Negotiating Socially Optimal Allocations of Resources

Ulle Endriss

Institute for Logic, Language and Computation

University of Amsterdam

[joint work with Yann Chevaleyre, Sylvia Estivie, Jérôme Lang, Nicolas Maudet, Fariba Sadri, and Francesca Toni]

Introduction

- Multiagent systems may be thought of as *"societies of agents"*.
- Agents *negotiate* deals to exchange resources to benefit either themselves or society as a whole.
- Agents may use very simple rationality criteria to decide what deals to accept, but interaction patterns may be complex (*multilateral* deals).

Talk Overview

- Definition of the basic negotiation framework
- Fundamental results linking individual interests and social welfare
- Efficient negotiation in restricted domains
- Complexity of negotiating socially optimal allocations
- Alternative social welfare measures
- Conclusions

Definition of the Basic Negotiation Framework

Resource Allocation by Negotiation

- Finite set of *agents* $\mathcal{A}$ and finite set of indivisible *resources* $\mathcal{R}$.
- An *allocation* A is a partitioning of $\mathcal{R}$ amongst the agents in $\mathcal{A}$.
  Example: $A(i) = \{r_5, r_7\}$ — agent i owns resources r_5 and r_7
- Every agent $i \in \mathcal{A}$ has got a *utility function* $u_i : 2^{\mathcal{R}} \to \mathbb{R}$.
  Example: $u_i(A) = u_i(A(i)) = 577.8$ — agent i is pretty happy
- Agents may engage in negotiation to exchange resources in order to benefit either themselves or society as a whole.
- 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 utility. 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.

Individual Rationality

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* iff there exists a payment function p such that $u_i(A') - u_i(A) > p(i)$ 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 utility (or money) that strictly outweighs a possible loss in money (or utility).

Utilitarian Social Welfare

The *social welfare* associated with an allocation of resources A is defined as follows:

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

This is the so-called *utilitarian* definition of social welfare, which measures the "sum of all pleasures" (Jeremy Bentham, ∼1820).

▶ Observe that maximising this function amounts to maximising the *average utility* enjoyed by agents in the system.

Example

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

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

Furthermore, suppose the initial allocation of resources 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* resource 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* $\{chair, table\}$.

Fundamental Results

Linking the Local and the Global Perspectives

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

Lemma 1 (Rationality and social welfare) *A deal $\delta = (A, A')$ with side payments is individually rational iff $sw(A) < sw(A')$.*

Proof. "$\Rightarrow$": Rationality means that overall utility gains outweigh overall payments (which are $= 0$). "$\Leftarrow$": Using side payments, the social surplus can be divided amongst all deal participants. □

▶ We can now prove a first result on negotiation processes:

Lemma 2 (Termination) *There can be no infinite sequence of individually rational deals, i.e. 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 3 (Sandholm, 1998) *Any sequence of individually rational deals will eventually result in an allocation with maximal social welfare.*

▶ Agents can act *locally* and need not be aware of the global picture (convergence towards a global optimum is guaranteed by the theorem).

T. Sandholm. *Contract types for satisficing task allocation: I Theoretical results*. AAAI Spring Symposium 1998.

U. Endriss, N. Maudet, F. Sadri and F. Toni. *On optimal outcomes of negotiations over resources*. AAMAS-2003.

Multilateral Negotiation

Optimal outcomes can only be guaranteed if the negotiation protocol allows for deals involving *any number of agents* and *resources:*

Theorem 4 (Necessity of complex deals) *Any deal $\delta = (A, A')$ may be necessary, i.e. there are utility functions and an initial allocation such that any sequence of individually rational deals leading to an allocation with maximal social welfare would have to include δ (unless δ is "independently decomposable").*

The proof involves the systematic definition of utility 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 concerning distinct sets of agents.

Efficient Negotiation in Restricted Domains

Negotiation in Restricted Domains

Most work on negotiation in multiagent systems is concerned with bilateral negotiation or auctions. $\leadsto$ Multilateral negotiation is difficult!

Maybe we can guarantee convergence to a socially optimal allocation for structurally simpler types of deals if we restrict the range of utility functions that agents can use? First, two negative results:

- Theorem 4 continues to hold even when all agents are required to use *monotonic* utility functions. $[R_1 \subseteq R_2 \Rightarrow u_i(R_1) \leq u_i(R_2)]$
- Theorem 4 continues to hold even when all agents are required to use *dichotomous* utility functions. $[u_i(R) = 0 \vee u_i(R) = 1]$

Modular Domains

A utility function u_i is called *modular* iff it satisfies the following condition for all bundles $R_1, R_2 \subseteq \mathcal{R}$:

$$u_i(R_1 \cup R_2) \quad = \quad u_i(R_1) + u_i(R_2) - u_i(R_1 \cap R_2)$$

That is, in a modular domain there are no synergies between items; you can get the utility of a bundle by adding up the utilities of the items in that bundle.

