Tutorial on Fair Division

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

Table of Contents

Introduction

Why Fair Division?

Fair division is the problem of dividing one or several goods amongst two or more agents in a way that satisfies a suitable fairness criterion.

Fair division has been studied in *philosophy*, *political science*, *economics*, and *mathematics* for a long time, but is also relevant to *computer science* and *multiagent systems:*

- *Resource allocation* is a central topic in MAS: it is either itself the application or agents need resources to perform tasks.
- Agents are *autonomous*. A solution needs to respect and balance their individual *preferences* $\rightsquigarrow$ requires definition of *fairness*.
- Once we have a well-defined fair division problem, we require an *algorithm* to solve it. And we might want to study its *complexity*.
- And there are many *applications*.

The Problem

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?

Tutorial Outline

This tutorial consists of three parts:

- Part 1. *Fairness and Efficiency Criteria* —
  What makes a good allocation? We will review and compare several proposals from the literature for how to define "fairness" and the related notion of economic "efficiency".

- Part 2. *Cake-Cutting Procedures* —
  How should we fairly divide a "cake" (a single *divisible good*)? We will review several algorithms and analyse their properties.

- Part 3. *Combinatorial Optimisation* —
  The fair division of *indivisible goods* gives rise to a combinatorial optimisation problem. We will cover centralised approaches (similar to auctions) and a distributed negotiation approach.

Fairness and Efficiency Criteria

What is a Good Allocation?

In this part of the tutorial we are going to give an 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 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 $u_i(A) \leq u_i(A')$ 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).

Social Welfare

Given the utilities of the individual agents, we can define a notion of social welfare and aim for an agreement that maximises social welfare.

The definition of *social welfare* commonly found in the MAS literature:

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

That is, social welfare is defined as the sum of the individual utilities. Maximising this function amounts to maximising *average utility*.

This is a reasonable definition, but it does not capture everything ...

▶ We need a systematic approach to defining social preferences.

Social Welfare Orderings

A *social welfare ordering* (SWO) $\preceq$ is a binary relation over $\mathbb{R}^n$ that is *reflexive*, *transitive*, and *complete*.

Intuitively, if $u, v \in \mathbb{R}^n$, then $u \preceq v$ means that v is socially preferred over u (not necessarily strictly).

We also use the following notation:

- $u \prec v$ iff $u \preceq v$ but not $v \preceq u$ (*strict social preference*)
- $u \sim v$ iff both $u \preceq v$ and $v \preceq u$ (*social indifference*)

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.

Every CUF *induces* an SWO: $u \preceq v \Leftrightarrow \mathrm{SW}(u) \leq \mathrm{SW}(v)$

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:

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

So this is what we have called "social welfare" a few slides back.

Remark: We define CUFs and SWOs 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).

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

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

Utilitarianism versus Egalitarianism

- In the MAS literature the utilitarian viewpoint (that is, social welfare = sum of individual utilities) is often taken for granted.
- In philosophy, economics, political science not.
- John Rawls' *"veil of ignorance"* (*A Theory of Justice*, 1971):

> *Without knowing what your position in society (class, race, sex, ...) will be, what kind of society would you choose to live in?*

- Reformulating the *veil of ignorance for multiagent systems:*

> *If you were to send a software agent into an artificial society to negotiate on your behalf, what would you consider acceptable principles for that society to operate by?*

- Conclusion: worthwhile to investigate egalitarian (and other) social principles also in the context of multiagent systems.

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 ($2 \cdot 6 < 4 \cdot 4$).

Ordered Utility Vectors

For any $u \in \mathbb{R}^n$, the *ordered utility vector* u^* is defined as the vector we obtain when we rearrange the elements of u in increasing order.

Example: Let $u = \langle 5, 20, 0 \rangle$ be a utility vector.

- $u^* = \langle 0, 5, 20 \rangle$ means that the weakest agent enjoys utility 0, the strongest utility 20, and the middle one utility 5.
- Recall that $u = \langle 5, 20, 0 \rangle$ means that the first agent enjoys utility 5, the second 20, and the third 0.

