# Computational Social Choice
# Session 1: Fair Allocation of Resources

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

Forum for Artificial Intelligence Research 2014, Stellenbosch
http://www.illc.uva.nl/~ulle/teaching/fair-2014/

# Table of Contents

# Introduction

# Computational Social Choice

*Computational social choice* deals with formal and algorithmic aspects of *collective decision making*. Each day will be devoted to one subfield:

- Session 1: *Fair Allocation of Resources* — Given the preferences of agents over resources, find a socially optimal allocation.
- Session 2: *Voting Theory* — Given the preferences of agents over candidates, pick the socially most preferred candidate.
- Session 3: *Judgment Aggregation* — Given the judgments of agents on certain statements, infer a collective judgment.

So, we will move from specific to more general aggregation problems. For a high-level view on the storyline, see the paper cited below.

U. Endriss. Social Choice Theory as a Foundation for Multiagent Systems. Proc. MATES-2014.

# Social Choice and AI (1)

Social choice theory, while originally developed in Economics, has a number of natural applications in AI. Examples:

- *Multiagent Systems:* to aggregate the beliefs + to coordinate the actions of groups of autonomous software agents
- *Search Engines:* to determine the most important sites based on links ("votes") + to aggregate the output of several search engines
- *Recommender Systems:* to recommend a product to a user based on earlier ratings by other users
- *AI Competitions:* to determine who has developed *the best* trading agent / SAT solver / RoboCup team

But not all of the classical assumptions will fit these new applications. So AI needs to develop *new models* and *ask new questions*.

## Social Choice and AI (2)

*Vice versa*, techniques from AI, and computational techniques in general, are useful for advancing the state of the art in social choice:

- *Algorithms and Complexity:* to develop algorithms for (complex) aggregation rules + to understand the hardness of "using" them
- *Knowledge Representation:* to compactly represent the preferences of individual agents over large spaces of alternatives
- *Logic and Automated Reasoning:* to formally model problems in social choice + to automatically verify (or discover) theorems

Indeed, you will find many papers on social choice at AI conferences (e.g., IJCAI, AAAI, ECAI, AAMAS, KR) and many AI researchers participate in events dedicated to social choice (e.g., COMSOC).

F. Brandt, V. Conitzer, and U. Endriss. Computational Social Choice. In G. Weiss (ed.), *Multiagent Systems*. MIT Press, 2013.

# Fair Allocation of Resources

Consider a set of agents and a set of goods. Each agent has their own preferences regarding the allocation of goods to agents to be selected.

▶ What constitutes a good allocation and how do we find it?

What goods? One or several goods? Available in single or multiple units? Divisible or indivisible? Can goods be shared? Are they static or do they change properties (e.g., consumable or perishable goods)?

What preferences? Ordinal or cardinal preference structures? Are monetary side payments possible, and how do they affect preferences? How are the preferences represented in the problem input?

Y. Chevaleyre, P.E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J.A. Rodríguez-Aguilar and P. Sousa. Issues in Multiagent Resource Allocation. *Informatica*, 30:3–31, 2006.

# Plan for Today

We will discuss three aspects of fair allocation problems:

- Part 1: *Fairness and Efficiency Criteria* —
  We will review several proposals from the literature for how to define "fairness" and the related notion of economic "efficiency".
- Part 2: *Combinatorial Optimisation and Convergence* —
  The allocation of *indivisible goods* gives rise to a combinatorial optimisation problem. We will focus on a distributed approach.
- Part 3: *Cake-Cutting Algorithms* —
  How should we fairly divide a "cake" (a single *divisible good*)? We will review several algorithms and analyse their properties.

Most of this material is covered in my lecture notes cited below.

U. Endriss. *Lecture Notes on Fair Division*. Institute for Logic, Language and Computation, University of Amsterdam, 2009/2010.

# Fairness and Efficiency Criteria

## What is a Good Allocation?

We are now going to give a (partial) overview of criteria that have been proposed for deciding what makes a "good" allocation:

- Of course, there are application-specific criteria, e.g.:
  - "*the allocation allows the agents to solve the problem*"
  - "*the auctioneer has generated sufficient revenue*"

  Here we are interested in general criteria that can be defined in terms of the individual agent preferences (*preference aggregation*).

