Bilateral Negotiation

An overview of the theory of negotiation & negotiating software agents

Almen, HTI Winterschool
Koen Hindriks
January 29, 2009

Outline

- A Classic Tale: Negotiating About an Orange
- Preferences and Utility Theory
- Bargaining Models and Solution Concepts
- Strategies and Concession Tactics
- Opponent Modeling & Learning
- Other Topics, or: What I Did Not Talk About
- The Future of Automated Negotiation
- Negotiation Session

A Classic Tale: Two Sisters Arguing Over an Orange*

Two sisters arguing over a *single available* orange:

Key Concepts:
- **Conflict of interests**
- **Distributional negotiation problem**
- **Positional bargaining**

* The story about the sisters' conflict over the orange has been attributed to Mary Parker Follett. See Deborah M. Kolb, 1995, The Love for Three Oranges, Or: What Did We Miss about Ms. Follett in the Library? 11 Negotiation J. 339. Also see: Roger Fisher & Danny Ertel, 1995, Getting Ready to Negotiate.

A Classic Tale: Two Sisters Arguing Over an Orange

Two sisters arguing over a *single available* orange:

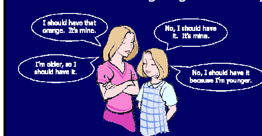

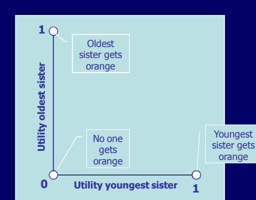

Key Concepts:
- **Outcome space (Ω)**
- **Status Quo**
- **Utility**

A Classic Tale: Two Sisters Arguing Over an Orange

The arguments are not decisive, but there is an alternative:

Key Concepts:
- **Making a concession, compromise**
- **Expand outcome space**
- **Fair outcome**

A Classic Tale: Two Sisters Arguing Over an Orange

The arguments are not decisive, but there is an alternative:

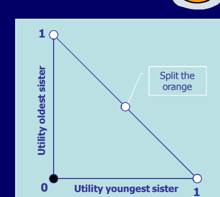

Key Concepts:
- **Pareto optimal outcome**
- **Pareto frontier**

A Classic Tale: Two Sisters Arguing Over an Orange

But why do both sisters want the orange?

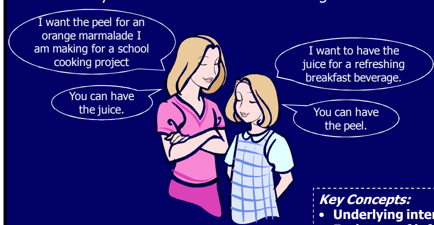

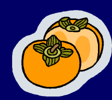

Key Concepts:
- **Underlying interests**
- **Exchange of information**

A Classic Tale: Two Sisters Arguing Over an Orange

But why do both sisters want the orange?

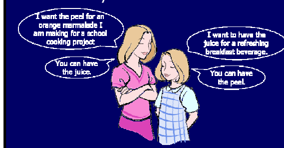

Key Concepts:
- **Inefficient outcome**
- **Pareto dominated**
- **Utopia, or Ideal point**

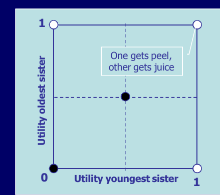

Summary

- A negotiation is aimed at resolving a **conflict of interests**.
- Each negotiation is defined by its **outcome space**. Ideally, an agreement is reached that is Pareto optimal.
- During a negotiation, the outcome space is **explored**, and sometimes, may be **expanded** to find creative agreements based on the underlying interests.

Literature:
- H. Raiffa, 1982, *The Art and Science of Negotiation*, Belknap Press
- L.L. Thompson, 2005 (3rd ed.), *The Mind and Heart of the Negotiator*, Prentice Hall

Preferences and Utility Theory

Von Neumann-Morgenstern Utility Theory

December 11, 2007

Where Do Utilities Come From?

- It is assumed that agent have preferences for outcomes, and thus are not completely indifferent about various outcomes.
- If preferences of an agent are both *rational* as well as *well-behaved* in a precise sense, then it can be shown that a utility function can be constructed that can be used for guiding the decision-making of that agent.

Predictable versus Unpredictable Preferences

Predictable preferences
- *Most people* prefer to have more rather than less money.

