

Using Gaussian Processes to Optimise Concession in Complex Negotiations against Unknown Opponents

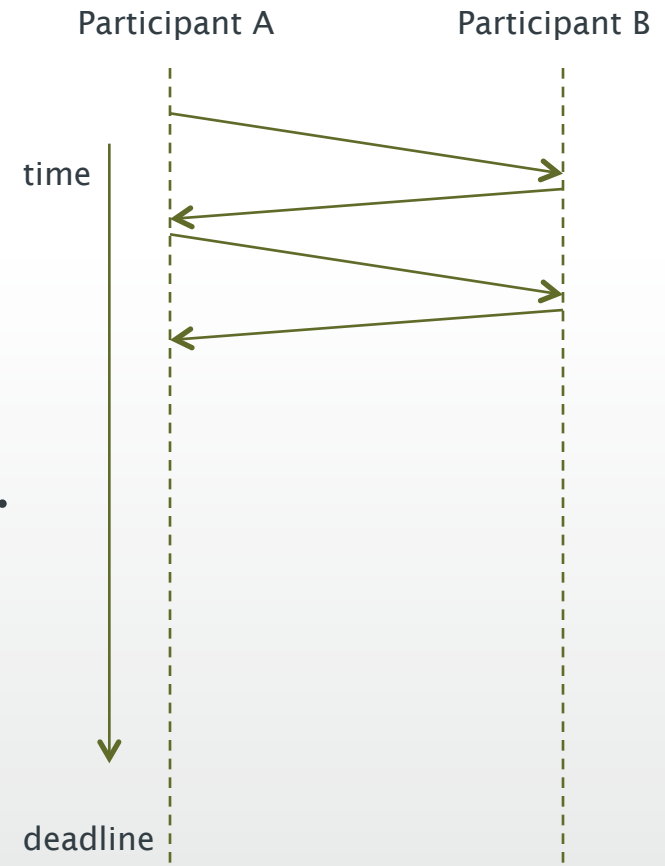
Colin R. Williams, Valentin Robu, Enrico H. Gerding, Nicholas R. Jennings
19 July 2011

Motivation

- Goal is to design a negotiation agent which can negotiate in complex environments.
- Can be used in a variety of scenarios:
 - Retail
(eg. Buying/Selling a car)
 - Service oriented negotiations
(eg. Provision of computational resources)

Complex Negotiations

- Many negotiation rounds
- Time constraints
 - Deadline, discounting factor
 - Both measured in **real-time**
- **Multiple** negotiation issues
 - High dimensionality of domains.
- Incomplete information
 - Opponent's utility function is **unknown**.
 - Opponent's behaviour is **unknown**.



Existing Approaches

- Tend to consider **less complex environments**:
 - Commonly, time constraints based on number of negotiation rounds. [Rubinstein, 1982; Coehoorn and Jennings, 2004; Nguyen and Jennings, 2003; Fatima et al., 2007]
 - Single issue negotiation. [Sandholm and Vulkan, 1999]
 - One sided uncertainty. [An et al., 2010]
 - Assume opponent belongs to a particular type.

Our Approach

- Uses a principled, decision-theoretic approach
 - Sets behaviour as a best response to the negotiation environment and opponent behaviour.

Design of IAMhaggler2011

Concession Strategy

- Used to select utility of offer to make at a given time.
- Aim is to concede:
 - **slowly** enough that the opponent does not take advantage.
 - **quickly** enough that:
 - an agreement is reached before the negotiation deadline.
 - losses due to discounting are minimised.

Concession Strategy

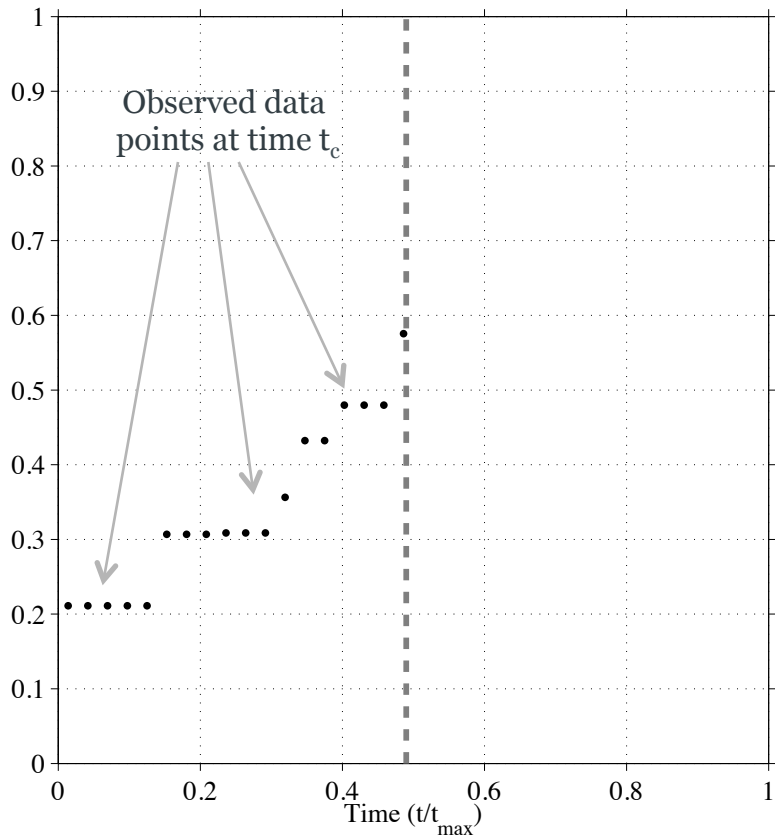
- To do this effectively, we consider:
 - The **discounting factor**
 - known
 - The **remaining time**
 - known
 - The effect of the opponent's **future concession** on our utility
 - can be estimated by learning from offers made during the current negotiation.

