**QUEEN'S
UNIVERSITY
BELFAST**

Fairness and efficiency in cake-cutting with single-peaked preferences

Bhardwaj, B., Kumar, R., & Ortega, J. (2020). Fairness and efficiency in cake-cutting with single-peaked preferences. *Economics Letters*, 190, [109064]. <https://doi.org/10.1016/j.econlet.2020.109064>

Published in:
Economics Letters

Document Version:
Peer reviewed version

Queen's University Belfast - Research Portal:
[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights

Copyright 2020 Elsevier.

This manuscript is distributed under a Creative Commons Attribution-NonCommercial-NoDerivs License

(<https://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits distribution and reproduction for non-commercial purposes, provided the author and source are cited.

General rights

Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

Fairness and Efficiency in Cake-Cutting with Single-Peaked Preferences

Bhavook Bhardwaj^a, Rajnish Kumar^b, Josué Ortega^b

^a*Indian Statistical Institute, Delhi, India.*

^b*Queen's Management School, Queen's University Belfast, UK.*

Abstract

We study the cake-cutting problem where agents have single-peaked preferences over the cake. We show that a recently proposed mechanism by [Wang and Wu \(2019\)](#) to obtain envy-free allocations can yield large welfare losses. Using a simplifying assumption, we characterize all Pareto optimal allocations, which have a simple structure: are peak-preserving and non-wasteful. Finally, we provide simple alternative mechanisms that Pareto dominate that of Wang–Wu and achieve envy-freeness or Pareto optimality.

Keywords: cake-cutting, fair division, single-peaked preferences.

JEL Codes: C78.

Email addresses: bhavook17r@isid.ac.in, rajnish.kumar@qub.ac.uk, j.ortega@qub.ac.uk

1. Introduction

The problem of how to fairly divide a heterogeneous good among n agents is one of the oldest problems in fair division (Steinhaus, 1948). Known as cake-cutting, this problem has spanned a long interdisciplinary literature in economics, computer science, mathematics and political science that includes multiple real-life applications, such as the division of land, inheritances, and computing resources.

One of the main goals of the cake-cutting literature is the construction of division methods that produce envy-free allocations. Such methods exist, but they require a large amount of information regarding agents' preferences (Brams and Taylor, 1995; Aziz and Mackenzie, 2016). However, when agents' preferences are well-behaved, an envy-free allocation can be computed with few queries about agents' preferences. In particular, when agents have single-peaked preferences, an envy-free allocation can be achieved with only $2n$ queries, using a clever procedure recently proposed by Wang and Wu (2019).¹ Single-peaked preferences are particularly interesting to study for two reasons. First, they can be described concisely: agents are specified as peak points, who prefer the pieces of cake that are close to their own locations. Second, single-peaked preferences can realistically describe preferences over land resources, clean water and mineral deposits; and have been thoroughly studied in other fair division problems.²

Although the Wang–Wu procedure achieves envy-freeness with low informational requirements, it may produce allocations that are Pareto dominated. In this note, we show that the Wang–Wu procedure may waste up to, but not more than, $\frac{n-1}{n}$ of the total achievable utility (Proposition 1). Obtaining a Pareto optimal allocation is easy only when single-peaked preferences have a common slope; in this case they have a simple structure and are completely characterized by being peak-preserving and non-wasteful (Proposition 2). Under those assumptions, we can construct two simple mechanisms that Pareto dominate the Wang–Wu mechanism, while attaining envy-freeness or Pareto optimality. Alternatively, the utilitarian mechanism that returns the allocation that maximizes the sum of agents' utilities is another reasonable option that is efficient and also easy to describe (Proposition 3, Table 1).

¹In the more general class of piecewise linear preferences, an envy-free allocation can be computed using $\Omega(n^6)$ queries (Kurokawa et al., 2013). A connected envy-free allocation cannot be obtained by a mechanism with finite queries (Stromquist, 2008).

²E.g. Moulin (1984); Sprumont (1991); Ehlers et al. (2002); Kasajima (2013); Long (2019), among others. Single-peaked preferences have also been (implicitly) used in cake-cutting (Weller, 1985, p. 17).