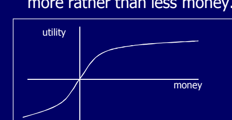

Monotonic preference
Logarithmic curve for money

- **Higher > Lower Grade**
- **Less > More Work**

Unpredictable preferences
- I like blue more than red, you like red more than blue.

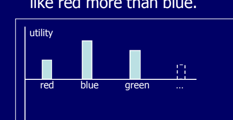

Varies from person to person
Not quite as well-behaved

- **Peugeot > Ford**
- **Swim > Cycle Ride**

Rational Preferences

***Preferences can be modeled by a (binary)* preference relation $\succeq$*. A preference relation is called* rational *if it is* complete *and* transitive.**

Preferences are **complete**:
- It is assumed that an agent can state his preferences when confronted with two arbitrary potential outcomes.
- Formally, for all $\omega,\omega' \in \Omega$: either $\omega \succeq \omega'$, or $\omega' \succeq \omega$, or $\omega \sim \omega'$.

Preferences are **transitive**:
- If an outcome ω_1 is preferred over ω_2, and ω_2 is preferred over ω_3, then ω_1 is preferred over ω_3.
- Formally, for all $\omega_1, \omega_2, \omega_3 \in \Omega$: $\omega_1 \succeq \omega_2$ and $\omega_2 \succeq \omega_3 \Rightarrow \omega_1 \succeq \omega_3$.

Intransitive Preferences & Money Pumps

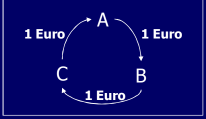

TU Delft

Modeling Risky Outcomes: Lotteries

***Outcomes that involve risk can be modeled by* lotteries. *For each lottery, it is assumed there is a* certainty-equivalent outcome.**

- A lottery associates chances or probabilities with outcomes. Notation: L = [p_1, ω_1 ; … ; p_n, ω_n].

Preferences are **continuous**:
- If an outcome ω′ is between ω and ω″ in preference, then there is a probability *p* such that an agent is indifferent between the lottery [p, ω; 1-p, ω″] and getting ω′ for sure.*
- Formally, if $\omega \succ \omega' \succ \omega''$, then there is a *p* s.t. [p, ω; 1-p, ω″] ~ ω′.

Example: Deal or No Deal.
- For a lottery involving monetary prizes, e.g.
  [¼, 20; ¼, 100; ¼, 20.000; ¼, 40.000]
  there is an amount of money, e.g. 5.000, such that:
  [¼, 20; ¼, 100; ¼, 20.000; ¼, 40.000]~5.000.

* ω′ is called the **certainty-equivalent outcome** for lottery [p, ω; 1-p, ω″].

TU Delft

Preferences are Monotone

Preferences are assumed to be* monotone. *Intuitively, this assumption expresses that getting more of what you want is preferred over less.

Preferences are m**onotone**:
- If an outcome ω is preferred over ω′, then an agent prefers a lottery with outcomes ω, ω′ that has a higher chance of getting ω over one with a lower chance of getting ω.
- Formally, [p, ω; (1-p), ω′] $\succeq$ [q, ω; (1-q), ω′] whenever p≥q.

Exceptions to monotonicity?
Most people prefer life to death, but still some people like to exercise dangerous sports, such as BASE jumping, heli-skiing, diving, mountaineering, big wave surfing and bull riding. The chances of getting killed are significantly higher than staying at home listening to music, drinking a beer, or reading a good book.

TU Delft

Equally Valued Outcomes

Equally valued outcomes may be* substituted *for each other.

Equally Valued Outcomes can be **substituted**:
- If an agent is indifferent between ω and ω′, then that agent is indifferent between a lottery that involves outcome ω over that same lottery with ω′ substituted for ω.
- Formally, [p, A; 1-p, C] ~ [p, B; 1-p, C] whenever A~B.

The intuition here is that indifference with respect to two outcomes is independent from preferences regarding other outcomes.

TU Delft

No Fun in Gambling

Lotteries that involve other lotteries can be reduced to simpler, non-composed lotteries.

Outcomes can be **decomposed:**
- Compound lotteries can be reduced using the laws of probability. E.g., [p, A; [q, B; 1-q, C]] ~ [p, A; (1-p)q, B; (1-p)(1-q), C]

The axiom has also been paraphrased as stating that there can be "no fun in gambling". That is, an agent is indifferent between:
- Playing a lottery in which prizes are tickets for a 2^{nd} lottery,
- Playing a single lottery that directly pays out cash

provided, off course, chances of winning cash are the same.

