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: 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*.

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?

Ordinal Preferences

- The *preference relation* of agent i over alternative agreements:

$$x \preceq_i y \quad \Leftrightarrow \quad \text{agreement } y \text{ is at least as good as } x \text{ (for agent } i)$$

- We shall also define the following bits of notation:
  - $x \prec_i y$ iff $x \preceq_i y$ but not $y \preceq_i x$ (*strict preference*)
  - $x \sim_i y$ iff both $x \preceq_i y$ and $y \preceq_i x$ (*indifference*)

Cardinal Preferences

- A *utility function* u_i (for agent i) is a mapping from the space of agreements to the reals.
- Example: $u_i(x) = 10$ means that agent i assigns a value of 10 to agreement x.
- A utility function u_i *representing* the preference relation $\preceq_i$:

$$x \preceq_i y \quad \Leftrightarrow \quad u_i(x) \leq u_i(y)$$

- Remark: In these slides, we are going to use the term *valuation* to model preferences over goods (allocations/bundles), while *utility* is used to model preferences over agreements, which may include a monetary component. If monetary side payments are not possible then we use valuation and utility interchangeably.

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 and Negotiation* —
  The fair division of a set of *indivisible goods* gives rise to a combinatorial optimisation problem. We will concentrate on an approach based on distributed negotiation.

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

- Let $\mathcal{A} = \{1..n\}$ be our agent society throughout.
- We have to decide on an *agreement*. This may be an allocation of goods, possibly coupled with monetary side payments (but much of this part of the tutorial is not specific to resource allocation).
- Each agent i has a *utility function* u_i over alternative agreements, which also induces a *preference ordering* $\preceq_i$.
- An agreement x gives rise to a *utility vector* $\langle u_1(x), \ldots, u_n(x) \rangle$
- 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

An agreement x is *Pareto-dominated* by another agreement y iff:

- $x \preceq_i y$ for all members i of society; and
- $x \prec_i y$ for at least one member i of society.

An agreement is called *Pareto efficient* iff it is not Pareto-dominated by any other feasible agreement (named so after Vilfredo Pareto, Italian economist, 1848–1923).

Pareto efficiency is very often considered a minimum requirement for any agreement/allocation.

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:

$$sw(u) \quad = \quad \sum_{i \in \mathcal{A}gents} u_i$$

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 *connected*.

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*)

Terminology: In the (economics) literature, connectedness is usually referred to as "completeness". Furthermore, many authors use the letters R, P and I instead of $\preceq$, $\prec$ and $\sim$.

Collective Utility Functions

- A *collective utility function* (CUF) is a function $W : \mathbb{R}^n \to \mathbb{R}$ mapping utility vectors to the reals.
- Intuitively, if $u \in \mathbb{R}^n$, then $W(u)$ is the utility derived from u by society as a whole.
- Every CUF *represents* an SWO: $u \preceq v \Leftrightarrow W(u) \leq W(v)$

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:

$$sw_u(u) \quad = \quad \sum_{i \in \mathcal{A}gents} u_i$$

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

Egalitarian Social Welfare

The *egalitarian* CUF measures social welfare as follows:

$$sw_e(u) \quad = \quad min\{u_i \mid i \in \mathcal{A}gents\}$$

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 usually taken for granted.
- In philosophy/sociology/economics 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 collective utility function* sw_N is defined as the product of individual utilities:

$$sw_N(u) \quad = \prod_{i \in \mathcal{A}gents} u_i$$

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

We now need some more notation . . .

For any $u \in \mathbb{R}^n$, the *ordered utility vector* $\vec{u}$ 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.

- $\vec{u} = \langle 0, 5, 20 \rangle$ 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 $k \in \mathcal{A}$ is mapping utility vectors to the utility enjoyed by the k-poorest agent:

$$sw_k(u) \quad = \quad \vec{u}_k$$

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 $k = \lfloor n/2 \rfloor$ 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* $\preceq_\ell$ is defined as follows:

$$u \preceq_\ell v \;\Leftrightarrow\; \vec{u} \text{ lexically precedes } \vec{v} \text{ (not necessarily strictly)}$$

That means: $\vec{u} = \vec{v}$ or there exists a $k \leq n$ such that

- $\vec{u}_i = \vec{v}_i$ for all $i < k$ and
- $\vec{u}_k < \vec{v}_k$

Example: $u \prec_\ell v$ for $\vec{u} = \langle 0, 6, 20, 29 \rangle$ and $\vec{v} = \langle 0, 6, 24, 25 \rangle$

Remark: The top agreement according to the *leximin*-ordering is also known as the *Kalai-Smorodinsky solution* (for "normalised" utilities).

