Fair Cake-Cutting Algorithms with Real Land-Value Data

Extended Abstract

Itay Shtechman
Dept. of Mathematics and Computer Science,
The Open University of Israel
itay.shtechman@gmail.com

Rica Gonen
Dept. of Management and Economics,
The Open University of Israel
ricagonen@gmail.com

Erel Segal-Halevi
Dept. of Computer Science,
Ariel University
erelsgl@gmail.com

ABSTRACT

Fair division of land is an important practical problem that is commonly handled either by hiring assessors or by selling and dividing the proceeds. A third way to divide land fairly is via algorithms for *fair cake-cutting*. However, the current theory of fair cake-cutting is not yet ready to optimally share a plot of land and such algorithms are seldom used in practical land-division.

We attempt to narrow the gap between theory and practice by performing extensive simulations of a classic cake-cutting algorithm on real land-value data. We improve the practical performance of this algorithm using heuristics we developed, and show their effectiveness on real land-value maps compared to actual assessment and sale data on various performance metrics. The cake-cutting algorithms perform better in most metrics.

We further examined the cake cutting algorithm with respect to strategic gain of an agent relative to a truthful agent. The strategic gain was found to be insignificant effect in cake-cutting algorithms.

KEYWORDS

fair allocation; cake-cutting; land division; simulation

ACM Reference Format:

Itay Shtechman, Rica Gonen, and Erel Segal-Halevi. 2020. Fair Cake-Cutting Algorithms with Real Land-Value Data. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 3 pages.

1 INTRODUCTION

Fair division occurs in inheritance cases, partnership resolutions, and public land allocations. If the partners do not trust an assessor, they often sell the land in the market and split the revenues.

In the last 70 years, economists and computer scientists have developed various algorithms for fair *cake-cutting* - fair division of a continuous heterogeneous resource among agents with different preferences. Using such automatic methods is not only cheaper than renting a human assessor — it is also theoretically fairer, since it guarantees to each agent a fair share by his/her personal value function. Our goal is to bring theory and practice closer together, and this paper describes several steps we made in this direction.

First, we constructed two-dimensional instances of cake-cutting problems based on real land-value data of New Zealand and of most of Israel. To simulate agents with different but correlated valuations, we created n different maps — each map is based on the original map and some random noise. We experimented with values of n between 4 and 128, two noise models, and noise levels, ranging between 20% and 60% of the base value.

Second, we adapted the classic cake-cutting algorithms of Even and Paz [3] and Steinhaus [5] to divide a two-dimensional map into rectangular land-plots. This adaptation can be done in many ways, since each cut of the interval made by the original algorithm can be converted to a horizontal cut or to a vertical cut of the map. The cut direction can be fixed in advance, or it can be decided dynamically based on various heuristics. Most previous works that we know of ignored these questions since they assumed that the cake is one-dimensional. However, when dividing a two-dimensional land, such decisions may affect the quality of the allocation.

Third, we conducted extensive experiments with different numbers of agents and three different maps (the above two and a randomly - generated map). In each setting, we compared the cake-cutting algorithm to the two baseline methods commonly used today for dividing land: (1) *assessor division*, where the land is partitioned into pieces with the same base value (ignoring the "noise") and each agent receives one piece; (2) *market sale*, where the land is sold for its total market value and each agent receives $1/n$ of the proceeds. We compared the methods using several metrics: *utilitarian welfare* (the sum of agents' utilities), *egalitarian welfare* (the smallest agent utility), *envy*, and a new metric particularly important in land division — the *length/width ratio* of the resulting pieces.

For each metric, one or two heuristics were superior to the others. This shows the importance of modelling land in two dimensions rather than reducing it to a 1-dimensional interval. The cake-cutting algorithms fared similar or better than assessor division in all metrics, and better than market sale in the two welfare metrics.