2. Model

We consider the standard cake-cutting problem (Nicolò and Yu, 2008; Procaccia, 2016). The cake is the $[0, 1]$ interval and the set of agents is $N = \{1, \dots, n\}$. A piece of cake P is a finite set of disjoint subintervals of $[0, 1]$. Each agent $i \in N$ has a integrable, non-negative value density function v_i . For each piece of cake P , an agent's value or utility on P is denoted by $V_i(P)$ and defined by the integral of its density function, i.e., $V_i(P) := \int_P v_i(x)dx$. Thus, $V_i : 2^{[0,1]} \rightarrow \mathbb{R}^+$ is called the valuation or utility function of agent i . The definition of the valuation function implies that it is additive and non-atomic. We simplify notation by writing $V_i(x, y)$ instead of $V_i([x, y])$ for each interval $[x, y]$. We normalize agents' valuations so that $V_i(0, 1) = 1$ for each $i \in N$. A *cake-cutting problem* is a triple $([0, 1], N, (v_i)_{i \in N})$.

We assume that agents' preferences over the cake are single-peaked. Agent i has a *single-peaked valuation* when her value density function v_i satisfies that there is a number $p_i \in [0, 1]$ such that $v_i(p_i) = \max_{x \in [0,1]} v_i(x)$ and $v_i(x) = \max\{0, v_i(p_i) - k_i |x - p_i|\}$ for all $x \in [0, 1]$ and some coefficient $k_i > 0$. We call p_i and $v_i(p_i)$ the peak location and peak density of agent i , respectively. Notice that the coefficient k_i as well as the density function is uniquely determined by the peak location and peak density as the normalized assumption says $\int_0^1 v_i(x)dx = 1$. We use the notation U_i to denote the maximum subinterval such that agent i has positive value on every point, i.e. $v_i(x) > 0$ for all $x \in U_i$ and $v_i(x) = 0$ for all $x \in [0, 1] \setminus U_i$. The points l_i and r_i denote the left and right endpoints of U_i . We assume without loss of generality that $\bigcup_{i \in N} U_i = [0, 1]$.³

An *allocation* $X = (X_1, \dots, X_n)$ is a partition of the cake among the n agents, i.e. each agent i receives the piece of cake X_i , the pieces are disjoint, and $\bigcup_{i \in N} X_i = [0, 1]$. An allocation X is *proportional* if $V_i(X_i) \geq 1/n$ for all $i \in N$, and *envy-free* (EF) if $V_i(X_i) \geq V_i(X_j)$ for all $i, j \in N$. Envy-freeness implies proportionality, and they are equivalent for $n = 2$. None of these fairness properties, however, imply efficiency in the sense of Pareto optimality. An allocation X is *Pareto optimal* (PO) if there is no other allocation X' that Pareto dominates it, i.e. such that, for all $i \in N$, $V_i(X'_i) \geq V_i(X_i)$; and for some i , $V_i(X'_i) > V_i(X_i)$.

A mechanism M is a function mapping a series of queries to agents about their valuations into an allocation. In the standard Robertson and Webb (1998) framework, in which we focus, only two types of queries are allowed: either an agent is asked to *cut* the cake at an specific point so that the left

³Pieces of cake that are not desired by anybody can be dropped or assigned arbitrarily, and play no role with regards to fairness and efficiency considerations.

piece has a specific valuation, or an agent is asked to *evaluate* a piece of cake, i.e. reveal her valuation for such a piece. Mechanism M is said to be envy-free (resp. proportional, Pareto optimal) if it always produces an envy-free (resp. proportional, Pareto optimal) allocation.

3. Results

Wang and Wu (2019) provide an ingenious mechanism to obtain envy-free allocations with only $2n$ queries. Their mechanism asks two cut queries to each agent (in the standard Robertson–Webb model) to find the l_i, p_i, r_i points (see their Algorithm 1). Then, all those points are ordered in increasing sequence, together with the 0 and 1 endpoints. Between any two of these points, the cake is divided into $2n$ equidistant pieces, and agent 1 receives pieces 1 and $2n$, agent 2 receives pieces 2 and $2n - 1$, and so on until agent n , who receives pieces n and $n + 1$ (see Figure 1 for an example). Because valuations are single-peaked, their procedure computes an envy-free allocation with only $2n$ queries (i.e. substantially less queries than those required in more general preference domains, see Kurokawa et al. (2013)).

