Computational Social Choice: Spring 2019

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

Plan for Today

The main purpose of today's lecture is to give you enough information to allow you to decide whether you want to take this course.

- What is Computational Social Choice? Why study COMSOC?
- Organisational Matters: planning, expectations, assessment, ...
- First Topic: Cake Cutting (Fair Division of a Continuous Resource)

What is Computational Social Choice?

Social choice theory is about methods for *collective decision making*, such as *political* decision making by groups of *economic* agents.

Its methodology ranges from the *philosophical* to the *mathematical*. It is traditionally studied in *Economics* and *Political Science* and it is a close cousin of both *decision theory* and *game theory*.

Its findings are relevant to multiple *applications*, such as these:

- How to fairly allocate resources to the members of a society?
- How to fairly divide computing time between several users?
- How to elect a president given people's preferences?
- How to combine the website rankings of multiple search engines?
- How to aggregate the views of different judges in a court case?
- How to extract information from noisy crowdsourced data?

Computational social choice, the topic of this course, emphasises the fact that any method of decision making is ultimately an *algorithm*.

Relationship with AI

Ideas from Economics entered AI when it became clear that we can use them to study interaction between agents in a *multiagent system*. Nowadays, the study of "*economic paradigms*" is all over AI.

The influential *One Hundred Year Study on Artificial Intelligence* (2016) singles out the following eleven "*hot topics*" in AI:

large-scale machine learning | deep learning | reinforcement learning
robotics | computer vision | natural language processing
collaborative systems | crowdsourcing and human computation
algorithmic game theory and *computational social choice*
internet of things | neuromorphic computing

And indeed, while COMSOC transcends several disciplines, about half of it gets published in AI conference proceedings and journals.

P. Stone et al. "Artificial Intelligence and Life in 2030". One Hundred Year Study on Artificial Intelligence. Stanford, 2016.

Course Overview

We will discuss three major scenarios of collective decision making:

- *Fair Allocation of Goods to Agents*

  Scenario: Several agents have individual preferences over which goods to receive. You need to compute a good allocation.

- *Voting and Preference Aggregation*

  Scenario: Several agents have individual preferences over alternative "states of affairs" (could be election outcomes, but also allocations).

- *Judgment Aggregation*

  Scenario: Several agents make judgments regarding the truth of certain statements (which could be of the form "A is better than B").

Thus: We will move from more specific to more general scenarios.

Remark: This is not an exhaustive list of topics studied in COMSOC. The main subarea omitted is *coalition formation* ($\hookrightarrow$ Game Theory).

Nature of the Course

This is an advanced *research-oriented* course: we'll move fast and often touch upon recent research. The focus is on *theory*.

Our methodology will be broad-ranging, with a special emphasis on the interplay between between *axiomatic* concerns ($\hookrightarrow$ Economics) and *algorithmic* concerns ($\hookrightarrow$ Computer Science).

Your main resource for this course will be the *Handbook of COMSOC*. Additional literature will get posted on the course website:

`http://www.illc.uva.nl/~ulle/teaching/comsoc/2019/`

Be ready to invest ∼20h/week (lectures, tutorials, readings, homework).

F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A.D. Procaccia (eds), *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

Prerequisites

I expect *mathematical maturity* (working out and writing up proofs), but little in terms of specific *mathematical knowledge:* some basic concepts from combinatorics, probability theory, and logic.

Some background in the following is useful but not strictly required:

- *Game Theory:* We'll often reason about agents being strategic. Prior exposure to game theory helps with this kind of thinking.
- *Complexity Theory:* Required to analyse social choice mechanisms from an algorithmic point of view. For those who need it, there'll be a tutorial to allow you to pick up the basics.
- *Programming:* For one homework assignment some very modest programming skills will be required (in Python). Help is available.

Assessment

Two parts: four pieces of *homework* (75%) and a final *exam* (25%).

