Report Number: WUCS-97-11

1997-01-01

Dialogue and Deliberation

Ronald P. Loui and Diana M. Moore

Formal accounts of negotiation tend to invoke the strategic models of conflict which have been impressively developed by game theorists in this half-century. For two decades, however, research on artificial intelligence (AI) has produced a different formal picture of the agent and of the rational deliberations of agents. AI's models are not based simply on intensities of preference and quantities of probability. AI's models consider that agents use language in various ways, that agents use and convey knowledge, that agents plan, search, focus, and argue. Agents can choose their language, apply their knowledge, change their plans, continue their search,... Read complete abstract on page 2.

Dialogue and Deliberation

Ronald P. Loui and Diana M. Moore

Complete Abstract:

Formal accounts of negotiation tend to invoke the strategic models of conflict which have been impressively developed by game theorists in this half-century. For two decades, however, research on artificial intelligence (AI) has produced a different formal picture of the agent and of the rational deliberations of agents. AI's models are not based simply on intensities of preference and quantities of probability. AI's models consider that agents use language in various ways, that agents use and convey knowledge, that agents plan, search, focus, and argue. Agents can choose their language, apply their knowledge, change their plans, continue their search, shift their focus, and rebut another's arguments. Especially for negotiation practitioners, AI's model may be more helpful than game theory's models inframing situations. AI offers a broader picture of the phenomenon of negotiated agreement, and is descriptive of cooperative phases of settlement.

Dialogue and Deliberation

Ronald P. Loui and Diana M. Moore

WUCS-97-11

September 1997

Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
St. Louis MO 63130

Draft Intended for Negotiation Journal

Dialogue and Deliberation

Ronald P. Loui[1]
Diana M. Moore[2]

I

Formal accounts of negotiation tend to invoke the strategic models of conflict which have been impressively developed by game theorists in this half-century. For two decades, however, research on artificial intelligence (AI) has produced a different formal picture of the agent and of the rational deliberations of agents. AI's models are not based simply on intensities of preference and quantities of probability. AI's models consider that agents use language in various ways, that agents use and convey knowledge, that agents

---

1. R. P. Loui B.A. Applied Mathematics 1982 Harvard, M.S. Computer Science 1985 and Ph.D. Cognitive Science (CS & Philosophy) 1988 University of Rochester. Associate Professor of Computer Science and Affiliate in the Program on Legal Studies, Affiliate in the Center for Control and Optimization, Washington University. This author would like to acknowledge Professor Roger Fisher for posing the important problems in his course on international conflict, eighteen years ago. Box 1045, Washington University, St. Louis, MO 63130. *loui@ai.wustl.edu*
2. Diana Moore (Neuman), B.Sc. Computer Science, Washington University 1996, currently National Security Agency Graduate Fellow at UNM and a principal of En Garde, Inc. The current form of this work has benefitted from discussions with Venkatesh Reddy (Harvard), Mark Foltz (MIT), Fernando Tohme (U. Nacional del Sur, Argentina), Katia Sycara (CMU), Kay Hashimoto (Harvard), and Anne Jump (Harvard).

plan, search, focus, and argue. Agents can choose their language, apply their knowledge, change their plans, continue their search, shift their focus, and rebut another's arguments.[3]

Especially for negotiation practitioners, AI's model may be more helpful than game theory's models in framing situations. AI offers a broader picture of the phenomenon of negotiated agreement, and is descriptive of cooperative phases of settlement.

II

To preview the merits of a new picture of the agent and a broader understanding of negotiation, consider some illustrative, simple exchanges that might take place during negotiation.[4]

3. Historically, the description of the agent as a repository of measurable preference intensities arises with Jeremy Bentham, while the view of the agent as an information processor arises with Alan Turing and Norbert Wiener. A mathematics of continuous, real-valued quantities naturally tends toward results on what behaviors are consistent, what can be controlled, and what states are in equilibrium. Computational models are different. A study of what information, processed in what way, would suffice to exhibit complex behaviors, naturally leads to results on how complexity can be designed, and how complexity of behavior might be understood.

Although this paper aims to be as technically simple as possible, it would be a mistake to think that game theory is more rigorous, more formal, or more precise than a computational model from AI. The implementation of an AI theory as a computer program is a complete demonstration of rigor, formality, and precision, if not often of elegance. The acknowledged disadvantage of a computational model is that complexity is usually introduced to make a computational model more expressive. In contrast, in a mathematical model, complexity is limited in order to make possible the proving of theorems.

4. The notation of proposals is chosen to conform to game theory's idea of a payoff matrix, but a formal commitment to such an ontology is not intended here.

1. Deliberating for a time.

A: I propose (3, 5)

B: Let me think about it.

Sometimes an agent's view of a proposal is not pre-computed, not summarized as a single static valuation. B might have to perform a computation to evaluate a proposal. The proposal describes an allocation of resources under a hypothetical agreement, and B might have to deliberate to determine what to do with such an allocation.

Game theory only permits deliberation upon a response, such as whether to accept a proposal when some proposals have not been made, and others have been made and have been rejected. Game theory does not consider deliberation upon the relative preference for a proposal, which is assumed to be known at the outset and unchanging. Here, B might need to think because B might not know what (3, 5) entails. B might not even be able to determine whether (3, 5) is preferred, for example, to (4, 4), much less be able to determine whether it is the offer to accept.[5]

---

5. Deliberation taking time is a main theme of AI. It is a pre-Fregean theme that disappears in "modern logic" with Bertrand Russell and Alfred Whitehead. John Maynard Keynes protested the turn away from processes of deliberation but had no followers (except Wittgenstein, who mocked the uniqueness of the formal logical game but unlike Keynes, did not deride the static nature of it). Although most of AI and most of computer science considers deliberation to take time, the most celebrated proponent of the idea in this half-century has been Herbert Simon. Simon in 1988, as in prior decades has:

> *Economics has largely been preoccupied with the results of rational choice rather than the process of choice. ... In the past twenty years, there have been important advances in our understanding of procedural rationality, particularly as a result of research in artificial intelligence and cognitive psychology. The importation of these theories of the process of choice into economics could provide immense help... .*

AI has gone forward since Simon's heyday, but economics continues to ignore it, as do disciplines colonized by economics.

2. Arguing for a proposal.

A: I propose (3, 5) because that is the industry standard.

B: Yes, but that standard has been exceeded for years.

Sometimes an agent produces an argument for a proposal, especially an argument from precedent. Sometimes there is rebuttal: in this case, an exclusion is sought, since any precedent is subject to posterior revision.[6]

Negotiation can embed argument as a component. But such an embedding raises questions of protocol: Must an agent accept a proposal if the argument is lost? What sanction does the agent risk with insubordination to rationality? What is the nature of the burden to establish the merits of the proposal? Does the scope of rational argumentation extend so far as to include arguments based on position, such as the existence of a threat? Negotiation is usually not considered enough constrained by rules to be called a game. But it seems to include subdialogues in which argument proceeds according to well-established rules determining what is a counterargument, what reasons have authority, and what is a sufficient rebuttal.

3. Informing.

A: I propose (3, 5).

---

6. Technically, AI regards arguments from precedent as "case-based reasoning" (e.g., David Skalak and Edwina Rissland, Katia Sycara), a variety of analogical reasoning, which appears to be a variety of defeasible reasoning (H.L.A. Hart), which in turn, is the most actively investigated current form of nondemonstrative reasoning (Aristotle). Roger Fisher's recommendation of "negotiating on principle" presupposes argumentation.
The formalization of argument games and arbitration games is the principal work of one of the authors (Loui).