However, a problem with the Wang–Wu mechanism is that it may produce an allocation that is not Pareto optimal, as they point out. In fact, the Wang–Wu procedure only produces a Pareto optimal allocation in a very special case: when $p_i = p_j$ and $k_i = k_j$ for all $i, j \in N$.

Moreover, the welfare losses than can occur in the Wang–Wu mechanism can be substantially large, as the example in Figure 1 shows. In our example, the Wang–Wu procedure gives to each agent only $1/n$ of the utility that they obtain in the unique PO allocation. The example can be extended to any number of agents to show the severity of the welfare losses in the Wang–Wu procedure *vis-a-vis* Pareto optimal allocations. To formalise this observation, let $\mathcal{X}^{PO}$ be the set of all PO allocations and X^{WW} be the allocation suggested by the Wang–Wu procedure. Let $T^{PO} := \min_{X \in \mathcal{X}^{PO}} \frac{\sum_{i \in N} V(X_i)}{n}$ and $T^{WW} := \frac{\sum_{i \in N} V(X_i^{WW})}{n}$.⁴ The welfare losses in the Wang–Wu mechanism are defined as $WL := T^{PO} - T^{WW}$. Our first Proposition provides a tight upper bound for WL .

⁴The choice of the Pareto optimal allocation with the minimum sum of utilities is arbitrary. Because in the proof we construct a cake-cutting problem with a unique Pareto optimal allocation, our result can be rephrased using $T^{PO} := \max_{X \in \mathcal{X}^{PO}} \frac{\sum_{i \in N} V(X_i)}{n}$. Our definition of welfare losses is related to the *price of fairness* (Bertsimas et al., 2011).

Figure 1: Agents 1, 2, and 3 receive the grey, white, and black pieces respectively. In the unique PO allocation (which is also EF), the sum of individual valuations is 3, whereas in the Wang–Wu allocation, the sum is only 1.

Proposition 1. *For any n , the welfare losses in the Wang–Wu mechanism WL can be up to, but no more than, $\frac{n-1}{n}$.*

Proof. To show that the upper bound of $\frac{n-1}{n}$ is achievable, construct a cake-cutting problem such that $\bigcap_{i \in N} U_i = \emptyset$. Then, there is a unique Pareto optimal allocation X^{PO} such that $V_i(X_i^{PO}) = 1$ for all $i \in N$. In contrast, in the Wang–Wu procedure $[l_i, p_i]$ is divided into $2n$ equidistant pieces, of which agent i receives 2 that she jointly values at $1/2n$; similarly she receives 2 of the $2n$ equidistant pieces in which the interval $[p_i, r_i]$ is divided, which she also jointly values at $1/2n$. Thus, $V_i(X_i^{WW}) = 1/n$ for all $i \in N$, and thus $T^{PO} - T^{WW} = \frac{n-1}{n}$.

To show that this upper bound is tight, note that the Wang–Wu procedure is envy-free, and thus proportional, so that $V_i(X_i^{WW}) \geq 1/n$, and thus, for any cake-cutting problem, $T^{PO} - T^{WW} \leq \frac{n-1}{n}$. $\square$

The trade-off between efficiency and envy-freeness is not specific to the Wang–Wu procedure, and is present in several fair division problems (e.g. [Maniquet and Sprumont, 2000](#)), including cake-cutting.⁵ To understand this trade-off better, we provide a characterization of all Pareto-optimal mecha-

⁵Large welfare losses also arise in matching problems ([Ortega, 2018, 2019](#)).

nisms. This is substantially easier to do when all valuation density functions v_i have the same slope, i.e. $k_i = k$ for all $i \in N$, an assumption that we impose from now.⁶ We focus on the generic case in which all peaks are different, i.e. $p_i \neq p_j$ for all $i, j \in N$.⁷