Cake-cutting algorithms' potential problem, that does not exist in market sale or assessor division, is that agents may try to manipulate it by misrepresenting their preferences. We measured the *strategic gain* — the amount by which a strategic agent with complete information can increase his/her utility. The average strategic gain of an agent was less than 1.5%. This implies that strategic manipulation is not a major factor in cake-cutting algorithms.

The advantage of the cake-cutting algorithms was much more pronounced in the two maps based on land values than in the random map. This illustrates the importance of using real value data for evaluating fair division algorithms. This is in contrast to previous work on simulation of fair division algorithms, which mostly used artificially-generated data [1, 2, 4, 6].

The improvements provided by the cake-cutting algorithm are encouraging. They imply that the algorithmic "do it yourself" approach to fair division, which is cheaper than the common approaches of selling or employing an expert assessor, is also better in terms of social welfare and reducing envy.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

2 ALGORITHMS AND METRICS

A land-estate C ("cake") has to be divided among n agents. Each agent i has a value-density function v_i, mapping each point of C to its monetary value for i. C should be partitioned into n disjoint contiguous pieces, $X_1, \ldots, X_n$, one piece per agent. C is assumed to be a rectangle, and each piece should be a rectangle too.

There are infinitely many land partitions, and there are various metrics by which a partition can be evaluated. We compared the methods using several metrics:

Utilitarian Value (UV): the sum of the agents' relative values.

Egalitarian Value (EV): the smallest relative value of an agent.

Largest Envy (LE): the largest amount by which an agent considers another agent's share as better than his/her own share.[1]

Average/Smallest Face-Ratio (AFR/SFR): the average/smallest of the calculated face-ratio of all allocated land-plots.

There are many algorithms for finding a proportional cake partition. In this experiment we focus on the *Even-Paz* algorithm [3], since we believe it is the algorithm most likely to be used in a cake-cutting problem with a large number of agents. This is due both to its simplicity and to its optimal run-time complexity: it runs in time $O(n \log n)$, which is provably the best possible [7]. We compare the results of Even-Paz to another classic proportional cake-cutting algorithm we examined - *Last-Diminisher* [5].

Both Even-Paz and Last-Diminisher are well-defined for a one-dimensional cake. However, when C is two dimensional, in each iteration, the agents can make their query-marks in many different directions. We applied 10 different heuristic for choosing the cut direction (horizontal or vertical) in each iteration and compared those heuristics empirically. We tested two sets of heuristics: pre-defined heuristics (all $n-1$ cuts are defined in advanced) and greedy heuristics (at each iteration level the cut direction is selected in order to maximize a certain quality function).

Assessor division. We compare our algorithms' results to the performance of an assessor. The division is done using the following method. For some integers k_1, k_2 with $k_1 \cdot k_2 = n$, the assessor first partitions the land using vertical cuts into k_1 parallel strips of value $1/k_1$, and then partitions each such strip using horizontal cuts into k_2 plots of value $(1/k_1)/k_2 = 1/n$. We take $k_1 = 2^{\lceil (\log_2 n)/2 \rceil}$ and $k_2 = 2^{\lfloor (\log_2 n)/2 \rfloor}$, so that the pieces have a balanced aspect ratio.

3 EXPERIMENTS AND RESULTS

We constructed three land-value maps. The first map was of New Zealand — based on the Forest Profit Expectations Dataset (U10073). The second map was of Israel — based on a commercial website for classified real-estate ads (http://madlan.co.il). For comparison, the third map was generated uniformly at random.

From each land-value map we created datasets, each dataset containing n variants, one variant per agent. The subjective preference of each agent was captured by a random noise that is added to the land-value map. For each map and n in 4, 8, 16 ,32, 64, 128, we ran 50 experiments with different randomly-generated noise. In most experiments the noise-ratio was $r = 0.6$.

[1]Note that if the land is sold and the money is divided among the agents, the relative value of each agent is exactly $1/n$, so $UV = EV = LE = 1$ (we defined the metrics such that their baseline is 1).