TU Delft

Utility Theory

The axioms of utility theory guarantee the existence of a utility function $U: \Omega \rightarrow \Re$ that represents a rational, well-behaved preference relation.

Theorem.
Suppose a preference relation $\succeq$ satisfies all conditions listed above. Then there exists a utility function $U: \Omega \rightarrow \Re$ such that:

$$U(\omega) \geq U(\omega') \text{ iff } \omega \succeq \omega$$

Example: Deal or No Deal:
- Suppose [¼, 20; ¼, 100; ¼, 20.000; ¼, 40.000]~5.000 and the preference relation satisfies all conditions above. Then the existence of a utility function *U* is guaranteed such that:
- *¼U(20)+¼U(100)+ ¼U(20.000)+ ¼U(40.000)= U(5.000).*

TU Delft

Buying a Supercomputer

In practice, in complex negotiations, it is hard to construct a good utility function that matches the preferences of a party.

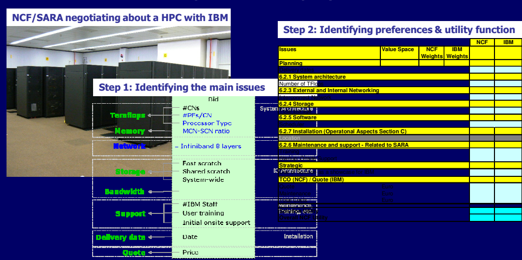

TU Delft

Summary

- A **preference relation** is a binary relation over an outcome space that models which outcomes are preferred over others.
- Preferences are **rational** if they are transitive and complete.
- A unique (Von Neumann-Morgenstern) utility function exists that models rational preferences that additionally satisfy continuity, monotonicity, substitutability, and decomposability.

Literature:

- R.D. Luce and H. Raiffa, 1957, *Games and Decisions: Introduction and a Critical Survey*, Dover Publications
- R.L. Keeney and H. Raiffa, 1976, *Decisions with Multiple Objectives*, Cambridge University Press

TU Delft

Bargaining Models and Solution Concepts

Axiomatic Foundations for Negotiation

December 11, 2007

TU Delft
Delft University of Technology

Abstract Model of Negotiation

Two parties that negotiate need to make a joint decision to settle on a final outcome. Can we predict the outcome that such agents will agree on?

- A **bargaining problem** is a pair $\langle \Omega, d \rangle$ with Ω an outcome space, and d the disagreement point.
- A **solution** for bargaining problems of the form $\langle \Omega, d \rangle$ is an outcome $\omega = F(\Omega, d) \in \Omega$.
- $F(\Omega, d)$ can be interpreted as a **prediction**, or **recommendation**, for the problem $\langle \Omega, d \rangle$.

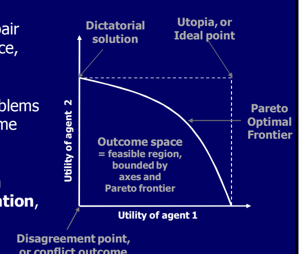

TU Delft

Which Solution Should We Select?

- Characterize desirable properties of solution.
- Ideally, but not necessarily, this would give a *unique* solution point for each bargaining problem.
- Axiomatic approaches to define properties.
- We discuss the axioms proposed by Nash (1950).

TU Delft

Axiom 1: All Gains Should Be Exhausted

- **Pareto-Optimality**: A solution $\omega = F(\Omega, d)$ should be Pareto-optimal, i.e. such that there is no $\omega' \in \Omega$ and $\omega' > \omega$.

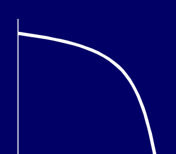

*Most solution concepts proposed satisfy this property.**

* The Egalitarian solution satisfies *weak Pareto optimality*, i.e. there is no outcome that is preferred by all negotiators to the Egalitarian solution. A natural extension of this solution concept that does satisfy Pareto optimality exists.

TU Delft

Axiom 2: Fully Symmetric Bargaining Problems

- **Symmetry**: If negotiators have completely symmetric roles, the solution should yield them equal utility, i.e. the solution $F(\Omega,d)=\langle x_1, x_2\rangle$ should satisfy $x_1=x_2$.

