Negotiation Amongst Agents
Principles and Techniques

Nicolas Maudet
nicolas.maudet@lip6.fr

Université Pierre et Marie Curie

2015 ESSENCE Summer School, Edinburgh

Outline of the Course

- Part I: Bilateral Negotiation (axiomatics, protocols, heuristics)
- Part II: Multilateral Negotiation (contract-based negotiation, networks, etc.)
- Part III: Negotiating the Meaning (?) (naming games, ontology alignment)

General References

Negotiation is a huge topic, studied in many fields for many years.

Raiffa. *The art and science of negotiation*. 1982.

Some general AI/MAS books, notes, with nice chapters on negotiation:

M. Wooldridge. *An Introduction to Multiagent Systems*. MIT Press-2004.

J. Vidal. *Fundamentals of Multiagent Systems*. 2007.

And these are three books dedicated to the subject:

J. Rosenschein & Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Agents*. 1994.

S. Kraus. *Strategic Negotiation in Multiagent Environments*. 2001.

S. Fatima, S. Kraus, and M. Wooldridge.. *Principles of Automated Negotiation. Cambridge University Press*. 2014.

Approaches to Negotiation

It is also common to find the following distinction:

- Game-theoretic—use of mathematical tools, as developed in game-theory, to analyze strategical interaction. Provable properties, strong assumptions.
- Heuristic-based—design of good strategies in practice, sometimes in specific domains of negotiation. More realistic assumptions, more difficult to guarantee properties.
- Argument-based—allows the exchange of arguments during negotiation.

Outline of the lecture

1. Bilateral Negotiation
   - The setting
   - What are good outcomes?
   - Axiomatics of negotiation
   - Protocols and Game-theoretical analysis
   - Heuristics
2. Multilateral Negotiation
   - A Mediated Protocol
   - Contract-Based Negotiation
   - Outcomes on Networks
3. Negotiation on Meaning
   - Naming Games
   - Negotiated Ontology Alignment

Bilateral Negotiation: The Setting

We first describe the outcome set $\mathcal{O}$. This set may have different characteristics.
Compare the following scenario:

1. we must decide on the next location for the summer school.

$$o_1 = \langle \text{bali} \rangle$$

2. we must divide a chocolate-vanilla cake
division of a continuous resource.

$$o_1 = \langle 1/3, 2/3 \rangle$$

3. there are 4 candies, we must decide on a complete allocation of resources to children.
allocation of indivisible resources.

$$o_1 = \langle \{c_1, c_4\}, \{c_2, c_3\} \rangle$$

The Setting: The Outcome Set

The outcome set may be very large, even in the discrete case:

- allocations of indivisible resources
  g goods, so $|\mathcal{O}| = |\mathcal{A}|^g$ outcomes
- choice in a multi-issue domain
  p issues, with D_i the domain of the issue i, so $|\mathcal{O}| = \Pi_i |D_i|$

Example: Choosing the next holiday package:

- $D_d = \{\text{1week}, \text{2weeks}\}$
- $D_c = \{\text{bali}, \text{lisboa}, \text{moscow}, \text{dakar}\}$
- $D_h = \{\text{pension}, \text{hotel1}, \text{hotel2}, \text{hotel3}, \text{hotel4}\}$
- $D_t = \{\text{plane}, \text{bike}, \text{car}\}$

This yields $2 \times 4 \times 5 \times 3 = 120$ outcomes.

The Setting: Agents' Preferences

Next we have discuss how agents express their preferences.

A preference structure represents an agent's preference over the set of outcomes $\mathcal{O}$. There are different types of preference structures:

Roughly speaking, preferences can be ordinal or cardinal.

- an ordinal preference structure is a binary relation over the outcomes $\mathcal{O}$, which is reflexive, transitive (and often complete).

$$o_1 \succeq o_2$$

"o_1 *is at least as good as* o_2"

The Setting: Agents' Preferences

Next we have discuss how agents express their preferences.

A preference structure represents an agent's preference over the set of outcomes $\mathcal{O}$. There are different types of preference structures:

Roughly speaking, preferences can be ordinal or cardinal.

- an ordinal preference structure is a binary relation over the outcomes $\mathcal{O}$, which is reflexive, transitive (and often complete).

$$o_1 \succeq o_2$$

"o_1 *is at least as good as* o_2"

- a cardinal preference structure is expressed as a valuation function

$$v : \mathcal{O} \mapsto Val$$

where *Val* can be a totally ordered scale of qualitative values ("very good", "good", ...), or some quantitative values.

The Setting: Agents' Preferences

Very often quantitative values are used. But beware of the exact interpretation of this "value".
Following (Luce and Raiffa, 1957), make sure to distinguish:

1. values are in utility terms, no interpersonal comparison of utility are permitted, and no side payments are allowed
2. values are in utility terms, interpersonal comparison is meaningful, and no side payment are allowed
3. values are in monetary terms, utility is linear in money, interpersonal comparisons are meaningful, and monetary side payments are allowed.

From now, we denote by $u_i(o)$ the utility of agent i for the outcome o.

Luce & Raiffa. *Games and Decisions*. 1957.

Negotiation Domains: Resource Allocation

In the context where agents seek to agree on an allocation of indivisible resources (or tasks), the following distinction is useful:

- task-oriented domains—the utility function is common to all agents (and commonly sub-additive), and agents are only concerned with the tasks it gets
- state-oriented domains—the utility function is common to all agents, but agents can value the state in general (not only its bundle of resources)
- worth-oriented domains—the utility function may be different for the different agents

Rosenschein & Zlotkin. *Rules of Encounter.* 1994.

Outcome Sets: example

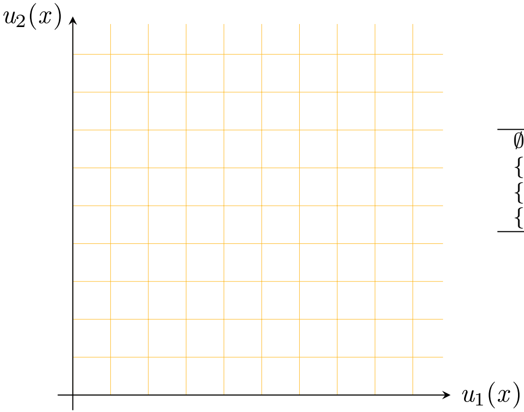

	u_1	u_2
$\emptyset$	0	0
$\{r_1\}$	1	3
$\{r_2\}$	3	3
$\{r_1, r_2\}$	7	8

What are the outcomes? Can you place them on this figure?

Outcome Sets: example

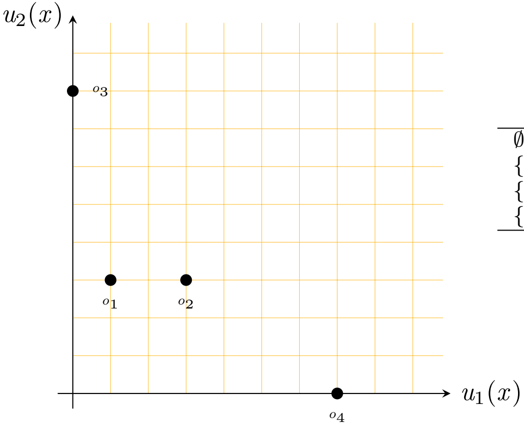

	u_1	u_2
$\emptyset$	0	0
$\{r_1\}$	1	3
$\{r_2\}$	3	3
$\{r_1, r_2\}$	7	8

What are the outcomes? Can you place them on this figure?

Convex Outcome Sets

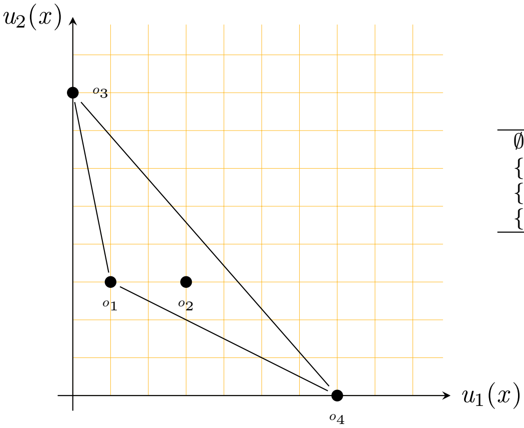

	u_1	u_2
$\emptyset$	0	0
$\{r_1\}$	1	3
$\{r_2\}$	3	3
$\{r_1, r_2\}$	7	8

Randomization between possible outcomes defines a new outcome. For instance, any point on the segment $o_3 - o_4$ is a randomized outcome. But then the outcome set becomes a convex region.

Some (important) assumptions and remarks

The following remarks are useful:

1. ordinal preferences do not allow interpersonal comparison
2. ordinal preferences cannot represent intensities, cardinal preferences can
3. ordinal preferences can handle incomparabilities, but cardinal preferences cannot
4. explicit representation of cardinal and ordinal preferences require space complexity of $O(|\mathcal{O}|)$ and $O(|\mathcal{O}|^2)$

In the following, we make some assumptions:

1. preferences of agents are common knowledge among all agents (we come back to this later)
2. agents can provide explicit representation of their preferences (more compact way of representing preferences are possible)

Good outcomes: Pareto-optimality

- An outcome o_1 Pareto-dominates another outcome o_2 if o_1 is at least as good as o_2 for all agents, and strictly better for at least one.
- An outcome is Pareto-optimal if no other outcome dominates it.