Gaussian Process Regression

- Use a Gaussian process regression technique in an attempt to learn the opponent's concession.
 - Mean prediction
 - Confidence measure
- Set concession rate according to this prediction.

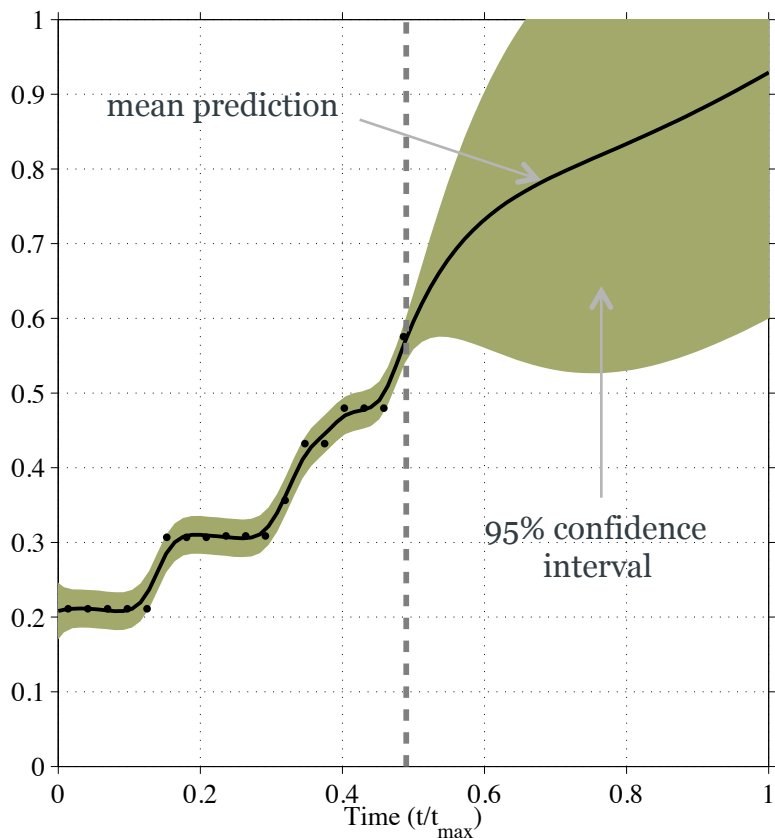
Gaussian Process Regression

- Use a Gaussian process regression technique in an attempt to learn the opponent's concession.



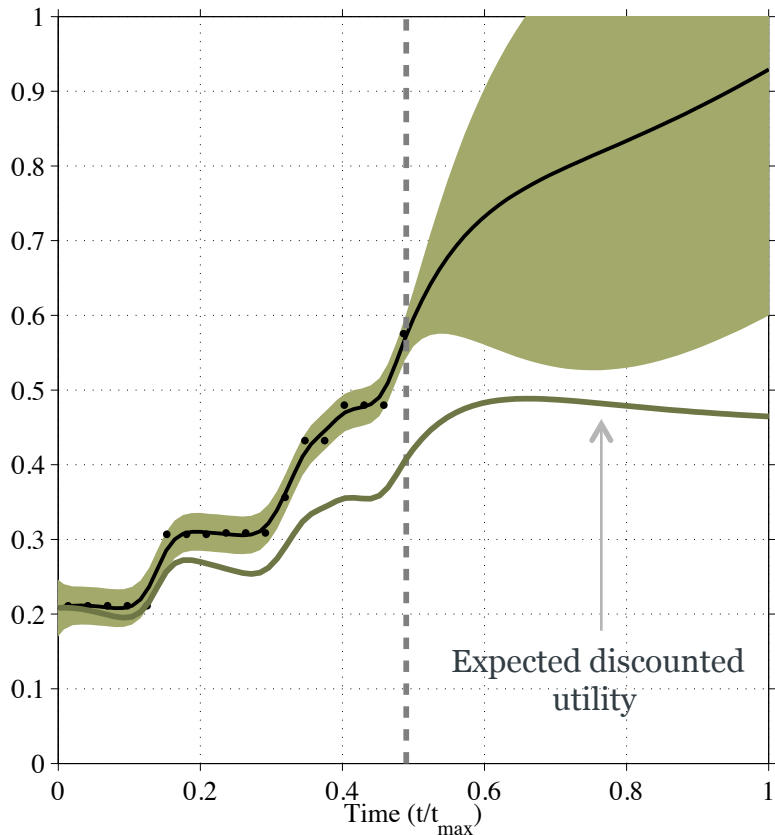
Gaussian Process Regression

- Use a Gaussian process regression technique in an attempt to learn the opponent's concession.