B: I strongly prefer (4, 4). But that's not a proposal.

Sometimes an agent chooses to convey information. Here, the information concerns relative degrees of preference; other information might be less precise characterizations of the the agent's situation, preferences, and desires. A threat, for example, indirectly conveys information about the agent's situation.

A: I threaten to play row 3, i.e., (3, .), in absence of agreement.

Under a broad range of assumptions about agents' willingness to threaten, useful characterizations of the the threatening party might be adduced by the party threatened.[7]

A: I will not accept any proposal in column 4, i.e., (. , 4).

Similarly, announcements of what is not acceptable convey prima facie information about the agent's contour of utilities. Some would warn that prima facie information cannot be taken seriously because of the possible mendacity of agents. More interestingly, an agent might be veridically reporting only the best perception of her or his utilities at a particular time, and this perception might be both fallible and corrigible: the agent's view could be revised with more evidence or deliberation; the agent's view might even be grossly mistaken.

---

7. Technically, such inferences from (linguistic) behavior to preferences requires the "principle of charity" (which is ironic in this context): the attribution to the agent that he or she is behaving under some principle of rationality. This principle is described, for example, by W.V.O. Quine when we assume that speakers are using logic and we wish to discover either the beliefs that they hold or the nature of the language that they use. "Charity" is also at the heart of the David Kahneman-Amos Tversky studies of allegedly non-Bayesian probability judgement. The charity appears in an extreme and possibly contorted form in the game theory literature when a solution concept is assumed for both players, for example, when it is assumed that both players are seeking Nash equilibria. In AI, all of these kinds of charity are included in what has been called "plan recognition" in the analysis of discourse.

Game theorists classify all locutions that are not proposals as "cheap talk." Such locutions are (usually) nonbinding and nonverifiable.

Game theorists are unwilling to assume the mechanisms of enforcement or self-regulation that would permit taking information exchanges seriously. More bluntly, they are interested in results that hold for unconstrained agents in weaker social settings.

Game theorists do however assume something about the social framework; they assume an institutional setting that permits enforcement of agreements and disallows withdrawal of proposals. Some might be willing to assume more elaborate institutional settings or to assume that agents are designed (for example, negotiating programs in electronic settings) in such a way that they cannot or do not lie all the time in all possible ways. There is no such thing as negotiation taking place in the absence of all social convention. Clearly the issue is one of modeling at this point, a question of emphasis, not an issue of a priori merit.

4. Querying and inviting.

A: Do you prefer (3, 5) to (4, 4)?

B: I'm not saying. But I invite a proposal in row 3.

Two other kinds of "speech acts," among others, are requests for information and invitations of proposals.[8] Protocols for negotiation are rarely so well defined that responses to such locutions or rewards for such locutions are compulsory. But the possibility exists in

8. The idea of speech acts clearly arises with John Austin, and their categorization begins in earnest with John Searle. AI models of discourse are so wedded to the Austin/H.P. Grice/Searle view that it is difficult to imagine an information-processing model of pragmatics without speech acts. A complete list of speech acts might include such moves as insulting and praising, explicitly taunting, annoying and confusing, filibustering, grandstanding, annointing and dubbing, incanting, focusing, clarifying, and even agreeing. The non-propositional uses of language are numerous. Surely some have no place in the model, while others, besides proposing and accepting, do.

artificial societies to advance such protocols, and conformity to strict social norms can be observed in many negotiations. It may even be useful to characterize the degree of cooperation achieved in many phases of negotiation by the severity of the linguistic and social regulations that the parties observe.

As with other speech acts, game theory can make no comment on such locutions, unless the model includes data representing the strategic relative benefit of all discourse options.[9]

---

9. The legitimate difference of views on modeling strategy, ontology, and regard for social convention might be found in the analysis of:

A: Let's play UNO.
B: Ok, but you'll have to teach me.
A: Actually, let's not.

We would say that a proposal was made, but that it was not maximally specific; its details were not fully specified, and it is a formalist's fiction to suppose that any proposal is ever fully specified. Anyone who has negotiated an important contract is aware of the various specificities of proposals and agreements.

It may be that every proposal includes implicitly, from a background or contextually implicit understanding, that interpretations of proposals will be resolved and clarifications made through some kind of procedure. So there is a conflict-resolving or detail-producing procedure proposed with every substantive proposal concerning distributions and allocations. Sometimes the negotiation appears to be principally concerned with the details of the allocation, and the insufficient specificity of language comes as a surprise. At other times, the negotiation is explicitly and consciously about the parameters of the proposal through which the details of the substantive agreement will be reached. Of course, the latter is a negotiation about a negotiation, a meta-negotiation, and the preliminary interaction can be clearly delineated from the ensuing one. The point remains, though, that it is a formalist's bane, and bust, always to assume (or perhaps ever to assume) that a proposal is maximally specific.

We would say that the proposal was accepted with clarification. Then the offer rescinded (a minor violation of a social norm). A classical mathematical economist would have to say that B's locution was a counterproposal. B's counterproposal raised an exception to the default (a linguistic default or communication convention), that an offer to play does not normally include a lesson. A rejected B's counterproposal, since it was not the same as A's original proposal.

5. Searching for oneself / Revising utilities (continuing dialogue 1).

A: I prefer (3, 5) to (4, 4).

B: Hmm.

A: Wait. No, actually I'm indifferent.

Sometimes agents revise their views of outcomes, settlements, distributions, or proposals. Agents might have a procedure for constructing utilities from more basic data. The data (that they use to construct utilities) may increase with time and with computational effort; for example, as more deliberation is performed, more scenarios are considered.

Real agents do not forsee future scenarios with infinite horizons. Their deliberations do not take place within Bayesian "small worlds" no enlargement of which can change evaluations. Agents think about acts and plans at varying levels of detail and with horizons of varying extent. Perhaps they are solving hard optimization problems, such as mixed non-linear programs, or they are solving constraint-satisfaction problems. Perhaps they have multiattribute utility expressions, but choose to introduce utility attributes one at a time, in order of importance.

In AI, outcomes are evaluated heuristically. Chess positions, in the usual example, are evaluated at different depths, with a scoring function. This scoring function or heuristic function provides a criterion for preference that is independent of expected utility or maximin. It does not conform to the theoretical presuppositions of utility, and no effort is made to reconcile it with those theories. This is not optimal (in any useful sense), but is heuristic.

Heuristic evaluation leads to the possibility that different deliberations can result in different preferences.[10]

In economics' models and all classical decision models, utilities can shift only because of (1) new evidence on which an agent's probabilities can be conditioned, which changes an agent's expectations, and (2) spurious (inexplicable) shifts in an agent's tastes.

6. Searching for another.

A: I can't meet Wednesday at noon.

B: You can if you end your lecture early.