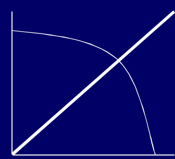

Does not allow for external factors that are not modeled.

TU Delft

Axiom 3 Scale Invariance

- **Scale Invariance**: The solution should be invariant under "scale transformations", i.e. for all positive linear transformations λ we have $\lambda(F(\Omega,d))=F(\lambda(\Omega), \lambda(d))$

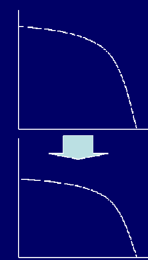

Does not allow for interpersonal comparison of utilities.

TU Delft

Axiom 4 Independence of Irrelevant Alternatives

- **Contraction Independence**: If $\Omega' \subseteq \Omega$ and $F(\Omega,d) \in \Omega'$, then $F(\Omega',d)=F(\Omega,d)$.

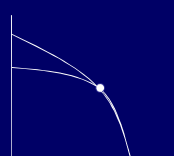

Does not allow to take into account that one negotiator has more or less options available to choose from in his favor.

TU Delft

The Nash Solution to a Bargaining Problem

- It can be shown that there is a *unique* outcome that satisfies all axioms proposed by Nash, called the **Nash solution**.
- The **Nash solution** is that outcome $\omega=\langle x_1, x_2\rangle \in \Omega$ that maximizes $x_1 \cdot x_2$.

TU Delft

Many Alternative Solutions Proposed

Best Known Solutions:	Many Other Solutions:
• **Nash Solution:** outcome $\omega=\langle x_1, x_2\rangle$ that maximizes $x_1 \cdot x_2$	• Dictatorial Solution • Raiffa Solution • Kalai-Rosenthal Solution
• **Kalai-Smorodinsky Solution:** Pareto-optimal outcome ω on line that connects disagreement and ideal point	• Lexicographic Egalitarian • Perles-Maschler Solution • The Equal Area Solution
• **Egalitarian Solution:** Weak Pareto-optimal outcome $\omega=\langle x_1, x_2\rangle$ such that $x_1=x_2$	• The Equal Loss Solution • The Utilitarian Solution • The Yu Solution

TU Delft

Summary

- A **bargaining problem** consists of an outcome space and a disagreement point
- A **solution** to a bargaining problem recommends a unique outcome as the final agreement
- Nash proposed to **axiomatically characterize** solution concepts
- Many solution concepts proposed in literature

Literature:

- R.D. Luce and H. Raiffa, 1957, *Games and Decisions: Introduction and a Critical Survey*, Dover Publications
- W. Thomson, 1994, *Cooperative Models of Bargaining*, in: Handbook of Game Theory, Volume 2, Elsevier

TU Delft

Strategies and Concession Tactics

Alternating Offers Protocol with Incomplete Information

December 11, 2007

Joint Exploration of Outcome Space

The Nash solution, and, more importantly, the Pareto frontier, cannot be computed when both parties only have incomplete preference information.

- Axiomatic approaches typically characterize possible outcomes of a negotiation, assuming **complete information**.
- In **closed negotiations**, negotiators at best have partial information about their opponent's preferences.
- For a good reason: Revealing information to an opponent may result in **exploitation**.
- Reaching an agreement requires negotiators with incomplete information to jointly **explore** the outcome space by **exchanging offers**, i.e. outcomes proposed to the other party.

Fundamental Problem of Negotiation

A dilemma faced by any negotiator is how to simultaneously achieve two typically conflicting goals: maximize own outcome and chance of a deal.

What is the optimal strategy to achieve the following:

- **Maximize own outcome**.
  Requires outcome with high own utility, i.e. that is individually rational. *Rule of Thumb*: Start negotiation and check whether offer that is best for self is acceptable.
- **Maximize chance of reaching an agreement**.
  Requires outcome with acceptable utility for opponent, i.e. resolving the conflict of interest.

Negotiation Decision functions

The complexity of negotiation with incomplete information does not allow for establishing game-theoretic optimal equilibrium strategy results. Instead negotiation heuristics are proposed in the literature.

Different families of negotiation heuristics:

- Time-dependent tactics
- Behaviour-dependent tactics

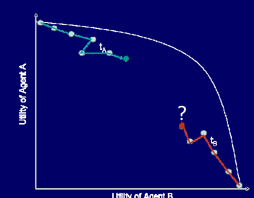

A tactic is used to determine a next offer.