- As we shall see, such criteria can be roughly divided into *fairness* and (economic) *efficiency* criteria.

# Notation and Terminology

- Let $\mathcal{N} = \{1, \ldots, n\}$ be a set of *agents* (or *players*, or *individuals*) who need to share several *goods* (or *resources*, *items*, *objects*).
- An *allocation* $A$ is a mapping of agents to *bundles* of goods.
- Most criteria will not be specific to allocation problems, so we also speak of *agreements* (or *outcomes*, *solutions*, *alternatives*, *states*).
- Each agent $i \in \mathcal{N}$ has a *utility function* $u_i$ (or *valuation function*), mapping agreements to the reals, to model their preferences.
  - Typically, $u_i$ is first defined on bundles, so: $u_i(A) = u_i(A(i))$.
  - Discussion: preference intensity, interpersonal comparison
- An agreement $A$ gives rise to a *utility vector* $\langle u_1(A), \ldots, u_n(A)\rangle$.
- Sometimes, we are going to define social preference structures directly over utility vectors $u = \langle u_1, \ldots, u_n\rangle$ (elements of $\mathbb{R}^n$), rather than speaking about the agreements generating them.

## Pareto Efficiency

Agreement $A$ is *Pareto dominated* by agreement $A'$ if $u_i(A) \leqslant u_i(A')$ for all agents $i \in \mathcal{N}$ and this inequality is strict in at least one case.

An agreement $A$ is *Pareto efficient* if there is no other feasible agreement $A'$ such that $A$ is Pareto dominated by $A'$.

The idea goes back to Vilfredo Pareto (Italian economist, 1848–1923).

Discussion:

- Pareto efficiency is very often considered a minimum requirement for any agreement/allocation. It is a very weak criterion.
- Only the ordinal content of preferences is needed to check Pareto efficiency (no preference intensity, no interpersonal comparison).

## Collective Utility Functions

A *collective utility function* (CUF) is a function $\mathrm{SW} : \mathbb{R}^n \to \mathbb{R}$ mapping utility vectors to the reals.

## Utilitarian Social Welfare

One approach to social welfare is to try to maximise overall profit. This is known as classical utilitarianism (advocated, amongst others, by Jeremy Bentham, British philosopher, 1748–1832).

The *utilitarian* CUF is defined as follows:

$$\mathrm{SW}_{\mathrm{util}}(u) \quad = \quad \sum_{i \in \mathcal{N}} u_i$$

Remark: We define CUFs on utility vectors, but the definitions immediately extend to allocations:

$$\mathrm{SW}_{\mathrm{util}}(A) \quad = \quad \mathrm{SW}_{\mathrm{util}}(\langle u_1(A), \ldots, u_n(A) \rangle) \quad = \quad \sum_{i \in \mathcal{N}} u_i(A(i))$$

# Egalitarian Social Welfare

The *egalitarian* CUF measures social welfare as follows:

$$\mathrm{SW}_{\mathrm{egal}}(u) \quad = \quad min\{u_i \mid i \in \mathcal{N}\}$$

Maximising this function amounts to improving the situation of the weakest member of society.

The egalitarian variant of welfare economics is inspired by the work of John Rawls (American philosopher, 1921–2002) and has been formally developed, amongst others, by Amartya Sen since the 1970s (Nobel Prize in Economic Sciences in 1998).

Refinement: *leximin-ordering*

J. Rawls. *A Theory of Justice*. Oxford University Press, 1971.

A.K. Sen. *Collective Choice and Social Welfare*. Holden Day, 1970.

## Nash Product

The *Nash* CUF is defined via the product of individual utilities:

$$\mathrm{SW}_{\mathrm{nash}}(u) \quad = \quad \prod_{i \in \mathcal{N}} u_i$$

This is a useful measure of social welfare as long as all utility functions can be assumed to be positive. Named after John F. Nash (Nobel Prize in Economic Sciences in 1994; Academy Award in 2001).