Sometimes agents report possibilities to each other that may not have been obvious, that may have been overlooked, or that may have been postponed for later consideration. This is different from situations wherein the other party is informing. In this example, the information (about the possibility) in this case is available to the deliberator, but has not yet figured in the deliberation.[11]

Searching for another requires some understanding of the problem that the other is trying to solve. It would be ludicrous to assume that one agent's picture of the other is perfect. We have already assumed that an agent's picture of her or his own position is not fully in focus. Here, it might be possible for B to know some of the parameters of A's scheduling problem, without knowing others. This knowledge

---

10. Heuristic search was already a common idea in operations research when Herbert Simon appropriated the idea to explain chess-playing programs and to theorize about bounded rationality. The idea that Bayesian small worlds (models of limited detail) should be equivalent to large worlds appears in decision theory with Leonard Savage, but is implicit in any axiomatic treatment of preference.
It is possible to view the ex post effect of a heuristic as an optimization of an unspecified objective function. This loses the important nuance of process. The nuance lost is in the move from nondeterministic potential behavior, ex ante, to observed behavior ex post. It is always possible to eliminate the actual procedures of deliberation on actual data and replace them with a historical fiction, of some other manner of choosing, based on invented data. The philosophical difference between ex ante implementation of a process and its ex post interpretation is reported by this author elsewhere.
An interestng paper on constructive decision theory is Glenn Shafer and Amos Tversky. Clearly John Pollock, who views probability arguments as constructive, foresaw constructive decision theory.

suffices to explain why B's rejoinder is relevant and calculated, rather than the result of chance or clairvoyance.

A could respond, for example,

A: I would suffer great penalties if I did that.

A might not fully agree with B's perception of A's valuations of outcomes. A's locution could arise in many other situations: for example, when B does not carry out the computation as fully as A, or vice versa, or even when A is being deceptive.

A could even respond,

---

11. Game theorists deal with information states. Hybrid game-theorists model agents' beliefs with sets of propositions in a logical language, with the intention that they define information states. There is a possible confusion here. The effect of deliberation upon fixed data is to alter the information states in the game theorist's model, in effect, to re-frame or re-model the problem. If deliberation and computation are not considered, then the additional step, from logical epistemic state to informational epistemic state, is a nuisance for game theorists and has no important modeling significance (for example, John Genakopoulos asked Joe Halpern in this author's presence, "why use logic?"). However, we are enlarging the picture of the agent considerably here. The epistemic state consists not just of the set of possibilities, contingencies, and fixed measures of probability and utility. It contains information about how those sets and measures might be derived from other data. This additional information guides and constrains the shifts of epistemic state that ensue as new information is adopted, or as further calculation is made with existing information. This kind of information (about the dynamics of beliefs) cannot simply be captured in information states.
A different model of the agent would restrict to a minimum the inferences that an agent makes and mandate that agents achieve this minimum. Epistemic states would then consist of propositions of which the deliberator has certain awareness. Drawing a conclusion would correspond to a revision of epistemic state, like adoption of new beliefs. No special status would be given to inference. Forgetting, for example, is modeled no differently from deliberate rejection of a belief in both the game theorist's model and the usual AI model.
Of course, theorists can model what they like. The present authors cannot conceive of a serious study of negotiation that does not model the interplay of agents' dialogue and inference.

A: No, I can't end that lecture early.

This reports that B's understanding of A's decision problem(s) is not only incomplete, but is in fact in error. It is a factual matter, not a question of incomplete or imprecise deliberation.

The response that B seeks from A is acknowledgement of the discovery that B has made. Presumably, if the discovery is important enough, resolution of the disagreement, follows:

A: How about that? You're right. So I can accept your proposal.

Even if B's discovery does not lead to settlement, at least it leads to a superior common understanding of the problem.[12]

7. Requesting problem-solving.

A: I propose to open the window.

B: I'll get cold.

Parties to negotiations do not just bombard each other with random facts. One kind of speech act is an indirect request. The surface-form of this request is not a question, but a declaration.[13]

---

12. Practicing negotiators might complain at this point that the analysis is shallow and aims to explain only mundane negotiator behaviors. They are correct, but the AI method requires focusing on simple and typical examples. The aim is to say what formal theory would be required to generate those behaviors, to write a program that would exhibit a range of appropriate behaviors, or to give an analysis detailed enough that such a range of behavior could be understood in terms of data and process. A computational model seeks to describe what data, processed in what way, would suffice for a behavior. Specifying the desired inputs and outputs is often the most important part of the modeling. Mathematical models seek minimal ontologies on which deep determinate analysis can be done, under the assumption that the model fits the situation. In contrast, we seek the largest useful ontologies with which to fit the situation.
13. For comparison, "Can you pass the salt?" is rarely intended to be understood its surface-form, as a request for a "yes" or "no" response. It is an indirect request for action.

Here, B is objecting to the proposal in a particular way. B announces a consideration or constraint which prevents B from accepting the proposal. B has thus invited A to address the problem of removing the constraint, mitigating or alleviating its impact, or otherwise compensating for its existence.

What is the extent to which B is bound to accept A's proposal if A can satisfactorily solve B's stated problem? The subdialogue takes the form of rational argumentation. A will be arguing that proposed compensation should suffice, and B will be arguing that there is a remaining inadequacy. One side has private access to counterarguments. The other side appeals to public standards of how problems might reasonably be solved.[14]

8. Using knowledge to find sidepayments.

B: I'll get cold.

A: I have a sweater.

Inventing relevant sidepayments requires knowledge. B's locution has meaning; it contains information about what kind of sidepayment might be worth pursuing. A's prospects of reaching an agreement with the locution "I have a sweater" are better than with the retort "I have an anvil."

The knowledge that A brought to bear need not have been about B particularly. A might know that a reasonable and usual solution to the problem of someone being cold is the wearing of a sweater. B might not understand how sweaters help, might hate sweaters, or might not even really be cold. But A has responded to B's locution in a rational, relevant, and cooperative way.

This is the main behavior that can be modeled well with the AI picture of agents and their intercourse. The AI picture will be superior to the game-theoretic picture in the modeling of interactions in direct proportion to the extent to which knowledge informs creativity.[15]

14. Katia Sycara has written extensively about this view of negotiation as problem-solving and problem-restructuring. Other authors in the AI study of dialogue have had similiar, if less precise views. Sycara's models are most relevant because they take seriously the adversarial relation between parties. Her models however do not import much of the theory of discourse; nor does she connect explicitly to the utility-based strategic view of negotiation.

AI's investigation of conversational dialogue is primarily concerned with coordinated problem-solving. This is something of an accident in intellectual history. In the tradition of H.P. Grice and David Lewis, AI's theorists were primarily interested in understanding conventional transformations of surface linguistic phenomena: agreements of referents for pronouns, tenses, narrative point of view, and the coordination of other meanings.

Problem-solving behavior naturally arises in discourse. Two speakers can enter a dialogue the goal of which is to come to an understanding about the meaning of some locution in the original dialogue. This is a *clarification subdialogue*. For example, one party might simply ask "What does that word mean?", "What do you mean by that?", or "Whom do you mean?" Such a dialogue is invoked in the same way that a subroutine is called by a calling routine. Its relevance to the larger dialogue is highly localized, encapsulated, and context-dependent.

Since the goals of communicating agents run the gamut of abstraction from low-level phonetics to high-level social coordination, it is not surprising that AI investigations of discourse soon grew from speech acts to *mixed-initiative planning* and negotiation. The discourse community supposes a rich and complex logic-based representation of agents' mental states. The community cares deeply about search and its effect on plans. If this essay were purveyed to an AI audience, most of the observations would be intuitive. The advances claimed would be (1) depicting the inescapable interplay of adversarial and cooperative behavior; (2) observing the connections between discourse (social search) and deliberation (individual search); and (3) intepreting game-theoretic constructs and modeling assumptions in the AI model and vice versa.

Regarding (1), existing mainstream AI research is too quick to assume a hierarchical relation (e.g., master-slave) between the negotiating parties, as part of the pragmatics of the situation. Regarding (2), no known work seriously addresses these issues. With respect to (3), mainstream AI is either ignorant of game theory's ontological assumptions, thus incapable of modifying its foundations, or else too quick to adopt the established tenets, thus unwilling to modify foundations.

