Computational Social Choice: Autumn 2012

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

Plan for Today

We shall continue our study of *fair allocation of indivisible goods*.

But instead of devising algorithms for computing a socially optimal allocation given agent preferences, we now want agents to be able to do this in a *distributed* manner. Main question addressed today:

- Under what circumstances will a system in which agents negotiate autonomously and contract *local* deals *converge* to a state considered optimal from a *global* point of view?

Notation and Terminology

We refine our formal framework for the allocation of indivisible goods to be able to model deals and monetary side payments:

- 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 $\mathcal{N}$.
- A *deal* $\delta = (A, A')$ is a pair of allocations (before/after).

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

  Example: $p(i) = 5$ and $p(j) = -5$ means that agent i *pays* €5, while agent j *receives* €5.
- Each agent $i \in \mathcal{N}$ has got a *valuation function* $v_i : 2^{\mathcal{G}} \to \mathbb{R}$.

  If agent i receives bundle B and the sum of her payments so far is x, then her *utility* is $u_i(B, x) = v_i(B) - x$ ("quasi-linear utility").

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

The general research agenda is outlined in the paper cited below.

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 is *myopic*, i.e., does not plan ahead) will only accept deals that improve her individual welfare:

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

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

The Global/Social Perspective

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

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

Observe that there is no need to include the agents' monetary balances into this definition, because they would 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 $\mathcal{N} = \{ann, bob\}$ and $\mathcal{G} = \{chair, table\}$ and suppose our agents use the following utility functions:

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

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

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

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

Convergence

The good news:

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

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

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

So why does this work?

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

- **Lemma 1** *A deal $\delta = (A, A')$ is individually rational if and only if $\mathrm{SW_{util}}(A) < \mathrm{SW_{util}}(A')$.*

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

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

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

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

Necessity of Multilateral Negotiation

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

Theorem 2 *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.

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

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, there are no synergies between items; you can get the value of a bundle by adding up the values of its elements (+ a constant for $\emptyset$).

- Negotiation in modular domains *is* feasible (the proof is easy):

Proposition 1 *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*: for no larger class can we still get the same convergence property.

Y. Chevaleyre, U. Endriss, and N. Maudet. Simple Negotiation Schemes for Agents with Simple Preferences. *JAAMAS*, 20(2):234–259, 2010.

Comparing Negotiation Policies

While we know from Proposition 1 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:

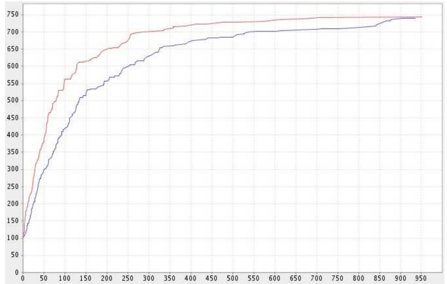

The graph shows how utilitarian social welfare (y-axis) develops as agents attempt to contract more an more deals (x-axis) amongst themselves. Graph generated using the MADRAS platform of Buisman *et al.* (2007).

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

Path Length

How many deals do we require to reach the optimal allocations?

Here are some simple results on *upper bounds* on the path length (here "socially optimal" = maximal utilitarian social welfare):

Proposition 2 *The shortest path of IR deals to a socially optimal allocation consists of at most 1 deal.*

Proposition 3 *The longest path of IR deals to a socially optimal allocation consists of at most* $|\mathcal{N}|^{|\mathcal{G}|} - 1$ *deals.*

Proposition 4 *In modular domains, the shortest path of IR 1-deals to a socially optimal allocation consists of at most* $|\mathcal{G}|$ *1-deals.*

Proposition 5 *In modular domains, the longest path of IR 1-deals to a socially optimal allocation consists of at most* $|\mathcal{G}| \cdot (|\mathcal{N}| - 1)$ *1-deals.*

U. Endriss and N. Maudet. On the Communication Complexity of Multilateral Trading: Extended Report. *JAAMAS*, 11(1):91–107, 2005.

Fairness

So far, the results have been all about *efficiency*. In fact, we have seen that the local criterion of *individual rationality* perfectly fits our global efficiency criterion of *maximal utilitarian social welfare* (⤳ Lemma 1).

If we take the individual agent behaviour as a given (we do, today), then we cannot possibly hope to always achieve *fair* outcomes.

The remainder of this lecture is about exploring how far we can get nevertheless, for one particular interpretation of fairness ...

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

Envy-freeness

Recall that an allocation is called *envy-free* if nobody wants to change bundle with any of the others:

Definition 1 (EF allocations) *An allocation* A *is envy-free iff* $v_i(A(i)) \geqslant v_i(A(j))$ *for all agents* $i, j \in \mathcal{N}$.

If we require all goods to be allocated, then envy-free allocations may not always exist. Example: 2 agents, 1 good, liked by both agents

There has been some work on the computational complexity of checking whether a given scenario admits an envy-free solution: it's NP-hard in the simplest cases. So, it's pretty difficult.

But above definition does not take *money* into account ...

S. Bouveret and J. Lang. Efficiency and Envy-freeness in Fair Division of Indivisible Goods: Logical Representation and Complexity. *JAIR*, 32:525–564, 2008.