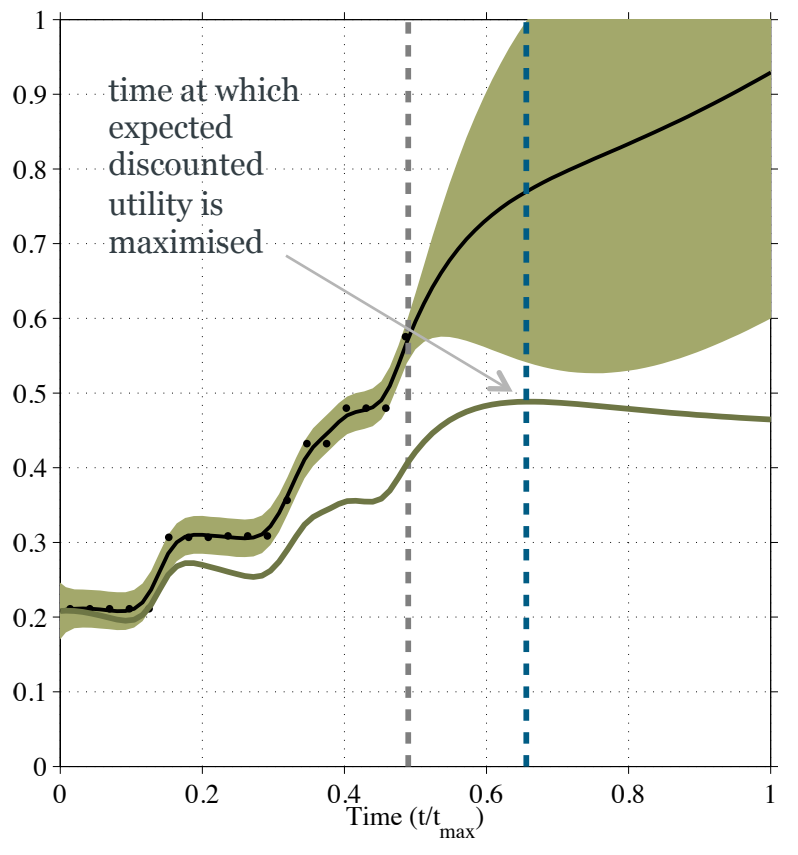
Maximise Expected Utility

- Apply **discounting** to determine expected discounted utility of our opponent's offer at time t .



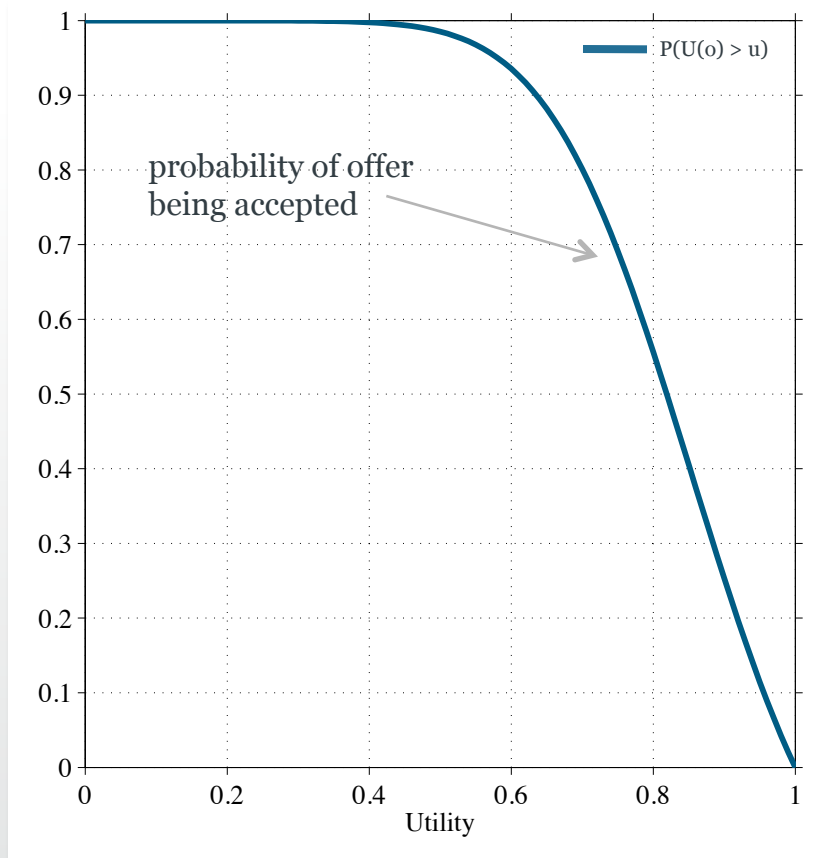
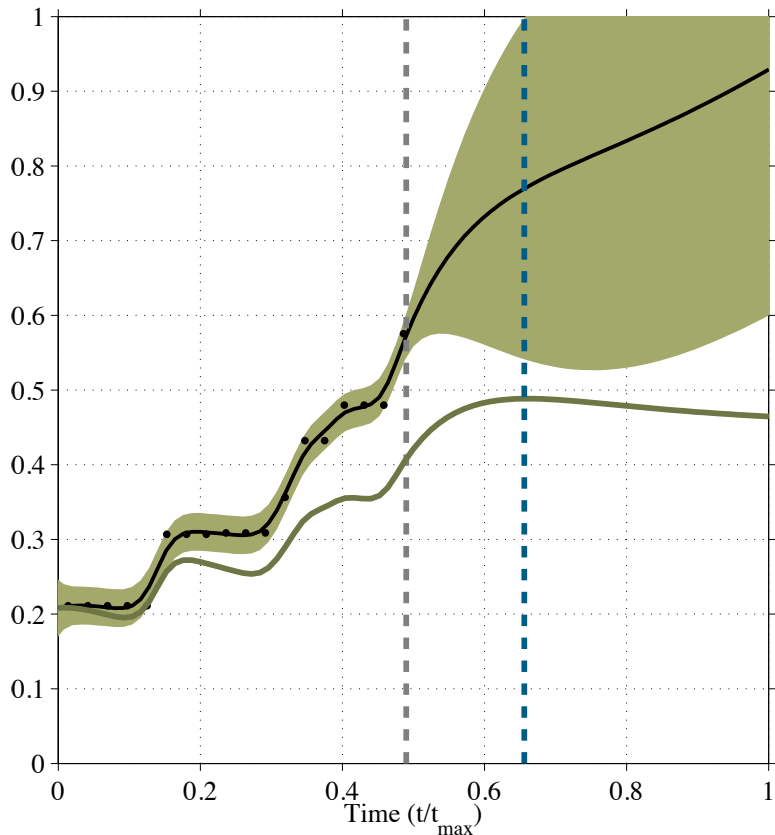
Maximise Expected Utility