9. Pursuing a subdialogue.

B: Sweaters don't keep my ears warm.

A: You can have my earmuffs, too.

B: I don't like the way earmuffs look.

A: No one will look at you.

A subdialogue can be pursued, perhaps indefinitely, as a party attempts to buttress a proposal. Sidepayment after sidepayment can be proposed. Subdialogues of subdialogues can be entered.

Agents who are inventing sidepayments are employing default models of the other party. Knowledge explicitly conveyed helps to identify the other agent within a taxonomy of types of individuals: persons, cold persons, locally cold persons, locally cold and socially

---

15. Fisher has emphasized the importance of creativity in negotiation. Creativity breaks deadlocks. Creativity is what separates good negotiators from mere hagglers, what distinguishes value-creating interactions from tugs-of-war. How can one make creativity possible, theorize about, explain, or implement systems that exhibit it? AI's way is to provide generative procedures which take data and transform implicit possibilities into something more explicit. In the same way that a chess-playing program generates its search space and creates its analysis, inegration of data and inference from multiple sources of knowledge generates conclusions and creates hypotheses.
Game theory's alternatives are (1) to include all possible potentially relevant sidepayments in the rows and columns of the payoff matrix; and (2) to permit unmotivated arbitrary redescriptions of the game matrix or its dimensionality. Note that under (1), game theory would require fixed payoffs to be given for a complete, high-dimensional payoff (bi)matrix with perhaps billions of billions of cells. Under (2), game theory could not distinguish the expansion of a game from a shift to a different and unrelated game.
Dean Pruitt wrote about integrative agreements, where negotiating agents begin to cooperate to solve each others' problems. The principal kind of integrative agreement was logrolling, where one agent makes a concession on one attribute in a multiattribute setting, in exchange for satisfaction on another. AI models of knowledge and planning permit a richer form of integrative agreement.

sensitive persons, etc.

Problems are mentioned which further identify the agent's situation. As proposed solutions to the problems are offered and rejected, the proposer's model must increasingly accomodate the other party as an unusual person, as someone who is an exception to the default, as a person who has a complaint in response to every reasonable accomodation. At some point, the proposer simply quits accomodating.

The failure to come to an agreement through problem-solving is one of the things risked when an agent attempts to be deceptive. What is an agent's response when offered a sweater, if the agent is not in fact cold? Acceptance of the sidepayment is almost impossible, since it is probably irrelevant to the actual preferences of the deceiving agent. The deceiver must simply endure the problem-solving subdialogue, continuing to deceive, at the risk of appearing particular, uncooperative, stubborn, stupid, or recalcitrant.

10. Rotating the focus.

A: You can have my sweater all evening.

B: Can I return it next week?

Sometimes both agents are so committed to trying to make a proposal acceptable that the focus of the negotiation changes. The parties who began negotiating the status of the window are now negotiating the time of possession of a sweater.

Thinking of dimensions in a payoff matrix, the negotiation begins with two dimensions; negotiation takes place in a 2d plane. Potential sidepayments are orthogonal directions, third, fourth, and further additional dimensions. These dimensions are discovered as the negotiation proceeds.

The traditional strategic picture of the situation can thus be altered to depict what happens when the negotiation proceeds in this way, pushing into new dimensions. But it cannot explain why it proceeds one way and not another. The AI model of the agents (specifically, of their search and their knowledge) is required to explain why things happen.

In the AI picture, the dimensions are not all there in the beginning, and fixed. This is mainly because agents are computationally limited and perform search. The ability to generate new space to search is made possible by agents' knowledge. An agent's decision to investigate some dimensions and not others is informed by the agent's inference about his or her situation, the other party's situation, and the nature of their impasse.

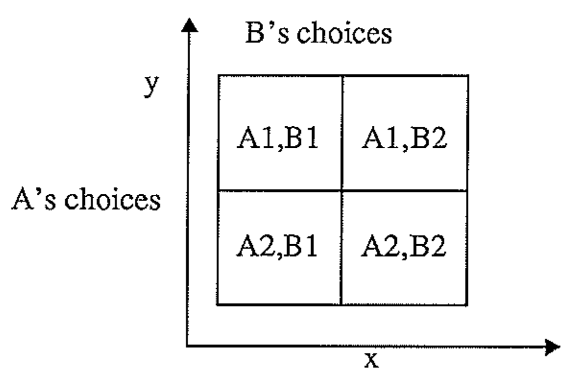

Two dimensions for the start of negotiation.

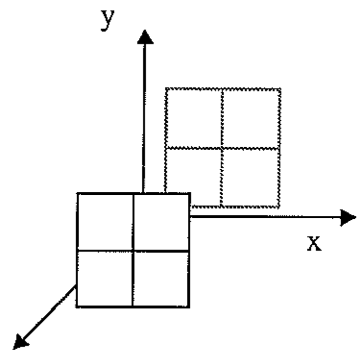

In which new dimension(s) to negotiate?

A party's perceived payoffs define a payoff surface over this space. When both parties have payoff surfaces with great cliffs, agreement can be hard to find.

No matter how many dimensions are used to describe proposals, all possible agreements can be graphed according to A's and B's payoffs (or their perception of their payoffs, at some level of analysis). In this way, 2d, 3d, 4d or higher-dimensional strategy space is viewed in a 2d payoff space. Each point in payoff space is a proposal. There are maximal points in this payoff space, and these form the Pareto frontier (though it is not generally useful to talk about a Pareto frontier when the payoff information is private).

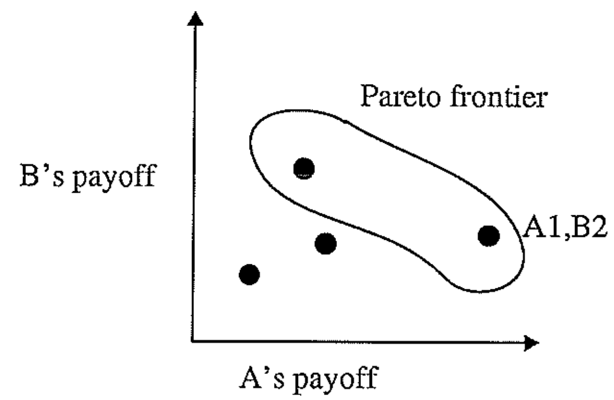

A familiar view of the (apparently) four possible proposals.

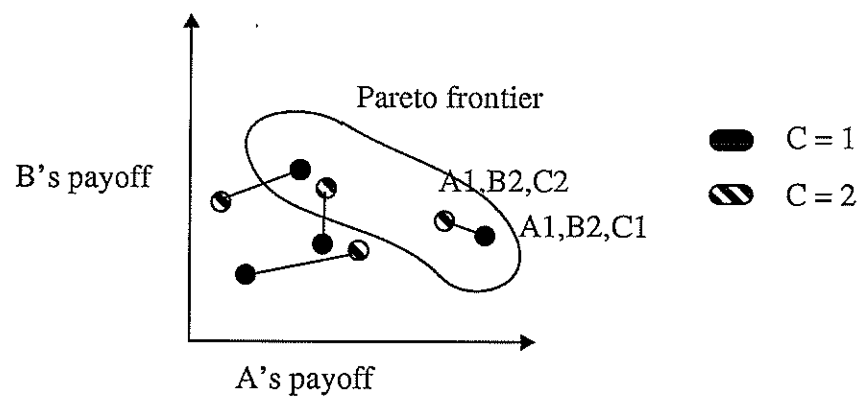

Each added dimension splits the possible agreements.