Time Dependent Tactics

Time dependent tactics take deadlines and discounting of future rewards into account.

- A time dependent tactic determines the next offer to be proposed as a function $\alpha^a(t)$ of time:

$$x^t_{a\rightarrow b} = \begin{cases} min^a + \alpha^a(t)(max^a - min^a) & \text{if } V^a \text{ decreasing} \\ min^a + (1-\alpha^a(t))(max^a - min^a) & \text{if } V^a \text{ increasing} \end{cases}$$

Various time-dependent functions are used:

- A polynomial function: $\alpha^a(t) = k^a + (1-k^a)\left(\frac{\min(t, t^a_{\max})}{t^a_{\max}}\right)^{\frac{1}{\beta}}$
- An exponential function: $\alpha^a(t) = e^{(1-\frac{\min(t, t^a_{\max})}{t^a_{\max}})^{\beta}\ln k^a}$

Time Dependent Functions

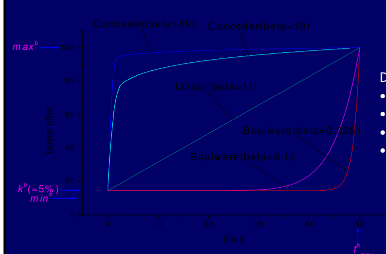

Decision Parameters:
- Reservation values: min^s, max^b
- Initial offer level: k
- Deadline (time available): t_{max}
- Rate of concession: $\beta/\mu/\gamma/\delta$,…

Behavior Dependent Tactics

Behavior dependent tactics take the behavior of the opponent into account measured relative to the agent's own utility space, giving rise to tit-for-tat strategies.

- A behavior dependent tactic determines the next offer to be proposed as a function $\psi(t_n)$ of the last proposed offer t_n:

$$x_{a \to b}^{t_{n+1}} = \min(\max(\psi(t_n), min^a), max^a)$$

Various functions $\psi(t_n)$ are used:

- Average tit-for-tat: $\psi(t_n) = \frac{x_{b \to a}^{t_{n-2\gamma}}}{x_{b \to a}^{t_n}} x_{a \to b}^{t_{n-1}}$
- Relative tit-for-tat: $\psi(t_n) = \frac{x_{b \to a}^{t_{n-2\delta}}}{x_{b \to a}^{t_{n-2\delta+2}}} x_{a \to b}^{t_{n-1}}$

Monotone Concession Protocol

A monotone concession protocol requires agent's to make concessions at each step, i.e. increase their opponent's utility at each negotiation move.

- Rule: Players are not allowed to make offers which have a lower utility for their opponent than their last offer.
- If the minimum concession is required to be >0, then a negotiation is guaranteed to terminate.
- Agents **accept** offer on table if utility exceeds that of own last offer or that of a viable counter-offer.
- *Example*: Zeuthian Strategy (cf. literature).

Note:
- In multi-issue negotiations with incomplete information about opponent preferences, it is not always possible to make monotonic concessions (due to "compatible issues")
- Neither is it best to monotonically make concessions, as we will see later.

Smart Meta-Strategy

Using domain knowledge it may be possible to improve the efficiency of a strategy to the benefit of the agent itself as well as of its opponent.

- *Key idea*: To better meet the demands of an opponent it is not always required to concess utility oneself.
- Propose alternative offer that is more *similar* to that of opponent's last offer with same own utility.
- Only concess if not possible.

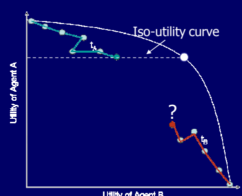

- *Result*: Issue trade-offs. *Example:* Propose ⟨Medium Quality ↓, Low Price ↑, 10 days =⟩ instead of ⟨High Quality, Medium Price, 10 days⟩
- See: Trade-Off Strategy (cf. literature)

Evaluating Strategies

A number of criteria are available to judge the quality of a strategy.

Process-Oriented:
- Cost-Effective: How many steps are required to reach agreement?
- Robustness: Is it possible to exploit the strategy?
- Negotiation move types: Is strategy sensitive to opponent?

Outcome-Oriented:
- Successful: Is an agreement reached?
- Efficiency: Is agreement Pareto optimal?
- Fairness: Is agreement close to Nash? Kalai-Smorodinsky?