- Find the **time**, t^* , at which expected discounted utility of our opponent's offer is maximised.



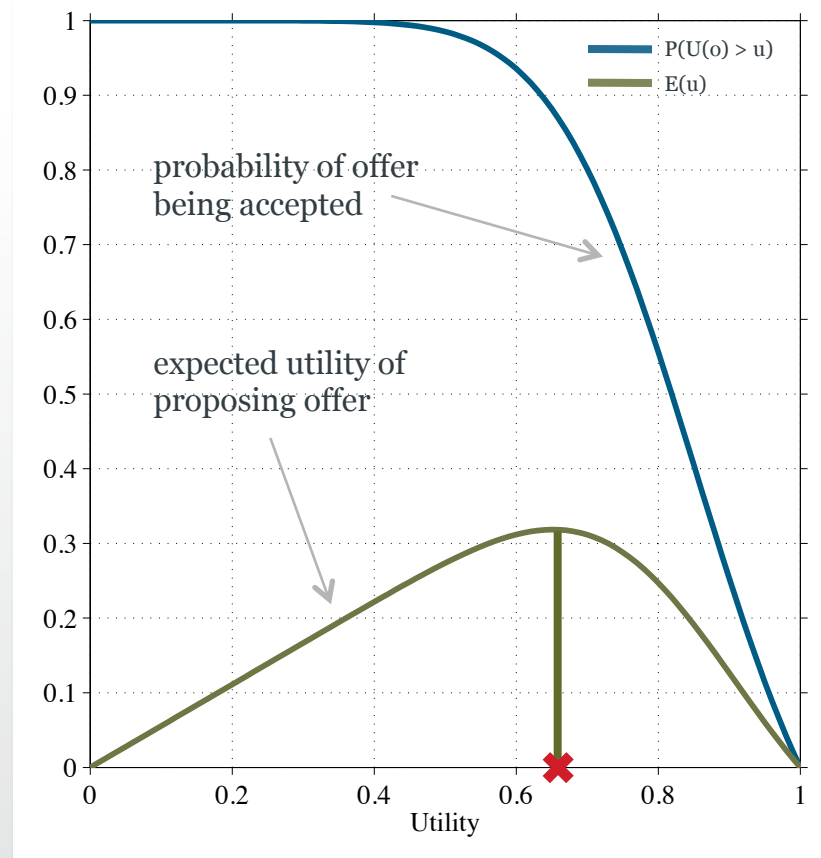

Maximise Expected Utility

- Find the **utility**, u^* , at which expected discounted utility of our offer is maximised.





Maximise Expected Utility

- Find the **utility**, u^* , at which expected discounted utility of our offer is maximised.