Each additional dimension causes a multiplication of points in payoff space, a multiplication of potential agreements.

The most mundane kind of sidepayment simply fills the gaps between points that are too widely spaced. In this way, sidepayments provide a smoothing of the payoff surfaces of each player, decreasing the contrast between alternatives, eroding the cliffs. Deadlocks are broken by providing increased levels of compromise. They have no more effect on the payoff space than the introduction of probabilistic mixtures of proposals (though they do so without increasing risk).

A more interesting kind of sidepayment is genuine creativity, which introduces new points beyond the Pareto frontier. Agents who search for ways that their reciprocal and coordinated actions can be mutually beneficial can usually find such points. This is the raison d'etre for social interaction. It happens in a static way when social search seeks the Pareto frontier; it happens in a dynamic way when knowledge pushes possible agreements beyond the Pareto edge.

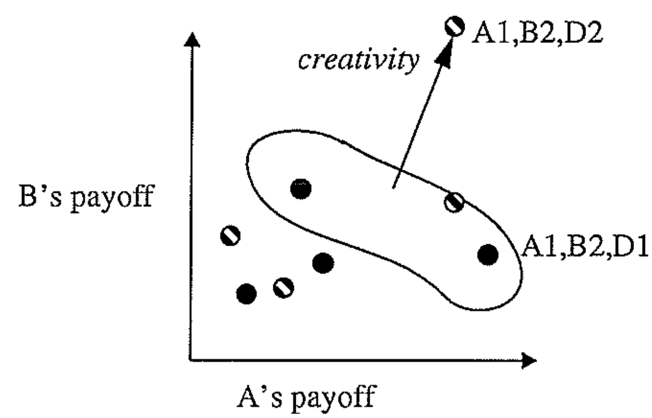

Some dimensions permit breakthrough proposals.
The trick is to find, or to construct, those dimensions.

But agents usually do not introduce new dimensionality without collapsing existing ranges of options. A subdialogue begins with a suggestion to fix attention on a particular proposal. Agents fix upon a single cell in the game matrix: a single row and a single column, while varying other dimensions. The negotiation that began in x-y space is settled in u-v space, not in x-y-u-v space. The four-dimensional space is never considered because it is entered only for fixed values of x and y. x-y-u-v space is thus projected onto u-v space for the purposes of the subdialogue.

In this way, additional considerations cause the focus of negotiation to rotate into a new plane. It is absurd to conceive of creativity as an unchecked ballooning of multiple independent options.

In the sweater example, negotiation begins with the 2x2 strategic form: rows describe A's choice of window position, up or down; columns describe B's choice of window position, up or down. Only the diagonal elements are reasonable outcomes in this game. This is because unresolved disagreement presumably is extremely costly for both parties (at any level of analysis). The new dimensions are *A's choice to give a sweater* and *B's choice to use it*: either can be true or false. This is a new 2x2 plane in which the negotiation subdialogue proceeds. It is not a 2x2x2x2 space; this four-dimensional space has been projected onto a plane with the *row value = A chooses the window up*; *column value = B chooses the window up*.[16]

There is no reason to suppose that negotiation always focuses on two dimensions, thus taking place in a plane. But there is reason to model dialogue as restricting attention to some dimensions, occupying subspaces at times.

---

16. It is easier to view this as rotation if one considers an example of three dimensions being projected to two, instead of four dimensions projected to two. We begin in x-y space and switch to x-z space.
Imagine a cube. Negotiation originally takes place in horizontal plane that slices through the cube at a fixed height. Every proposal is a point in that plane, and that point has a depth. When a subdialogue begins in response to a proposal, the focus switches to the vertical plane that cuts the cube at that depth. Negotiation now takes place in the vertical plane. Thinking just of the two rectilinear planes, the change in focus has caused the plane of attention to rotate 90 degrees forward.

This restriction to subspaces is one example of a larger class of search patterns that help to induce agreement for normal agents who employ non-pathological search procedures. This is because agents normally focus their personal deliberative search where the dialogue has focused its attention.

Some kinds of search have an inherent bias; they are characterized by goal-satisfaction or monotone improvement of an approach to a problem. Heuristic strategies for optimization are examples of this kind of search. The more time granted to analyze a constrained nonlinear optimization problem, the better the solution. Search can only uncover improvements; it cannot make earlier discoveries disappear nor can it taint them.[17]

When two parties have proposals on the table, some search is normally allocated to the analysis of what might be done with the other person's proposal. There is no guarantee that this happens. There is also no guarantee that the search improves their estimation of a proposal. If however, the effect of search is to improve the estimation of the other's proposal, the focus of deliberation naturally leads to a smoothing of payoff surfaces. This contributes to the finding of agreement.

The effect on topology can be even more dramatic. Suppose the usual game in which each player's payoff surface slopes downward from the favorite corner of the strategy matrix. The preferred settlements are at opposing corners. Along the off-diagonal is a valley between the two agents' utility surfaces, where neither party finds much prima facie value. In this valley, however, is where agreement is likely to be found. If both parties search to improve their views in this area of likely agreement, the valley pushes upward as utilities are revised. With enough time and serendipity, search can even result in both players valuing settlements in the

17. Some kinds of search can be described as neither optimistic nor pessimistic. The breadth-first expansion of a search tree in a game like chess is an example. Additional search does not inherently buttress the case for what currently appears to be the best move.

center more than the best values to be found at the corners.[18] The focus of dialogue induces a focus of deliberation. The effect on utilities is like the meeting of landmasses, a geological process that can create mountains.

One party might even ask that the other party expend computational resources on a promising proposal.

> *A: Can you think more about what my present proposal, (3, 4), might mean to you?*

The problem with such a request is that it can easily be rejoined:

> *B: Well, can you think more about what my proposal, (4, 3), might mean to you?*

Explicit requests to deliberate upon a proposal are not usually useful because it is already in the interest of an agent to deliberate on the other's proposals. Without deliberating on a proposal, it is hard to see what would move an agent to accept or reject it.

It seems simply counterproductive to dwell on rejected proposals[19]:

> *B: I reject the proposal (3, 4).*

---

18. It is an interesting question whether a player should accept a proposal for which there ought to be great value, when the player lacks the understanding of how to achieve that value. A player might know for certain that value increases monotonically with row and column, assuming there is unbounded computation available to identify that value. The player might, however, have computed a better solution for a theoretically inferior proposal. For example, the row and column values might be parameters in highly discontinuous optimization problems. The objective function might be increasing in those values. Nevertheless, the computation of the solution for, e.g., (10, 10) might be considerably more difficult than the computation of the solution for (4, 5). Solutions are likely to be heuristic. Familiarity with the (4, 5) problem might explain why it is more easily solved or why a better solution for it is known. As another example, suppose that a player's value depends on solving large scheduling problems. Value increases with the size of the schedule, but solvability decreases. Clearly, introducing a probability of finding an ideal solution or a probability of improving a heuristic valuation reduces the problem, but does not improve understanding.

A: I've been thinking more about how great (3, 4) would have been for me.

11. Popping from a subdialogue.

A: We agree that my sweater will keep you warm.

B: Yes, but I still reject the proposal.

Sometimes agents finish a subdialogue without success. They then return to the original dialogue. In geometric terms, they have aborted their social search for agreement in the rotated dimensions and have returned to the original plane.

In AI models of discourse, the agents are said to have popped the stack. Positing a stack as the appropriate data structure means that the subdialogue has no further ramifications once it is ended. Agents do not switch between a few active subdialogues, for example.[20]

There are several ways to terminate a subdialogue.