Rank Dictators

The *k-rank dictator* CUF for $k \in \mathcal{N}$ is mapping utility vectors to the utility enjoyed by the k-poorest agent:

$$\mathrm{SW}_k(u) \quad = \quad u^*_k$$

Interesting special cases:

- For $k = 1$ we obtain the *egalitarian* CUF.
- For $k = n$ we obtain an *elitist* CUF measuring social welfare in terms of the happiest agent.
- For $k = \lfloor \frac{n+1}{2} \rfloor$ we obtain the *median-rank-dictator* CUF.

The Leximin Ordering

We now introduce an SWO that may be regarded as a refinement of the SWO induced by the egalitarian CUF.

The *leximin ordering* $\preceq_{\text{lex}}$ is defined as follows:

$$u \preceq_{\text{lex}} v \;\Leftrightarrow\; u^* \text{ lexically precedes } v^* \text{ (not necessarily strictly)}$$

That means: $u^* = v^*$ or there exists a $k \leq n$ such that

- $u^*_i = v^*_i$ for all $i < k$ and
- $u^*_k < v^*_k$

Example: $u \prec_{\text{lex}} v$ for $u^* = \langle 0, 6, 20, 35 \rangle$ and $v^* = \langle 0, 6, 24, 25 \rangle$

Axiomatic Approach

So far we have simply defined some SWOs and CUFs and informally discussed their attractive and less attractive features.

Next we give a couple of examples for *axioms* — properties that we may or may not wish to impose on an SWO.

Interesting results are then of the following kind:

- A given SWO may or may not satisfy a given axiom.
- A given (class of) SWO(s) may or may not be the only one satisfying a given (combination of) axiom(s).
- A given combination of axioms may be *impossible* to satisfy.

The Pigou-Dalton Principle

A fair SWO will encourage inequality-reducing welfare redistributions.

Axiom 1 (PD) *An SWO $\preceq$ respects the Pigou-Dalton principle if, for all $u, v \in \mathbb{R}^n$, $u \preceq v$ holds whenever there exist $i, j \in \mathcal{N}$ such that:*

- $u_k = v_k$ *for all* $k \in \mathcal{N} \setminus \{i, j\}$ *— only* i *and* j *are involved;*
- $u_i + u_j = v_i + v_j$ *— the change is mean-preserving; and*
- $|u_i - u_j| > |v_i - v_j|$ *— the change is inequality-reducing.*

The idea is due to Arthur C. Pigou (British economist, 1877-1959) and Hugh Dalton (British economist and politician, 1887-1962).

Example: The leximin ordering satisfies the Pigou-Dalton principle.

Zero Independence

If agents enjoy very different utilities before the encounter, it may not be meaningful to use their absolute utilities afterwards to assess social welfare, but rather their relative gain or loss in utility. So a desirable property of an SWO may be to be independent from what individual agents consider "zero" utility.

Axiom 2 (ZI) *An SWO $\preceq$ is zero independent if $u \preceq v$ entails $(u+w) \preceq (v+w)$ for all $u, v, w \in \mathbb{R}^n$.*

Example: The SWO induced by the utilitarian CUF is zero independent, while the egalitarian SWO is not.

In fact, an SWO satisfies ZI iff it is represented by the utilitarian CUF. See Moulin (1988) for a precise statement of this result.

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

Scale Independence

Different agents may measure their personal utility using different "currencies". So a desirable property of an SWO may be to be independent from the utility scales used by individual agents.

Assumption: Here, we use positive utilities only: $u \in (\mathbb{R}^+)^n$.

Notation: Let $u \cdot v = \langle u_1 \cdot v_1, \ldots, u_n \cdot v_n \rangle$.

Axiom 3 (SI) *An SWO $\preceq$ over positive utilities is scale independent if $u \preceq v$ entails $(u \cdot w) \preceq (v \cdot w)$ for all $u, v, w \in (\mathbb{R}^+)^n$.*

Example: Clearly, neither the utilitarian nor the egalitarian SWO are scale independent, but the Nash SWO is.

By a similar result as the one mentioned before, an SWO satisfies SI iff it is represented by the Nash CUF.

Independence of the Common Utility Pace

Another desirable property of an SWO may be that we would like to be able to make social welfare judgements without knowing what kind of tax members of society will have to pay.

Axiom 4 (ICP) *An SWO $\preceq$ is independent of the common utility pace if $u \preceq v$ entails $f(u) \preceq f(v)$ for all $u, v \in \mathbb{R}^n$ and for every increasing bijection $f : \mathbb{R} \to \mathbb{R}$.*

For an SWO satisfying ICP only interpersonal comparisons ($u_i \leq v_j$ or $u_i \geq v_j$) matter, but no the (cardinal) intensity of $u_i - v_j$.