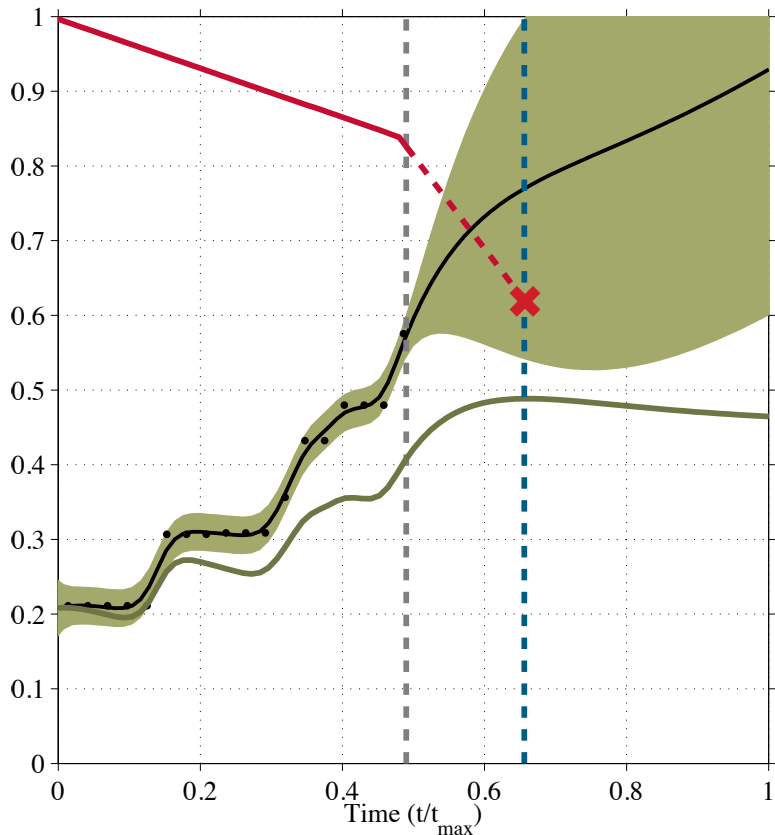
Choose Target Utility

- Concede towards $[t^*, u^*]$.




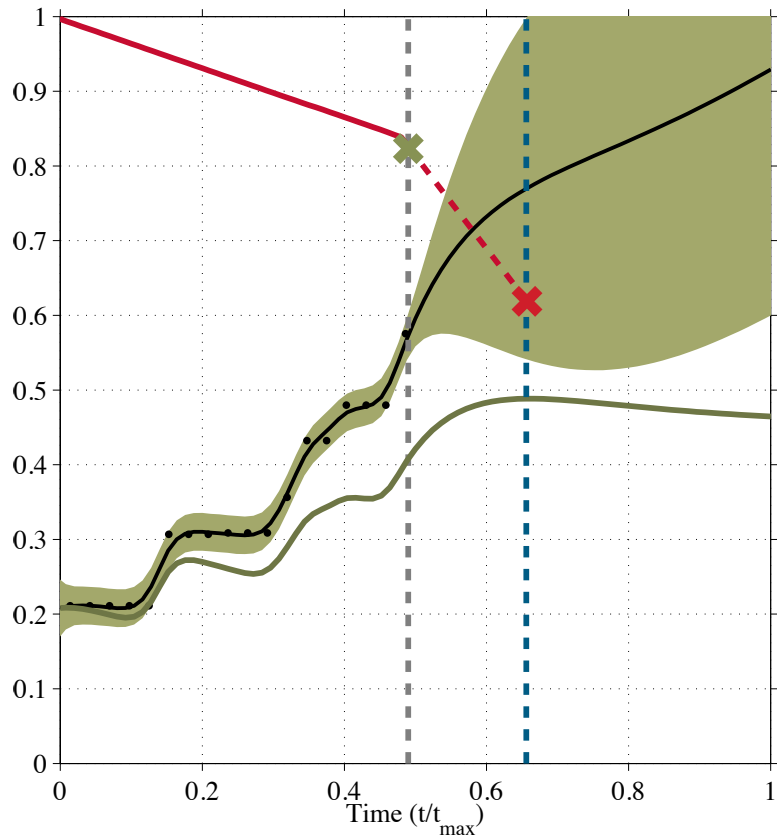
Choose Target Utility

- Concede towards $[t^*, u^*]$.



Choose Target Utility

- Concede towards $[t^*, u^*]$.

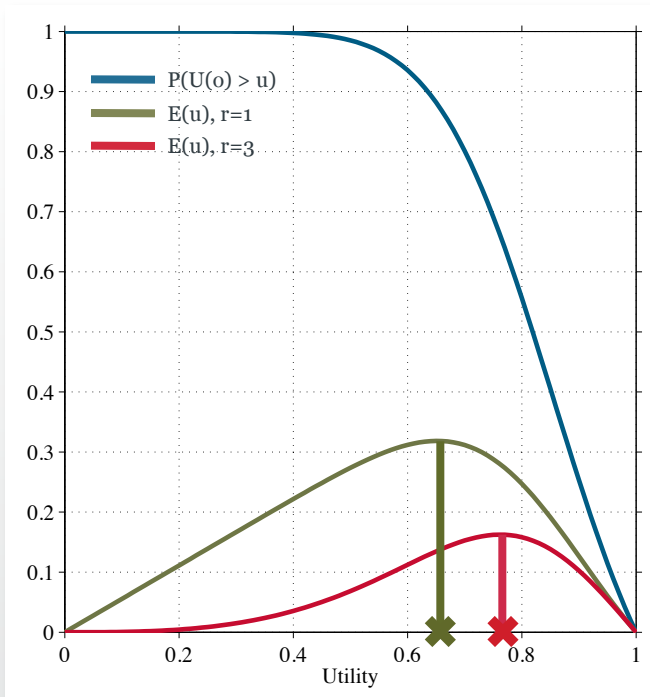


Dealing with Multiple Issues

- Select a random package, with utility close to target (according to concession strategy).
 - Fast process allows many offers to be made.
 - Encourages exploration of outcome space.

Risk Attitude

- Maximise risk-adjusted utility instead of expected utility.
- Particularly relevant in tournament, where aim is to beat opponents rather than to simply achieve high utility.