Remark: The Nash (like the utilitarian) CUF favours increases in overall utility, but also inequality-reducing redistributions ($2 \cdot 6 < 4 \cdot 4$).

## Envy-Freeness

An allocation is called *envy-free* if no agent would rather have one of the bundles allocated to any of the other agents:

$$u_i(A(i)) \quad \geqslant \quad u_i(A(j))$$

Recall that $A(i)$ is the bundle allocated to agent $i$ in allocation $A$.

Remark: Envy-free allocations do not always *exist* (at least not if we require either complete or Pareto efficient allocations).

## Summary: Fairness and Efficiency Criteria

- The quality of an allocation can be measured using a variety of fairness and efficiency criteria.
- We have seen Pareto efficiency, collective utility functions (utilitarian, egalitarian, Nash), leximin ordering, and envy-freeness.
- Understanding the structure of different notions of social welfare is in itself an interesting research area (so-called *axiomatic method*).

## Literature: Fairness and Efficiency Criteria

Moulin (1988) provides an excellent introduction to welfare economics. Moulin (2003) offers a less technical version of the material.

The "MARA Survey" (Chevaleyre *et al.*, 2006) discusses many social welfare concepts and their relevance to multiagent resource allocation.

H. Moulin. *Axioms of Cooperative Decision Making*. Econometric Society Monographs, Cambridge University Press, 1988.

H. Moulin. *Fair Division and Collective Welfare*. MIT Press, 2003.

Y. Chevaleyre, P.E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J.A. Rodríguez-Aguilar and P. Sousa. Issues in Multiagent Resource Allocation. *Informatica*, 30:3–31, 2006.

# Indivisible Goods: Combinatorial Optimisation and Convergence

## Allocation of Indivisible Goods

Next we will consider the case of allocating indivisible goods. We can distinguish two approaches:

- In the *centralised approach* (e.g., combinatorial auctions), we need to devise an optimisation algorithm to compute an allocation meeting our fairness and efficiency requirements.
- In the *distributed approach*, allocations emerge as a consequence of the agents implementing a sequence of local deals. What can we say about the properties of these emerging allocations?

# Formal Framework: Indivisible Goods

We will work in the following formal framework:

- Set of *agents* $\mathcal{N} = \{1, \ldots, n\}$ and finite set of indivisible *goods* $\mathcal{G}$.
- An *allocation* $A$ is a partitioning of $\mathcal{G}$ amongst the agents in $\mathcal{N}$.
  Example: $A(i) = \{a, b\}$ — agent $i$ owns items $a$ and $b$
- Each agent $i \in \mathcal{N}$ has got a *valuation function* $v_i : 2^{\mathcal{G}} \to \mathbb{R}$.
  Example: $v_i(A) = v_i(A(i)) = 577.8$ — agent $i$ is pretty happy
- If agent $i$ receives bundle $B$ and the sum of her payments is $x$, then her *utility* is $u_i(B, x) = v_i(B) - x$ ("quasi-linear utility").

For fair allocation of indivisible goods *without money*, assume that payment balances are always equal to 0 (and utility = valuation).

▶ How can we find a socially optimal allocation of goods?

# Preference Representation

Example: Allocating 10 goods to 5 agents means $5^{10} = 9765625$ allocations and $2^{10} = 1024$ bundles for each agent to think about.

So we need to choose a good *language* to compactly represent preferences over such large numbers of alternative bundles, e.g.:

- Logic-based languages (weighted goals)
- Bidding languages for combinatorial auctions (OR/XOR)
- Program-based preference representation (straight-line programs)
- CP-nets and CI-nets (for ordinal preferences)

The choice of language affects both *algorithm design* and *complexity*.

See our *AI Magazine* article for an introduction to the problem of preference modelling in combinatorial domains.

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Preference Handling in Combinatorial Domains: From AI to Social Choice. *AI Magazine*, 29(4):37–46, 2008.