An allocation X is *non-wasteful* if $X_i \subseteq U_i$ for all i . A connected allocation is completely described by $n - 1$ cutting points $c_1 \leq \dots \leq c_{n-1}$ such that $X_1 = [0, c_1]$, $X_2 = (c_1, c_2]$, and so on. A connected allocation is *peak-preserving* if the interval $[0, c_1]$ is allocated to the agent with the smallest peak point, $(c_1, c_2]$ is allocated to the agent with the second smallest peak point, and so on, until the piece $(c_{n-1}, 1]$ is allocated to the agent with the largest peak point. In our second Proposition, we show that non-wastefulness and peak-preservingness completely characterize Pareto optimal allocations, which implies that all Pareto optimal allocations are connected.

Proposition 2. *When $k_i = k$ for all $i \in N$, an allocation is Pareto optimal if and only if it is a non-wasteful and peak-preserving allocation.*

Proof. First, we show (by contradiction) that if X is Pareto optimal, then it must be non-wasteful and peak preserving. If X was wasteful, some agent i would receive a part of the cake that she does not value, but which by assumption is valued by someone else. A clear Pareto improvement exists, and thus X is not Pareto optimal, a contradiction.

If X was not peak-preserving, it means it is either not connected or it is connected but it does not respect peaks. In both cases, there must exist two agents i, j such that $p_i < p_j$, $[a, b] \subseteq X_j$ and $(b, c] \subseteq X_i$, where $a < b < b' < c$ (see subfigure 2.a).⁸ We show geometrically in Figure 2 that a reassignment of the interval $[a, c]$ such that i receives $[a, b']$ and j receives $[b', c]$ leads to a Pareto improvement, where b' is chosen so that $V_j(b', c) = V_j(a, b)$.⁹

Second, we prove that if X is non-wasteful and peak-preserving, then it must be Pareto optimal. By contradiction, suppose X is Pareto dominated

⁶One may consider the following broader preference domain: if $l_i \leq l_j$, then $r_i \leq r_j$ for all $i, j \in N$. Our results do not extend to this more general preference domain.

⁷Our results extend to the case when some agents have the same peaks *mutatis mutandis*. In such a case, an allocation may be disconnected only between agents who have the same peaks.

⁸That agents i, j as above must exist is trivial when X does not respect peaks. If X is not connected, there is an agent α who obtains at least two disconnected pieces of cake $[\alpha^0, \alpha']$ and $(\alpha'', \alpha''']$, and agents β and γ with allocations $(\alpha', \beta']$ and $(\gamma', \alpha'']$, such that $\alpha' < \beta' \leq \gamma' < \alpha''$. If $p_\alpha > p_\beta$, then $i = \beta$ and $j = \alpha$. Otherwise, if $p_\alpha < p_\beta$ and $p_\alpha < p_\gamma$, then $i = \alpha$ and $j = \gamma$. And if $p_\alpha < p_\beta$ and $p_\alpha > p_\gamma$, then $i = \gamma$ and $j = \beta$.

⁹A similar argument is used to show that, when two agents divide several objects, at most one object is divided in all Pareto optimal allocations (Moulin, 2004, p. 255). See also Corollary 2.4 in Sandomirskiy and Segal-Halevi (2019).

(a) A non peak-preserving allocation.

(b) A Pareto improvement.

Figure 2: In subfigure 2.a, we have a non peak-preserving allocation in which agent j gets $[a, b]$ and agent i gets $(b, c]$ (the sum of utilities is the shaded area). In subfigure 2.b, we show how such allocation is Pareto dominated by a peak-preserving allocation. We reassign the interval $[a, c]$ so that j gets an interval $(b', c]$ such that $V_j(b', c) = V_j(a, b)$. The difference $V_i(a, b') + V_j(b', c) - V_j(a, b) - V_i(b, c)$ is equal to the area of polygon 1–2–6–5 (coordinates $(a, v_j(a)) - (a, v_i(a)) - (b', v_i(b')) - (b', v_j(b'))$) minus the area of the polygon 3–4–8–7 (coordinates $(b, v_j(b)) - (b, v_i(b)) - (c, v_i(c)) - (c, v_j(c))$). The former is strictly larger than the latter because $b' > b$, due to the fact that $c \leq p_j$. Since by construction $V_j(b', c) = V_j(a, b)$, we must have that $V_i(a, b') > V_i(b, c)$, showing that the allocation in subfigure 2.a is Pareto dominated. The argument is completely symmetrical when the allocation endpoints a, b, c are located elsewhere.