Regarding homework:

- Each assignment will be graded on the usual 1–10 scale.
- Homework should be submitted in pairs (via Canvas).
- Collaboration is subject to common-sense rules (see Canvas).
- Regrading within one week only and in exceptional cases only (mapping mistakes to points is subjective, so not up for discussion)

To pass the course, you must get $\geqslant 5.5$ both in the exam and overall.

Resit exam in June (maybe oral exam if small number of candidates).
No resit opportunity for the homework component.

Requirements for Homework Solutions

Most questions will be of the problem-solving sort, requiring:

- a good understanding of the topic to see what the question is
- some creativity to find the solution
- mathematical maturity, to write it up correctly, often as a proof
- good taste, to write it up in a reader-friendly manner

Solutions must be *typed* up professionally (LaTeX strongly preferred).

Of course, solutions should be *correct*. But just as importantly, they should be *short* and *easy to understand*. (This is the advanced level: it's not anymore just about you getting it, it now is about your reader!)

Common mistakes will be discussed during *tutorials*.

Read the Homework Guidelines on Canvas. Ask next time if unclear.

What to Expect at the Exam

The exam will assess your *understanding of the concepts* introduced in the course (so: less focus on mathematical problem solving).

This will be a closed-book exam, but you may bring one piece of paper (A4, double-sided) of *handwritten notes* with you.

Further Activities

In June I plan to offer a project course on Advanced Topics in Computational Social Choice (6EC). More information in May.

Consider attending relevant seminar talks ($\hookrightarrow$ COMSOC Seminar).

Plan for the Rest of Today

We will discuss methods for dividing a single divisible (heterogeneous) resource (the "*cake*") between several agents.

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

- Definition of the problem: how can you cut a cake fairly?
- Presentation of several protocols for cutting a cake
- Complexity analysis: how many cuts do you need?

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*. A.K. Peters, 1998.

U. Endriss. *Lecture Notes on Fair Division*. ILLC, University of Amsterdam, 2009.

A.D. Procaccia. Cake Cutting Algorithms. In F. Brandt et al. (eds.), *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

The Model

The *cake* is the interval $[0, 1]$ of the real numbers from 0 to 1:

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

We need to divide the cake between n *agents* (with $n = 2, 3, 4, 5, \dots$).
A *piece* of cake is a finite union of disjoint subintervals of $[0, 1]$.

Each agent i has a *valuation function* v_i to measure how much she likes any given piece of cake. Assumptions:

- Normalisation: $v_i(\texttt{full\_cake}) = 1$ and $v_i(\texttt{nothing}) = 0$
- Additivity: $v_i(A \cup B) = v_i(A) + v_i(B)$ if A and B don't overlap
- Continuity: small increases in cake $\Rightarrow$ small increases in value

Proportional Fairness

We want to design protocols that are "fair". What does that mean?

One possible definition:

> *An allocation of pieces of cake to agents is proportionally fair, if every agent's subjective value for her piece is at least $\frac{1}{n}$.*

Other options: *envy-freeness* (discussed later), *equitability* (not today)

But more precisely, we want this:

> *A cake-cutting protocol is proportionally fair, if every agent can ensure she gets a piece that she values at at least $\frac{1}{n}$.*

For all proportionally fair protocols we will see, agents can in fact guarantee their fair share by answering all questions *truthfully*.

Cut-and-Choose Protocol

For the case of *two agents*, you all know how to do this:

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

This clearly is *proportionally fair*!

Remark: Truthfully answering the questions ("where is the middle?" and "which one do you prefer?") is the best you can do. But if the cutter knows the valuation of the chooser, she can do even better.

Exercise: *What about three agents? Or more?*

The Steinhaus Protocol

This *proportional* protocol for *three agents* was 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. ✓

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

The Dubins-Spanier Moving-Knife Protocol

Dubins and Spanier (1961) proposed this protocol (for *any* n):

(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 agents, until just one agent is left (who takes the rest).

This is *proportionally fair*! (Of course: right-to-left works as well.)

