Games and Economic Behavior 77 (2013) 284–297

Contents lists available at SciVerse ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Truth, justice, and cake cutting ☆

Yiling Chen [a], John K. Lai [a], David C. Parkes [a,*], Ariel D. Procaccia [b]

[a] School of Engineering and Applied Sciences, Harvard University, 33 Oxford Street, Cambridge, MA 02138, United States
[b] Computer Science Department, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, United States

ARTICLE INFO

Article history:
Received 11 October 2010
Available online 1 November 2012

JEL classification:
C78
D63
D82

Keywords:
Fair division
Envy-freeness
Cake cutting
Maximum flows
Strategy-proofness
Mechanism design

ABSTRACT

Cake cutting is a common metaphor for the division of a heterogeneous divisible good. There are numerous papers that study the problem of fairly dividing a cake; a small number of them also take into account self-interested agents and consequent strategic issues, but these papers focus on fairness and consider a strikingly weak notion of truthfulness. In this paper we investigate the problem of cutting a cake in a way that is truthful, Pareto-efficient, and fair, where for the first time our notion of dominant strategy truthfulness is the ubiquitous one in social choice and computer science. We design both deterministic and randomized cake cutting mechanisms that are truthful and fair under different assumptions with respect to the valuation functions of the agents.

Superman: "I'm here to fight for truth, justice, and the American way."
Lois Lane: "You're gonna wind up fighting every elected official in this country!"

Superman (1978)

1. Introduction

Cutting a cake is often used as a metaphor for allocating a divisible good. The difficulty is not cutting the cake into pieces of equal size, but rather that the cake is not uniformly tasty: different agents prefer different parts of the cake, depending, e.g., on whether the toppings are strawberries or cookies. The goal is to divide the cake in a way that is "fair"; the definition of fairness is a nontrivial issue in itself, which we discuss in the sequel. The cake cutting problem dates back to the work of Steinhaus in the 1940s (Steinhaus, 1948), and for over sixty years has attracted the attention of mathematicians, economists, political scientists, and more recently, computer scientists (Edmonds and Pruhs, 2006a, 2006b; Procaccia, 2009).

☆ A significantly shorter version of this paper appeared in the proceedings of AAAI 2010. Here we fill in missing proofs, discuss the Lorenz dominance of our allocation, and provide an extended discussion of related work and future directions. The paper was also presented in the Third International Workshop on Computational Social Choice in Düsseldorf, Germany, September 2010, at a workshop on prior-free mechanism design in Guanajuato, Mexico, May 2010, and in the Harvard EconCS seminar, February 2010.

* Corresponding author.
E-mail addresses: yiling@eecs.harvard.edu (Y. Chen), jklai@post.harvard.edu (J.K. Lai), parkes@eecs.harvard.edu (D.C. Parkes), arielpro@cs.cmu.edu (A.D. Procaccia).

0899-8256/$ – see front matter
http://dx.doi.org/10.1016/j.geb.2012.10.009

Slightly more formally, the cake is represented by the interval $[0,1]$. Each of n agents has a valuation function over the cake, which assigns a value to every given piece of cake and is additive. The goal is to find a partition of the cake among the agents (while possibly throwing a piece away) that satisfies one or several fairness criteria. In this paper we consider the two most prominent criteria. A *proportional* allocation is one where the value each agent has for its own piece of cake is at least $1/n$ of the value it assigns to the entire cake. An *envy-free (EF)* allocation is one where the value each agent assigns to its own piece of cake is at least as high as the value it assigns to any other agent's piece of cake. There is a rather large body of literature on fairly cutting a cake according to these two criteria (see, e.g., the books by Robertson and Webb, 1998 and Brams and Taylor, 1996).