by an alternative allocation X' so that for some agent i , $V_i(X'_i) > V_i(X_i)$ and $V_j(X'_j) \geq V_j(X_j)$ for all $j \in N$. X' can be assumed to be Pareto optimal, otherwise there exists another allocation X'' that Pareto dominates both X' and X (we can reach the Pareto optimal allocation by trading among agents). So let X' be Pareto optimal, which by our previous argument must be non-wasteful and peak preserving. Both X and X' can be completely described by a sequence of $n-1$ cut points, called $c = (c_1, \dots, c_{n-1})$ and $c' = (c'_1, \dots, c'_{n-1})$, with endpoint $c_n = c'_n = 1$. Let s be the smallest index for which $c'_s < c_s$, if such exists. Then, because X is non-wasteful, $V_s(X_s) > V_s(X'_s)$, and thus X' does not Pareto dominate X . If no such index exist, $V_n(X_n) > V_n(X'_n)$, and thus X' does not Pareto dominate X . We conclude that X must be Pareto optimal. $\square$

The assumption of common slopes is critical to obtain such a simple characterization of Pareto optimality: when each agent's valuation density has a different slope, Pareto optimal allocations may be disconnected, as Figure 3 shows.

4. Three different mechanisms that improve on Wang–Wu

We turn to the design of a mechanism that returns an allocation that is both Pareto optimal and envy-free. Although such an allocation exists

Figure 3: An cake-cutting problem with single-peaked valuations and different slopes. The allocation in which agent i obtains the grey parts, and agent j obtains the black parts is Pareto optimal, but not connected.

(Weller, 1985), constructing an algorithm that returns it (in the Robertson–Webb model) remains an open question when preferences are single-peaked (and is impossible in more general preference domains (Kurokawa et al., 2013)). Nevertheless, we provide three deterministic mechanisms that achieve both properties individually knowing only agents’ peaks and endpoints, and that, at the same time, improve on the original Wang–Wu procedure by either Pareto dominating it or generating a higher total utility. We note that the three mechanisms we describe below are in the Robertson–Webb model, since the peak and endpoints of valuations’ densities can be recovered using Algorithm 1 in Wang and Wu (2019). Figure 4 exemplifies how our three different solutions differ and improve on the Wang–Wu mechanism.

1. *Utilitarian Mechanism* (UM). Our first mechanism computes the allocation that maximizes the sum of agents’ utilities, i.e. the area under the upper envelope of the valuation density functions. Computing such allocation is easy: it is the peak-preserving allocation described by $n - 1$ cut points $c_1^u, \dots, c_{n-1}^u$, each located at the intersection of the valuation density functions, i.e. c_i^u is such that $v_i(c_i^u) = v_{i+1}(c_{i+1}^u)$. The utilitarian mechanism is Pareto optimal (and arguably the best one among all Pareto optimal ones, since it maximizes the sum of utilities), but it may create envy.¹⁰

¹⁰The utilitarian mechanism maximizes the sum of utilities, but maximizing other aggregate utility functions may also be of interest, such as the Nash product (Segal-Halevi and Sziklai, 2019). See Juarez and Kumar (2013) for further applications of the utilitarian mechanism.