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).

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 1 (ZI) *An SWO $\preceq$ is zero independent iff $u \preceq v$ 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 CUF 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, *i.e.* $u \in (\mathbb{R}^+)^n$.

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

Axiom 2 (SI) *An SWO $\preceq$ over positive utilities is scale independent iff $u \preceq v$ 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 CUF are scale independent, but the Nash CUF 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 3 (ICP) *An SWO $\preceq$ is independent of the common utility pace iff $u \preceq v$ 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 ($u_i \leq v_i$ or $u_i \geq v_i$) matter, but no the (cardinal) intensity of $u_i - v_i$.

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

All-or-Nothing Criteria

Next we are going to introduce two criteria for "good" allocations that either are or are not satisfied, but there is no further differentiation amongst allocations of varying quality:

- *proportionality* (a.k.a. proportional fairness)
- *envy-freeness*

Remark 1: Pareto efficiency may also be considered such an "all-or-nothing" criterion: an allocation either is or is not Pareto efficient. However, the notion of Pareto-dominance does permit a more fine-grained ranking.

Remark 2: We could interpret proportionality and envy-freeness also as SWOs, although this is not commonly done.

Proportionality

Let $\hat{u}_i$ be the utility that agent i would get for the most attractive agreement (think of this as the utility for obtaining all the goods).

The number of agents is n. So an agent may feel that they are entitled to $\frac{1}{n}$th of the overall value of the goods under discussion.

An agreement is called *proportional* iff $u_i \geq \frac{1}{n} \cdot \hat{u}_i$ for each agent i.

Envy-Freeness

The next definition is specific to allocations (so we cannot continue working with the more abstract notion of "agreement").

An *allocation* A is a function mapping each agent i to the bundle of goods it receives in that allocation. Suppose utility functions are now declared over such bundles.

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

$$A(i) \quad \succeq_i \quad A(j)$$

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

Note that envy-free allocations do not always *exist* (at least not if we require either complete or Pareto optimal 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 even without $\max$]
- Degree of envy of a single agent:
  $0\text{-}1, \max, \text{sum}$
- Degree of envy of a society:
  $\max, \text{sum}$ [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), *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) is an easier read but of less mathematical depth.

The "MARA Survey" (Chevaleyre *et al.*, 2006) lists many SWOs and discusses the 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 will discuss the division of a single divisible good, commonly referred to as a *cake* (amongst n *players*). It's the sort of cake where you can cut off slices with a single cut (so not a round *tart*).

More abstractly, you may think of a cake as the unit interval $[0, 1]$:

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

Each player i has a *valuation* function v_i mapping finite unions of subintervals (slices) to the reals, satisfying the following conditions:

- Non-negativity: $v_i(X) \geq 0$ for all $X \subseteq [0, 1]$
- Additivity: $v_i(X \cup Y) = v_i(X) + v_i(Y)$ for disjoint $X, Y \subseteq [0, 1]$
- v_i is continuous (the Intermediate-Value Theorem applies) and single points do not have any value.
- $v_i([0, 1]) = 1$ (*i.e.* it's like a probability measure)

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.

Further Properties

We may also be interested in the following properties:

- *Equitability:* Under an equitable division, each player assigns the same value to the slice they receive.

  Discussion: Cut-and-choose clearly violates equitability. Furthermore, for $n > 2$, equitability is often in conflict with envy-freeness, and we shall not discuss it any further today.

- *Pareto efficiency:* Under an efficient division, no other division will make somebody better and nobody worse off.

  Discussion: Generally speaking, cut-and-choose violates Pareto efficiency: suppose player 1 really likes the middle of the cake and player 2 really like the two outer parts (then *no* one-cut procedure will be efficient). But amongst all divisions into two contiguous slices, the cut-and-choose division will be efficient.

Operational Properties

The properties discussed so far all relate to the fairness (or efficiency) of the resulting division of the cake. Beyond that 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.k.a. a *protocol*), requiring a finite number of specific actions from the participants (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 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*. However, observe that players 2 and 3 will not envy anyone. Only player 1 may envy one of the others in case the situation where 2 and 3 play cut-and-choose occurs.
- 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 for $n = 3$ 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* anymore and it requires the active 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 $1/n$).

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 cuts required to produce a proportional division for n *players*, without allowing either a moving knife or "virtual cuts" (marks).

