that the dynamic program runs in $O(n^2 |\widehat{P}|)$ time and only requires eval queries.

Overall, we can find an allocation with social welfare η close to the optimal in time $O(\text{poly}(n, \log \lambda, \log \frac{1}{\eta}))$. Since the precision parameter η can be driven exponentially close to zero, in time that is polynomial in $\log \frac{1}{\eta}$ (i.e., in the bit complexity of η) the stated claim follows. \square

²⁰As in the case of P , we include 0 and 1 in \widehat{P} for ease of presentation.

²¹By convention, the agents are indexed following the MLRP order.

F DISTRIBUTION FAMILIES WITH MLRP

This section highlights that various well-studied distribution families bear MLRP. Recall, that two probability density functions f_i and f_j are said to satisfy MLRP iff, for every $x \leq y$ in the domain, we have

$$\frac{f_j(x)}{f_i(x)} \leq \frac{f_j(y)}{f_i(y)}.$$

That is, the likelihood ratio $f_j(x)/f_i(x)$ is non-decreasing in the argument $x \in \mathbb{R}$. It is relevant to note that this property can be verified analytically. We support this observation through two examples: Binomial polynomials (and, hence, linear value densities) in Proposition F.1 and Gaussian densities in Proposition F.2. Similar analytic arguments can be used to show that many other classes of value densities satisfy MLRP [9, 22, 29].

Binomial Polynomials: The following proposition shows that every pair of binomial polynomials bear MLRP over $[0, 1]$, i.e., we have a total order over this family of density functions with respect to MLRP.²² Hence, the results developed in the work hold for cake-division instances in which the value densities are binomial. Also, setting the exponent parameters $s = 1$ and $t = 0$ in this proposition, we observe that linear functions form a special case of binomial polynomials.

PROPOSITION F.1. *With integer exponents $s > t$, let $f_i(x) = a_i x^s + b_i x^t$ and $f_j(x) = a_j x^s + b_j x^t$ be two binomial polynomials. Then, f_i and f_j bear MLRP iff $a_i b_j - a_j b_i \leq 0$.*

PROOF. For any two points $x, y \in [0, 1]$, such that $x \leq y$, the MLRP condition for binomials corresponds to the following inequality

$$\frac{a_j x^s + b_j x^t}{a_i x^s + b_i x^t} \leq \frac{a_j y^s + b_j y^t}{a_i y^s + b_i y^t}$$

Since f_i and f_j constitute value densities with full support, the values of these functions are positive over the cake. Hence, the previous equation can be rewritten as

$$(a_j x^s + b_j x^t)(a_i y^s + b_i y^t) \leq (a_j y^s + b_j y^t)(a_i x^s + b_i x^t)$$

This equation further simplifies to

$$(a_i b_j - a_j b_i) x^t y^s \leq (a_i b_j - a_j b_i) x^s y^t$$

Finally, we rewrite the last inequality to obtain that for binomials MLRP is equivalent to

$$x^t y^t (a_i b_j - a_j b_i) (y^{s-t} - x^{s-t}) \leq 0 \tag{14}$$

Since $y - x \geq 0$ and $(s - t) \in \mathbb{Z}_+$, we have that²³ $y^{s-t} - x^{s-t} \geq 0$. Furthermore, since $x^t y^t \geq 0$, inequality (14) holds iff $a_i b_j - a_j b_i \leq 0$. That is, f_i and f_j bear MLRP iff $a_i b_j - a_j b_i \leq 0$. \square

Gaussian distributions: The next proposition shows that Gaussian distributions with different means, but the same variance, bear MLRP.

PROPOSITION F.2. *Let f_i and f_j be two Gaussian density functions with the same variance σ^2 and means $\mu_i \leq \mu_j$, respectively. Then, f_i and f_j satisfy MLRP.*

²²As mentioned previously, MLRP is preserved under scaling and, hence, here we do not have to explicitly enforce normalization.

²³For $(s - t) \in \mathbb{N}$, write $y^{s-t} - x^{s-t} = (y-x)(y^{s-t-1} + y^{s-t-2}x + \dots + yx^{s-t-2} + x^{s-t-1})$. Hence, $(y-x)$ and $(y^{s-t} - x^{s-t})$ have the same signs, for $0 \leq x \leq y \leq 1$.