2. *Leftmost leaves* (LL). Our second mechanism is leftmost leaves and is also simple. Index agents according to their peaks, so that agent 1 is the one with the smallest peak, and let c_1^{ll} be the point for which $V_1[0, c_1^{\text{ll}}] = 1/n$ and $c_1^{\text{ll}} = \max\{c_1^{\text{ll}}, l_2\}$. Agent 1 is assigned the interval $[0, c_1^{\text{ll}}]$. Then, let c_2^{ll} be the point such that $V_2[c_1^{\text{ll}}, c_2^{\text{ll}}] = \frac{V_2[c_1^{\text{ll}}, 1]}{n-1}$ and $c_2^{\text{ll}} = \max\{c_2^{\text{ll}}, l_3\}$. Agent 2 is assigned the interval $(c_1^{\text{ll}}, c_2^{\text{ll}}]$ and so on, until the last agent who receives the interval $(c_{n-1}^{\text{ll}}, 1]$. Leftmost leaves generates a Pareto optimal allocation, guaranteeing all agents at least a utility of $1/n$ (Kyropoulou et al., 2019; Ortega and Segal-Halevi, 2019), in contrast with the Wang–Wu mechanism which guarantees exactly a $1/n$ utility to all agents; thus leftmost leaves Pareto dominates the Wang–Wu mechanism. However, leftmost leaves can generate envy.

3. *Modified Wang–Wu* (MWW). The third mechanism is a modification of the original Wang–Wu mechanism. As in the original version, the three points l_i, p_i, r_i for all agents are ordered in increasing sequence, together with the 0 and 1 endpoints. The cake is cut at each of these points, generating up to $3n+1$ cake pieces. The difference with the original Wang–Wu procedure is that each of these pieces is divided only among the agents who have a positive value for it. Formally, for each cake piece P_i let $n_i := \{i \in N : U_i \cap P_i \neq \emptyset\}$. Then, each piece P_i is divided into $2n_i$ equidistant pieces, which are shared into the interested n_i agents as in the Wang–Wu procedure: the first interested agent gets pieces 1 and $2n_i$, the second one gets pieces 2 and $2n_i - 1$ and so on. This procedure is envy-free (like the original Wang–Wu procedure, which it Pareto dominates), although is not Pareto optimal since it generates a disconnected allocation.

We summarize our findings in our third Proposition and in Table 1.

Table 1: Summary of the properties of each mechanism.

	Efficient	Envy-free	Max Utility	Dominates WW
Wang–Wu	No	Yes	No	
Utilitarian	Yes	No	Yes	No
LL	Yes	No	No	Yes
MWW	No	Yes	No	Yes

Proposition 3. *The utilitarian mechanism is Pareto optimal and maximizes the sum of agents’ utilities. Leftmost leaves is Pareto optimal and Pareto*

Figure 4: An example showing the differences between the four different mechanisms. Agents 1, 2, and 3 receive the grey, white, and black pieces respectively. Agents' utilities appear in parentheses.

dominates the original Wang–Wu mechanism. MWW is envy-free and also Pareto dominates the original Wang–Wu mechanism.

Proof. Proposition 2 implies that the utilitarian and leftmost leaves mechanisms are Pareto optimal. MWW is envy-free because, for any of the (up to) $3n + 1$ subintervals generated, any agent values the share received by of any other agent (among those who also receive a share) as much as her own (as in Theorem 2 in Wang and Wu (2019)). Leftmost leaves Pareto dominates Wang–Wu follows because leftmost leaves is proportional, and only assigns a utility of $1/n$ to every agent when all agents' have the same preferences (Kyropoulou et al., 2019), whereas Wang–Wu always gives a utility of $1/n$ to all agents. That MWW Pareto dominates Wang–Wu follows from the fact that, although both WW and MWW divide the cake into the same number of pieces ($3n + 1$), MWW only assigns to agents cake shares of the pieces they desire. $\square$

5. Conclusion

Dividing a cake fairly and efficiently is a difficult problem, even when agents have simple preferences. Nevertheless, we have proposed fair division mechanisms that are either Pareto optimal or envy-free, and which Pareto dominate the allocation generated by the envy-free procedure of Wang–Wu. Using any of these mechanisms instead is justified due to the large utility losses that the Wang–Wu mechanism can generate.

Acknowledgements

We are grateful to Chenhao Wang, Erel Segal-Halevi and an anonymous referee for their helpful comments. Sarah Fox proofread this paper. This work was completed while Bhavook Bhardwaj was a visiting student at Queen’s University Belfast; their hospitality is gratefully acknowledged. Rajnish Kumar acknowledges financial support provided by British Council Grant UGC-UKIERI 2016-17-059.