Example: The utilitarian SWO is not independent of the common utility pace, but the egalitarian SWO is. Any k-rank dictator SWO is.

Proportionality

If utility functions are *monotonic* ($B \subseteq B' \Rightarrow u(B) \leq u(B')$), then agents may want the *full* bundle and feel entitled to $1/n$ of its value.

In the context of monotonic utilities, this definition makes sense:

An allocation A is *proportional* if $u_i(A(i)) \geq \frac{1}{n} \cdot \hat{u}_i$ for every agent $i \in \mathcal{N}$, where $\hat{u}_i$ is the utility given to the full bundle by agent i.

Remark: Mostly used in the context of *additive* utilities.

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:

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

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).

Degrees of Envy

As we cannot always ensure envy-free allocations, another approach would be to try to *reduce* envy as much as possible.

But what does that actually mean?

A possible approach to systematically defining different ways of measuring the *degree of envy* of an allocation:

- Envy between two agents:
  $\max\{u_i(A(j)) - u_i(A(i)), 0\}$ or
  1 if $u_i(A(j)) > u_i(A(i))$ and 0 otherwise
- Degree of envy of a single agent:
  $\max$, sum
- Degree of envy of a society:
  $\max$, sum [or indeed any SWO/CUF]

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, Nash, egalitarian and other k-rank dictators), the leximin ordering, proportionality, and envy-freeness.
- All of these (and others) are interesting for multiagent systems. Which is appropriate depends on the application at hand, and some applications may even require the definition of new criteria.
- Understanding the structure of social welfare orderings is in itself an interesting research area (see discussion of axioms).

Literature

Moulin (1988) provides an excellent introduction to welfare economics. Much of the material from this part of the slides is taken from his book. Moulin (2003) offers a less technical version of the material.

The "MARA Survey" (Chevaleyre *et al.*, 2006) lists many SWOs and discusses their relevance to multiagent resource allocation in detail.

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.

Divisible Goods: Cake-Cutting Procedures

Cake-Cutting Procedures

- Cake-cutting as a metaphor for the fair division of a single divisible (and heterogeneous) good between n agents (called *players*).
- Studied seriously since the 1940s (Banach, Knaster, Steinhaus). Simple model, yet still many open problems.
- This part of the tutorial will be an introduction to the field:
  - Problem definition (proportionality, envy-freeness)
  - Classical procedures (Cut-and-Choose, Banach-Knaster, . . .)
  - Some open problems

Cakes

We want to divide a single divisible good, commonly referred to as a *cake* (amongst n *players*), by means of a series of parallel cuts.

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

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

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

- Non-negativity: $u_i(B) \geq 0$ for all $B \subseteq [0, 1]$
- Additivity: $u_i(B \cup B') = u_i(B) + u_i(B')$ for disjoint $B, B' \subseteq [0, 1]$
- Normalisation: $u_i([0, 1]) = 1$
- 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 players:*

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

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

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

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

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

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

Operational Properties

Beyond fairness, we may also be interested in the "operational" properties of the procedures themselves:

- Does the procedure guarantee that each player receives a single *contiguous* slice (rather than the union of several subintervals)?
- Is the *number of cuts* minimal? If not, is it at least bounded?
- Does the procedure require an active *referee*, or can all actions be performed by the players themselves?
- Is the procedure a proper algorithm (a *protocol*), i.e., can it be implemented as a discrete sequence of queries to the agents? (no need for a "continuously moving knife"—to be discussed)?

Cut-and-choose is ideal and as simple as can be with respect to all of these properties. For $n > 2$, it won't be quite that easy though ...

Proportionality and Envy-Freeness

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

Fact 1 *Any envy-free division is also proportional, but there are proportional divisions that are not envy-free.*

Over the next few slides, we are first going to focus on cake-cutting procedures that achieve proportional divisions.

The Steinhaus Procedure

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

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

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

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

(4) If neither player 2 or player 3 passed, then player 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. ✓

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

Properties

The Steinhaus procedure —

- Guarantees a *proportional* division of the cake (under the standard assumption that players are risk-averse: they want to maximise their payoff in the worst case).
- Is *not envy-free*.
- Is a discrete procedure that does not require a referee.
- Requires *at most 3 cuts* (as opposed to the minimum of 2 cuts). The resulting pieces do not have to be contiguous (namely if both 2 and 3 label the middle piece as "bad" and 1 takes it; and if the cut-and-choose cut is different from 1's original cut).