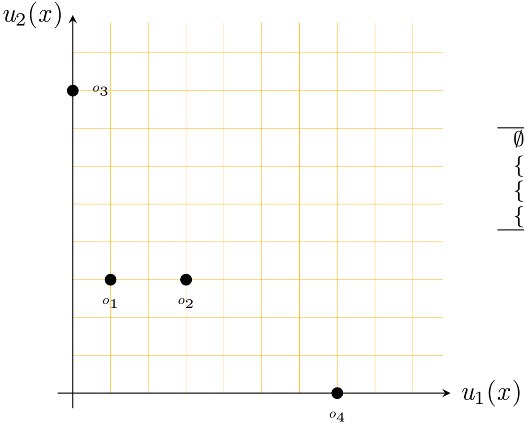

	u_1	u_2
$\emptyset$	0	0
$\{r_1\}$	1	3
$\{r_2\}$	3	3
$\{r_1, r_2\}$	7	8

Good outcomes: more on Efficiency and Fairness

There may be many Pareto-optimal outcomes. Outcomes may also maximize some measure of social welfare:

- utilitarian— maximizes $\sum_i u_i(o)$
- egalitarian— maximizes $\min_i u_i(o)$
- Nash product— maximizes $\Pi_i u_i(o)$

Example:

	u_1	u_2
$\emptyset$	0	0
$\{r_1\}$	1	3
$\{r_2\}$	3	3
$\{r_1, r_2\}$	7	8

Good outcomes: more on Efficiency and Fairness

There may be many Pareto-optimal outcomes. Outcomes may also maximize some measure of social welfare:

- utilitarian— maximizes $\sum_i u_i(o)$
- egalitarian— maximizes $\min_i u_i(o)$
- Nash product— maximizes $\Pi_i u_i(o)$

Example:

	u_1	u_2
$\emptyset$	0	0
$\{r_1\}$	1	3
$\{r_2\}$	3	3
$\{r_1, r_2\}$	7	8

- Which outcome maximizes the utilitarian social welfare, the egalitarian social welfare, and the Nash product?

Good outcomes: more on Efficiency and Fairness

There may be many Pareto-optimal outcomes. Outcomes may also maximize some measure of social welfare:

- utilitarian— maximizes $\sum_i u_i(o)$
- egalitarian— maximizes $\min_i u_i(o)$
- Nash product— maximizes $\Pi_i u_i(o)$

Example:

	u_1	u_2
$\emptyset$	0	0
$\{r_1\}$	1	3
$\{r_2\}$	3	3
$\{r_1, r_2\}$	7	8

- Which outcome maximizes the utilitarian social welfare, the egalitarian social welfare, and the Nash product?
- Which of these notions imply Pareto-optimality?

Individual Rationality

- We denote by $o^{\downarrow}$ the disagreement (or conflict) point.
  It indicates the utility that each player gets if the negotiation fails.
  This needs not be the same for both agents.

Individual Rationality

- We denote by $o^{\downarrow}$ the disagreement (or conflict) point. It indicates the utility that each player gets if the negotiation fails. This needs not be the same for both agents.
- Individual rationality: agents should be better off engaging in the negotiation, that is, for all i, the outcome of the negotiation o must be such that:

$$u_i(o) \geq u_i(o^{\downarrow})$$

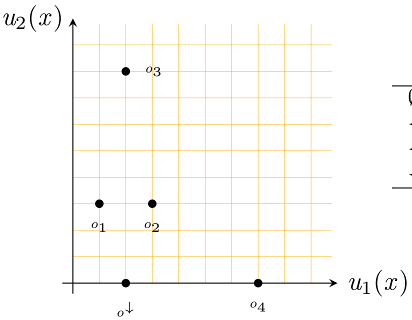

	u_1	u_2
$\emptyset$	2	0
$\{r_1\}$	1	3
$\{r_2\}$	3	3
$\{r_1, r_2\}$	7	8

The Negotiation Set

IR plus Pareto brings the negotiation set (Luce and Raiffa, 1957). However, this set contains many possible solutions. Can we restrict further the set of intuitively "fair" outcomes?

Nash (1950) takes an axiomatic approach, and under some assumptions (in particular that the outcome set is convex), shows that the unique solution to a bargaining problem must be the Nash product, provided we accept some "intuitive" axioms.

Nash Bargaining Solution

A bargaining problem is described as a pair $\langle \mathcal{O}, o^{\downarrow} \rangle$.
We write $o^* = NBS(\langle \mathcal{O}, o^{\downarrow} \rangle)$ for the outcome selected.
Basic axioms:

- Pareto— the solution should be on the Pareto-frontier
- IR— the outcome should be individually rational

Additional axioms:

- Symmetry
- Linear Invariance
- Independance of Irrelevant Alternatives

We discuss them in more details now.

Nash Bargaining Solution: Symmetry

Intuitively, symmetry says that agents should be treated the same when their initial situation is equivalent. Thus:

1. if $u_1(o^{\downarrow}) = u_2(o^{\downarrow})$, and
2. if $\forall o \in \mathcal{O}$: $\exists o' \in \mathcal{O}$ such that $u_1(o) = u_2(o')$ and $u_2(o) = u_1(o')$

then the outcome o^* must be such that $u_1(o^*) = u_2(o^*)$

Nash Bargaining Solution: Linear Invariance

Intuitively, linear invariance says two things:

- independance of scale—the outcome does not depend on the scale used by the agent to represent its utility.
  Suppose agent 1 uses a scale [0,10] to represent its utility, while agent 2 uses a scale [0,100]. The fact that agent 1 enjoys utility 9 and agent 2 utility 50 does not mean that agent 2 is more "happy".
- independance of zero—a translation of the scale of utilities does not affect the outcome.
  Suppose agent 1 uses a scale [0,9], while agent 2 uses a scale [1,10]. The scale of agent 2 can be translated to [0,9] without any consequence on the outcome.

Nash Bargaining Solution: IIA

Intuitively, Independence of Irrelevant Alternatives (IIA) says that if the outcome o^* of the negotiation lies in some sub-region of the outcome set, then the negotiation should still select o^* if we restrict the outcome set to this sub-region.

So, removing "irrelevant outcomes" should not affect the result.

More precisely, for any $O \subseteq \mathcal{O}$,
if $NBS(\langle \mathcal{O}, o^{\downarrow} \rangle) = o^* \in O$ then $NBS(\langle O, o^{\downarrow} \rangle) = o^*$.

Properties of Protocols

- A protocol specifies the rules of interaction (who can say what?). For instance, we may allow simultaneous moves, or sequential moves.
- A strategy specifies the behavior of the agent (which move to select among all the legal ones?)

Properties of Protocols

- A protocol specifies the rules of interaction (who can say what?). For instance, we may allow simultaneous moves, or sequential moves.
- A strategy specifies the behavior of the agent (which move to select among all the legal ones?)

We usually require the following properties of protocols+strategies:

- termination—the negotiation will terminate
- guaranteed agreement—the negotiation will end on an agreement (not on the conflict point)
- efficiency—upon termination, the negotiation provides an efficient (eg. Pareto-optimal) outcome
- equilibrium—captures a notion of stability. In particular:
  - symmetric Nash equilibrium: assuming agent 1 uses strategy s, agent 2 cannot be better off using a different strategy than s.
  - subgame perfect equilibrium: in the case of sequential protocol.

Price of Anarchy

Can we provide worst-case guarantees on the loss of social welfare in a state at equilibrium?

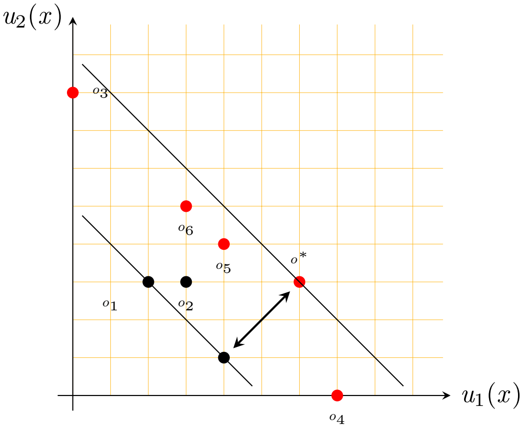

Price of Anarchy

Monotonic Concession Protocol

The protocol proceeds in rounds where agents make simultaneous offers. Offers are assumed to be in the negotiation set to start with.
Let o_i^t and o_j^t be the offers made by agent i and agent j, at round t.
In the initial round, agents make the offer they like, then in the following rounds, each agent must either:

- stick to their previous offer, or
- make a concession (an offer which gives the other more utility)

Monotonic Concession Protocol

The protocol proceeds in rounds where agents make simultaneous offers. Offers are assumed to be in the negotiation set to start with.
Let o_i^t and o_j^t be the offers made by agent i and agent j, at round t.
In the initial round, agents make the offer they like, then in the following rounds, each agent must either:

- stick to their previous offer, or
- make a concession (an offer which gives the other more utility)

An agreement is found when, for at least an agent, the offer made by the other agent is at least as good as its own current offer. That is:

$$u_i(o_j^t) \geq u_i(o_i^t) \text{ or } u_j(o_i^t) \geq u_j(o_j^t)$$

(Flip a coin if both agents agree).