References

- AZIZ, H. AND S. MACKENZIE (2016): “A discrete and bounded envy-free cake cutting protocol for any number of agents,” in *IEEE 57th Annual Symposium on Foundations of Computer Science*, 416–427.
- BERTSIMAS, D., V. F. FARIAS, AND N. TRICHAKIS (2011): “The price of fairness,” *Operations Research*, 59, 17–31.
- BRAMS, S. J. AND A. D. TAYLOR (1995): “An envy-free cake division protocol,” *The American Mathematical Monthly*, 102, 9–18.
- EHLERS, L., H. PETERS, AND T. STORCKEN (2002): “Strategy-proof probabilistic decision schemes for one-dimensional single-peaked preferences,” *Journal of Economic Theory*, 105, 408–434.
- JUAREZ, R. AND R. KUMAR (2013): “Implementing efficient graphs in connection networks,” *Economic Theory*, 54, 359–403.
- KASAJIMA, Y. (2013): “Probabilistic assignment of indivisible goods with single-peaked preferences,” *Social Choice and Welfare*, 41, 203–215.
- KUROKAWA, D., J. K. LAI, AND A. D. PROCACCIA (2013): “How to cut a cake before the party ends,” in *Proceedings of the 27th AAAI Conference on Artificial Intelligence*.
- KYROPOULOU, M., J. ORTEGA, AND E. SEGAL-HALEVI (2019): “Fair cake-cutting in practice,” in *Proceedings of the 2019 ACM Conference on Economics and Computation*, ACM, 547–548.
- LONG, Y. (2019): “Strategy-proof group selection under single-peaked preferences over group size,” *Economic Theory*, 68, 579–608.
- MANIQUET, F. AND Y. SPRUMONT (2000): “On resource monotonicity in the fair division problem,” *Economics Letters*, 68, 299–302.
- MOULIN, H. (1984): “Generalized Condorcet-winners for single peaked and single-plateau preferences,” *Social Choice and Welfare*, 1, 127–147.
- (2004): *Fair division and collective welfare*, MIT press.
- NICOLÒ, A. AND Y. YU (2008): “Strategic divide and choose,” *Games and Economic Behavior*, 64, 268 – 289.
- ORTEGA, J. (2018): “Social integration in two-sided matching markets,” *Journal of Mathematical Economics*, 78, 119–126.

- (2019): “The losses from integration in matching markets can be large,” *Economics Letters*, 174, 48–51.
- ORTEGA, J. AND E. SEGAL-HALEVI (2019): “Obvious manipulations in cake-cutting,” *ZEW Discussion Papers 19-056*.
- PROCACCIA, A. (2016): “Cake cutting algorithms,” in *Handbook of Computational Social Choice*, ed. by F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, Cambridge University Press, 311–330.
- ROBERTSON, J. AND W. WEBB (1998): *Cake-cutting algorithms: Be fair if you can*, AK Peters/CRC Press.
- SANDOMIRSKIY, F. AND E. SEGAL-HALEVI (2019): “Fair division with minimal sharing,” *arXiv preprint arXiv:1908.01669*.
- SEGAL-HALEVI, E. AND B. R. SZIKLAI (2019): “Monotonicity and competitive equilibrium in cake-cutting,” *Economic Theory*, 68, 363–401.
- SPRUMONT, Y. (1991): “The division problem with single-peaked preferences: a characterization of the uniform allocation rule,” *Econometrica*, 59, 509–519.
- STEINHAUS, H. (1948): “The problem of fair division,” *Econometrica*, 16, 101–104.
- STROMQUIST, W. (2008): “Envy-free cake divisions cannot be found by finite protocols,” *The Electronic Journal of Combinatorics*, R11–R11.
- WANG, C. AND X. WU (2019): “Cake-cutting with single-peaked valuations,” in *Proceedings of the International Conference on Combinatorial Optimization and Applications*, Springer, 507–516.
- WELLER, D. (1985): “Fair division of a measurable space,” *Journal of Mathematical Economics*, 14, 5–17.