The results for the New Zealand map and Israel map were very similar, but the results for the random map were substantially different. For all tested metrics we found that Even-Paz is at least as good, and often better, than Last-Diminisher.

Below, we discuss the Even-Paz results for New Zealand map averaged over the 50 runs and a 95% confidence interval.

The best-performing heuristic is **MostValuableMargin** (MVM), a greedy heuristic that chooses a cut direction resulting in a more valuable margin between the actual cut proposed by the algorithm and the two closest cuts proposed by agents. Intuitively, the reason is that this heuristic takes the most advantage of the differences between the agents' valuations.

The *UV* of MVM is significantly better than the market sale ($p < 0.001$) and the advantage grows with n. For $n = 128$, the advantage is about 2.9%. In contrast, the UV of an assessor division is 1 – the same as market sale.

The EV of Even-Paz is always higher than the EV of selling the land and with MVM it is significantly better ($p < 0.001$). With $n = 128$, the advantage is about +1%. The EV of an assessor division is always *lower* than of selling the land. This means that, in an assessor division, there is always at least one person who receives less than his/her fair share.

Besides market sale which trivially attains the minimum LE, the best heuristic is MVM: it scores better than the other heuristics we tested, but the advantage is not statistically significant ($p = 0.2$). The LE of Even-Paz is better (smaller) than of the assessor-division. For $n = 128$, MVM is significantly smaller than assessor division in about 3% ($p < 0.001$).

The best heuristic for both AFR and SFR, is **SquarePiece (SP)** ($p < 0.001$), a greedy heuristic that chooses a cut direction resulting in a higher face-ratio. In contrast to the previous metrics, we did not find a significant difference between SP and an assessor division. The AFR of SP is always above 1/3, and sometimes above 1/2. However, the SFR (for $n = 128$) is very small — around 0.02. This means that at least one agent might get a very thin plot.

We ran experiments to check how much an agent may gain by being untruthful (assuming all other agents are truthful). The *largest* strategic gain attained by an agent was 6.54%. However, the *average* statistic gain in all experiments was less than 1.5%.

4 CONCLUSION AND DISCUSSION

Cake-cutting algorithms may be a viable alternative to the common methods for land division, namely assessor division and market sale. In particular, Even-Paz attains higher social welfare than both these methods, and lower envy than assessor division. The effect is larger when there are more agents and when there is more variation in the agents' valuations. Even-Paz performs better than the other cake-cutting algorithm we checked (Last Diminisher).

When adapting cake-cutting algorithms to two dimensions, in terms of social welfare and envy MVM is superior, while SP is superior in terms of face-ratio. This shows a trilateral trade-off between run-time, social welfare and geometric shape as well as illustrating the importance of metric driven practical decision making.

Strategic manipulation may improve the welfare of an agent with complete information, but the improvement is relatively small.

REFERENCES

[1] Ruggiero Cavallo. 2012. Fairness and welfare through redistribution when utility is transferable. In *Proceedings of the Conference on Artificial Intelligence (AAAI)*.

[2] John Dickerson, Jonathan Goldman, Jeremy Karp, Ariel Procaccia, and Tuomas Sandholm. 2014. The computational rise and fall of fairness.. In *Proceedings of the Conference on Artificial Intelligence (AAAI)*.

[3] Shimon Even and Azaria Paz. 1984. A Note on Cake Cutting. *Discrete Applied Mathematics* 7, 3 (March 1984), 285–296. https://doi.org/10.1016/0166-218x(84)90005-2

[4] Erel Segal-Halevi, Haris Aziz, and Avinatan Hassidim. 2017. Fair allocation based on diminishing differences. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence.* AAAI Press, 1254–1261.

[5] Hugo Steinhaus. 1948. The problem of fair division. *Econometrica* 16, 1 (1948), 101–104.

[6] Toby Walsh. 2011. Online Cake Cutting. *Algorithmic Decision Theory* 6992 (2011), 292–305.

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