A: I can think of no way to solve the problem of someone seeing you wearing earmuffs.

The various kinds of termination have various effects on the parties' willingess to cooperate. An analysis of these effects requires precise specification of the meta-games that the players are inheriting and constructing.

---

19. We do not deny that such a locution can be used for tactical purposes, whether the deliberation actually took place or not. It might be useful to explain what one party has been denied, in detail, in some social settings. For instance, it might behoove agents to argue convincingly that they perceive a large amount of regret, pain, or self-sacrifice because their proposals are rejected. Or it might be useful, in support of a threat, to convince the other party that one has an unusual willingness to forego agreement, if the negotiation is not largely settled in the threatening party's favor. These seem to be easily abused social conventions and are perhaps better analyzed with fully adversarial models.

12. Adhering to protocol.

A: I threaten row 3 under breakdown.

B: I thought we agreed to negotiate without threats.

Sometimes parties agree to the language that they will use, or to particular conventions. They might agree to restrict their talk to substantive proposals. They might agree not to threaten. They might agree to a limit on the number of proposals, the total time to deliberate, or the minimum time between proposals. They might agree to match concessions with concessions. They might subject themselves to the results of adversarial argument from precedent. They might agree to respond to all informational queries truthfully.

---

20. Stack behavior is imposed on dialogues mainly to give them a simple form; the structure conforms to a depth-first traversal of a tree with linearly ordered nodes. Anything that can be outlined can be given this latter form, and anyone who reads an outline top-to-bottom is performing a depth-first traversal of the structure. The appropriate data structure for keeping track of such a traversal is a stack.
A terminated subdialogue might actually to be reentered (thus, violating stack behavior):

> *A: (considerably later) Remember the possibility of my giving you a sweater? Suppose I can find a hat, too. Would you accept that?*

This recalling of a subdialogue is better considered a queer way of introducing a new subdialogue.
Barbara Grosz's work is considered the main AI investigation of the structure of dialogue. There, the desire to eliminate the temporary data associated with a subdialogue upon its termination is easily understood. Consider pronoun reference. Conversation begins with things in focus: principal actors, objects, goals, times, and places. A subdialogue temporarily suspends this frame of reference, substituting objects that serve the subdialogue's goal. In the case of a clarification goal, the subdialogue refers to the word or sentence, not to the actors and actions. When the subdialogue ends, the objects go away; the original frame is restored. This is exactly how subroutines work in most programming languages. The more limited the focus and the stronger the social conventions that govern subdialogues, the more sense it makes to view dialogue as a stack.

To a game theorist, protocol is simply a pruning of the extensive form of the game. There is no need to specify the protocol as a set of rules, especially not as a set of rules that might be the objects of negotiation, interpretation and dispute. The reverse is true of an AI model, wherein sets of rules that govern procedures, expressed in a precise symbolic language, are central objects of theorizing.

Many negotiations are characterized by ascent to meta-negotiation: negotiation over the protocol and ad hoc conventions, not over substantive issues. It would be desirable to theorize about protocol, its description, adoption, parties' adherence to it, and penalties for failing to adhere to it. While AI's models have little to say about this at the moment, they do at least provide ways of posing the question because they admit rule-making, rule-following, and language into their models.

III

The AI view need not be as complex as the discussion might suggest. The essential features are that:

1) parties can individually deliberate upon outcomes, and they have the data that generate search spaces for this deliberation; the relevant AI idea is *planning*, which presupposes *knowledge*;

2) parties communicate in a logical language that is at least, under some assumptions (such as a probability of truthfulness) interpretable at face value; the idea here is that there is *inference* and *plan-recognition*;

3) the effect of conversation on the mental state of agents can be clearly represented, especially the effect on one party's model of the other; the idea here

is *default reasoning*;

4) there is a model of rational *argument*; and

5) there is a specification of protocol that defines well-formedness of the dialogue; this is just a matter of defining a *language game*.

AI considers each of those topics, *planning*, *knowledge representation, automated reasoning,* and *argumentation* as a separate field of study. Obviously, we can only sketch a few of the ideas in the space that remains.

The mental state of an agent consists of a set of beliefs, or *knowledge base*, KB. KB contains different kinds of knowledge: knowledge about what is possible, knowledge about what are the probabilities of the possibilities, knowledge about which effects are had by which acts; knowledge about the average person, knowledge about what is average for certain kinds of exceptional persons; knowledge about what has value for the agent and for others.

This may sound like a great indulgence, but it is actually compliance. It is much less than what the game theorists ask for. AI supposes that there may be knowledge of different kinds, but does not suppose that there is necessarily a lot of knowledge. The knowledge assumed for any agent is always permitted to be incomplete. The agent might even be pathetically ignorant; that ignorance might yet be heterogeneous. What little knowledge there is might still be useful if it is recognized for the kind of knowledge that it is. AI aims to show how incomplete information of various kinds bears on deliberation. It aims to produce models that are robust to the lack of information.

In contrast to mathematical approaches of prior decades, we do not assume that there must be probabilistic knowledge about the benefits of search.

Knowledge is usually represented in a logical language with an appropriate *ontology*. Referring to an earlier example, the knowledge of the agent might include[21]:

Has(agent2,sweater,Do(give-sweater, current-situation))

FORALL s
IF Has(agent2,sweater,s) AND True(window-open,s)
THEN False(cold(agent2),s)

An agent's *plan* is a tree which determines what acts should be taken in what order, under various scenarios of possible futures. Those scenarios might not be exhaustive. No particular horizon for action or evaluation is predetermined. Agents who have planned little will have short horizons; more deliberative agents will have long horizons. One of the main points of this paper is that the AI model permits plans to be extended during the negotiation.

The plan also includes *ramifications* of actions and events. Not all ramifications need be considered. Agents value a scenario by evaluating their like or dislike for the ramifications of a scenario. Ramifications are properties that can be proved to hold at some resultant state of a scenario. The more scenarios considered, the more branching the tree. The more ramifications contemplated, the more the agent can say about each leaf in the tree.

An AI plan is essentially a decision tree from decision theory (e.g., Raiffa), with the allowance that it can be usefully deepened with more search.

---

21. The predicate, *True*, is introduced so that meta-language and object-language are kept separate. This way of representing situations, the situation calculus, arises with Patrick Hayes and John McCarthy. It is but one ontology that AI has investigated for its computational and expressive properties. It is perhaps one of the least practical ontologies, but perhaps the most theoretically useful one.
To represent knowledge is to define a symbol system and a set of rewriting rules that govern the shorthands of that symbol system. Logics are one class of knowledge representation system: their syntax defines what symbols may combine, while their inference rules define what sets of symbols can be rewritten as others.
When the logical construct, *IF* ... *THEN* ... is permitted to be nonstandard (defeasible or fuzzy), the specification of an agent's knowledge is considerably improved.

Game Theory:

if player a does a3, player b can do b4

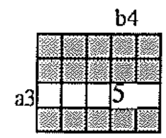

Bargaining Theory:

a's and b's proposals are close

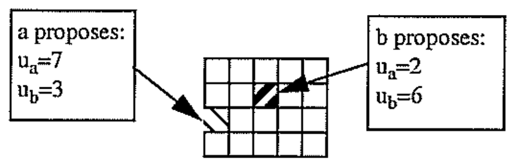

Planning Theory:

do more search? do more analysis of the outcome?