# Complexity Results

Before we look into the "how", here are some complexity results:

- Finding an allocation with maximal *utilitarian* social welfare is NP-hard. If all valuations are *modular* (additive) then it is polynomial.
- Finding an allocation with maximal *egalitarian* social welfare is also NP-hard, even when all valuations are modular.
- Checking whether an allocation is *Pareto efficient* is coNP-complete.
- Checking whether an *envy-free* allocation exists is NP-complete; checking whether an allocation that is both Pareto efficient and envy-free exists is even $\Sigma_2^p$-complete.

References to these results may be found in the "MARA Survey".

Y. Chevaleyre, P.E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J.A. Rodríguez-Aguilar and P. Sousa. Issues in Multiagent Resource Allocation. *Informatica*, 30:3–31, 2006.

# Algorithms for Finding an Optimal Allocation

If our goal is to find an allocation with maximal *utilitarian* social welfare, then the allocation problem is equivalent to the winner determination problem in *combinatorial auctions:*

- valuation of agent $i$ for bundle $B \sim$ price offered for $B$ by bidder $i$
- utilitarian social welfare $\sim$ revenue (1st price auction)

Winner determination is a hard problem, but empirically successful algorithms are available. See Sandholm (2006) for an introduction.

For other optimality criteria, much less work has been done on algorithms. An exception is the work of Bouveret and Lemaître (2009).

T. Sandholm. Optimal Winner Determination Algorithms. In P. Cramton *et al.* (eds.), *Combinatorial Auctions*, MIT Press, 2006.

S. Bouveret and M. Lemaître. Computing Leximin-optimal Solutions in Constraint Networks. *Artificial Intelligence*, 19(2):343–364, 2009.

## Distributed Approach

Instead of devising algorithms for computing a socially optimal allocation in a centralised manner, we now want agents to be able to do this in a distributed manner by contracting deals locally.

- A *deal* $\delta = (A, A')$ is a pair of allocations (before/after).
- A deal may come with a number of side payments to compensate some of the agents for a loss in valuation. A *payment function* is a function $p : \mathcal{N} \to \mathbb{R}$ with $p(1) + \cdots + p(n) = 0$.

  Example: $p(i) = 5$ and $p(j) = -5$ means that agent $i$ *pays* €5, while agent $j$ *receives* €5.

## Negotiating Socially Optimal Allocations

We are not going to talk about designing a concrete negotiation protocol, but rather study the framework from an abstract point of view. The main question concerns the relationship between

- the *local view:* what deals will agents make in response to their individual preferences?; and
- the *global view:* how will the overall allocation of goods evolve in terms of social welfare?

We will go through this for one set of assumptions regarding the local view and one choice of desiderata regarding the global view.

U. Endriss, N. Maudet, F. Sadri and F. Toni. Negotiating Socially Optimal Allocations of Resources. *Journal of AI Research*, 25:315–348, 2006.

## The Local/Individual Perspective

A rational agent (who does not plan ahead) will only accept deals that improve her individual welfare:

- A deal $\delta = (A, A')$ is called *individually rational* (IR) if there *exists* a payment function $p$ such that $v_i(A') - v_i(A) > p(i)$ for all $i \in \mathcal{N}$, except possibly $p(i) = 0$ for agents $i$ with $A(i) = A'(i)$.

That is, an agent will only accept a deal if it results in a gain in value (or money) that strictly outweighs a possible loss in money (or value).

## The Global/Social Perspective

Suppose that, as system designers, we are interested in maximising *utilitarian social welfare:*

$$\mathrm{SW}_{\mathrm{util}}(A) \quad = \quad \sum_{i \in \mathcal{N}} v_i(A(i))$$

Observe that there is no need to include the agents' monetary balances into this definition, because they'd always add up to $0$.

While the local perspective is driving the negotiation process, we use the global perspective to assess how well we are doing.