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. ✓

Fact 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) has come up with an envy-free procedure for $n = 3$ producing contiguous pieces, albeit requiring the use of 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, and the other two pieces are allocated in the obvious manner (ensuring proportionality). ✓

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.

Overview of Procedures

Procedure	Players	Type	Division	Pieces	Cuts
Cut-and-choose	$n = 2$	protocol	envy-free $(*)$	contiguous	minimal
Steinhaus	$n = 3$	protocol	proportional	not contig.	min.+1
Banach-Knaster (last-diminisher)	any n	protocol	proportional	not contig. (but could be)	bounded
Dubins-Spanier	any n	1 knife	proportional	contiguous	minimal
Even-Paz (divide-and-conquer)	any n	protocol	proportional	contiguous	$O(n \log n)$
Selfridge-Conway	$n = 3$	protocol	envy-free $(*)$	not contig.	≤ 5
Stromquist	$n = 3$	4 knives	envy-free $(*)$	contiguous	minimal

$(*)$ Recall that envy-freeness entails proportionality.

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 and Negotiation

Allocation of Indivisible Goods

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

- In the *centralised* approach, we need to devise an optimisation algorithm to compute an allocation meeting our fairness and efficiency requirements.
  - We will briefly mention complexity results,
  - and point out connections to combinatorial auctions.
- 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* $\mathcal{A} = \{1..n\}$ and finite set of indivisible *goods* $\mathcal{G}$.
- An *allocation* A is a partitioning of $\mathcal{G}$ amongst the agents in $\mathcal{A}$.
  Example: $A(i) = \{r_5, r_7\}$ — agent i owns resources r_5 and r_7
- Each agent $i \in \mathcal{A}$ 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

Later we will define (quasi-linear) utility functions over these valuations (to account for payments). For now think of valuation as utility.

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

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 a Pareto efficient envy-free allocation 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 (2007).

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

S. Bouveret and M. Lemaître. *New Constraint Programming Approaches for the Computation of Leximin-Optimal Solutions in Constraint Networks*. IJCAI-2007.

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 $p : \mathcal{A} \to \mathbb{R}$ with $\sum_{i \in \mathcal{A}} p(i) = 0$.

  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 resources evolve in terms of social welfare?

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) iff there exists a payment function p such that $v_i(A') - v_i(A) > p(i)$ for all $i \in \mathcal{A}$, except possibly $p(i) = 0$ for agents i with $A(i) = A'(i)$.

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

The Global/Social Perspective

For now, suppose that as system designers we are interested in maximising *utilitarian social welfare:*

$$sw_u(A) \quad = \quad \sum_{i \in \mathcal{A}gents} v_i(A)$$

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* $\{chair, table\}$.

Linking the Local and the Global Perspectives

It turns out that individually rational deals are exactly those deals that increase social welfare:

Lemma 3 (Rationality and social welfare) *A deal* $\delta = (A, A')$ *with side payments is individually rational iff* $sw_u(A) < sw_u(A')$.

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:

$$p(i) \;=\; v_i(A') \;-\; v_i(A) \;-\; \underbrace{\frac{sw_u(A') - sw_u(A)}{|\mathcal{A}|}}_{>0}$$

✓

Discussion: The lemma confirms that individually rational behaviour is "appropriate" in utilitarian societies.

Termination

We can now prove a first result on negotiation processes:

Lemma 4 (Termination) *There can be no infinite sequence of IR deals; that is, negotiation must always terminate.*

Proof: Follows from the first lemma and the observation that the space of distinct allocations is finite. ✓

Convergence

It is now easy to prove the following *convergence* result (originally stated by Sandholm in the context of distributed task allocation):

Theorem 5 (Sandholm, 1998) *Any sequence of IR deals will eventually result in an allocation with maximal social welfare.*

Proof: Termination has been shown in the previous lemma. So let A be the terminal allocation. Assume A is *not* optimal, *i.e.* there exists an allocation A' with $sw_u(A) < sw_u(A')$. Then, by our first lemma, $\delta = (A, A')$ is individually rational $\Rightarrow$ contradiction. ✓

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.

Multilateral Negotiation

On the downside, outcomes that maximise utilitarian social welfare can only be guaranteed if the negotiation protocol allows for deals involving *any number of agents* and *resources:*

Theorem 6 (Necessity of complex deals) *Any deal* $\delta = (A, A')$ *may be necessary, i.e. there are valuation functions and an initial allocation such that any sequence of individually rational 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* iff 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 7 (Modular domains) *If all valuation functions are modular, then individually rational 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 strictly larger class could still guarantee the same convergence property.