PROOF. Here, the two density functions are $f_i(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu_i)^2}{2\sigma^2}}$ and $f_j(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu_j)^2}{2\sigma^2}}$, for $x \in \mathbb{R}$. Note that the likelihood ratio of f_j and f_i is equal to

$$\frac{e^{-(x-\mu_j)^2/2\sigma^2}}{e^{-(x-\mu_i)^2/2\sigma^2}} = \exp\left(\frac{1}{2\sigma^2} \left((x-\mu_i)^2 - (x-\mu_j)^2\right)\right)$$

Write $g(x) := \frac{1}{2\sigma^2} \left((x-\mu_i)^2 - (x-\mu_j)^2\right)$ and note that the derivate of this function $g'(x) = \frac{(\mu_j-\mu_i)}{\sigma^2} \geq 0$ if and only if $\mu_i \leq \mu_j$. Hence, g is an increasing function in x —and so is $\exp(g(x))$ —iff $\mu_i \leq \mu_j$. That is, f_i and f_j bear MLRP iff $\mu_i \leq \mu_j$. \square

An analogous result holds for Gaussians with mean zero, but distinct variances.

G BIT COMPLEXITY OF CAKE DIVISION

This section provides a cake-division instance (with MLPR) in which the unique envy-free allocation has an irrational cut point. Notably the parameters that specify the value densities in this instance are rational. This example implies that, in general, one cannot expect an efficient algorithm that outputs an exact envy-free allocation. That is, when considering cake-division algorithms with bounded bit complexity, a precision loss is inevitable.

We will also show, through the example, that to obtain a nontrivial bound on the envy (between the agents) the bit complexity of the output has to be $\Omega(\log \lambda)$; here $\lambda \geq 1$ is the Lipschitz constant of the cut and eval queries. Hence, a runtime dependence of $\log \lambda$ is unavoidable as well.

Consider cake division between three agents with identical value densities, $f : [0, 1] \mapsto \mathbb{R}_+$. Such an instance trivially satisfies MLRP, since the likelihood ratio of the value densities is equal to 1 throughout the cake. In particular, the density function f is piecewise linear and is defined as follows

$$f(x) := \begin{cases} \frac{\lambda}{3(\lambda-1)} & \text{if } 0 \leq x \leq (1 - \frac{1}{\lambda}) \\ \frac{2\lambda^2}{3}x + \left(\lambda - \frac{2\lambda^2}{3}\right) & \text{if } (1 - \frac{1}{\lambda}) \leq x \leq 1 \end{cases}$$

Note that this function can be given as input using only rational parameters. In addition, f is discontinuous at $1 - \frac{1}{\lambda}$; recall that our results require the value densities to be integrable, and not necessarily continuous.

Here, the values of the agents are normalized, $\int_0^1 f(x)dx = 1$. Also, the following bounds hold for the density: $0 < \frac{\lambda}{3(\lambda-1)} \leq f(x) \leq \lambda$ for all $x \in [0, 1]$. Therefore, Proposition C.3 implies that the cut and eval queries in this instance are 3λ -Lipschitz. Since in this instance the three agents have identical value densities (with full support over $[0, 1]$), there exists a unique envy-free allocation wherein each agent receives an interval of value exactly equal to $1/3$. Write $0 = x_0^* \leq x_1^* \leq x_2^* \leq x_3^* = 1$ to denote the cut points of this (unique) envy-free allocation; in particular, agent $i \in [3]$ receives the i th interval $[x_{i-1}^*, x_i^*]$.

First, we will show that x_2^* is irrational. Note that the interval $[0, 1 - 1/\lambda]$ is of value $1/3$: $\int_0^{1-\frac{1}{\lambda}} f(x)dx = \frac{1}{3}$. Hence, the first cut point $x_1^* = 1 - \frac{1}{\lambda}$. The second cut point x_2^* now lies in the interval $([x_1^*, 1])$ of width $1/\lambda$ and it must satisfy $\int_{1-\frac{1}{\lambda}}^{x_2^*} f(x)dx = \frac{1}{3}$. Using the definition of f in this range, we get that x_2^* is a solution of the following quadratic equation

$$\lambda^2(x_2^*)^2 + (3\lambda - 2\lambda^2)x_2^* + (\lambda^2 - 3\lambda + 1) = 0 \tag{15}$$

With the constraint $x_2^* \in (1 - \frac{1}{\lambda}, 1)$, equation (15) gives us $x_2^* = 1 - \left(\frac{3-\sqrt{5}}{2}\right)\frac{1}{\lambda}$. That is, the cut point x_2^* is irrational, for $\lambda \in \mathbb{Q}_+$.