# Example

Let $\mathcal{A} = \{ann, bob\}$ and $\mathcal{G} = \{chair, table\}$ and suppose our agents use the following utility functions:

$$
\begin{array}{rclrcl}
v_{ann}(\emptyset) & = & 0 & v_{bob}(\emptyset) & = & 0 \\
v_{ann}(\{chair\}) & = & 2 & v_{bob}(\{chair\}) & = & 3 \\
v_{ann}(\{table\}) & = & 3 & v_{bob}(\{table\}) & = & 3 \\
v_{ann}(\{chair, table\}) & = & 7 & v_{bob}(\{chair, table\}) & = & 8
\end{array}
$$

Furthermore, suppose the initial allocation of goods is $A_0$ with $A_0(ann) = \{chair, table\}$ and $A_0(bob) = \emptyset$.

Social welfare for allocation $A_0$ is $7$, but it could be $8$. By moving only a *single* good from agent $ann$ to agent $bob$, the former would lose more than the latter would gain (not individually rational).

The only possible deal would be to move the whole *set* $\{chair, table\}$.

## Convergence

The good news:

**Theorem 1 (Sandholm, 1998)** *Any sequence of IR deals will eventually result in an allocation with maximal social welfare.*

Discussion: Agents can act *locally* and need not be aware of the global picture (convergence is guaranteed by the theorem).

T. Sandholm. Contract Types for Satisficing Task Allocation: I Theoretical Results. Proc. AAAI Spring Symposium 1998.

## So why does this work?

The key to the proof is the insight that IR deals are exactly those deals that increase social welfare:

- **Lemma 2** *A deal $\delta = (A, A')$ is individually rational if and only if $\mathrm{SW}_{\mathrm{util}}(A) < \mathrm{SW}_{\mathrm{util}}(A')$.*

  Proof: ($\Rightarrow$) Rationality means that overall utility gains outweigh overall payments (which are $= 0$).

  ($\Leftarrow$) The social surplus can be divided amongst all agents by using, say, the following payment function:

$$p(i) \; = \; v_i(A') \; - \; v_i(A) \; - \; \underbrace{\frac{\mathrm{SW}_{\mathrm{util}}(A') - \mathrm{SW}_{\mathrm{util}}(A)}{|\mathcal{N}|}}_{> 0} \qquad \checkmark$$

Thus, as SW increases with every deal, negotiation must *terminate*. Upon termination, the final allocation $A$ must be *optimal*, because if there were a better allocation $A'$, the deal $\delta = (A, A')$ would be IR.

# Necessity of Multilateral Negotiation

The bad news is that outcomes that maximise utilitarian social welfare can only be guaranteed if the negotiation protocol allows for deals involving *any number of agents* and *goods:*

**Theorem 3** *Any deal $\delta = (A, A')$ may be necessary: there are valuations and an initial allocation such that any sequence of IR deals leading to an allocation with maximal utilitarian social welfare would have to include $\delta$ (unless $\delta$ is "independently decomposable").*

The proof involves the systematic definition of valuation functions such that $A'$ is optimal and $A$ is the second best allocation.

Independently decomposable deals (to which the result does not apply) are deals that can be split into two subdeals involving distinct agents.

U. Endriss, N. Maudet, F. Sadri and F. Toni. Negotiating Socially Optimal Allocations of Resources. *Journal of AI Research*, 25:315–348, 2006.

# More Convergence Results

For any given fairness or efficiency criterion, we would like to know how to set up a negotiation framework so as to be able to guarantee convergence to a social optimum. Some existing results:

- Pareto efficient outcomes via rational deals without money
- Outcomes maximising the egalitarian or the Nash CUF via specifically engineered deal criteria
- Envy-free outcomes via IR deals with a fixed payment function, for supermodular valuations (also on social networks)

U. Endriss, N. Maudet, F. Sadri and F. Toni. Negotiating Socially Optimal Allocations of Resources. *Journal of AI Research*, 25:315–348, 2006.

S. Ramezani and U. Endriss. *Nash Social Welfare in Multiagent Resource Allocation*. Proc. AMEC-2009.

Y. Chevaleyre, U. Endriss, and N. Maudet. *Allocating Goods on a Graph to Eliminate Envy*. Proc. AAAI-2007.

## Summary: Allocating Indivisible Goods

We have seen that finding a fair/efficient allocation in case of indivisible goods gives rise to a combinatorial optimisation problem.