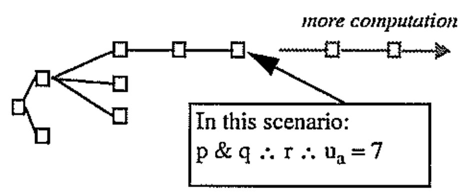

Negotiation in AI:

evaluation of proposals depends on plans, which can be deepened

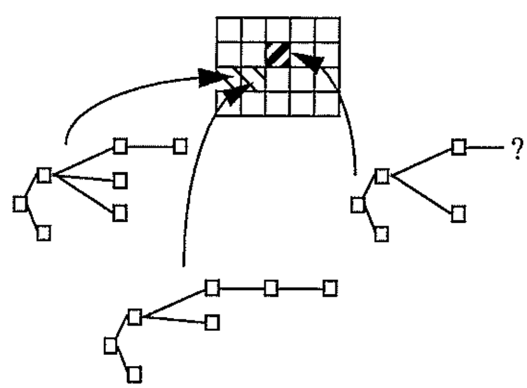

A Change in View

During negotiation, each party develops a plan for each proposed agreement. Each proposal might require a separate plan. The plan is what the agent intends to do if that agreement is made. A proposal defines a distribution of resources, and this distribution affects the relevant scenarios, probabilities, and valuations. Under one settlement, an act might be possible for an agent that might be impossible under a different settlement. Or the settlement might change the cost of performing the act. Or it might change the ramifications, thereby changing the valuations.

An agent may refuse to plan deeply in evaluating a proposal. Nevertheless, to perform any evaluation, at least a degenerately trivial plan must be developed.

When parties communicate, they sometimes make descriptive statements about themselves or about the joint situation. The main interpretive challenge of the hearer is to move from the speaker's *surface form* of an utterance to its *pragmatic meaning*.

Even when statements are veridical and detailed, there are problems with interpretation. A statement uttered in the context of the speaker's background knowledge does not have the same implications for the hearer, when that statement is adopted by the hearer in the context of the hearer's knowledge.

Mathematical puzzlers have focused on two pathological issues, each minimally relevant to the discovery of meaning: (1) the nesting of speaker accomodation of hearer and vice versa (the problem of his knowing what she knows he knows and so forth); (2) the proclivity of speakers to truth and hearers to trust. Each can be made to be important or found to be important, but can skew one's understanding of ordinary negotiation dialogues.

Not all language is meaningful in an AI model, but at least the declarative statements are. Anything that can be expressed (at face value) as a logical sentence can be adopted by the hearer, and new conclusions drawn from it. The AI problem of *natural language understanding* is certainly not relevant here in the broadest sense of natural language. But the narrow use of language to convey information, resulting in an integration of what each party knows, is usable. Each of the speaker's statements is converted into its *propositional form* and added to the hearer's knowledge base.

A: **propose!** (up, up)

B: **propose!** (down,down)

A: **why?** (down, down)

B: **declare:** cold(B,now)

A: **so-what?**

B: **declare:** ~cold(B,see-to-it-that(down,now))

A: **so-what?**

B: **declare:** prefers(B,~cold(B,s),cold(B,s))

A: **declare:** has(A,sweater,now)

Speech Acts
and
Propositional Content

Understanding the speaker's meaning is difficult even when speakers are truthful and relevant. Four speech acts are illustrated (bold). These are the surface forms. A's repeated query of ***so-what*** *might be the surface form of some implicit act, other than a request for information.*

More generally, the meaning of an utterance is decided through *plan-recognition*. The general idea of plan-recognition supposes that the hearer has a library of goals, any of which might be ascribed to the speaker. The usual example is the utterance "Do you know the time?" which is a surface request for information (is the knowledge of the time possessed?) but is really a request for the performance of an action (informing what time it is). Depending on the context (e.g., a polite social engagement that has lasted too long, or a chance interaction on a busy street), the utterance is used to identify the probable goals or intentions of the speaker. In the example above, the declaration "I'm cold" was taken to be an invitation for problem-solving. Recognizing what it is that the speaker would like the utterance to achieve depends on what plans the hearer believes the speaker possesses. AI's idea of plan-recognition requires the hearer to consider for each goal what the agent might be planning to achieve in the current situation. If the speaker's utterance can make sense as a part of a plan, then that plan can be attributed to the speaker. Many plans might be attributed, but some are more probable than others. An agent brings knowledge of the situation, knowledge of typical individuals, and knowledge about the particular speaker, as evidence to the ascertainment of probabilities.

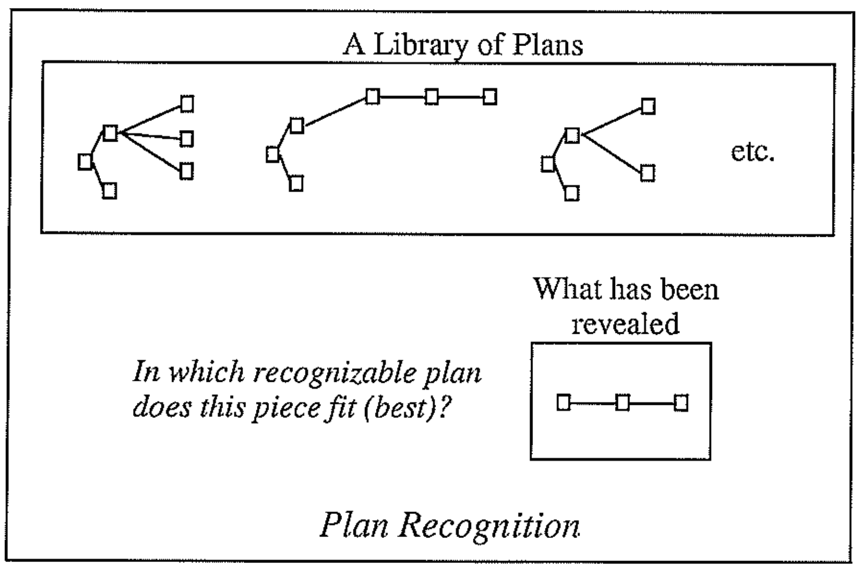

Game theory cannot model language richly because meanings ae limited by the representation of an agent as a bag of utilities. The meaning of an utterance can only be what it reveals about private utilities. Linguists refer to two distinct levels of meaning: *semantics* (what does the expression entail; what are the equivalent ways of wording it?) and *pragmatics* (what is the speaker trying to achieve:

what goals, through what plans?). AI has discovered that in most cases, it is the pragmatics that matters: inferring the speaker's purpose. A good exercise for negotiators and AI researchers would be to enumerate in detail what range of purposes speakers have for their locutions. In some sense, this is what the bulk of the qualitative analysis of negotiation attempts. It could further attempt to relate the analysis to formal models of knowledge, planning, and preference.

Default reasoning is possible when information is organized in a robust way. It might be held, by default, that a person wants the status quo; a cold person wants to be warmer; a cold person in front of an open window wants the window down; a cold person in front of an open window who has a sweater wants the status quo. In the simplest case, objects are arranged in hierarchies of classes, subclasses, sub-subclasses, and so forth. A property that is generally true of the class is *inherited* by any member of the subclass, unless there is a different default for the subclass. If there is a different default, then the subclass is distinguished from the class. This is a preferred way to organize knowledge because it permits parsimonious expression of discontinuity. It is just as good for simple ignorance as it is for complex understanding of variation.

As information is exchanged, parties to a negotiation improve their knowledge of each other's situation. The default inferences that they make are the result of this constant re-classification under defaults.