We next establish a lower bound on the bit complexity of the output. Consider, in the above-mentioned instance, any allocation $\mathcal{I} = \{I_1, I_2, I_3\}$ wherein the envy between the agents is, say, less than $\eta = 1/2$, i.e., $v_i(I_i) \geq v_i(I_j) - 1/2$ for all $i, j \in [n]$.²⁴ For any such allocation \mathcal{I} one of the cut points must lie in $[1 - \frac{1}{\lambda}, 1]$. Indeed, this interval is of value $2/3$ to each agent, $\int_{(1-1/\lambda)}^1 f = 2/3$. The bit complexity of such a cut point is $\Omega(\log \lambda)$. Hence, in general, the bit complexity of the any algorithm that finds an allocation (equivalently, outputs cut points) with bounded envy is $\Omega(\log \lambda)$.

Welfare-maximizing allocations induced by irrational cuts: One can also construct cake-division instances (with rational and MLRP value densities) in which the welfare-maximizing allocations have irrational cut points.

In particular, consider cake division between two agents with identical value densities, $f(x) = x + 1/2$. Here, the (unique) allocation $\mathcal{I} = \{I_1, I_2\}$ that maximizes egalitarian welfare consists of the intervals $I_1 = [0, \frac{\sqrt{5}-1}{2}]$ and $I_2 = [\frac{\sqrt{5}-1}{2}, 1]$. Allocation \mathcal{I} is also the (unique) equitable, proportional, and envy-free allocation in this instance.

In addition, consider a cake-division instance with two agents and the following value densities: $f_1(x) = 1$ and $f_2(x) = 3x^2$. Note that these densities satisfy MLRP. In this instance, the (unique) social welfare maximizing allocation $\mathcal{S} = \{S_1, S_2\}$ is obtained by an irrational cut; specifically, $S_1 = [0, \frac{1}{\sqrt{3}}]$ and $S_2 = [\frac{1}{\sqrt{3}}, 1]$. The cut point $\frac{1}{\sqrt{3}}$ is the switching point (as defined in Section 7) between the two densities f_1 and f_2 .

H ROBUSTNESS OF MLRP

This section shows that our framework extends to the class of non-full-support value densities considered in [1]. In particular, Alijani et al. [1] established that a contiguous envy-free cake division can be efficiently computed if every agent $i \in [n]$ uniformly values a single interval $[\ell_i, r_i] \subseteq [0, 1]$ and these intervals satisfy the following ordering property OP: for all $i, j \in [n]$ we have $\ell_i \leq \ell_j$ iff $r_i \leq r_j$. This ordering property OP ensures that, in particular, any interval $[\ell_i, r_i]$ is not a strict subset of $[\ell_j, r_j]$ for $i, j \in [n]$.

Let $\mathcal{K} = \langle [n], \{f_i\}_{i \in [n]} \rangle$ denote a cake-division instance with such value densities and note that, for each $i \in [n]$,

$$f_i(x) := \begin{cases} \frac{1}{r_i - \ell_i} & \text{if } x \in [\ell_i, r_i] \\ 0 & \text{otherwise} \end{cases} \quad (16)$$

Write $h_i := 1/(r_i - \ell_i)$. The work of Alijani et al. [1] shows that an envy-free allocation of \mathcal{K} can be computed efficiently.

Indeed, the value densities in \mathcal{K} do not have full support over the cake. However, we will show that we can perturb these densities f_i s (in a structured manner) to obtain an instance $\widehat{\mathcal{K}} = \langle [n], \{\widehat{f}_i\}_{i \in [n]} \rangle$ such that (i) the value densities \widehat{f}_i s in $\widehat{\mathcal{K}}$ have full support and bear MLRP (Claim 1) and (ii) any envy-free allocation in $\widehat{\mathcal{K}}$ forms an envy-free allocation, up to a small precision loss, in the original instance \mathcal{K} (Claim 2).

Index the agents such that $0 \leq \ell_1 \leq \ell_2 \leq \dots \leq \ell_n \leq 1$. Property OP ensures that, under this indexing, the right endpoints r_i s are also sorted: $r_1 \leq r_2 \leq \dots \leq r_n$. To construct the densities \widehat{f}_i , for each agent i , we will define counting functions $c_i : [0, 1] \mapsto \{0, 1, \dots, i\}$ and $d_i : [0, 1] \mapsto \{0, 1, 2, \dots, n-i\}$ as follows: $c_i(x) := \sum_{k=1}^i \mathbb{1}\{\ell_k \leq x\}$ and $d_i(x) := \sum_{k=i}^n \mathbb{1}\{r_k \leq x\}$, i.e., $c_i(x)$ is the

²⁴Here, the choice of $\eta = 1/2$ is essentially for ease of exposition; by scaling down the density in the range $[0, (1-1/\lambda)]$, we can drive η close to one.

number of left endpoints in $\{\ell_1, \ell_2, \dots, \ell_i\}$ that appear before x and $d_i(x)$ is equal to the number of right endpoints in $\{r_i, r_{i+1}, \dots, r_n\}$ that appear before x .