Two approaches:

- *Centralised:* Give a complete specification of the problem to an optimisation algorithm (related to combinatorial auctions).
- *Distributed:* Try to get the agents to solve the problem. For certain fairness criteria and certain assumptions on agent behaviour, we can predict convergence to an optimal state.

# Literature: Allocating Indivisible Goods

Besides listing *fairness and efficiency criteria*, the "MARA Survey" also gives an overview of *allocation procedures* for indivisible goods. (It also covers *applications*, *preference* languages, and *complexity* results.)

We have largely neglected strategic (and have been brief on algorithmic) aspects, which are better developed in the *combinatorial auction* literature The handbook edited by Cramton *et al.* (2006) is a good starting point.

To find out more about *convergence* in distributed negotiation you may start by consulting the JAIR 2006 paper cited below.

Y. Chevaleyre, P.E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J.A. Rodríguez-Aguilar and P. Sousa. Issues in Multiagent Resource Allocation. *Informatica*, 30:3–31, 2006.

P. Cramton, Y. Shoham, and R. Steinberg (eds.). *Combinatorial Auctions*. MIT Press, 2006.

U. Endriss, N. Maudet, F. Sadri and F. Toni. Negotiating Socially Optimal Allocations of Resources. *Journal of AI Research*, 25:315–348, 2006.

# Divisible Goods: Cake-Cutting Algorithms

## Cake-Cutting Algorithms

Next we discuss methods for dividing a single *divisible* (heterogeneous) good between several agents, often referred to as the "*cake*".

Studied seriously since the 1940s (Banach, Knaster, Steinhaus). Simple model, yet still many open problems.

This will be an introduction to *cake-cutting algorithms:*

- Problem definition (formal model, proportionality, envy-freeness)
- Classical procedures (cut-and-choose, Banach-Knaster, . . . )
- Some open problems

# Formal Framework: Cake-Cutting

We want to divide a single divisible good, the *cake*, amongst $n$ *agents* by means of a series of parallel cuts.

The cake is represented by the unit interval $[0, 1]$:

```
|----------------------|
0                      1
```

Each agent $i$ has a *utility function* $u_i$ (or *valuation*, *measure*) mapping finite unions of subintervals of $[0, 1]$ to the reals, satisfying:

- Non-negativity: $u_i(B) \geqslant 0$ for all $B \subseteq [0, 1]$
- Normalisation: $u_i(\emptyset) = 0$ and $u_i([0, 1]) = 1$
- Additivity: $u_i(B \cup B') = u_i(B) + u_i(B')$ for disjoint $B, B' \subseteq [0, 1]$
- $u_i$ is continuous: the Intermediate-Value Theorem applies and single points do not have any value.

# Cut-and-Choose

The classical approach for dividing a cake between *two agents:*

- *One agent cuts the cake in two pieces (she considers to be of equal value), and the other chooses one of them (the piece she prefers).*

The cut-and-choose procedure satisfies two important properties:

- *Proportionality:* Each agent is guaranteed at least one half (general: $1/n$) according to her own valuation.

  Discussion: In fact, the first agent (if she is risk-averse) will receive exactly $1/2$, while the second will usually get more.

- *Envy-freeness:* No agent will envy (any of) the other(s).

  Discussion: Actually, for two agents, proportionality and envy-freeness amount to the same thing (in this model).

What if there are more than two agents?

# The Steinhaus Procedure

This procedure for *three agents* has been proposed by Steinhaus around 1943. Our exposition follows Brams and Taylor (1995).

(1) Agent 1 cuts the cake into three pieces (which she values equally).

(2) Agent 2 "passes" (if she thinks at least two of the pieces are $\geqslant 1/3$) or labels two of them as "bad". — If agent 2 passed, then agents 3, 2, 1 each choose a piece (in that order) and we are done. ✓

(3) If agent 2 did not pass, then agent 3 can also choose between passing and labelling. — If agent 3 passed, then agents 2, 3, 1 each choose a piece (in that order) and we are done. ✓

(4) If neither agent 2 or agent 3 passed, then agent 1 has to take (one of) the piece(s) labelled as "bad" by both 2 and 3. —
The rest is reassembled and 2 and 3 play cut-and-choose. ✓

Properties: *proportional*, *not* (always) *envy-free*, *not* (always) *contiguous*

S.J. Brams and A.D. Taylor. An Envy-free Cake Division Protocol. *American Mathematical Monthly*, 102(1):9–18, 1995.