Envy-freeness in the Presence of Money

We refine our negotiation framework as follows ...

- Associate each allocation A with a *balance function* $\pi : \mathcal{N} \to \mathbb{R}$, mapping agents to the sum of payments they've made so far.
- A *state* (A, π) is a pair of an allocation and a payment balance.
- Each agent $i \in \mathcal{N}$ has got a (quasi-linear) *utility function* $u_i : 2^{\mathcal{G}} \times \mathbb{R} \to \mathbb{R}$, defined as follows: $u_i(B, x) = v_i(B) - x$.

... and adapt the definition of envy-freeness:

Definition 2 (EF states) *A state* (A, π) *is envy-free iff* $u_i(A(i), \pi(i)) \geqslant u_i(A(j), \pi(j))$ *for all agents* $i, j \in \mathcal{N}$.

Note: In the presence of money, *Pareto efficiency* (defined over states, not allocations) = *maximal utilitarian social welfare*.

An *efficient envy-free* (EEF) state is an EF state that is Pareto efficient. Does such a state exist under all circumstances?

Existence of EEF States

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

Theorem 3 (Alkan et al., 1991) *If all valuations are supermodular, then an EEF state always exists.*

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

$$v(B_1 \cup B_2) \geqslant v(B_1) + v(B_2) - v(B_1 \cap B_2)$$

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

A. Alkan, G. Demange and D. Gale. Fair Allocation of Indivisible Goods and Criteria of Justice. *Econometrica*, 59(4):1023–1039, 1991.

Proof of Theorem 3

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:

$$\pi^*(i) \quad := \quad v_i(A^*) - \mathrm{SW}_{\mathrm{util}}(A^*)/n$$

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

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

$$v_i(A^*(i)) + v_j(A^*(j)) \geqslant v_i(A^*(i) \cup A^*(j))$$

Now apply the supermodularity condition ... and rewrite:

$$v_i(A^*(i)) + v_j(A^*(j)) \geqslant v_i(A^*(i)) + v_i(A^*(j))$$

$$v_i(A^*(i)) - [v_i(A^*) - \mathrm{SW}_{\mathrm{util}}(A^*)/n] \geqslant v_i(A^*(j)) - [v_j(A^*) - \mathrm{SW}_{\mathrm{util}}(A^*)/n]$$

$$u_i(A^*(i), \pi^*(i)) \geqslant u_i(A^*(j), \pi^*(j))$$

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 good

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

Suppose agent 1 owns g to begin with.

The efficient allocation would be where agent 2 owns g.

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

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

Compromise: We shall enforce an *initial equitability payment* $\pi_0(i) := v_i(A_0) - \mathrm{SW}_{\mathrm{util}}(A_0)/n$ before beginning negotiation ...

Globally Uniform Payments

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!):

Definition 3 (GUPF) *Let $\delta = (A, A')$ be an IR deal. The payments as given by the globally uniform payment function are defined as:*

$$p(i) := [v_i(A') - v_i(A)] - [\mathrm{SW}_{\mathrm{util}}(A') - \mathrm{SW}_{\mathrm{util}}(A)]/n$$

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

Convergence in Supermodular Domains

We obtain a surprising result:

Theorem 4 *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):

$$\pi(i) = v_i(A) - \mathrm{SW}_{\mathrm{util}}(A)/n \quad (*)$$

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

$$\begin{aligned} \pi'(i) &= \pi(i) + [v_i(A') - v_i(A)] - [\mathrm{SW}_{\mathrm{util}}(A') - \mathrm{SW}_{\mathrm{util}}(A)]/n \\ &= v_i(A') - \mathrm{SW}_{\mathrm{util}}(A')/n \quad \rightsquigarrow (*) \text{ holds by induction} \end{aligned}$$

Theorem 1 shows that the system must *converge* to an *efficient* allocation A^* (whatever the payment function). Then the proof of Theorem 3 (existence of EEF states, by constructing precisely π above) demonstrates that condition $(*)$ implies that (A^*, π^*) must be an *EEF* state. ✓

Social Networks

Imagine our agents are nodes in a social network and only deals in fully connected groups are possible.

- Clearly, our first convergence result (Theorem 1 will break, as there will now be some deals δ that are not allowed anymore (cf. Theorem 2). No obvious way of fixing this.
- But: There is a very natural adaptation of envy-freeness to this setting, namely that agents only envy those agents they are connected. Theorem 4 now generalises beautifully.

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

Summary

We have discussed a framework for the allocation of *indivisible goods* where agents agree on deals to exchange goods in a *distributed* way.

Results concerning *convergence* to a socially optimal allocation :

- About maximising *utilitarian social welfare:*
  - Convergence via IR deals can be guaranteed
  - But might need structurally complex deals
  - Simple deals might work for restricted valuations (modularity)
  - Path to optimum can always be short, but might be very long
- About *envy-freenes:*
  - Convergence via IR deals under restrictive assumptions
  - Generalisation to fair division on social networks
- Other fairness criteria have also been studied

What next?

We will complete the section of fair division with a lecture on *cake-cutting algorithms*, where the “cake” stands for a single *divisible* resource to be divided amongst the agnets.