Y. Chevaleyre, U. Endriss, and N. Maudet. On Maximal Classes of Utility Functions for Efficient one-to-one Negotiation. IJCAI-2005.

Comparing Negotiation Policies

While we know from Theorem 7 that 1-deals (blue) guarantee an optimal result, an experiment (20 agents, 200 resources, modular utilities) suggests that general bilateral deals (red) achieve the same goal in fewer steps:

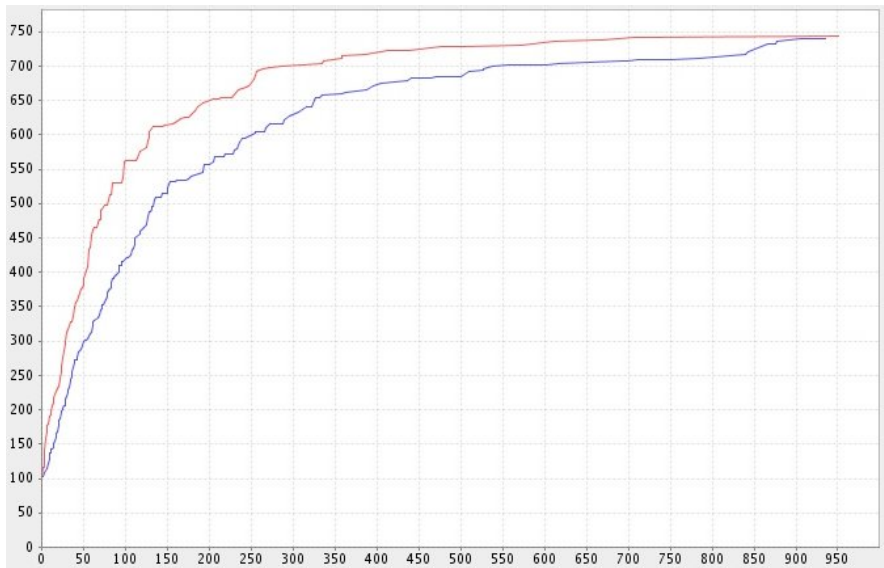

The graph shows how utilitarian social welfare (y-axis) develops as agents attempt to contract more an more deals (x-axis) amongst themselves. Graph generated using the MADRAS platform of Buisman *et al.* (2007).

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

There are some known results of this sort, notably for Pareto efficiency and maximal egalitarian social welfare.

Next we will consider the case of *envy-freeness*. Guaranteeing convergence to an envy-free allocation is particularly difficult:

- Envy-free allocations do not always exist.
- A local deal in one part of society can make another agent somewhere else envious.

Y. Chevaleyre, U. Endriss, S. Estivie and N. Maudet. *Reaching Envy-free States in Distributed Negotiation Settings*. IJCAI-2007.

Envy-freeness in the Presence of Money

Unfortunately, there are cases where envy-free allocations do not *exist*.

Example: 2 agents, 1 good desired by both

We can try to circumvent this problem by taking the balance of past side payments into account when defining envy-freeness:

- Associate each allocation A with a payment *balance* $\pi : \mathcal{A} \to \mathbb{R}$, mapping agents to the sum of payments they have made so far.
- A *state* (A, π) is a pair of an allocation and a payment balance.
- Each agent $i \in \mathcal{A}$ has got a (quasi-linear) *utility function* $u_i : 2^{\mathcal{G}} \times \mathbb{R} \to \mathbb{R}$, defined as follows: $u_i(R, x) = v_i(R) - x$.
- A state (A, π) is *envy-free* iff $u_i(A(i), \pi(i)) \geq u_i(A(j), \pi(j))$ for all agents $i, j \in \mathcal{A}$. An *efficient envy-free* (EEF) state is an envy-free state maximising utilitarian social welfare.

Existence of EEF States

Unlike for the case without money, EEF states always exist. There's a simple proof for supermodular valuations (proof on next slide):

Theorem 8 (Existence of EEF states) *If all valuations are supermodular, then an EEF state always exists.*

Note that *supermodular* valuations are valuations satisfying the following condition for all bundles $B_1, B_2 \subseteq \mathcal{G}$:

$$v(B_1 \cup B_2) \;\geq\; v(B_1) + v(B_2) - v(B_1 \cap B_2)$$

For ease of presentation (not technically required), from now on we assume that all valuations are *normalised:* $v(\{\,\}) = 0$.