A disagreement occurs when both agents stick to their current offer.

Monotonic Concession Protocol: Zeuthen strategy

How should agents play this game? Zeuthen proposes the following:

The willingness to risk conflict (denoted Z_i^t), intuitively captures "how bad" would be a conflict for agent i at round t. It is given by the following formula (assuming $(0,0)$ for the conflict):

$$Z_i^t = \begin{cases} 1 & \text{if} \quad u_i(o_i^t) = 0 \\ \frac{u_i(o_i^t) - u_i(o_j^t)}{u_i(o_i^t)} & \text{otherwise} \end{cases}$$

Monotonic Concession Protocol: Zeuthen strategy

How should agents play this game? Zeuthen proposes the following:

The willingness to risk conflict (denoted Z_i^t), intuitively captures "how bad" would be a conflict for agent i at round t. It is given by the following formula (assuming $(0,0)$ for the conflict):

$$Z_i^t = \begin{cases} 1 & \text{if} \quad u_i(o_i^t) = 0 \\ \frac{u_i(o_i^t) - u_i(o_j^t)}{u_i(o_i^t)} & \text{otherwise} \end{cases}$$

From this, the Zeuthen strategy is specified as follows, for agent i:

- compute your willingness to risk conflict Z_i^t and that of your partner
- the one with the smallest value should concede
- make the minimal concession making Z_j^t become smaller than Z_i^t

Monotonic Concession Protocol: Example

Two robots need to collect items at different sites:

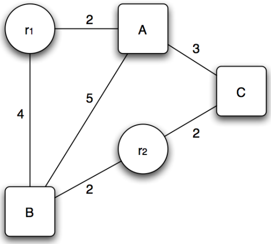

Monotonic Concession Protocol: Example

round	offer a_1	offer a_2	$u_1(o^t_{a_1}), u_1(o^t_{a_2})$	$u_2(o^t_{a_1}), u_2(o^t_{a_2})$	Z_1	Z_2
1	$\langle \emptyset, \{a, b, c\} \rangle$	$\langle \{a, b, c\}, \emptyset \rangle$	9,0	0,9	1	1

Monotonic Concession Protocol: Example

round	offer a_1	offer a_2	$u_1(o^t_{a_1}), u_1(o^t_{a_2})$	$u_2(o^t_{a_1}), u_2(o^t_{a_2})$	Z_1	Z_2
1	$\langle \emptyset, \{a,b,c\} \rangle$	$\langle \{a,b,c\}, \emptyset \rangle$	9,0	0,9	1	1
2	$\langle \{a\}, \{b,c\} \rangle$	$\langle \{a,c\}, \{b\} \rangle$	7,4	3,7	$\frac{3}{7}$	$\frac{4}{7}$

Monotonic Concession Protocol: Example

round	offer a_1	offer a_2	$u_1(o^t_{a_1}), u_1(o^t_{a_2})$	$u_2(o^t_{a_1}), u_2(o^t_{a_2})$	Z_1	Z_2
1	$\langle\emptyset, \{a, b, c\}\rangle$	$\langle\{a, b, c\}, \emptyset\rangle$	9,0	0,9	1	1
2	$\langle\{a\}, \{b, c\}\rangle$	$\langle\{a, c\}, \{b\}\rangle$	7,4	3,7	$\frac{3}{7}$	$\frac{4}{7}$
3	$\langle\{a, c\}, \{b\}\rangle$	$\langle\{a, c\}, \{b\}\rangle$	4,4	7,7	stop	stop

Monotonic Concession Protocol: Properties

The properties of the MCP + Zeuthen strategy are as follows:

- termination is guaranteed, as well as agreement upon termination (there is always at least an agent willing to concede)
- because the offers considered are in the Negotiation Set to start with, Pareto-optimality is obvious

Now a stronger result:

- The outcome maximizes the Nash product.

Monotonic Concession Protocol: Properties

The properties of the MCP + Zeuthen strategy are as follows:

- termination is guaranteed, as well as agreement upon termination (there is always at least an agent willing to concede)
- because the offers considered are in the Negotiation Set to start with, Pareto-optimality is obvious

Now a stronger result:

- The outcome maximizes the Nash product.

Did we know this already by Nash result?

Monotonic Concession Protocol: Properties

The properties of the MCP + Zeuthen strategy are as follows:

- termination is guaranteed, as well as agreement upon termination (there is always at least an agent willing to concede)
- because the offers considered are in the Negotiation Set to start with, Pareto-optimality is obvious

Now a stronger result:

- The outcome maximizes the Nash product.

Did we know this already by Nash result?

Warning: the domain is not convex here.

Monotonic Concession Protocol: Properties

The properties of the MCP + Zeuthen strategy are as follows:

- termination is guaranteed, as well as agreement upon termination (there is always at least an agent willing to concede)
- because the offers considered are in the Negotiation Set to start with, Pareto-optimality is obvious

Now a stronger result:

- The outcome maximizes the Nash product.

Did we know this already by Nash result?

Warning: the domain is not convex here.

How about stability?

- the Zeuthen strategy is not in symmetric equilibrium

Monotonic Concession Protocol: Properties

The properties of the MCP + Zeuthen strategy are as follows:

- termination is guaranteed, as well as agreement upon termination (there is always at least an agent willing to concede)
- because the offers considered are in the Negotiation Set to start with, Pareto-optimality is obvious

Now a stronger result:

- The outcome maximizes the Nash product.

Did we know this already by Nash result?

Warning: the domain is not convex here.

How about stability?

- the Zeuthen strategy is not in symmetric equilibrium

Explanation: The problem comes from the last step of the protocol. If both agents have the same Z, both are willing to concede, and so one agent can exploit this and deviate to get a better outcome.

Monotonic Concession Protocol

Some final remarks on MCP.

- it is possible to extend the Zeuthen strategy (by allowing a mixed strategy in the last step) to retrieve stability
- a more simple one-step protocol is possible!

The one-step protocol is as follows:

- agents simultaneously make a single offer
- select the one maximizing the product of utilities

What is the best strategy for an agent given this protocol?

Monotonic Concession Protocol

Some final remarks on MCP.

- it is possible to extend the Zeuthen strategy (by allowing a mixed strategy in the last step) to retrieve stability
- a more simple one-step protocol is possible!

The one-step protocol is as follows:

- agents simultaneously make a single offer
- select the one maximizing the product of utilities

What is the best strategy for an agent given this protocol?
Given this protocol, the strategy for an agent is to select, among the outcomes maximizing, the one giving him the best utility.

Rosenschein & Zlotkin. *Rules of Encounter*. 1993.

Alternating Offers

We now discuss a sequential protocol.

- an agent starts by making an offer. In the next round, the other agent can either accept or make a counter-offer.
- the protocol integrates a discount factor λ_i to capture the fact that negotiation is time constrained. An offer accepted at round t by agent i brings utility $u_i(o^t) \times (\lambda_i)^t$.

The sequential nature of this protocol allows backward induction solving.

Rubinstein. *Perfect equilibrium in a bargaining model.* Econometrica-1982.

Alternating Offers

Set $\lambda = 1$ for agents (they are patient).

- suppose the number of rounds is known in advance. But then the last agent to make an offer gets all the "power". What is his best strategy?

Alternating Offers

Set $\lambda = 1$ for agents (they are patient).

- suppose the number of rounds is known in advance. But then the last agent to make an offer gets all the "power". What is his best strategy?
  Always refuse the offers of the other, then make an offer $\langle 1-\epsilon, \epsilon \rangle$ *in the last round (this last step is actually an ultimatum game: more on this later)*

Alternating Offers

Set $\lambda = 1$ for agents (they are patient).

- suppose the number of rounds is known in advance. But then the last agent to make an offer gets all the "power". What is his best strategy?
  Always refuse the offers of the other, then make an offer $\langle 1 - \epsilon, \epsilon \rangle$ *in the last round (this last step is actually an ultimatum game: more on this later)*
- suppose the number of rounds is not known in advance
  Suppose a_1 *uses this strategy: Always propose* $(1 - \epsilon, \epsilon)$*, and always refuse the offer of the other. What is* a_2 *best response to this?*

Alternating Offers

Set $\lambda = 1$ for agents (they are patient).

- suppose the number of rounds is known in advance. But then the last agent to make an offer gets all the "power". What is his best strategy?
  Always refuse the offers of the other, then make an offer $\langle 1-\epsilon, \epsilon \rangle$ in the last round (this last step is actually an ultimatum game: more on this later)
- suppose the number of rounds is not known in advance
  Suppose a_1 uses this strategy: Always propose $(1-\epsilon, \epsilon)$, and always refuse the offer of the other. What is a_2 best response to this? Always refusing yields the conflict outcome. So a_2 must accept at some point, no reason to postpone: accept in the first round.
  Immediate acceptance of any offer is a Nash equilibrium, given that a_2 knows a_1's strategy.

Alternating Offers: Example

Take $\mathcal{O} = \{o_2, o_3, o_6\}$, with $o_2 = \langle 7, 3\rangle$, $o_3 = \langle 5, 4\rangle$, and $o_6 = \langle 4, 7\rangle$.