For a sufficiently large parameter $H \geq 1$ (which we will set in Claim 2), we define $\widehat{f}_i : [0, 1] \mapsto \mathbb{R}_+$ as follows²⁵

$$\widehat{f}_i(x) := h_i H^{(c_i(x)-i-d_i(x))} \quad (17)$$

In other words, we initialize $\widehat{f}_i(0) = h_i H^{-i}$ and scale up this density by a multiplicative factor of H whenever an interval $[\ell_k, r_k]$ starts, for any agent $k \leq i$. Note that $c_i(x) = i$ and $d_i(x) = 0$ for all $x \in [\ell_i, r_i]$ and, hence, in this interval $\widehat{f}_i(x) = h_i = f_i(x)$. Beyond this range, i.e., for $x \geq r_i \geq \ell_i$, we have $c_i(x) = i$. Therefore, once the interval $[\ell_i, r_i]$ ends, we scale down \widehat{f}_i by a multiplicative factor of H , whenever an interval $[\ell_j, r_j]$ ends, for any agent $j \geq i$. Furthermore, for all $x \notin [\ell_i, r_i]$, we have $\widehat{f}_i(x) \leq h_i/H$. Also, as required, \widehat{f}_i s have full support over $[0, 1]$.

Applying Proposition C.3, we obtain that the Lipschitz constant of cut and eval queries in $\widehat{\mathcal{K}}$ is $\lambda = H^n \max_{1 \leq i \leq n} h_i$. The next claim shows that the the value densities in $\widehat{\mathcal{K}}$ satisfy MLRP.

CLAIM 1. *Let $\mathcal{K} = \langle [n], \{f_i\}_{i \in [n]} \rangle$ be a cake-division instance in which the value densities satisfy equation (16) and the ordering property OP. Then, for parameter $H \geq 1$, the value densities \widehat{f}_i s (as defined in equation (17)) satisfy MLRP.*

PROOF. In this constructed instance $\widehat{\mathcal{K}} = \langle [n], \{\widehat{f}_i\}_{i \in [n]} \rangle$ the agents are indexed such that $0 \leq \ell_1 \leq \ell_2 \leq \dots \leq \ell_n \leq 1$ and, by OP, we have $r_1 \leq r_2 \leq \dots \leq r_n$. Fix any two agents $i, j \in [n]$ such that $i < j$. The indexing among the agents ensures that $\ell_i \leq \ell_j$ and $r_i \leq r_j$.

To prove the stated claim it suffices to show that the likelihood ratio $\frac{\widehat{f}_j(x)}{\widehat{f}_i(x)}$ is nondecreasing throughout $[0, 1]$. We establish the monotonicity of this likelihood ratio by considering three different ranges in the cake $[0, \ell_i]$, $[\ell_i, r_j]$, and $[r_j, 1]$.

Initially, the likelihood ratio $\frac{\widehat{f}_j(0)}{\widehat{f}_i(0)} = \left(\frac{h_j}{h_i}\right) \frac{1}{H^{-j}}$. For all $x \in [0, \ell_i]$ whenever the density $\widehat{f}_i(x)$ experiences a multiplicative increase so does $\widehat{f}_j(x)$, by the same factor H . In particular, if c_i increases then so does c_j throughout the interval $[0, \ell_i]$: if at a point $x \in [0, \ell_i]$ the density \widehat{f}_i is scaled up by a multiplicative factor H (i.e., c_i increases by one), then it must have been the case that an interval $[\ell_k, r_k]$, with $k \leq i$, starts at $\ell_k = x$. In such a case, by definition, \widehat{f}_j also increases by a multiplicative factor H (i.e., c_j increases by one as well). Therefore, for all $x \in [0, \ell_i]$, the likelihood ratio stays constant at $\frac{\widehat{f}_j(0)}{\widehat{f}_i(0)}$.