Proof of Theorem 8

Of course, there's always an *efficient* allocation; let's call it A^*.

We'll try to fix a payment balance π^* such that (A^*, π^*) is EEF:

$$\pi^*(i) \quad = \quad v_i(A^*) \, - \, sw(A^*)/n$$

Note: the $\pi^*(i)$ add up to 0, so it's a *valid* payment balance. ✓

Now let $i, j \in \mathcal{A}$ be any two agents. As A^* is efficient, giving both $A^*(i)$ and $A^*(j)$ to i won't increase social welfare any further:

$$v_i(A^*(i)) \, + \, v_j(A^*(j)) \;\geq\; v_i(A^*(i) \cup A^*(j))$$

Now apply the supermodularity condition ... and rewrite:

$$\begin{aligned} v_i(A^*(i)) + v_j(A^*(j)) \;&\geq\; v_i(A^*(i)) + v_i(A^*(j)) \\ v_i(A^*(i)) - [v_i(A^*) - sw(A^*)/n] \;&\geq\; v_i(A^*(j)) - [v_j(A^*) - sw(A^*)/n] \\ u_i(A^*(i), \pi^*(i)) \;&\geq\; u_i(A^*(j), \pi^*(j)) \end{aligned}$$

That is, i does not envy j. Hence, (A^*, π^*) is *envy-free* (and EEF). ✓

Envy-freeness and Individual Rationality

Now that we know that EEF states always exist, we want to find them by means of *rational* negotiation. Unfortunately, this is *impossible*.

Example: 2 agents, 1 resource

$$v_1(\{r\}) = 4 \qquad v_2(\{r\}) = 7$$

Suppose agent 1 owns r to begin with.

The efficient allocation would be where agent 2 owns r.

An individually rational deal would require a payment within $(4, 7)$.

But to ensure envy-freeness, the payment should be in $[2, 3.5]$.

Compromise: We shall enforce an *initial equitability payment* $\pi_0(i) = v_i(A_0) - sw(A_0)/n$ before beginning negotiation.

Discussion: If we have $v_i(A_0) = 0$ or if we think of A_0 as random (agents cannot derive any entitlements), then this is ok. Also note that the payments do not achieve either envy-freeness or efficiency.

Globally Uniform Payments

Again, realise just how unlikely it seems that our goal of guaranteeing EEF outcomes for distributed negotiation amongst self-interested (IR) agents could succeed ("non-local effects of local deals") ...

We will have to restrict the freedom of agents a little by fixing a specific payment function (still IR!):

- Let $\delta = (A, A')$ be an IR deal. The payments as given by the globally uniform payment function (GUPF) are defined as:

$$p(i) \;=\; [v_i(A') - v_i(A)] \;-\; [sw(A') - sw(A)]/n$$

That is, we evenly distribute the (positive!) social surplus to *all* agents.

Convergence in Supermodular Domains

After having deciphered all the acronyms, this should be rather surprising:

Theorem 9 (Convergence) *If all valuations are supermodular and if initial equitability payments have been made, then any sequence of IR deals using the GUPF will eventually result in an EEF state.*

Proof: First try to show that this invariant holds for all states (A, π):

$$\pi(i) \;=\; v_i(A) - sw(A)/n \quad (*)$$

True initially by definition (initial equitability payments). Now let $\delta = (A, A')$ be a deal, with payment balances π and π'. Compute:

$$\begin{aligned} \pi'(i) &= \pi(i) + [v_i(A') - v_i(A)] - [sw(A') - sw(A)]/n \\ &= v_i(A') - sw(A')/n \quad \leadsto \; (*) \text{ holds by induction} \end{aligned}$$

Theorem 5 shows that the system must *converge* to an *efficient* allocation A^* (whatever the payment function). Then the proof of Theorem 8 demonstrates that $(*)$ implies that (A^*, π^*) must be an *EEF* state. ✓

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.
  - maximal utilitarian social welfare
  - envy-free states (allocation + side payments)

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 algorithmic and strategic 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).
- Distributed approach based on negotiation for indivisible goods.

These slides will remain available on the tutorial website, and more extensive material may be obtained from the website of my course on Computational Social Choice given at the ILLC in Amsterdam:

- `http://www.illc.uva.nl/~ulle/teaching/aamas-2008/`
- `http://www.illc.uva.nl/~ulle/teaching/comsoc/`