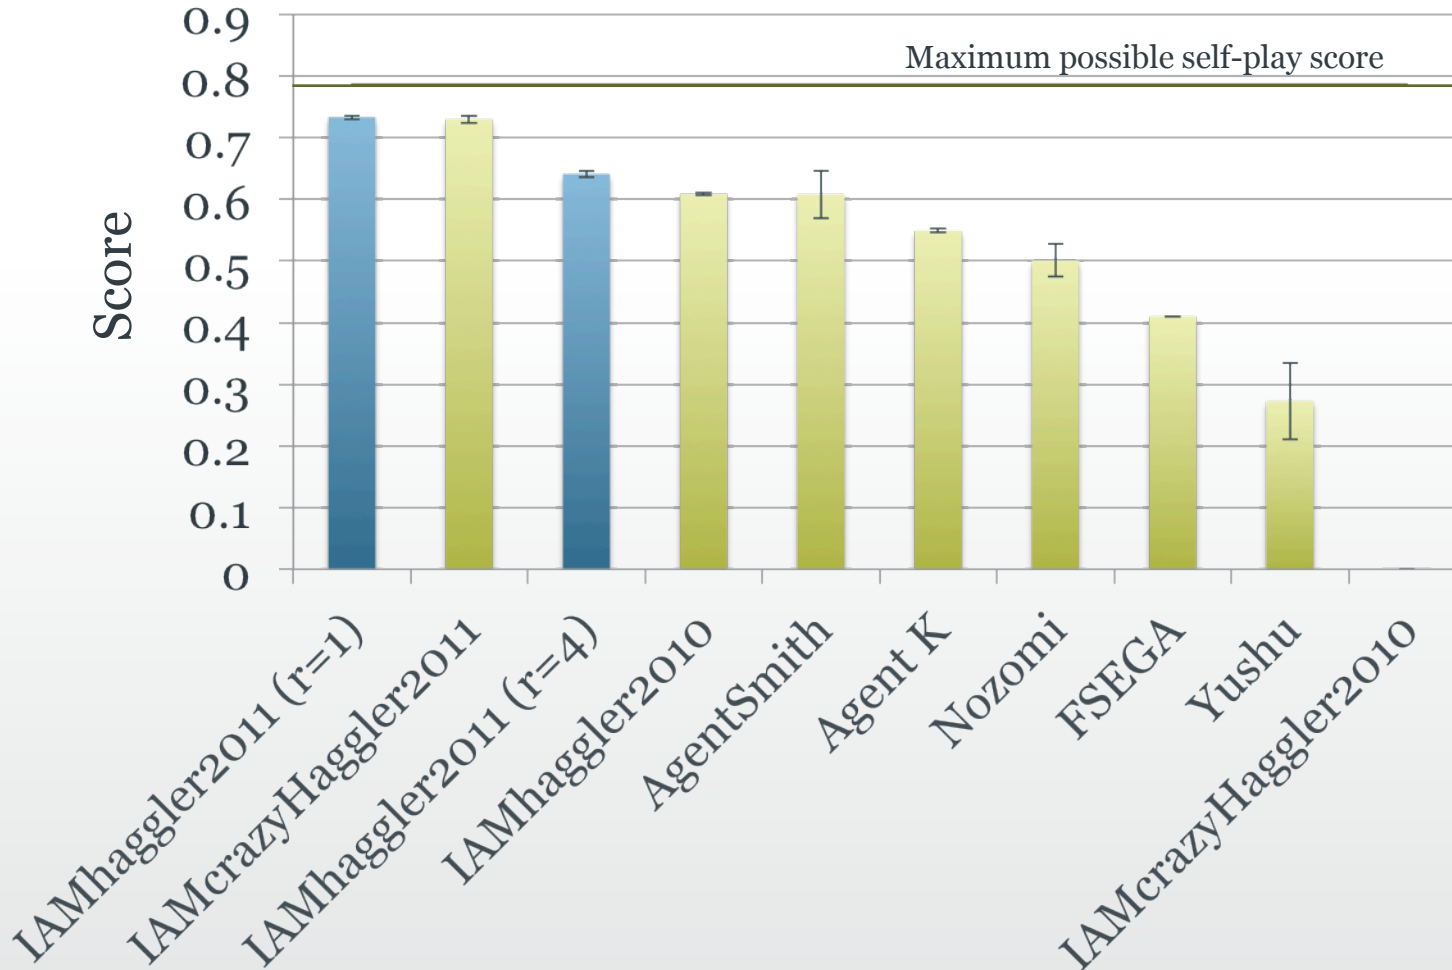
Maximisation of expected utility for different risk attitudes.