# The Dubins-Spanier Procedure

Dubins and Spanier (1961) proposed a *proportional* procedure for $n$ *agents*, producing *contiguous* slices.

(1) A referee moves a knife slowly across the cake, from left to right. Any agent may shout "stop" at any time. Whoever does so receives the piece to the left of the knife.

(2) When a piece has been cut off, we continue with the remaining $n-1$ agents, until just one agent is left (who takes the rest). ✓

Observe that this is also *not envy-free*.

Note that this is not a discrete "protocol" (i.e., *not an algorithm*) in the narrow sense of the word (you cannot actually implement this!).

L.E. Dubins and E.H. Spanier. How to Cut a Cake Fairly. *American Mathematical Monthly*, 68(1):1–17, 1961.

# The Banach-Knaster Last-Diminisher Procedure

In the first ever paper on fair division, Steinhaus (1948) reports on a *proportional* procedure for $n$ *agents* due to Banach and Knaster.

(1) Agent 1 cuts off a piece (that she considers to represent $1/n$).

(2) That piece is passed around the agents. Each agent either lets it pass (if she considers it too small) or trims it down further (to what she considers $1/n$).

(3) After the piece has made the full round, the last agent to cut something off (the "last diminisher") is obliged to take it.

(4) The rest (including the trimmings) is then divided amongst the remaining $n-1$ agents. Play cut-and-choose once $n = 2$. ✓

Requires $O(n^2)$ cuts (though arguably fewer cuts "on average").

May not be *contiguous* (unless you always trim "from the right").

H. Steinhaus. The Problem of Fair Division. *Econometrica*, 16:101–104, 1948.

# The Even-Paz Divide-and-Conquer Procedure

Even and Paz (1984) introduced the *divide-and-conquer* procedure:

(1) Ask each agent to cut the cake at her $\lfloor \frac{n}{2} \rfloor : \lceil \frac{n}{2} \rceil$ mark.

(2) Associate the union of the leftmost $\lfloor \frac{n}{2} \rfloor$ pieces with the agents who made the leftmost $\lfloor \frac{n}{2} \rfloor$ cuts, and the rest with the others.

(3) Recursively apply the same procedure to each of the two groups, until only a single agent is left. ✓

Each agent is guaranteed a *proportional* piece. Takes $O(n \log n)$ cuts.

Woeginger and Sgall (2007) later showed that we cannot do much better: $\Omega(n \log n)$ is a lower bound on the number of queries for any proportional procedure producing contiguous pieces.

S. Even and A. Paz. A Note on Cake Cutting. *Discrete Applied Mathematics*, 7(3):285–296, 1984.

G.J. Woeginger and J. Sgall. On the Complexity of Cake Cutting. *Discrete Optimization*, 4(2):213–220, 2007.

## Proportionality and Envy-Freeness

Except for cut-and-choose, we have not yet seen any procedure that would guarantee envy-freeness.

**Fact 4** *Under additive preferences, a division for two agents is proportional if and only if it is envy-free.*

For $n \geqslant 3$, proportionality and envy-freeness are not the same properties anymore (unlike for $n = 2$):

**Fact 5** *Under additive preferences, any envy-free division is also proportional, but there are proportional divisions that are not envy-free.*

# Envy-Free Procedures

Achieving *envy-freeness* is much harder than achieving proportionality:

- For $n = 2$ the problem is easy: cut-and-choose does the job.
- For $n = 3$ we will see two solutions. They are already quite complicated: either the number of cuts is *not minimal* (but $> 2$), or *several simultaneously moving knives* are required.
- For $n = 4$, to date, no procedure producing *contiguous pieces* is known. Barbanel and Brams (2004), for example, give a moving-knife procedure requiring up to 5 cuts.
- For $n \geqslant 6$, to date, only procedures requiring an *unbounded* number of cuts are known (see, e.g., Brams and Taylor, 1995).