Exercise 1: *You love strawberries. There is a single large strawberry on the right end of the cake. Do you prefer left-to-right or right-to-left?*

Exercise 2: *How would you program a computer to play for you?*

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

Problem

Each agent has to continuously monitor the knife as it moves over all the real numbers from 0 to 1. For each number, the agent has to evaluate the piece to the left of the knife. This is impossible.

The Robertson-Webb Model

What counts as a "protocol"? — A reasonable protocol should be implementable in terms of just two types of queries:

- $\mathrm{Cut}_i(x, \alpha) \mapsto y$: Ask agent i to cut off a piece of value α, starting from point x (she cuts at point y).
- $\mathrm{Eval}_i(x, y) \mapsto \alpha$: Ask agent i to indicate her value for the piece between points x and y (she answers α).

Now we can count queries and compare the complexity of protocols.

J. Robertson and W. Webb. *Cake-Cutting Algorithms*. A.K. Peters, 1998.

Simulating the Moving-Knife Protocol

We can "discretise" the moving-knife protocol to solve our problem:

(1) Ask each agent to *mark* the cake where she *would* shout "stop". Then *cut* the cake at the *leftmost mark* and give the resulting piece to the agent who made that mark.

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

Formally, the marks are cut-queries. No evaluation-queries needed.

Exercise: *How complex is this (how many queries do we need)?*

Complexity Analysis: Number of Marks

In each round, each participating agent makes one mark. The number of participating agents goes down from n to 2. Thus:

$$n + (n-1) + (n-2) + \cdots + 3 + 2 \;=\; \frac{n \cdot (n+1)}{2} - 1 \;\approx\; \frac{1}{2} \cdot n^2$$

Proof:

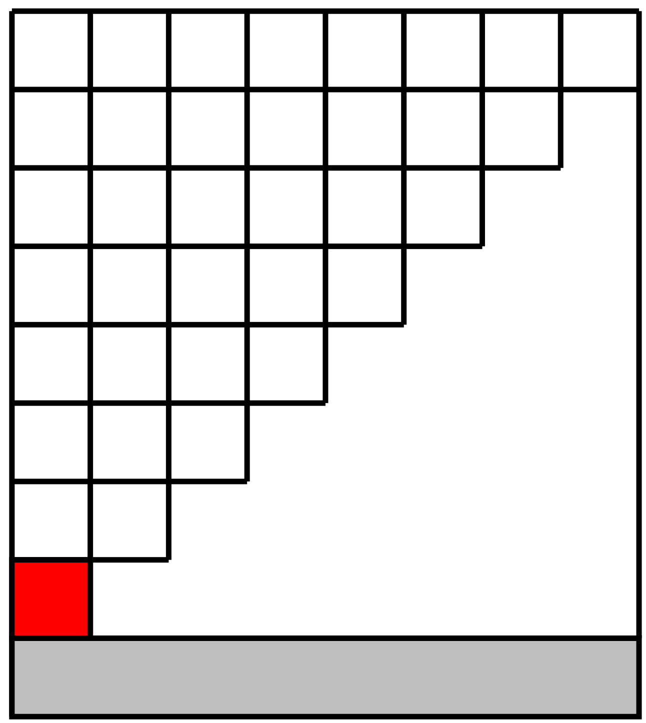

Can we do better?

The Even-Paz Divide-and-Conquer Protocol

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

(1) Ask each agent to put a *mark* on the cake.

(2) *Cut* the cake at the $\lfloor \frac{n}{2} \rfloor$*th mark* (counting from the left).

Associate the agents who made the *leftmost* $\lfloor \frac{n}{2} \rfloor$ *marks* with the *lefthand part*, and the *remaining agents* with the *righthand part*.

(3) *Repeat* for each group, until only one agent is left.

This also is *proportionally fair*! Again, we only require cut-queries.