Evaluation of IAMhaggler2011

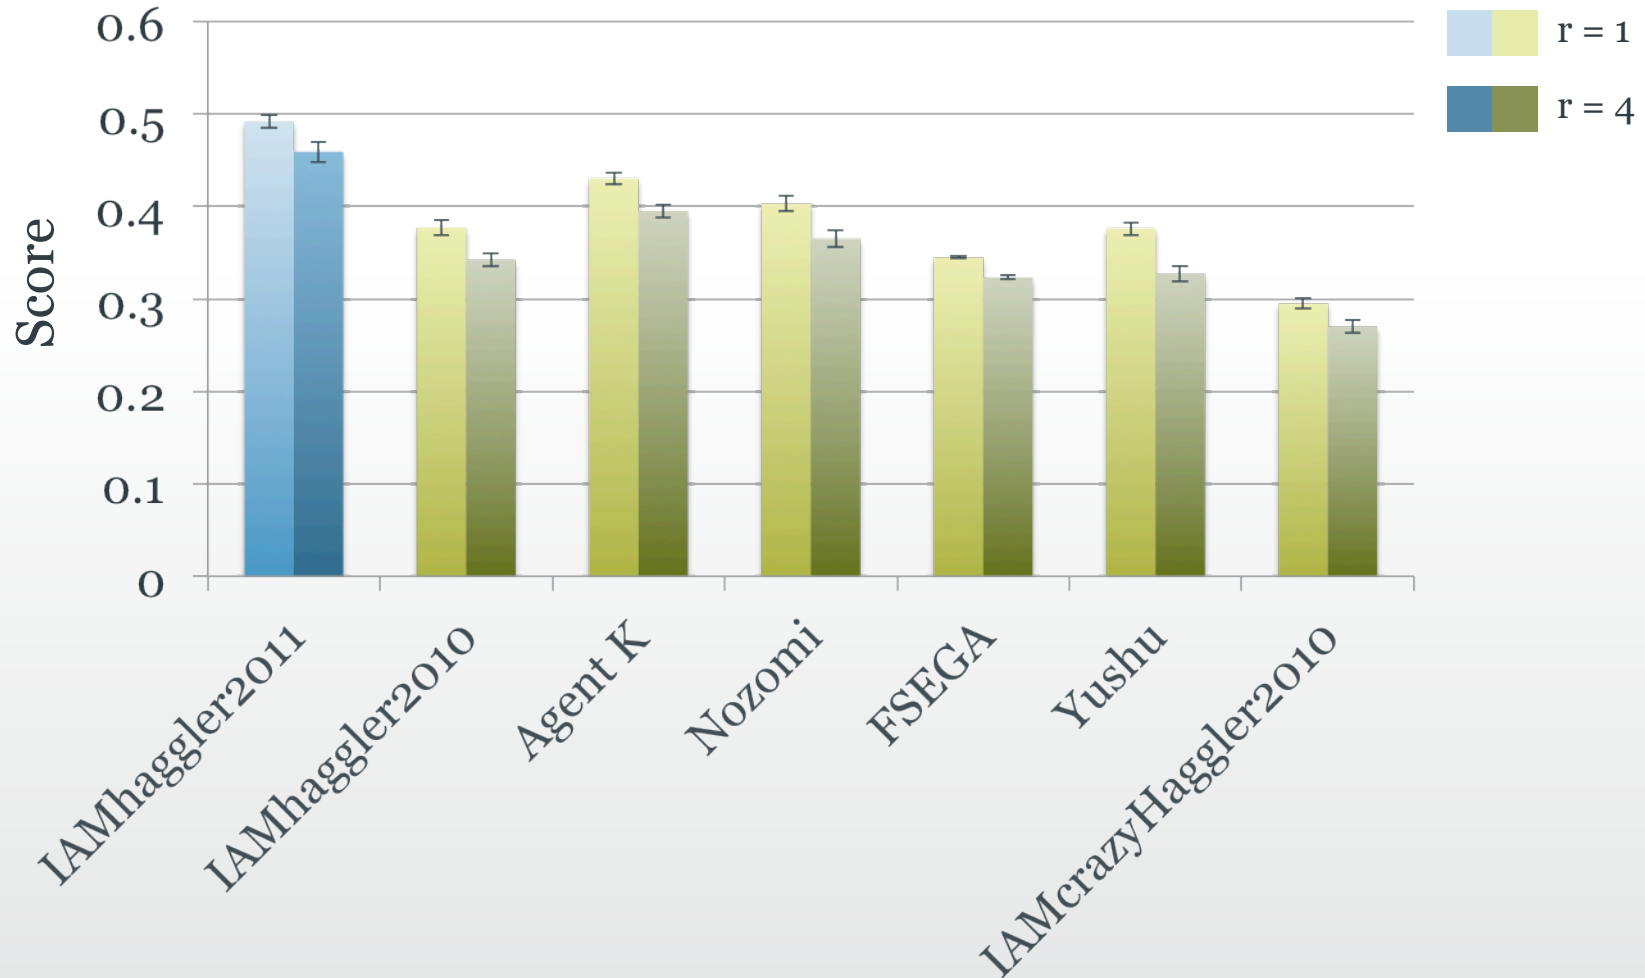
Automated Negotiating Agents Competition (ANAC 2010)

- Set up to encourage research into automated negotiation.
- Contains repository of:
 - Negotiation agents
 - Negotiation domains
- Evaluation against agents and domains submitted to ANAC 2010.
 - Plus discounted versions of ANAC 2010 domains.