▶ Negotiation in modular domains *is* feasible:

Theorem 5 (Modular domains) *If all utility functions are modular, then individually rational 1-deals (involving just one resource) suffice to guarantee outcomes with maximal social welfare.*

U. Endriss, N. Maudet, F. Sadri and F. Toni. *On optimal outcomes of negotiations over resources*. AAMAS-2003.

Sufficiency, Necessity, Maximality

- Theorem 5 says that the class of modular utility functions is *sufficient* for successful 1-deal negotiation.
- Is it also *necessary*?
  Answer: No. It's easy to construct examples.
- Is there *any* class of functions that is *sufficient* and *necessary*?
  Answer: No. Seems surprising at first, but for a proof it suffices to find two sufficient classes the union of which is not sufficient.
- As there can be no unique class of utility functions characterising all situations where 1-deal negotiation works, we have looked for *maximal* classes . . .

Maximality of Modular Utilities

We say that a class of utility functions $\mathcal{F}$ *permits 1-deal negotiation* iff any sequence of individually rational 1-deals will converge to a socially optimal allocation whenever all utility functions belong to $\mathcal{F}$.

Another surprising result:

Theorem 6 (Maximality) *Let $\mathcal{M}$ be the class of modular utility functions. Then for any class of utility functions $\mathcal{F}$ such that $\mathcal{M} \subset \mathcal{F}$, $\mathcal{F}$ does not permit 1-deal negotiation.*

▶ Are there other (interesting) classes of functions that are maximal?

Y. Chevaleyre, U. Endriss and N. Maudet. *On maximal classes of utility functions for efficient one-to-one negotiation*. IJCAI-2005.

Complexity Issues

Complexity of Maximising Social Welfare

Theorem 7 (Complexity) *Let $K \in \mathbb{Z}$. For a given scenario, deciding whether there exists an allocation A with $sw(A) > K$ is NP-complete.*

This is essentially a well-known result that is closely related to the winner determination problem in combinatorial auctions, but there are some interesting variations depending on how we represent utilities (bundle enumeration, k-additive form, straight-line programs, . . .).

M. H. Rothkopf, A. Pekeč and R. M. Harstad. *Computationally manageable combinational auctions*. Management Science, 44(8):1131–1147, 1998.

P. E. Dunne, M. Wooldridge and M. Laurence. *The complexity of contract negotiation*. Artificial Intelligence, 164(1–2):23–46, 2005.

Y. Chevaleyre, U. Endriss, S. Estivie and N. Maudet. *Multiagent resource allocation with k-additive utility functions*. DIMACS-LAMSADE Workshop 2004.

Proof of NP-completeness

NP-membership: We can check $sw(A) > K$ in polynomial time. ✓

NP-hardness: The following problem is known to be NP-complete:

WEIGHTED SET PACKING

Instance: Collection $\mathcal{C}$ of finite sets with positive weights.

Solution: Collection of disjoint sets $\mathcal{C}' \subseteq \mathcal{C}$.

Question: Does the sum of weights of the sets in $\mathcal{C}'$ exceed K?

This can be reduced to our problem as follows:

- For every set R in $\mathcal{C}$ with weight x, introduce an agent i and define $u_i(R) = x$ and $u_i(R') = 0$ for all bundles $R' \neq R$.
- "Free disposal": introduce an additional agent i^* with $u_{i^*} \equiv 0$.

Now any allocation A with $sw(A) > K$ corresponds to a set packing $\mathcal{C}'$ with a sum of weights exceeding K. Hence, our problem is at least as hard as WEIGHTED SET PACKING. ✓

Communication Complexity

- The NP-completeness result concerns the *computational* complexity of an *abstract* problem: finding a socially optimal allocation *somehow* (not necessarily by negotiation).
- What we are really interested in is the complexity of negotiation processes in our multilateral trading framework.
- We therefore consider also the *communication complexity* of negotiating socially optimal allocations of resources, i.e. we focus on the length of negotiation processes and the amount of information exchanged, rather than just on computational aspects.

U. Endriss and N. Maudet. *On the communication complexity of multilateral trading*. Journal of Autonomous Agents and Multiagent Systems, 11(1):91–107, 2005.

Aspects of Complexity

(1) How many *deals* are required to reach an optimal allocation?

- communication complexity as number of individual deals
- technical results to follow

(2) How many *dialogue moves* are required to agree on one such deal?

- affects communication complexity as number of dialogue moves

(3) How expressive a *communication language* do we require?

- Minimum requirements: performatives *propose*, *accept*, *reject* + content language to specify multilateral deals
- affects communication complexity as number of bits exchanged

(4) How complex is the *reasoning* task faced by an agent when deciding on its next dialogue move?