The Banach-Knaster Last-Diminisher Procedure

In the first ever paper on fair division, Steinhaus (1948) reports on his own solution and a generalisation to *arbitrary n* proposed by Banach and Knaster.

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

(2) That piece is passed around the players. Each player 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 player to cut something off (the "last diminisher") is obliged to take it.

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

The procedure's properties are similar to that of the Steinhaus procedure (proportional; not envy-free; not contiguous; bounded number of cuts).

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

The Dubins-Spanier Procedure

Dubins and Spanier (1961) proposed an alternative *proportional* procedure for *arbitrary* n. It produces *contiguous* slices (and hence uses a minimal number of cuts), but it is *not discrete* and requires the help of a *referee*.

(1) A referee moves a knife slowly across the cake, from left to right. Any player 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 $n-1$ players, until just one player is left (who takes the rest). ✓

Observe that this is also *not envy-free*. The last chooser is best off (she is the only one who can get more than $1/n$).

Remark: Discretisation is possible by asking players to mark the cake where they *would* call "stop" . . .

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

The Even-Paz Divide-and-Conquer Procedure

Even and Paz (1984) investigated *upper bounds* for the number of *queries* (cuts or marks) required to produce a proportional division for n *players*, without allowing a moving knife.

They introduced the following *divide-and-conquer* protocol:

(1) Ask each player 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 players who made the leftmost $\lfloor \frac{n}{2} \rfloor$ cuts (group 1), and the rest with the others (group 2).

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

Theorem 2 *The Even-Paz procedure requires* $O(n \log n)$ *cuts.*

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

Envy-Free Procedures

Next we discuss procedures for achieving *envy-free* divisions.

- 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 $n \geq 5$, 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 $n = 3$ has been discovered independently by Selfridge and Conway (around 1960). Our exposition follows Brams and Taylor (1995).

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

(2) Player 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 players 3, 2, 1 pick (in that order). ✓

(3) If player 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 players choose in this order: non-cutter, player 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 player is moving her own knife so that it would cut the righthand piece in half (wrt. her own valuation).
- The first player 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 player knifes. If neither of the other two players 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 Procedures

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

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

Literature

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.

Indivisible Goods: Combinatorial Optimisation

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?

Setting

For the remainder of today we will work in this 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* $v_i : 2^{\mathcal{G}} \rightarrow \mathbb{R}$.
  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 division 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:

- Checking whether an allocation is *Pareto efficient* is coNP-complete.
- 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 *envy-free* allocation exists is NP-complete; checking whether an allocation that is both Pareto efficient and envy-free exists is even Σ_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.

Maximising Egalitarian Social Welfare

An algorithm using (mixed) *integer programming* for maximising *egalitarian* social welfare with utilities represented in the *XOR language:*

XOR-language means: Agent i submits n_i atomic bids $\langle B_{ij}, u_{ij} \rangle$ with $B_{ij} \subseteq \mathcal{G}$ and $u_{ij} \in \mathbb{R}^+$. Then utility of $B \subseteq \mathcal{G}$ is $\max\{u_{ij} \mid B_{ij} \subseteq B\}$.

IP variables: $x_{ij} \in \{0, 1\}$ ("agent i gets jth bundle"); $y \geq 0$ $(= \mathrm{SW}_{\mathrm{egal}})$

IP algorithm: *maximise y subject to* three constraints:

(1) Each good gets allocated to at most one agent:

$$(\forall k \leq |\mathcal{G}|) \quad \sum_{i \in \mathcal{N}} \sum_{j=1}^{n_i} [k \in B_{ij}] \cdot x_{ij} \quad \leq \quad 1, \quad \text{where } [k \in B_{ij}] \in \{0, 1\}$$

(2) Each agent receives at most one bundle specified in their XOR-bid:

$$(\forall i \in \mathcal{N}) \quad \sum_{j=1}^{n_i} x_{ij} \quad \leq \quad 1$$

(3) Egalitarian social welfare is at most equal to any individual utility:

$$(\forall i \in \mathcal{N}) \quad y \quad \leq \quad \sum_{j=1}^{n_i} u_{ij} \cdot x_{ij}$$

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 $p : \mathcal{N} \rightarrow \mathbb{R}$ 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 its 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:

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

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

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 3 (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.