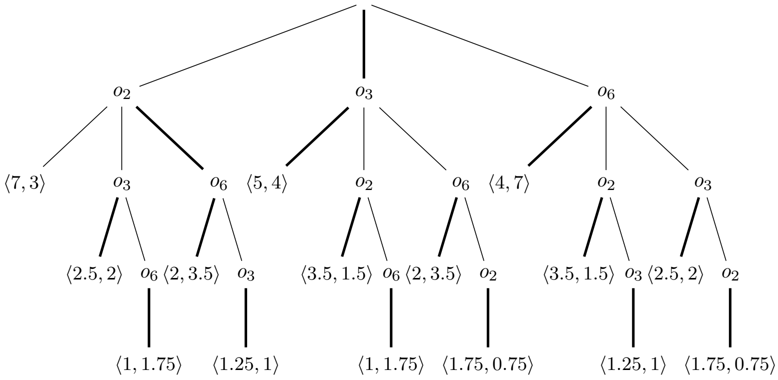

Figure: Backward Induction with the alternating-offer protocol

Some words of caution (I)

Back to the ultimatum game.

Remember: One agent proposes an offer (say a division of a pie), the offer may either accept or reject. If it accepts the offer is chosen outcome, otherwise the conflict outcome.

What do you think a human agent will propose in real life?

Some words of caution (I)

Back to the ultimatum game.

Remember: One agent proposes an offer (say a division of a pie), the offer may either accept or reject. If it accepts the offer is chosen outcome, otherwise the conflict outcome.

What do you think a human agent will propose in real life?

- many studies in economics
- usually offers more around a 60/40 division
- importance of social context, reputation, etc.

Some words of caution (II)

Now consider the following game, known as the centipede game.

There are 100 candies to share, and two agents. The protocol for negotiation is as follows. In each round:

- player i can either take 1 or 2 candies
- if he takes 2 candies, the protocol terminates, and agents keep the candies they have collected so far (the rest is wasted)
- if he takes 1 candy, the protocol continues, by giving the turn to the other agent, and so on.

Some words of caution (II)

Now consider the following game, known as the centipede game.

There are 100 candies to share, and two agents. The protocol for negotiation is as follows. In each round:

- player i can either take 1 or 2 candies
- if he takes 2 candies, the protocol terminates, and agents keep the candies they have collected so far (the rest is wasted)
- if he takes 1 candy, the protocol continues, by giving the turn to the other agent, and so on.

Can you analyze this game? (maybe with 4 candies ;-)

Some words of caution (III)

There is evidence from economics that agents decide based on reasons they have at their disposal. This violates many "irrelevant alternatives" assumptions.

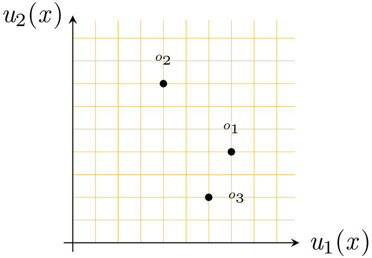

Dominance effect

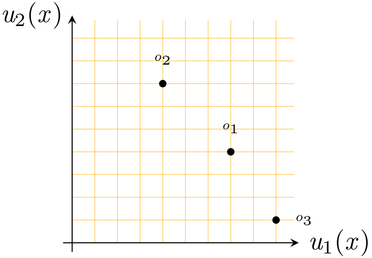

Compromise effect

Shafir, Simonson, & Tversky. *Reason-based Choice*. Cognition-1993.

Heuristic-Based Negotiation

Often, agents have not an exact knowledge about the preferences of others, nor of the strategy they use: this is a setting under incomplete information.

It is possible to conceive general profiles of agents, specifying high-level behavior. These agents will typically adapt to different parameters of the negotiation setting (time, proposals of the other, etc.). There is a large spectrum of techniques, up to very sophisticated opponent modelling.

Basic tactics based on deadlines:

- boulware agents—very slow concession until we get close to the deadline, then exponential increase
- conceder agents—prone to concede in the first rounds of negotiation and get close to reserve price, then slow increase

Faratin et al.. *Negotiation decision functions for autonomous agents*. Robotics and Autonomous Systems-1998.

Heuristic-Based Negotiation

How to be sure that a move is indeed a concession in the first place? Trying to guess/approximate an agent preference structure based on its negotiation behavior is very challenging!

Idea: seek the offer which is the "closest" from the other agent offer in the preceding move. To do this, compute similarity among offers.

Heuristic-Based Negotiation

How to be sure that a move is indeed a concession in the first place? Trying to guess/approximate an agent preference structure based on its negotiation behavior is very challenging!

Idea: seek the offer which is the "closest" from the other agent offer in the preceding move. To do this, compute similarity among offers.

- we can then compute similarity among offers (by summing similarity, taking weights into account)
- finally the agent seeks among all offers giving her the same utility the one which is most similar to the other agent's previous offer.

Faratin et al.. *Using similarity criteria to make issue trade-offs in automated negotiations*. AIJ-2002.

Testing Strategies: Agent Competitions

In recent years, competitions involving negotiating agents have emerged, allowing to test and compare various strategies on different problems.

- ANAC Competition: Automated Negotiating Agents Competition
- TAC: Trading Agent Competition (auctions, etc.)
- Genius platform (negotiation problems, library of agents' strategies) `http://mmi.tudelft.nl/negotiation/index.php/Genius`
- many papers and even books on analysis of the best strategies

Wellmann, Greenwald, & Stone. *Autonomous Bidding Agents: Strategies and Lessons from the TAC competition*. 2007.

Outline of the lecture

Single Text Mediated Protocol

A first possible approach is to use a mediator.
The protocol is as follows (K is fixed a priori):

```
for t:=1 to K do
begin
  the mediator proposes an offer o ;
  agents votes on o (accept/refuse);
  if all agents accept, then current := o;
end;
```

So the protocol returns the latest unanimously accepted offer.

Essentially, the protocol starts from an offer, and performs Pareto improvements.

- Is the protocol guaranteed to reach a Pareto-optimal outcome?
- Is the protocol guaranteed to stop when having reached a Pareto optimal outcome?

Single Text Mediated Protocol: Example

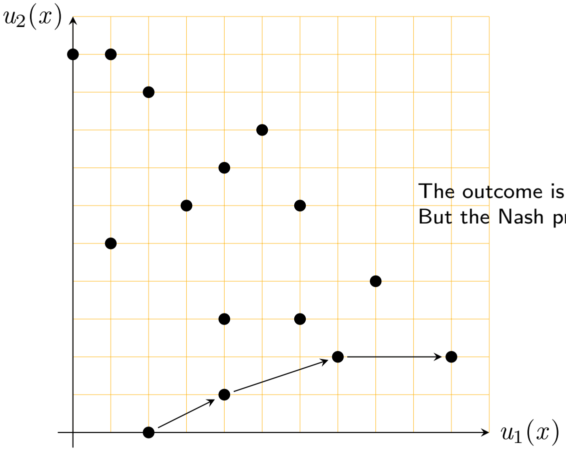

The outcome is $\langle 10, 2\rangle$
But the Nash product optimum is $\langle 5, 8\rangle$

Single Text Mediated Protocol

This protocol may be problematic in some contexts:

- it requires a mediator (not always possible)
- requires many rounds of communication from all agents to the mediator
- it can reach outcomes with very low social welfare

Some extensions have been proposed to try to address some of these limitations:

- use of meta-heuristic techniques to avoid local optima (eg. simulated annealing)
- learning of agents preferences to guide the offer proposal from the mediator

Klein et al.. *Protocols for negotiating complex contracts*. IEEE Intelligent Systems.

Contract-Based Negotiation

Instead of trying to deal with multilateral encounters, let us try to build on simple building blocks. We take inspiration from Contract-Net protocols.

- negotiation starts with an initial allocation
- agents asynchronously negotiate resources
- deals to move from one allocation to another, ie $\delta = (A, A')$
- deals can involve payments (utility transfer);
- agents accept deals on the basis of a rationality criterion, we assume myopic IR: $v_i(A') - v_i(A) > p(i)$

Contract-Based Negotiation

Different types of deals can be considered
"natural" restrictions on the type of exchanges allowed between agents, in particular:

- 1-deals: exchange of a single resource
- swap deals: swapping two resources among agents
- bilateral deal: exchange involving two agents
- clique deal: exchange among agents in a clique of neighbours

Different assumptions on the preference structures
"natural" restrictions/assumptions to be made on the preferences of all the agents of the system, in particular:

- monotonicity: $v_i(B_1) \leq v_i(B_2)$ when $B_1 \subseteq B_2$
- modularity: $v(S_1 \cup S_2) = v(S_1) + v(S_2) - v(S_1 \cap S_2)$

A very simple model of contract-based negotiation

A house market setting...

- n agents, n resources, each agent has to get one resource, *and initially holds one*,
- agents have preferences (linear order) over resources

... under a dynamic perspective:

- agents exchange resources thanks to mutually beneficial (rational) swap contracts (no money involved)
- ... until a *stable* allocation is reached (no more deal is possible).