Evaluation: Dynamics of Negotiation Process

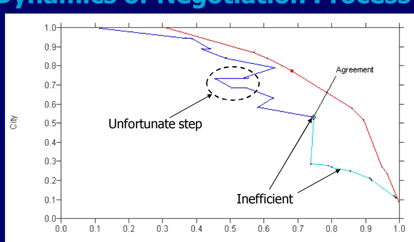

Trade-Off versus ABMP strategy on the City-vs-AMPO domain (Raiffa)

Evaluation: Dynamics of Negotiation Process

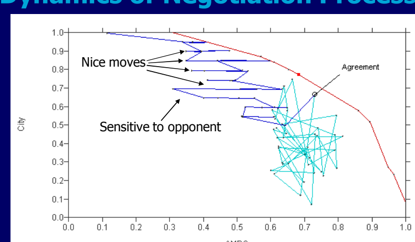

Trade-Off versus Random Walker strategy on the City-vs-AMPO domain

Summary

- Due to incomplete information about opponent preferences, it is not possible to define "optimal" strategies in any precise sense.
- Various negotiation heuristics have been proposed, including **time** and **behavior dependent tactics**.
- Efficiency can be improved by using domain knowledge and proposing an offer on a utility iso-curve more similar to the opponent's last offer.

Literature:

- P. Faratin, C. Sierra, N.R. Jennings, 1997, *Negotiation Decision Functions for Autonomous Agents*, in: *Int. Journal of Robotics and Autonomous Systems*, 24(3-4), 159-182
- P. Faratin, C. Sierra, N.R. Jennings, 2002, *Using Similarity Criteria to Make Issue Trade-Offs in Automated Negotiations*, in: Artificial Intelligence, 142, 205-237
- K.V. Hindriks, C.M. Jonker, D. Tykhonov, 2007, Analysis of Negotiation Dynamics, in: Cooperative Information Agents XI, LNCS 4676, 27-35
- J.S. Rosenschein, G. Zlotkin, 1994, Rules of Encounter, MIT Press

TU Delft

Opponent Modeling and Learning

Improving negotiated agreements by estimating opponent preferences.

December 11, 2007

TU Delft
Delft University of Technology

Learning Opponent Preferences

Various features of an opponent's strategy can be learned, e.g. type of concession tactic, reservation value. We discuss learning opponent preferences.

- The strategies and tactics discussed so far are blind for opponent preferences.
- Various techniques have been proposed to learn features and construct an **opponent model**.
- Can information about **opponent preferences** be learned in single-instance negotiation?
  - Single-instance negotiation typical for e-commerce.
  - Hard because only bids exchanged can be used, no previous history can be assumed.

TU Delft

Learning: Structural Assumptions

In order to be able to learn opponent preferences we have to exploit certain structural features of negotiation problems.

- Linear additive utility functions: $u(b_t) = \sum_{i=1}^{n} w_i e_i (x_i \in b_t)$
- Sufficient to learn ranking of weights: $w_i = 2\frac{r_i^j}{n(n+1)}$
- Shape of Evaluation functions (utility of issues):

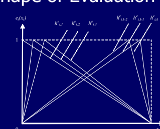

Downhill,
Uphill, and
Triangular

TU Delft

Learning: Rationality Assumptions

In order to be able to learn opponent preferences we have to assume agents are (at least to some extent) rational.

- Opponent starts by proposing its best bid.
- Opponent uses a concession-based tactic

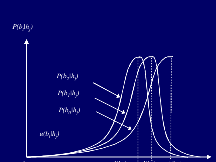

- Probability distributions are associated with possible tactics.

TU Delft

Bayesian Learning

Each (new) bid provides new information that is used to update the model of the opponent's preferences using Bayes' rule.

- Using Bayes' Rule the probabilities associated with hypotheses regarding weights and evaluation functions can be updated for every new bid b_t:

$$P(h_j|b_t) = \frac{P(h_j)P(b_t|h_j)}{\sum_{k=1}^{m} P(h_k)P(b_t|h_k)}$$

- An estimate of the utility an opponent associates with a bid can be computed using the computed probabilities:

$$\overline{u}(b_t) = \sum_{j=1}^{|H|} P(h_j) \sum_{i=1}^{n} w_i e_i (x_i \in b_t)$$

TU Delft

Using an Opponent Model

An opponent model can be used to choose a next offer in bilateral negotiation.