Why does this work?

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

- **Lemma 4** *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 deal participants 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.

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 5 *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 δ (unless δ 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.

The theorem holds even when valuation functions are restricted to be monotonic or dichotomous.

Modular Domains

A valuation function v_i is called *modular* if it satisfies the following condition for all bundles $B_1, B_2 \subseteq \mathcal{G}$:

$$v_i(B_1 \cup B_2) \quad = \quad v_i(B_1) + v_i(B_2) - v_i(B_1 \cap B_2)$$

That is, in a modular domain there are no synergies between items; you can get the value of a bundle by adding up the values of its elements.

▶ Negotiation in modular domains *is* feasible:

Theorem 6 *If all valuation functions are modular, then IR 1-deals (each involving just one item) suffice to guarantee outcomes with maximal utilitarian social welfare.*

We also know that the class of modular valuation functions is *maximal:* no larger class can guarantee the same convergence property.

Y. Chevaleyre, U. Endriss, and N. Maudet. Simple Negotiation Schemes for Agents with Simple Preferences. *JAAMAS*, 2009. In press.

Comparing Negotiation Policies

While we know from Theorem 6 that 1-deals (blue) guarantee an optimal result, an experiment (20 agents, 200 goods, modular valuations) suggests that general bilateral deals (red) achieve the same goal in fewer steps:

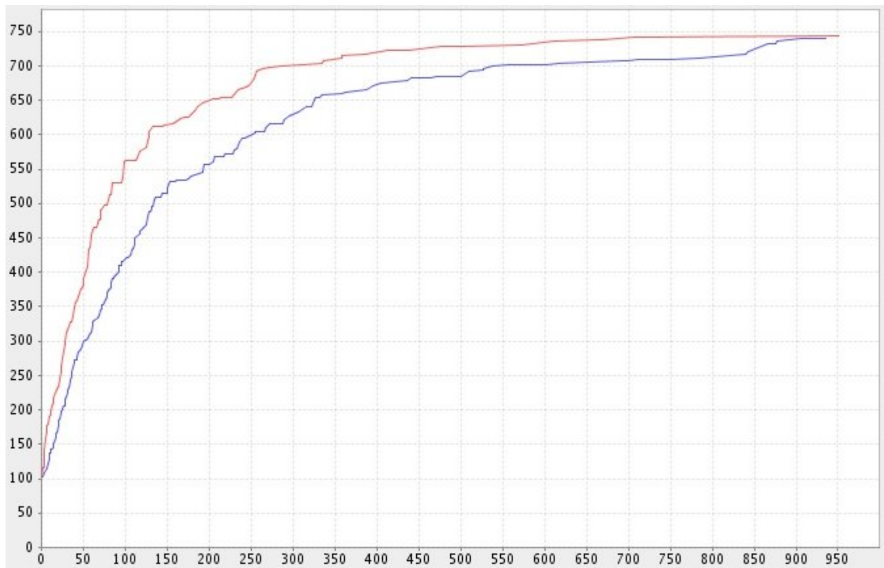

The graph shows how utilitarian social welfare (y-axis) develops as agents attempt to contract more an more deals (x-axis) amongst themselves. Graph generated using the MADRAS platform of Buisman *et al.* (2007).

H. Buisman, G. Kruitbosch, N. Peek, and U. Endriss. *Simulation of Negotiation Policies in Distributed Multiagent Resource Allocation*. Proc. ESAW-2007.

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 work:

- 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

Besides listing *fairness and efficiency criteria* (Part 1), 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.

Conclusion

Conclusion

Fair division is relevant to multiagent systems research. In this tutorial we have covered three topics:

- Fairness and efficiency defined in terms of individual preferences.
- Classical algorithms for the cake-cutting problem (divisible good).
- Combinatorial optimisation and negotiation for indivisible goods.

These slides and the lecture notes will remain available on the tutorial website, and more extensive material can be found on the website of my Amsterdam course on Computational Social Choice:

- `http://www.illc.uva.nl/~ulle/teaching/easss-2009/`
- `http://www.illc.uva.nl/~ulle/teaching/comsoc/`

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