We assess the quality of allocations on:

- Pareto-optimality
- utilitarian and egalitarian social welfare (giving utilities to ranks in preference)

Example

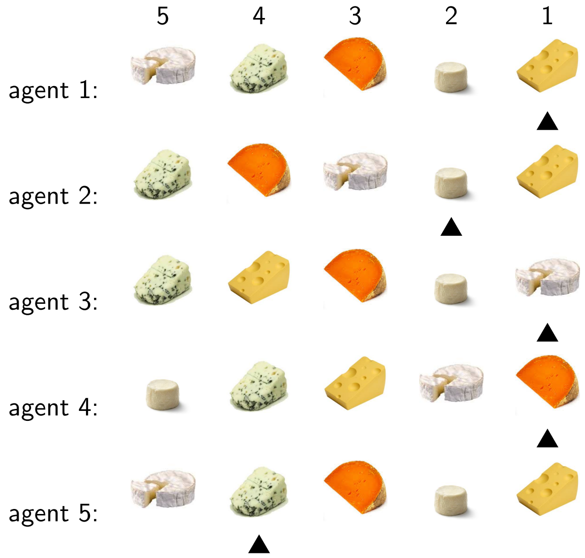

Example

5 4 3 2 1

agent 1:

agent 2:

agent 3:

agent 4:

agent 5:

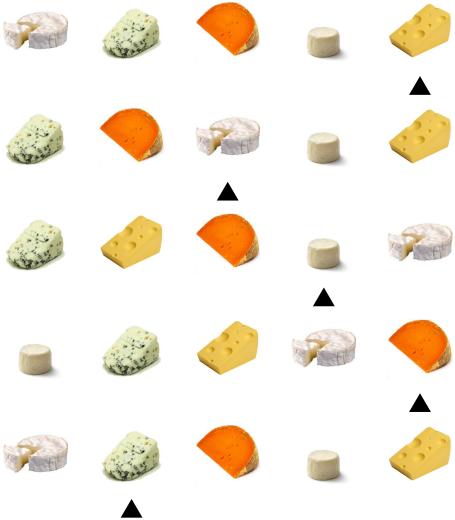

Example

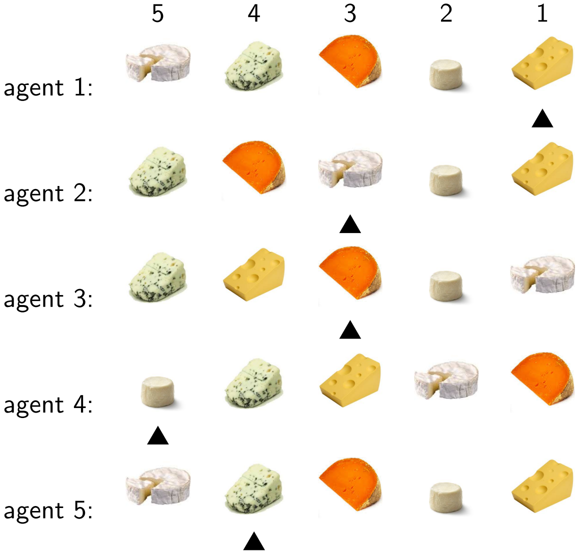

Example

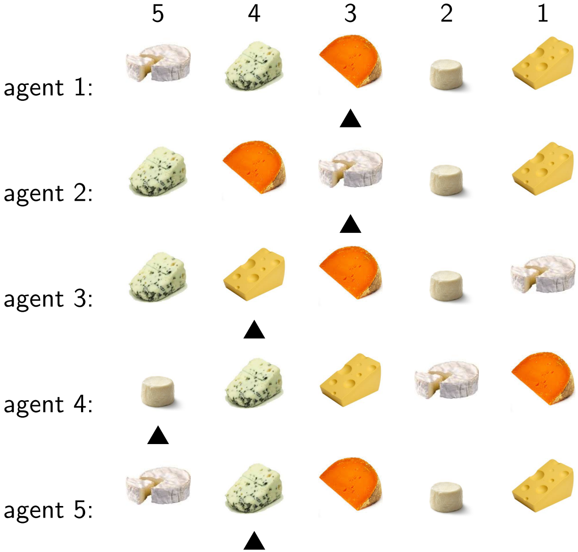

Example

agent 1:
agent 2:
agent 3:
agent 4:
agent 5:

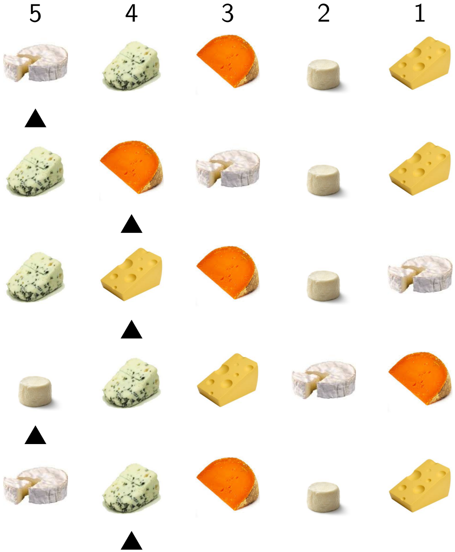

$\Rightarrow sw_u(A) = 22$

$\Rightarrow sw_e(A) = 4$

Another example

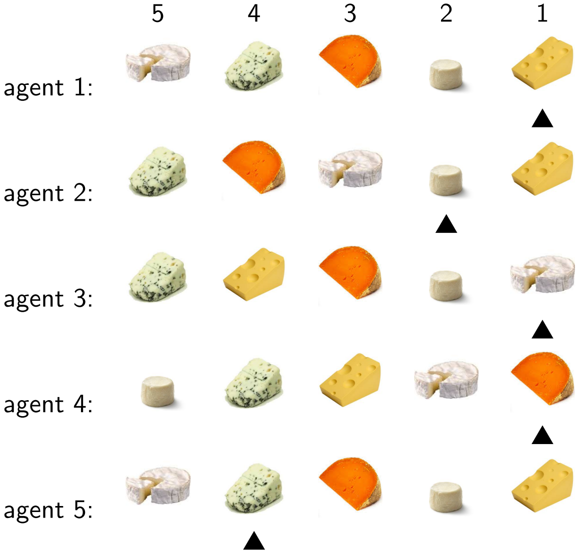

Another example

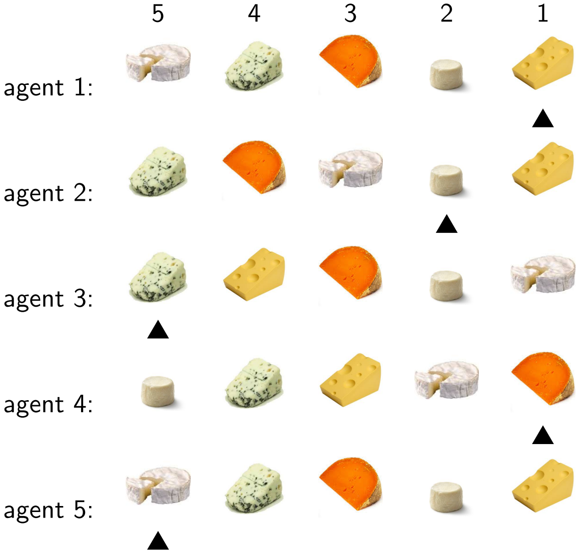

Another example

5 4 3 2 1

agent 1:

agent 2:

agent 3:

agent 4:

agent 5:

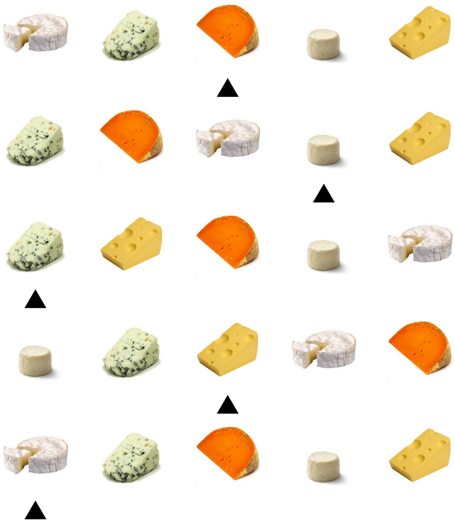

$\Rightarrow sw_u(A) = 18$

Lack of Pareto-optimality

First remark: swap deals do not guarantee convergence to Pareto-optimal allocations:

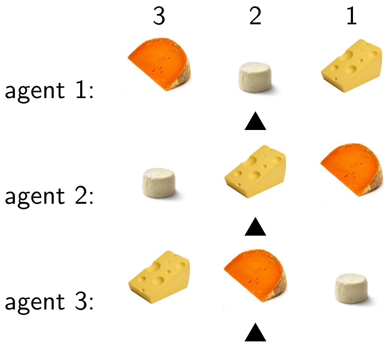

Price of Anarchy

Gap between socially optimal and stable allocations:

$$PoA = max_{I \in \mathcal{I}} \frac{max_{A \in I}\, sw_u(A)}{min_{A \in C_k(I)}\, sw_u(A)}$$

➤ Any IR procedure have PoA ≥ 2.

a_1	1	$\succ$	5	$\succ$	...	$\succ$	...	$\succ$	...
a_2	1	$\succ$	2	$\succ$	...	$\succ$	...	$\succ$	...
a_3	2	$\succ$	1	$\succ$	3	$\succ$	...	$\succ$	...
a_4	3	$\succ$	1	$\succ$	2	$\succ$	4	$\succ$	...
a_5	4	$\succ$	...	$\succ$	...	$\succ$	...	$\succ$	5

Price of Anarchy

Gap between socially optimal and stable allocations:

$$PoA = max_{I \in \mathcal{I}} \frac{max_{A \in I}\, sw_u(A)}{min_{A \in C_k(I)}\, sw_u(A)}$$

➤ Any IR procedure have PoA ≥ 2.

a_1	[1]	$\succ$	5 (highlighted)	$\succ$	$\dots$	$\succ$	$\dots$	$\succ$	$\dots$
a_2	1 (highlighted)	$\succ$	[2]	$\succ$	$\dots$	$\succ$	$\dots$	$\succ$	$\dots$
a_3	2 (highlighted)	$\succ$	1	$\succ$	[3]	$\succ$	$\dots$	$\succ$	$\dots$
a_4	3 (highlighted)	$\succ$	1	$\succ$	2	$\succ$	[4]	$\succ$	$\dots$
a_5	4 (highlighted)	$\succ$	$\dots$	$\succ$	$\dots$	$\succ$	$\dots$	$\succ$	[5]

But take a 2-stable allocation A: for each pair of agents (x, y), at least one agent ranks the resource of the other below her current...