J.B. Barbanel and S.J. Brams. Cake Division with Minimal Cuts. *Mathematical Social Sciences*, 48(3):251–269, 2004.

S.J. Brams and A.D. Taylor. An Envy-free Cake Division Protocol. *American Mathematical Monthly*, 102(1):9–18, 1995.

# The Selfridge-Conway Procedure

The first discrete protocol achieving *envy-freeness* for $n = 3$ agents has been discovered independently by Selfridge and Conway (around 1960). It doesn't ensure contiguous pieces. Our exposition follows Brams and Taylor (1995).

(1) Agent 1 cuts the cake in three pieces (she considers equal).

(2) Agent 2 either "passes" (if she thinks at least two pieces are tied for largest) or trims one piece (to get two tied for largest pieces). — If she passed, then let agents 3, 2, 1 pick (in that order). ✓

(3) If agent 2 did trim, then let 3, 2, 1 pick (in that order), but require 2 to take the trimmed piece (unless 3 did). Keep the trimmings unallocated for now (note: the partial allocation is envy-free).

(4) Now divide the trimmings. Whoever of 2 and 3 received the *un*trimmed piece does the cutting. Let agents choose in this order: non-cutter, agent 1, cutter. ✓

S.J. Brams and A.D. Taylor. An Envy-free Cake Division Protocol. *American Mathematical Monthly*, 102(1):9–18, 1995.

# The Stromquist Procedure

Stromquist (1980) found an *envy-free* procedure for $n = 3$ producing *contiguous* pieces, though requiring four simultaneously *moving knifes:*

- A referee slowly moves a knife across the cake, from left to right (supposed to cut somewhere around the $1/3$ mark).
- At the same time, each agent is moving her own knife so that it would cut the righthand piece in half (wrt. her own valuation).
- The first agent to call "stop" receives the piece to the left of the referee's knife. The righthand part is cut by the middle one of the three agent knifes. If neither of the other two agents hold the middle knife, they each obtain the piece at which their knife is pointing. If one of them does hold the middle knife, then the other one gets the piece at which her knife is pointing. ✓

W. Stromquist. How to Cut a Cake Fairly. *American Mathematical Monthly*, 87(8):640–644, 1980.

# Summary: Cake-Cutting Algorithms

We have discussed various procedures for fairly dividing a cake (a metaphor for a single divisible good) amongst several agents.

- Fairness criteria: *proportionality* and *envy-freeness* (but other notions, such as equitability, Pareto efficiency, strategy-proofness ... are also of interest)
- Distinguish *discrete* from hypothetical *moving-knife* procedures.
- The problem becomes non-trivial for more than two agents, and there are many open problems relating to finding procedures with "good" properties for larger numbers.

# Literature: Cake-Cutting Algorithms

Both the book by Brams and Taylor (1996) and that by Robertson and Webb (1998) cover the cake-cutting problem in great depth.

The paper by Brams and Taylor (1995) does not only introduce their procedure for envy-free division for more than three players (not covered in this tutorial), but is also very nice for presenting several of the classical procedures in a systematic and accessible manner.

S.J. Brams and A.D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.

J. Robertson and W. Webb. *Cake-Cutting Algorithms: Be Fair if You Can*. A.K. Peters, 1998.

S.J. Brams and A.D. Taylor. An Envy-free Cake Division Protocol. *American Mathematical Monthly*, 102(1):9–18, 1995.

# Conclusion

## Conclusion: Fair Allocation of Resources

Fair allocation is a concrete example for a social choice problem, with obvious applications in AI, particularly in multiagent systems.

We have covered three topics:

- Fairness and efficiency criteria
- Allocation of indivisible goods, particularly convergence
- Divisible goods: cake-cutting algorithms

Most of the material covered is included in my lecture notes.

Next: *voting theory*

U. Endriss. *Lecture Notes on Fair Division*. Institute for Logic, Language and Computation, University of Amsterdam, 2009/2010.