- computational complexity (local rather than global view)

Number of Deals

We have two results on *upper bounds* pertaining to the first variant of our negotiation framework (with side payments, general utility functions, and aiming at maximising utilitarian social welfare):

Theorem 8 (Shortest path) *A single rational deal is sufficient to reach an allocation with maximal social welfare.*

Proof. Use Lemma 1 [$\delta = (A, A')$ rational iff $sw(A) < sw(A')$]. □

Theorem 9 (Longest path) *A sequence of rational deals can consist of up to $|\mathcal{A}|^{|\mathcal{R}|} - 1$ deals, but not more.*

Proof. No allocation can be visited twice (same lemma) and there are $|\mathcal{A}|^{|\mathcal{R}|}$ distinct allocations ⇒ upper bound follows ✓
To show that the upper bound is *tight*, we need to show that it is possible that all allocations have distinct social welfare . . . □

Path Length in Modular Domains

If all agents are using modular utility functions and only negotiate 1-deals, then we obtain the following bounds:

- *Shortest path:* $\leq |\mathcal{R}|$
- *Longest path:* $\leq |\mathcal{R}| \cdot (|\mathcal{A}| - 1)$

Alternative Social Welfare Measures

Negotiation without Side Payments

- Problem: Agents may require *unlimited amounts of money* to get through a negotiation process.
- Without side payments, however, rational negotiation cannot guarantee outcomes with maximal social welfare.
  Example: *Would you give me your bike just because I value it more highly than you do? . . . note that this would be socially beneficial!*
- It is possible to show that *cooperatively rational* deals (only one agent requires a *strictly* positive payoff) without side payments are sufficient to negotiate *Pareto optimal* allocations (and multilateral deals are again necessary).
- Similar types of results as before on *1-deal negotiation* and on *communication complexity* . . .

Egalitarian Agent Societies

- The utilitarian sw is not the only *collective utility function* . . .
- The *egalitarian* collective utility function sw_e, for instance, measures social welfare as follows:

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

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

- We have defined a local rationality criterion ("equitable deals") for agents operating in egalitarian systems and proved convergence and necessity theorems similar to those we have seen earlier.

U. Endriss, N. Maudet, F. Sadri and F. Toni. *Resource allocation in egalitarian agent societies*. MFI-2003.

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.

Notions of Social Welfare

- *Utilitarian:* sum of utilities $sw_u(A) = \sum_{i \in \mathcal{A}} u_i(A)$
- *Nash product:* product of utilities $sw_N(A) = \prod_{i \in \mathcal{A}} u_i(A)$
- *Egalitarian:* utility of the weakest $sw_e(A) = min\{u_i(A) \mid i \in \mathcal{A}\}$
- *Elitist:* utility of the strongest $sw_{el}(A) = max\{u_i(A) \mid i \in \mathcal{A}\}$
- *Pareto optimality:* no other allocation is better for some agents without being worse for any of the others
- *Lorenz optimality:* the sum of utilities of the k weakest agents cannot be maintained for all and increased for some $k \leq |\mathcal{A}|$
- *Envy-freeness:* no agent would rather have the bundle allocated to one of the other agents $u_i(A(i)) \geq u_i(A(j))$
  - envy-free allocations are not always possible
  - could search for *envy-reducing* deals (for instance, with respect to the number of envious agents or the average degree of envy)

Welfare Engineering

- Choice (and possibly design) of *social welfare orderings* that are appropriate for specific agent-based applications.
  - Example: The *elitist* collective utility function sw_{el} seems unethical for human society, but may be appropriate for a distributed application where each agent gets the same task.
  - Slogan: "welfare economics for *artificial* agent societies"
- Design of suitable *rationality criteria* for agents participating in negotiation in view of different notions of social welfare.
  - Example: To achieve *Lorenz optimal* allocations in *0-1 domains without money*, ask agents to negotiate *cooperatively rational* or *inequality-reducing* deals over *one resource* at a time.
  - Slogan: "*inverse* welfare economics" ($\rightsquigarrow$ mechanism design)

U. Endriss and N. Maudet. *Welfare engineering in multiagent systems*. ESAW-2003.

Criteria for Social Welfare Choice

We have tried to identify criteria that determine what social welfare ordering is appropriate for which application (work in progress):

- What does the income of the system provider depend on?
  - *Utility-dependent* ("tax on gain") $\rightsquigarrow$ utilitarian
  - *Membership-dependent* ("joining fee") $\rightsquigarrow$ "fair" approach
  - *Transaction-dependent* ("pay as you go") $\rightsquigarrow$ not clear (but note the connections to communication complexity)
- Can agents join or leave the society *during* negotiation?
  Yes: review definitions (e.g. utilitarian welfare as average utility)
- Can agents participate in *more than one* negotiation?
  Yes: strong point for fair approaches (egalitarian, envy-reducing)

Y. Chevaleyre, U. Endriss, S. Estivie and N. Maudet. *Welfare engineering in practice: On the variety of multiagent resource allocation problems*. ESAW-2004.

Conclusions

- Negotiation is an exciting and fruitful area of research:
  - Knowledge transfer from economics to computer science and AI
  - Application of computational tools to problems in economics
- Many open problems and scope for new directions of research!
- For more information on the field in general, have a look at the MARA (Multiagent Resource Allocation) website:

  http://www.illc.uva.nl/~ulle/MARA/

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 (to appear).