➤ C_2 have PoA $\leq 2 \Rightarrow$ All cycle procedures have PoA $= 2$.

☞ *The size of the allowed cycles does not change anything regarding the social welfare loss (in the worst-case)*

Contract-Based Negotiation with Money

Some known results:

- a deal is IR (with money) iff it increases utilitarian social welfare (thus generates a surplus).
- allows to show that any sequence of IR deals converges to an allocation maximizing utilitarian social welfare
- however, may require very complex deals to be implemented during the negotiation (in fact, for any conceivable deal we may construct a scenario requiring exactly that deal).
- for modular domains, convergence is guaranteed for negotiations involving 1-deals only

Sandholm. *Contract types for satisficing task allocation*. IEEE Symposium-1998.

Endriss et al.. *Negotiating socially optimal allocation of resources*. JAIR-2006.

Contract-Based Negotiation

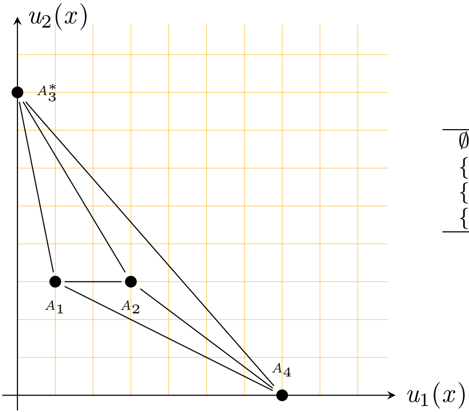

	u_1	u_2
$\emptyset$	0	0
$\{r_1\}$	1	3
$\{r_2\}$	3	3
$\{r_1, r_2\}$	7	8

Contract-Based Negotiation

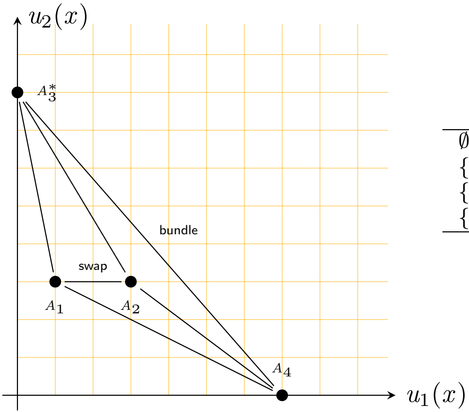

	u_1	u_2
$\emptyset$	0	0
$\{r_1\}$	1	3
$\{r_2\}$	3	3
$\{r_1, r_2\}$	7	8

Contract-Based Negotiation

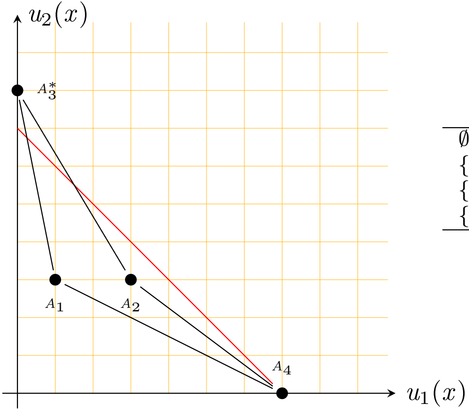

	u_1	u_2
$\emptyset$	0	0
$\{r_1\}$	1	3
$\{r_2\}$	3	3
$\{r_1, r_2\}$	7	8

Maximality of Modular wrt. Bilateral-deals

How far can we get with bilateral deals?
Assume (at least) 3 agents, and take an arbitrary non-modular valuation function:

$$v_1 = a + b.r_1 + c.r_2 + d.r_1.r_2$$

Maximality of Modular wrt. Bilateral-deals

How far can we get with bilateral deals?
Assume (at least) 3 agents, and take an arbitrary non-modular valuation function:

$$v_1 = a + b.r_1 + c.r_2 + d.r_1.r_2$$

We need to show that it is possible to construct two modular functions and select an initial allocation such that no bilateral deals would lead to optimal sw. Assuming $d > 0$ here, take:

$$v_2 = v_3 = (b + \frac{1}{3}d).r_1 + (c + \frac{1}{3}d).r_2$$

Initially (A_0), we allocate r_1 to agent 2 and r_2 to agent 3.
Hence, $sw(A_0) = a + b + c + \frac{2}{3}d < sw(A^*)$, where A^* is the allocation where agent 1 receives both objects.

Chevaleyre et al.. *Simple Negotiation Schemes for Agents with Simple Preferences*. JAAMAS-2010.

Balanced Outcomes

Network Exchange Theory: agents can only negotiate with neighbours.

- agents are now located on a graph G
- each agent can reach an agreement with at most one neighbour
- each pair of agents negotiate over the division of 1 euro

Example:

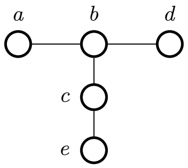

Can you guess how the negotiation will unfold?
Which agreements are met, how the money is divided?

Kleinberg & Tardos. *Balanced Outcomes in Social Exchange Networks*. STOC-08.

Balanced Outcomes

Network Exchange Theory: agents can only negotiate with neighbours.

- agents are now located on a graph G
- each agent can reach an agreement with at most one neighbour
- each pair of agents negotiate over the division of 1 euro

Example:

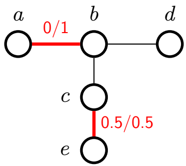

Intuition: b uses his "power" to have two potential agreements (a/d). c sees that b would not deal with him, so focus on e.

Kleinberg & Tardos. *Balanced Outcomes in Social Exchange Networks*. STOC-08.

Balanced Outcomes

More precisely, we can define:

- an outcome is a pair $\langle M, \alpha \rangle$, where M is a matching (which agents agree on a deal), and values α_x for each agent x, with:
  - $\alpha_x + \alpha_y = 1$ when $(x, y) \in M$,
  - $\alpha_x = 0$ when $x \notin M$.

Let β_x be the best alternative for x, that is, $max\{1 - \alpha_y | (x, y) \in G\}$

- an outcome is stable if $\alpha_x \geq \beta_x$, for all x.

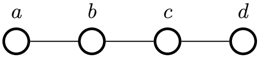

Balanced Outcomes

More precisely, we can define:

- an outcome is a pair $\langle M, \alpha \rangle$, where M is a matching (which agents agree on a deal), and values α_x for each agent x, with:
  - $\alpha_x + \alpha_y = 1$ when $(x, y) \in M$,
  - $\alpha_x = 0$ when $x \notin M$.

Let β_x be the best alternative for x, that is, $max\{1 - \alpha_y | (x, y) \in G\}$

- an outcome is stable if $\alpha_x \geq \beta_x$, for all x.

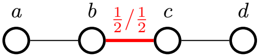

Is this multi-outcome stable?

Balanced Outcomes

More precisely, we can define:

- an outcome is a pair $\langle M, \alpha \rangle$, where M is a matching (which agents agree on a deal), and values α_x for each agent x, with:
  - $\alpha_x + \alpha_y = 1$ when $(x, y) \in M$,
  - $\alpha_x = 0$ when $x \notin M$.

Let β_x be the best alternative for x, that is, $max\{1 - \alpha_y | (x, y) \in G\}$

- an outcome is stable if $\alpha_x \geq \beta_x$, for all x.

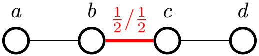

Is this multi-outcome stable? No. Eg. take b: $\alpha_b = \frac{1}{2}$, when $\beta_b = 1$.

Balanced Outcomes

More precisely, we can define:

- an outcome is a pair $\langle M, \alpha \rangle$, where M is a matching (which agents agree on a deal), and values α_x for each agent x, with:
  - $\alpha_x + \alpha_y = 1$ when $(x, y) \in M$,
  - $\alpha_x = 0$ when $x \notin M$.

Let β_x be the best alternative for x, that is, $max\{1 - \alpha_y | (x, y) \in G\}$

- an outcome is stable if $\alpha_x \geq \beta_x$, for all x.

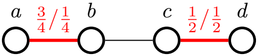

Is this multi-outcome stable?

Balanced Outcomes

More precisely, we can define:

- an outcome is a pair $\langle M, \alpha \rangle$, where M is a matching (which agents agree on a deal), and values α_x for each agent x, with:
  - $\alpha_x + \alpha_y = 1$ when $(x, y) \in M$,
  - $\alpha_x = 0$ when $x \notin M$.