An opponent model:

- enables choosing an offer that has utility x for the negotiator and maximizes the utility of the opponent (which helps to increase the chance of acceptance)
- enables to select an offer that increases the utility of both parties if previous offers have deviated from the Pareto frontier, which is typical. I.e. it allows making a **fortunate move**.

Defining an Opponent-based Strategy

Strategy using opponent model:

- Use a tit-for-tat tactic, i.e. match negotiation move that opponent just made.
- Propose Pareto-efficient offers, i.e. go "straight up" (agent A) to the Pareto frontier in order to maximize own utility.
- Many possible variants.

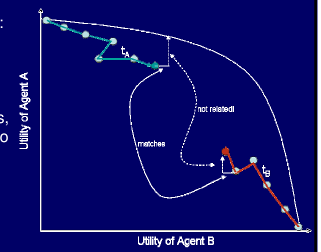

Summary

- It is possible to learn opponent preferences even in single-instance negotiations by making some rationality and structural assumptions.
- Efficiency can be improved by using learnt opponent preferences and by proposing an offer that increases utility of both parties.

Literature:

- D. Zeng, K. Sycara, 1998, *Bayesian Learning in Negotiation*, in: International Journal of Human–Computer Studies 48, 125–141
- R. Lin, S. Kraus, J. Wilkenfeld, J. Barry, 2006, *An Automated Agent for Bilateral Negotiation with Bounded Rational Agents with Incomplete Information,* ECAI 2006, 270-274
- K.V. Hindriks, D. Tykhonov, 2007, Opponent Modelling in Automated Multi-Issue Negotiation Using Bayesian Learning, Submitted.

Other Topics, or: What I Did Not Talk About

Automated negotiation is a very active and broad area of research.

December 11, 2007

Other Topics

Other Topics

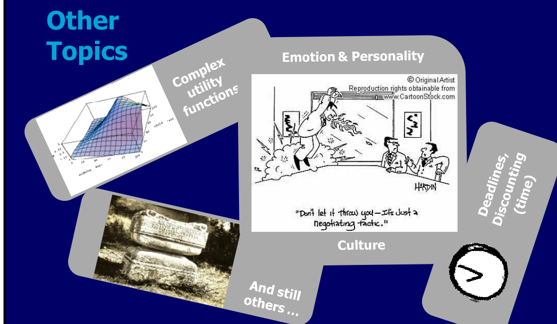

Summary

- **Topics discussed** include: Decentralized, Bilateral (one-to-one) & closed versus open negotiation. Multi-issue, additive utility functions. Alternating offers protocols without time pressure. Automated negotiating software agents. Human-machine negotiation.
- **Topics not discussed** include: Multi-lateral negotiation. Auctions. Human negotiation, emotion, personality & culture. Argumentation-based and qualitative models of negotiation. Deadlines. Complex, non-linear utility functions.

Literature:

- V. Krishna, 2002, *Auction Theory*, Academic Press
- I. Rahwan, S.D. Ramchurn, N.R. Jennings, P. McBurney, S. Parsons, L. Sonenberg, 2003, *Argumentation-based Negotiation*, in: Knowledge Engineering Review, 18(4), 343-375.
- L.L. Thompson, 2005 (3rd ed.), *The Mind and Heart of the Negotiator*, Prentice Hall

December 11, 2007 55

TU Delft

The Future of Automated Negotiation

Automated negotiation has potential applications in the near future in electronic commerce and marketplaces, and negotiation support systems.

December 11, 2007

56

TU Delft
Delft University of Technology

The Future of Automated Negotiation

Research Challenges

Strategies:
- Optimal strategies for negotiation with incomplete information.

Learning:
- negotiation tactics, opponent models, deadlines, …

Qualitative negotiation models:
- Qualitative preference models, extensions of the alternating offers protocol that allow for additional forms of information exchange.

Preference modeling:
- Elicitation techniques, utility spaces with issue dependencies.

December 11, 2007 57

TU Delft

The Future of Automated Negotiation

Applications

Negotiation support systems:
- Human-machine interaction,
- Integration of emotion & personality, negotiation styles, culture
- Analyzing domains and real-world cases

Electronic Marketplaces:
- Online negotiation, shopping bots (pre-negotiation)
- Recommender systems, personalized negotiation assistance, information selection & provision.
- Empirical analysis.

December 11, 2007 58

TU Delft