<u>Exercise:</u> *How complex is this?*

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

Complexity Analysis: Number of Marks

In each round, every agent makes one mark. Thus: n *marks per round*

But how many rounds?

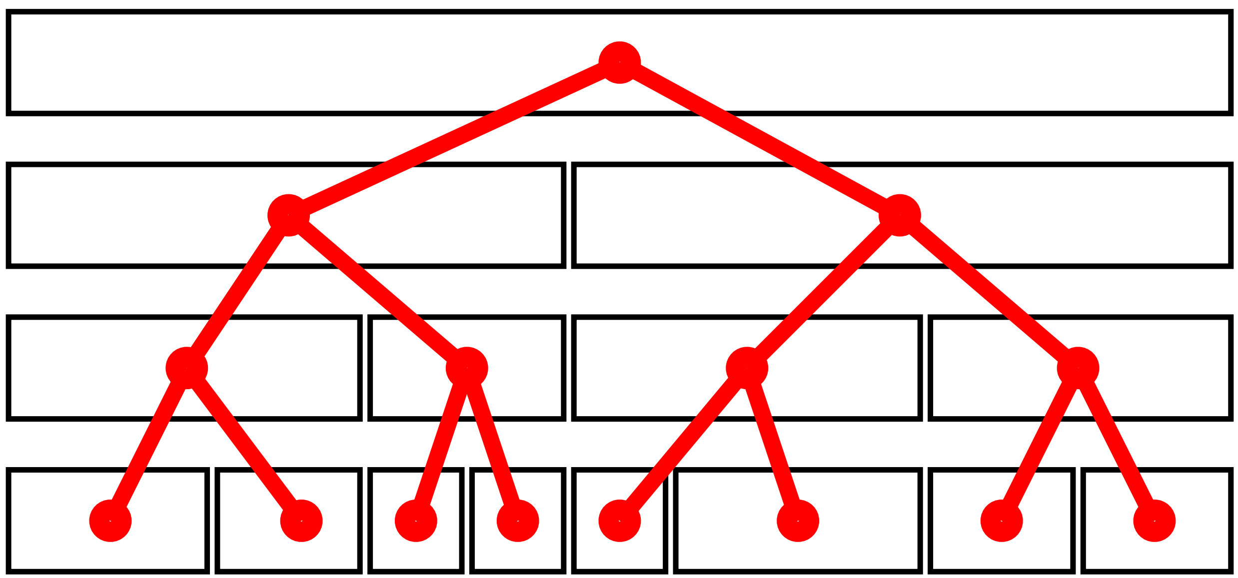

rounds $=$ number of times you can divide n by 2 before hitting $\leqslant 1$

$\approx \log_2 n$ (example: $\log_2 8 = 3$)

Thus: number of *marks* required $\approx n \cdot \log_2 n$

Comparison and Limitations

Recall: *simulated moving-knife* requires around $\frac{1}{2} \cdot n^2$ marks and *divide-and-conquer* requires around $n \cdot \log_2 n$ marks.

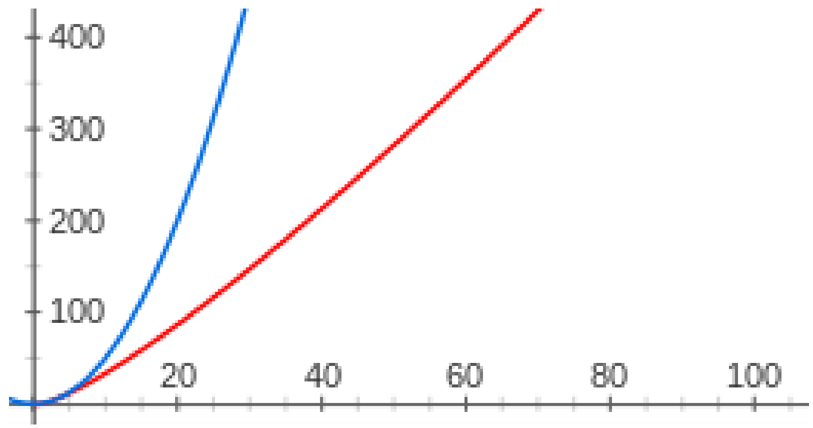

So: divide-and-conquer is much better (for large n, complexity-wise).