Let β_x be the best alternative for x, that is, $max\{1 - \alpha_y | (x, y) \in G\}$

- an outcome is stable if $\alpha_x \geq \beta_x$, for all x.

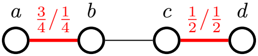

Is this multi-outcome stable? No. Eg. take b: $\alpha_b = \frac{1}{4}$, when $\beta_b = \frac{1}{2}$.

Balanced Outcomes

More precisely, we can define:

- an outcome is a pair $\langle M, \alpha \rangle$, where M is a matching (which agents agree on a deal), and values α_x for each agent x, with:
  - $\alpha_x + \alpha_y = 1$ when $(x, y) \in M$,
  - $\alpha_x = 0$ when $x \notin M$.

Let β_x be the best alternative for x, that is, $max\{1 - \alpha_y | (x, y) \in G\}$

- an outcome is stable if $\alpha_x \geq \beta_x$, for all x.

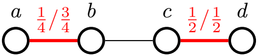

Is this multi-outcome stable?

Balanced Outcomes

More precisely, we can define:

- an outcome is a pair $\langle M, \alpha \rangle$, where M is a matching (which agents agree on a deal), and values α_x for each agent x, with:
  - $\alpha_x + \alpha_y = 1$ when $(x, y) \in M$,
  - $\alpha_x = 0$ when $x \notin M$.

Let β_x be the best alternative for x, that is, $max\{1 - \alpha_y | (x, y) \in G\}$

- an outcome is stable if $\alpha_x \geq \beta_x$, for all x.

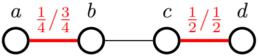

Is this multi-outcome stable? Yes!

Balanced Outcomes

More precisely, we can define:

- an outcome is a pair $\langle M, \alpha \rangle$, where M is a matching (which agents agree on a deal), and values α_x for each agent x, with:
  - $\alpha_x + \alpha_y = 1$ when $(x, y) \in M$,
  - $\alpha_x = 0$ when $x \notin M$.

Let β_x be the best alternative for x, that is, $max\{1 - \alpha_y | (x, y) \in G\}$

- an outcome is stable if $\alpha_x \geq \beta_x$, for all x.

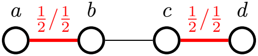

Is this multi-outcome stable?

Balanced Outcomes

More precisely, we can define:

- an outcome is a pair $\langle M, \alpha \rangle$, where M is a matching (which agents agree on a deal), and values α_x for each agent x, with:
  - $\alpha_x + \alpha_y = 1$ when $(x, y) \in M$,
  - $\alpha_x = 0$ when $x \notin M$.

Let β_x be the best alternative for x, that is, $max\{1 - \alpha_y | (x, y) \in G\}$

- an outcome is stable if $\alpha_x \geq \beta_x$, for all x.

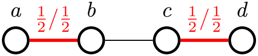

Is this multi-outcome stable? Yes!
But contradicted by experiments (b and c have more negotiation power)

Balanced Outcomes

The idea is to strengthen the notion of stability.

- A balanced outcome is an outcome such that, for all $(x, y) \in M$, (α_x, α_y) constitutes a Nash Bargaining Solution considering (β_x, β_y) as the disagreement outcome.

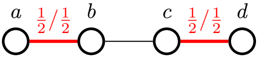

This is not a balanced outcome.

Balanced Outcomes

The idea is to strengthen the notion of stability.

- A balanced outcome is an outcome such that, for all $(x, y) \in M$, (α_x, α_y) constitutes a Nash Bargaining Solution considering (β_x, β_y) as the disagreement outcome.

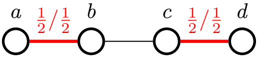

This is not a balanced outcome. Indeed take $(\alpha_a, \alpha_b) = (0.5, 0.5)$.
Given $(\beta_a, \beta_b) = (0, 0.5)$, the surplus $1 - 0.5$ should be evenly divided, yielding $(\alpha_a, \alpha_b) = (0.25, 0.75)$.
But now given $(\beta_c, \beta_d) = (0.25, 0)$, the values of (α_c, α_d) should be modified... $\Rightarrow$ fixed-point definition

- Can you guess the balanced outcome here?

Balanced Outcomes

The idea is to strengthen the notion of stability.

- A balanced outcome is an outcome such that, for all $(x, y) \in M$, (α_x, α_y) constitutes a Nash Bargaining Solution considering (β_x, β_y) as the disagreement outcome.

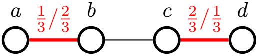

Gives rise to many questions:

- are balanced outcomes guaranteed to exist? (if not, when?)
- are these values rational?
- is it easy to compute these values?
- etc.

Kleinberg & Tardos. *Balanced Outcomes in Social Exchange Networks*. STOC-08.

Outline of the lecture

1. Bilateral Negotiation
   - The setting
   - What are good outcomes?
   - Axiomatics of negotiation
   - Protocols and Game-theoretical analysis
   - Heuristics
2. Multilateral Negotiation
   - A Mediated Protocol
   - Contract-Based Negotiation
   - Outcomes on Networks
3. Negotiation on Meaning
   - Naming Games
   - Negotiated Ontology Alignment

Is there anything like negotiation of meaning?

Two extracts from the literature:

Is there anything like negotiation of meaning?

Two extracts from the literature:

> *"By definition it will not be possible to define all the needed communication conventions and ontologies in advance and robots will have to build up and negotiate their own communication systems, situated and grounded in their ongoing activities" (A. Baronchelli, citing the work of L. Steels)*

Is there anything like negotiation of meaning?

Two extracts from the literature:

> *"By definition it will not be possible to define all the needed communication conventions and ontologies in advance and robots will have to build up and negotiate their own communication systems, situated and grounded in their ongoing activities" (A. Baronchelli, citing the work of L. Steels)*

> *"how can agents align ontologies that they do not want to disclose? [...] Agents need to agree on what correspondences they believe to be the most relevant to resolve ambiguous combinations, whilst attempting to reduce the number of messages communicated, and minimise the number of beliefs disclosed." (T. Payne and V. Tamma)*

Naming Games

Each agent maintains a inventory, i.e. a set of terms (initially empty) associated to objects. Agents meet in a pairwise fashion, one being the speaker and the other the hearer.

1. the speaker picks a word from his inventory (if the vocabulary is empty, it makes one's up)
2. the speaker communicates this word to the hearer
3. the hearer checks whether she has this word (associated to this object) in her inventory:
   - if this the case, then both the speaker and hearer only keep this word in their inventoty
   - otherwise, the hearer adds the word to her inventory

Baronchelli et al. *Sharp transitions towards shared vocabularies in multiagent systems.* Journal of statistical mechanics, 2006.

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

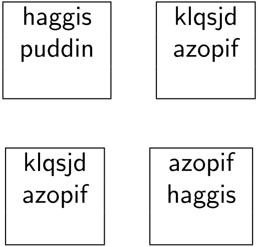

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

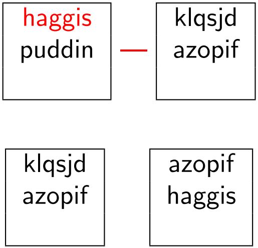

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

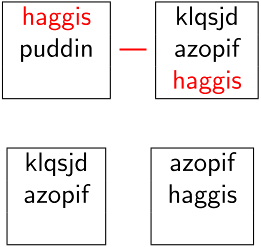

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

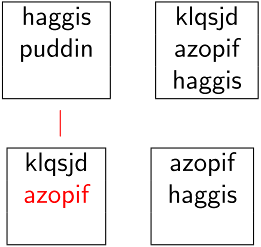

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

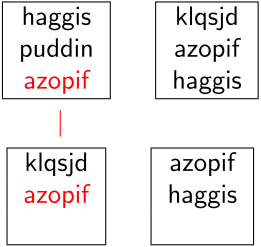

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

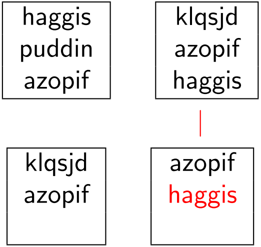

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

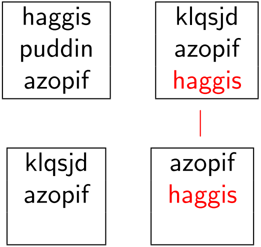

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

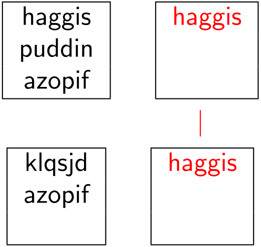

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

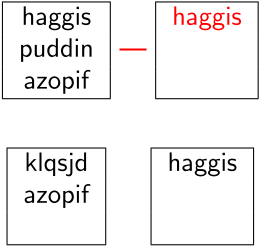

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

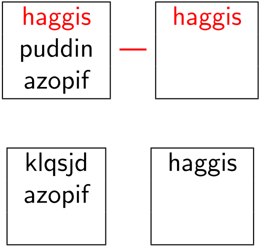

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

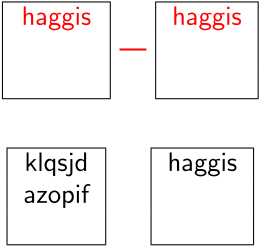

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

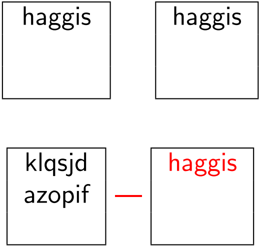

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

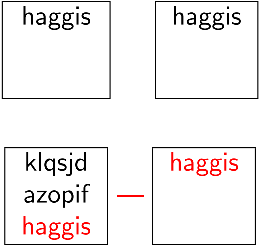

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

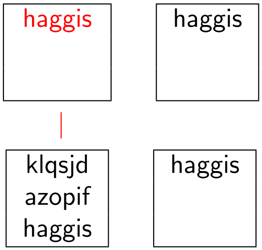

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

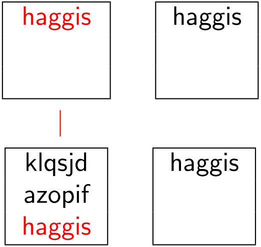

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

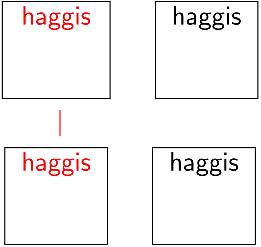

Naming Games (example)

Without loss of generality, assume we are concerned with a single object.