In the interval $[\ell_i, r_j]$, the density f_i does not increase, since here c_i saturates at i after ℓ_i and d_i increases from zero beyond r_i . At the same time, in this interval $[\ell_i, r_j]$ the density f_j is nondecreasing. Therefore, the likelihood ratio continues to be monotonic in the interval $[\ell_i, r_j]$ as well.

Finally, for $x \geq r_j$, we note that d_i and d_j increase synchronously. Also, for points $x \in [r_j, 1]$ we have $c_j(x) = j$ and $c_i(x) = i$. Therefore, in this range the likelihood ratio stays constant at $\frac{\widehat{f}_j(r_j)}{\widehat{f}_i(r_j)}$. Overall, we obtain the monotonicity of the likelihood ratio and the MLRP guarantee follows. \square

CLAIM 2. *Given a cake-division instance $\mathcal{K} = \langle [n], \{f_i\}_{i \in [n]} \rangle$ in which the value densities satisfy equation (16) and the ordering property OP. Let $\widehat{\mathcal{K}} = \langle [n], \{\widehat{f}_i\}_{i \in [n]} \rangle$ be the cake-division instance defined above, with parameter $H := \frac{2}{\eta} \max_{1 \leq i \leq n} h_i$, and suppose that allocation $I = \{I_1, \dots, I_n\}$ is*

²⁵The integrals of the value densities \widehat{f}_i s are not normalized to 1, though, as mentioned previously, MLRP is preserved under scaling and, hence, we do not have to explicitly enforce normalization.

envy-free up to an additive factor of η in $\widehat{\mathcal{K}}$ (i.e., $v_i(I_i) \geq v_i(I_j) - \eta$ for all $i, j \in [n]$). Then, allocation \mathcal{I} is envy-free up to 2η in \mathcal{K} .

PROOF. For agent $i \in [n]$, write v_i and \widehat{v}_i to denote the valuation function of agent i under f_i and \widehat{f}_i , respectively. By construction, the densities f_i and \widehat{f}_i coincide over the interval $[\ell_i, r_i]$; in particular, $f_i(x) = \widehat{f}_i(x) = h_i$ for all $x \in [\ell_i, r_i]$. Hence, for any interval I we have $v_i(I \cap [\ell_i, r_i]) = \widehat{v}_i(I \cap [\ell_i, r_i])$.

Furthermore, the construction (equation (17)) of \widehat{f}_i ensures that $\widehat{v}_i(I \setminus [\ell_i, r_i]) \leq \frac{h_i}{H} \leq \eta/2$. The last inequality follows from the choice of the parameter H . Note that $f_i(x) = 0$ for all $x \notin [\ell_i, r_i]$ and, hence, $v_i(I \setminus [\ell_i, r_i]) = 0$. These bounds imply that, for any agent $i \in [n]$, the difference in values of any interval I is at most $\eta/2$:

$$|\widehat{v}_i(I) - v_i(I)| \leq \frac{\eta}{2} \quad (18)$$

Hence, an allocation $\mathcal{I} = \{I_1, \dots, I_n\}$ that is envy-free up to an additive factor of η in $\widehat{\mathcal{K}}$ (i.e., $\widehat{v}_i(I_i) \geq \widehat{v}_i(I_j) - \eta$ for all $i, j \in [n]$) is envy-free up to 2η in \mathcal{K} . Specifically, for any $i, j \in [n]$,

$$\begin{aligned} v_i(I_i) &\geq \widehat{v}_i(I_i) - \eta/2 && \text{(using inequality (18) with interval } I_i) \\ &\geq \widehat{v}_i(I_j) - \frac{3\eta}{2} && \text{(by } \eta\text{-envy-freeness of } \mathcal{I}) \\ &\geq v_i(I_j) - 2\eta && \text{(using inequality (18) with interval } I_j) \end{aligned}$$

That is, \mathcal{I} is envy-free, up to 2η precision, in \mathcal{K} . This completes the proof. \square

Since the constructed instance $\widehat{\mathcal{K}}$ satisfies MLRP, we can use Algorithm 1 (Section 5) to efficiently compute, up to an arbitrary precision, an envy-free allocation \mathcal{I} in the instance $\widehat{\mathcal{K}}$. The previous claim ensures that \mathcal{I} is an envy-free allocation (up to an arbitrary precision) in the original instance \mathcal{K} as well. This observation highlights the fact that the ideas developed in this work are somewhat robust and extend to other value-density settings.