So far we have briefly discussed "justice", but have not yet mentioned "truth". Taking the game-theoretic point of view, an agent's valuation function is its private information, which is reported to a cake cutting mechanism. We would like a mechanism to be *truthful*, in the sense that agents are motivated to report their true valuation functions. Like fairness, this idea of truthfulness also lends itself to many interpretations. One variation, referred to as *strategy-proofness* in previous papers by Brams et al. (2006, 2008), assumes that an agent will report its true valuation rather than lie if there *exist* valuations of the other agents such that reporting truthfully yields at least as much value as lying. In the words of Brams et al., "the players are risk-averse and never strategically announce false measures if it does not guarantee them more-valued pieces. … Hence, a procedure is strategy-proof if no player has a strategy that dominates his true value function" (Brams et al., 2008, p. 362).

The foregoing notion is strikingly weak compared to the notion of truthfulness that is common in the social choice literature. Indeed, strategy-proofness is usually taken to mean that an agent can *never* benefit by lying, that is, *for all* valuations of the other agents reporting truthfully yields at least as much value as lying. Put another way, truth-telling is a dominant strategy. This notion is worst-case, in the sense that an agent cannot benefit by lying even if it is fully knowledgeable of the valuations of the other agents. In order to prevent confusion we will avoid using the term "strategy-proof", and instead refer to the former notion of Brams et al. as "weak truthfulness" and to the latter standard notion as "truthfulness".

To illustrate the difference between the two notions, consider the most basic cake cutting mechanism for the case of two agents, the *Cut and Choose* mechanism.[1] Agent 1 cuts the cake into two pieces that are of equal value according to its valuation; agent 2 then chooses the piece that it prefers, giving the other piece to agent 1. This mechanism is trivially proportional and EF.[2] It is also weakly truthful, as if agent 1 divides the cake into two pieces that are unequal according to its valuation then agent 2 may prefer the piece that is worth more to agent 1. Agent 2 clearly cannot benefit by lying. The mechanism is even truthful (in the strong sense) if the agents have the same valuation function. However, the mechanism is not truthful in general. Indeed, consider the case where agent 1 would simply like to receive as much cake as possible, whereas the single-minded agent 2 is only interested in the interval $[0,\epsilon]$ where ϵ is small (for example, it may only be interested in the cherry). If agent 1 follows the protocol it would only receive half of the cake. Agent 1 can do better by reporting that it values the intervals $[0,\epsilon]$ and $[\epsilon,1]$ equally.

In this paper we consider the design of truthful and fair cake cutting mechanisms. We will design mechanisms where the agents report valuations and the mechanism, taking the role of a direct revelation mechanism, determines the cuts of the cake. However, there is a major obstacle that must be circumvented: regardless of strategic issues, and when there are more than five agents, even finding a proportional and EF allocation in a bounded number of steps with a deterministic mechanism is a long-standing open problem! See Procaccia (2009) for an up-to-date discussion.[3] We shall therefore restrict ourselves to specific classes of valuation functions where efficiently finding fair allocations is a non-issue; the richness of our problem stems from our desire to additionally achieve truthfulness.

In particular, we will focus on the case where agents hold *piecewise uniform, piecewise constant, and piecewise linear* valuation functions. Piecewise uniform valuation functions capture the case where each agent is interested in a collection of subintervals of $[0,1]$ and has the same marginal value for each fractional piece in each desired subinterval. They can be thought of as dichotomous preferences since intervals can be separated into acceptable and unacceptable intervals, and the acceptable intervals all have the same marginal utility. While these valuations are restrictive, they are expressive enough to capture some realistic settings. Piecewise uniform valuations capture the case when some parts of the cake satisfy a certain property that an agent cares about and an agent desires as much of these parts as possible. An example of such a setting is a shared computing server. Each researcher has certain times during which the server is not useful because of scheduling constraints or waiting times for experiments, but outside of these times, the researcher would like as much time on the computing server as possible. Piecewise constant and linear valuation functions capture the case where an agent's marginal value is constant and linear respectively, and are significantly more expressive. In particular, these valuation functions can approximate any valuation function.

[1] This mechanism is described here with the agents taking actions; equivalently, the mechanism acts on behalf of agents using the reported valuations.

[2] Proportionality and envy-freeness coincide if there are two agents and the entire cake is allocated.

[3] To be precise, previous work assumed that the entire cake has to be allocated, but this does not seem to be a significant restriction in the context of fairness.