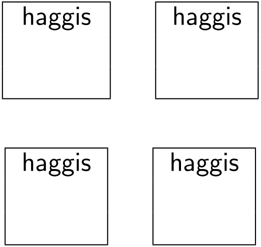

Naming Games: properties

What properties?

- coherent state: all agents have a single (same) word in their inventory.

The coherent state is an absorbing state, i.e. the system is stable once in this state (more interactions will not modify it).

Naming Games: properties

What properties?

- coherent state: all agents have a single (same) word in their inventory.

The coherent state is an absorbing state, i.e. the system is stable once in this state (more interactions will not modify it).
Are there other absorbing states?

Naming Games: properties

What properties?

- coherent state: all agents have a single (same) word in their inventory.

The coherent state is an absorbing state, i.e. the system is stable once in this state (more interactions will not modify it).
Are there other absorbing states?

- reachability: are we sure to attain the absorbing state (at some point)?

Naming Games: properties

What properties?

- coherent state: all agents have a single (same) word in their inventory.

The coherent state is an absorbing state, i.e. the system is stable once in this state (more interactions will not modify it).
Are there other absorbing states?

- reachability: are we sure to attain the absorbing state (at some point)?

This can be shown to occur with probability 1.
Argument: from any state, we can reach the absorbing state by $2(n-1)$: a single agent will talk with all the other agents twice, using the same word w (that is, after these two interactions, the other agent will only have w in his inventory). Since this sequence can occur with some probability, an absorbing state must be attained asymptotically.

Naming Games: properties

But can we get a more detailed understanding of the dynamics?

- $N_w(t)$: total number of words in the system
- $N_d(t)$: total number of different words in the system
- $S(t)$: success rate of interaction

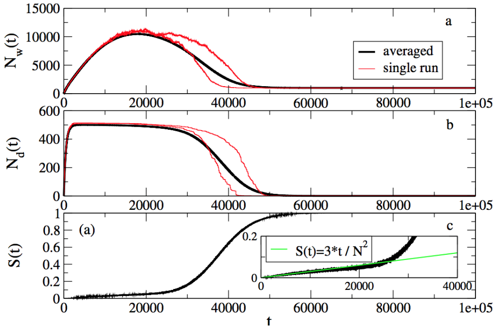

Figure (n=1000) from: Baronchelli et al. *Sharp transitions towards shared vocabularies in multiagent systems.* Journal of statistical mechanics, 2006.

Naming Games: properties

The process can be divided in three phases:

1. (very early) the number of different words increases in the system
2. building correlations
3. (close to $N_w(t)$ max): phase transition to coherence

Many other variants of this model have been studied, in particular:

- dynamics on various types of graphs (regular graphs, small worlds, ...)
- other parameters, like probability of successful update

See the work of A. Baronchelli, A. Barrat, L. Steels, K. Tuyls...

Negotiated Ontology Alignment

An example where ontology alignments are negotiated.

- A correspondence is a mapping between two entities, one in source ontology, and one in a target ontology. An alignment is a set of such correspondence.
- Each agent associates a degree of belief to each correspondence ("the likelihood of being included in some alignment").
- A protocol is proposed to align ontology without full disclosure of beliefs.

Terry Payne and Valentina Tamma. *Negotiating over Ontological Correspondences with Asymmetric and Incomplete Knowledge*. AAMAS-14.

See also:

Laera *et al.* *Argumentation over ontology correspondences in MAS*. AAMAS-07.

Negotiated Ontology Alignment

Each agent is equipped with a private knowledge base of correspondence. Each agent has an estimate of the others "which reflects the maximum degree of belief an agent has in its undisclosed correspondences"

c	K_c^{Alice}	K_c^{Bob}	joint(c)
$\langle a, x\rangle$	0.8	0.6	0.7
$\langle b, x\rangle$	0.5	0.8	0.65
$\langle b, w\rangle$	0.6	0.4	0.5
$\langle b, z\rangle$	0.9	–	0.45
$\langle c, y\rangle$	–	0.2	0.1
$\langle a, z\rangle$	0.1	–	0.05

Negotiated Ontology Alignment

Negotiation proceeds in rounds.
Agents exchange moves by means of communicative acts: assert, object, accept, reject.

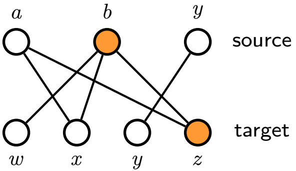

c	K_c^{Alice}	K_c^{Bob}	joint(c)
$\langle a, x\rangle$	0.8	0.6	0.7
$\langle b, x\rangle$	0.5	0.8	0.65
$\langle b, w\rangle$	0.6	0.4	0.5
$\langle b, z\rangle$	0.9	–	0.45
$\langle c, y\rangle$	–	0.2	0.1
$\langle a, z\rangle$	0.1	–	0.05

- Alice's estimate: 1, Bob's estimate: 1
- Alice picks the best correspondence $\langle b, z\rangle$ and asserts it

Negotiated Ontology Alignment

Negotiation proceeds in rounds.
Agents exchange moves by means of communicative acts:
assert, object, accept, reject.

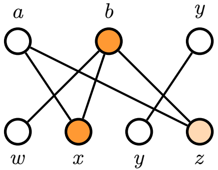

c	K_c^{Alice}	K_c^{Bob}	joint(c)
$\langle a, x\rangle$	0.8	0.6	0.7
$\langle b, x\rangle$	0.5	0.8	0.65
$\langle b, w\rangle$	0.6	0.4	0.5
$\langle b, z\rangle$	0.9	–	0.45
$\langle c, y\rangle$	–	0.2	0.1
$\langle a, z\rangle$	0.1	–	0.05

- Alice's estimate: 1, Bob's estimate: 0.9
- Bob computes that $joint(\langle b, z\rangle) = 0.45$, and thinks $\langle b, x\rangle$ may be better since $\frac{1}{2}(0.8 + 0.9) = 0.85$, thus $object((\langle b, z\rangle, 0.0), (\langle b, x\rangle, 0.8))$

Negotiated Ontology Alignment

Negotiation proceeds in rounds.
Agents exchange moves by means of communicative acts:
assert, object, accept, reject.

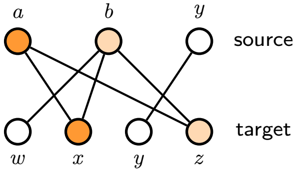

c	K_c^{Alice}	K_c^{Bob}	joint(c)
$\langle a, x\rangle$	0.8	0.6	0.7
$\langle b, x\rangle$	0.5	0.8	0.65
$\langle b, w\rangle$	0.6	0.4	0.5
$\langle b, z\rangle$	0.9	–	0.45
$\langle c, y\rangle$	–	0.2	0.1
$\langle a, z\rangle$	0.1	–	0.05

- Alice's estimate: 1, Bob's estimate: 0.9
- Alice computes that $joint(\langle b, z\rangle) = 0.65$, and thinks $\langle b, x\rangle$ may be better since $\frac{1}{2}(0.8 + 0.8) = 0.8$, thus $object((\langle b, x\rangle, 0.5), (\langle a, x\rangle, 0.8))$

Negotiated Ontology Alignment

Negotiation proceeds in rounds.
Agents exchange moves by means of communicative acts: assert, object, accept, reject.

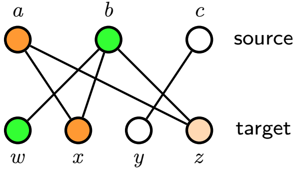

c	K_c^{Alice}	K_c^{Bob}	joint(c)
$\langle a, x\rangle$	0.8	0.6	0.7
$\langle b, x\rangle$	0.5	0.8	0.65
$\langle b, w\rangle$	0.6	0.4	0.5
$\langle b, z\rangle$	0.9	–	0.45
$\langle c, y\rangle$	–	0.2	0.1
$\langle a, z\rangle$	0.1	–	0.05

- Alice's estimate: 1, Bob's estimate: 0.9
- Bob computes that $joint(\langle b, w\rangle) = 0.75$, and thus $assert(\langle b, w\rangle, 0.4)$, and so on.

Thank you for your attention!