And in fact divide-and-conquer is the best you can do:

Theorem 1 (Edmonds and Pruhs, 2006) *Any proportionally fair protocol requires $\Omega(n \log n)$ queries in the Robertson-Webb model.*

J. Edmonds and K. Pruhs. Cake cutting really is not a piece of cake. SODA-2006.

Envy

Proportional fairness is but one formalisation of "fairness":

> *A cake-cutting protocol is called envy-free, if every agent can ensure that she will receive a subjectively largest piece.*

Connections between these two notions of fairness:

- Observe that for $n = 2$ *agents*, we have:

$$\textit{envy-freeness} \Longleftrightarrow \textit{proportional fairness}$$

- But for $n \geqslant 3$ *agents*, we only have:

$$\textit{envy-freeness} \Longrightarrow \textit{proportional fairness}$$

Indeed, of our protocols only *cut-and-choose* guarantees envy-freeness.

Exercise: *Give an example where divide-and-conquer violates EF.*

Four Simultaneously Moving Knives

Stromquist (1980) found this *envy-free* protocol for *three agents:*

- A referee slowly moves a knife across the cake, from left to right (supposed to eventually cut somewhere around the $\frac{1}{3}$ mark).
- At the same time, each agent is moving her own knife so that it *would* cut the righthand piece in half (w.r.t. 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.

The Selfridge-Conway Protocol

The first discrete protocol achieving *envy-freeness* for *three agents* was 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.

Beyond Three Agents

Regarding envy-free protocols, we saw:

- For $n = 2$ the problem is easy: cut-and-choose does the job.
- For $n = 3$ we saw two protocols, each with some drawbacks.

For *arbitrary* n, Brams and Taylor (1995) propose a protocol requiring an *unbounded* number of queries in the R-W model (so the number of queries required doesn't just depend on n but also on the valuations).

Achieving envy-freeness really is harder than achieving proportionality:

Theorem 2 (Procaccia, 2009) *Any envy-free protocol requires $\Omega(n^2)$ queries in the Robertson-Webb model.*

Recall: Proportionality only requires $O(n \log n)$ queries.

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

A.D. Procaccia. Thou Shalt Covet Thy Neighbor's Cake. IJCAI-2009.

Recent Advances

The problem of envy-free cake cutting was studied since the 1940s. Best result until 1995 for $n = 4$ has been the unbounded protocol of Brams and Taylor (contrast this with the lower bound of $\Omega(n^2)$).

Breakthrough results by Aziz and Mackenzie (2016):

- a protocol for *four agents* requiring at most 584 queries in the Robertson-Webb model
- a protocol for n *agents* requiring at most $n^{n^{n^{n^{n^{n}}}}}$ such queries.

Note that the latter is considerably better than the earlier bound of ∞.

H. Aziz and S. Mackenzie. A Discrete and Bounded Envy-free Cake Cutting Protocol for Four Agents. STOC-2016.

H. Aziz and S. Mackenzie. A Discrete and Bounded Envy-free Cake Cutting Protocol for Any Number of Agents. FOCS-2016.

Summary

This has been an introduction to cake cutting. We saw:

- usable protocols for guaranteeing *proportional fairness*
- severe limitations for protocols guaranteeing *envy-freeness*

In terms of methodology, we saw:

- how to define *fairness* in terms of guarantees for the agents
- how to formalise the concept of "*protocol*" (Robertson-Webb)
- how to analyse the *complexity* of a cake-cutting protocol

Read Chapter 1 of the Handbook to get an understanding of the nature and history of the field of computational social choice.

What next? Fair allocation of several indivisible goods, focusing first on the axiomatic and then on the algorithmic perspective.