I NONEXISTENCE OF CONTIGUOUS PERFECT DIVISIONS UNDER MLRP

A cake division $\mathcal{D} = \{D_1, D_2, \dots, D_n\}$ (consisting of connected or disconnected pieces) is said to be *perfect* if all the agents agree on the value of every piece, i.e., $v_i(D_j) = 1/n$ for all $i, j \in [n]$.

Perfect divisions are not guaranteed to exist under the contiguity requirement, i.e., there are cake-division instances that do not admit perfect allocations. In this section, we will show that the nonexistence continues to hold even with MLRP. That is, we will provide a cake-division instance with MLRP that does not admit a perfect allocation.

Consider a cake-division instance $\mathcal{C} = \langle [2], \{f_1, f_2\} \rangle$ with two agents. For a fixed constant $\alpha \in (0, 1)$, we define the value densities of the two agents $f_1, f_2 : [0, 1] \mapsto \mathbb{R}_+$ as follows

$$f_1(x) = \begin{cases} 1 + \alpha & \text{if } 0 \leq x \leq (1 - \alpha) \\ \alpha & \text{if } (1 - \alpha) < x \leq 1 \end{cases} \quad \text{and} \quad f_2(x) = \begin{cases} 1 - \alpha & \text{if } 0 \leq x \leq (1 - \alpha) \\ 2 - \alpha & \text{if } (1 - \alpha) < x \leq 1 \end{cases}$$

The agents' values for the cake $[0, 1]$ are normalized, $\int_0^1 f_1(x) dx = \int_0^1 f_2(x) dx = 1$. Also, the likelihood ratio of f_1 and f_2 satisfies

$$\frac{f_2(x)}{f_1(x)} = \begin{cases} \frac{1-\alpha}{1+\alpha} & \text{if } 0 \leq x \leq (1 - \alpha) \\ \frac{2-\alpha}{\alpha} & \text{if } (1 - \alpha) < x \leq 1 \end{cases}$$

Since $\frac{2-\alpha}{\alpha} > \frac{1-\alpha}{1+\alpha}$, for all $\alpha \in (0, 1)$, the likelihood ratio is nondecreasing over $[0, 1]$ and, hence, the densities bear MLRP.

We assume, towards a contradiction that there exists a perfect allocation in C . That is, there exists a point $x \in [0, 1]$ such that $v_1(0, x) = v_2(0, x) = 1/2$, and $v_1(x, 1) = v_2(x, 1) = 1/2$. The fact that the value of intervals $[0, x]$ and $[x, 1]$ is equal to $1/2$ ensures that the point x cannot be 0 or 1. We consider two complementary and exhaustive cases: Case (i) $0 < x \leq (1 - \alpha)$ and Case (ii) $(1 - \alpha) < x \leq 1$. The analysis is similar in both the cases. Hence, we only address Case (i) and omit the analysis for Case (ii).

For the first interval $[0, x]$ (with $0 < x \leq (1 - \alpha)$), we have $v_1(0, x) = x(1 + \alpha)$ and $v_2(0, x) = x(1 - \alpha)$. However, for any $\alpha \in (0, 1)$ and $x > 0$, the following strict inequality holds: $v_1(0, x) = x(1 + \alpha) > x(1 - \alpha) = v_2(0, x)$. This leads to a contradiction and proves that C does not admit a perfect division with connected pieces.

Even though we might not have a perfect allocation, the work of Alon [2] proves that a perfect division with $n(n - 1)$ cuts always exists. Hence, in the above-mentioned instance C with 2 agents, 2 cuts should suffice to form a perfect division. In particular, we note that the following division $\mathcal{D}^* = \{D_1^*, D_2^*\}$ is perfect in C ; here $D_1^* = [\frac{1}{2} - \frac{\alpha}{2}, 1 - \frac{\alpha}{2}]$ and $D_2^* = [0, \frac{1}{2} - \frac{\alpha}{2}] \cup [1 - \frac{\alpha}{2}, 1]$. Here,

$$v_1(D_2^*) = (1 + \alpha) \left(\frac{1}{2} - \frac{\alpha}{2} \right) + \alpha \left(\frac{\alpha}{2} \right) = \frac{1}{2}$$

and

$$v_2(D_2^*) = (1 - \alpha) \left(\frac{1}{2} - \frac{\alpha}{2} \right) + (2 - \alpha) \left(\frac{\alpha}{2} \right) = \frac{1}{2}$$

That is, both the agents value the piece D_2^* at $1/2$. Since the valuations are normalized, we additionally have $v_1(D_1^*) = v_2(D_1^*) = 1/2$. This shows that \mathcal{D}^* is a perfect division (with disconnected pieces) in C .