Self-play



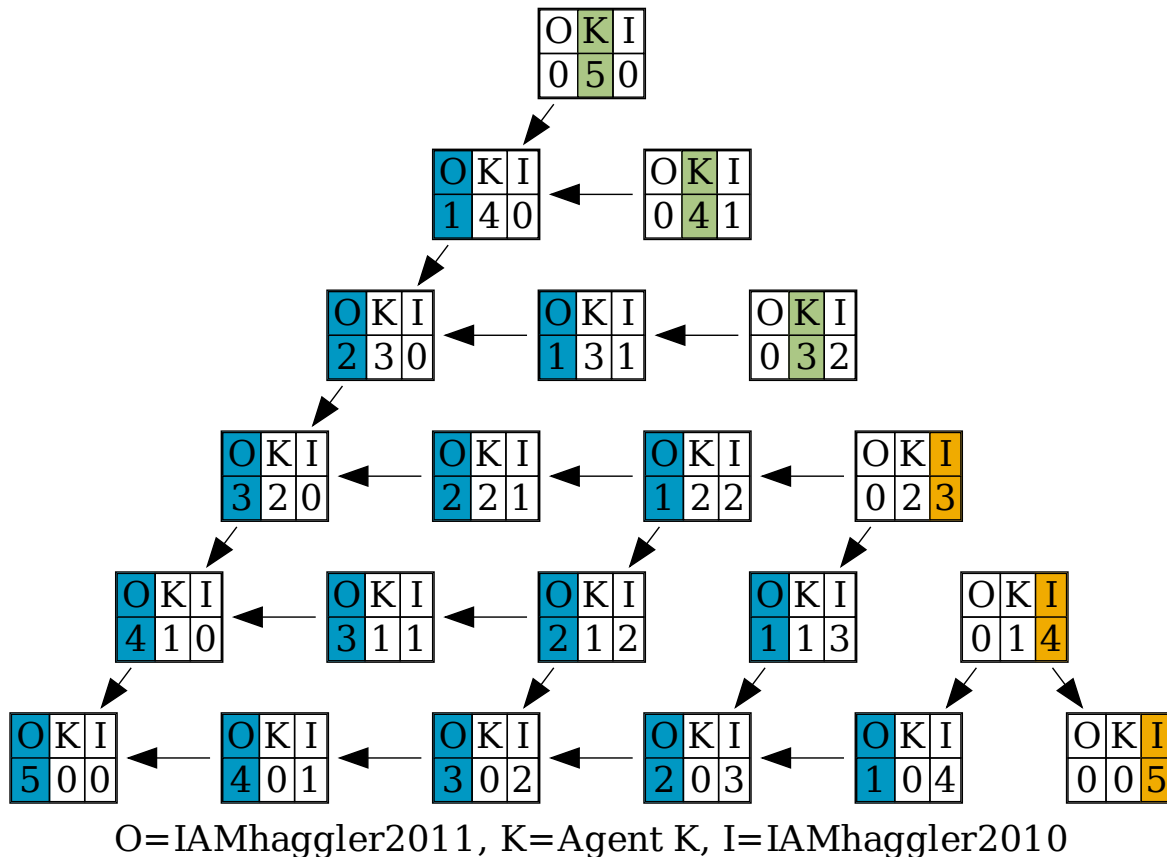
Tournament Results



Empirical Game Theoretic Analysis

- Consider tournaments with different mixtures of strategies.
- Identify incentives for agents to change strategy.
- Find Nash equilibria where no single agent can benefit by changing its strategy without others also changing theirs.
- Based on work by [Jordan et al., 2007]
 - analysis of Trading Agent Competition

Empirical Game Theoretic Analysis



Deviation in discounted version of Travel domain, for IAMhaggler2011 (with $r = 1$), Agent K and IAMhaggler2010. The shaded strategies are those which achieve the highest scores.

Conclusions and Future Work

- Developed principled concession strategy
 - for use in complex negotiation domains.
- Demonstrated performance of agent
 - against range of state-of-the-art opponents.
 - in a variety of domains.
- Planned extensions
 - further analysis against ANAC 2011 agents and domains.
 - improve strategy to avoid exploitation by non-concessive opponents.
 - improve multi-issue aspect by selecting packages which are believed to offer high utility to opponent.

References

- [An et al., 2010] B. An, V. R. Lesser, and K. M. Sim. Decommitment in multi-resource negotiation. *Proc. 7th Int. Joint Conf. on Autonomous Agents and Multiagent Systems*, 3:1553–1556.
- [Coehoorn and Jennings, 2004] R. M. Coehoorn and N. R. Jennings. Learning on opponent's preferences to make effective multi-issue negotiation trade-offs. *Proc. 6th Int. Conf. on Electronic Commerce*, pp. 59–68.
- [Fatima et al., 2007] S. Fatima, M. Wooldridge, and N. R. Jennings. On efficient procedures for multi-issue negotiation. *Lecture Notes in Computer Science*, 4452:31–45.
- [Jordan et al., 2007] P. R. Jordan, C. Kiekintveld, and M. P. Wellman. Empirical game-theoretic analysis of the TAC supply chain game. *Proc. 6th Int. Joint Conf. on Autonomous Agents and Multiagent Systems (2007)* pp. 1188-1195
- [Nguyen and Jennings, 2003] T. D. Nguyen and N. R. Jennings. Concurrent bi-lateral negotiation in agent systems. *14th Int. Workshop on Database and Expert Systems Applications*, pp. 844–849.
- [Rubinstein, 1982] A. Rubinstein. *Modeling bounded rationality*. The MIT Press.
- [Sandholm and Vulkan, 1999] T. Sandholm and N. Vulkan. Bargaining with deadlines. *Proc. 16th National Conf. on Artificial Intelligence*, 1:44–51.