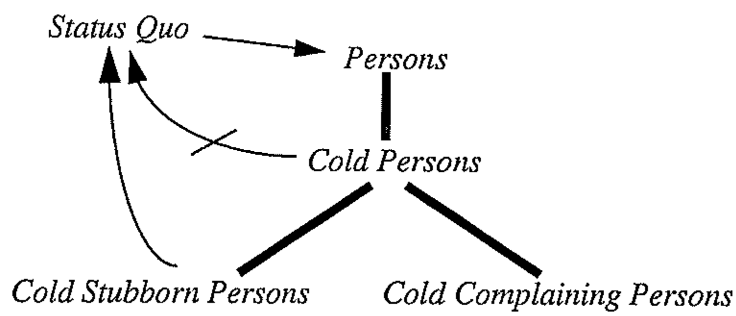

Default knowledge
expressed for a hierarchy of classes and their subclasses:
bold lines are for subsets; arrows for defaults

Note that the default rules and the taxonomies to which they apply are particular to each agent. Agents with different knowledge will represent their knowledge differently.

In the end, default reasoning can be seen to be just a special case of argumentation.[22] But we are trying to keep the pictures simple here.

Argumentation consists of a series of linguistic exchanges between parties, one of whom is trying to establish a proposition *p*, and the other of whom is trying to block *p*'s establishment. The simplest argument for a proposal refers to a prior settlement that both parties think is fair, or are obliged to think is fair. Arguments from

22. Lines of default (or defeasible) reasoning are arguments. There are arguments and counterarguments when there are multiple default paths of reasoning. Here, we are using defaults for the knowledge about individuals, and arguments for the dialogue about fair settlement. To push further, planning is really a form of argumentation. There is a limit to the usefulness of reduction. It would not be too useful to assert that a negotiation simply involves lots of arguing with oneself and arguing with another.

precedent may list the similarities between the current situation and the past situation. Counterarguments may cite dissimilarities between past and present.

The *features* of which there can be similarities and dissimilarities are objective properties of the distribution, such as $5/hour, or of the negotiators, such as one party having made a promise of good faith. The features might also be relative properties of the proposed distribution, such as equal division, or of the negotiators, such as one party being an employer of the other. The features might also be about the relative strategic position of parties, such as the existence of a threat for one and not the other.

Argument, Counterargument, and Rebuttal

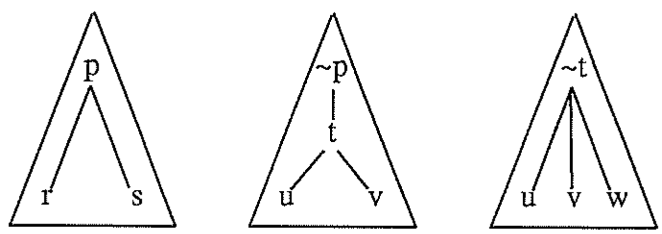

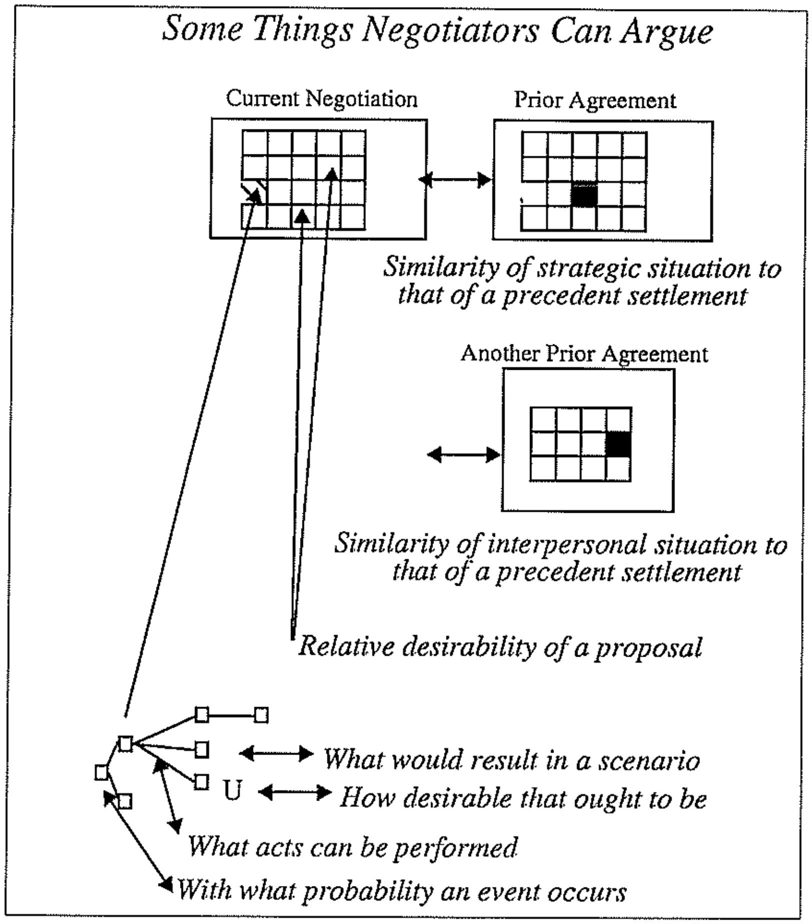

Argument is an example of a *language game* because it is dialogue that is regulated. A language game is defined by a set of rules that the speakers follow. Argument can be a highly regulated game in the sense that quite specific rules can be given that govern who must move when, and govern what are the legitimate moves. In the extreme, the arguments might be required to take a logical form, where claims are connected by reasons, and each successive argument *defeats* the other player's prior arguments.

Negotiators can enter and exit language games, sometimes explicitly and at other times implicitly. They enter by adopting and adhering to conventions for a period of time, and exit by violating those conventions. Those conventions define various layers of protocol, and protocols can be written formally.

A language might consist of several subclasses of sentences, $L = C_1 + C_2 + ... + C_k$. The simplest rules restrict locutions to ranges of L, e.g., $C_1 ... C_5$. Other rules might require that locutions in one class be met with replies in another class, e.g., letting *A.t* refer to speaker *A*'s locution at step *t*, *if X.t is in* C_1, *then Y.t must be in* C_2. Requiring the simple politeness of responding to questions with answers requires rules of this kind; requiring that concessions be met with concessions requires rules of a slightly more complex kind. If some locutions are particularly distinguished, such as threats or automatic responses, there might be rules that discourage their use: *if X.t is not in* $C_{threats}$, *then Y.t should not be in* $C_{threats}$. Of course, this rule is violated unilaterally as soon as a threat is made. If conversational rules have names, e.g., *rule7*, then we can write rules about rules: *if X complies with rule7 during [t1:t2], then Y shall comply with rule7 during [t2:t3] or until X violates rule7.*

AI does not currently advance any theses about these rules, though some are worth posing: does a particular subset of language and set of rules incline negotiators toward problem-solving and social search? Is there a minimal set of conventions that can be adopted, without which it is not worth calling the dialogue a negotiation? Are agreements on rules merely psychological preludes to substantive agreement, or is there some structural purpose that they serve?

Certainly anyone programming a machine to negotiate with a broad range of language will need a theory of how to regulate the use of that language.

IV

A formal model should be simple and memorable, and it should provide the right leverage of analysis. Game theory has and does all of these, but is a model of the wrong phenomenon. It is seeks to explain settlement in terms of agents' self-interests in a static problem formulation with static valuations of proposals. Negotiation includes more important phenomena: dialogue, argument, planning, thinking, focusing, and reformulating.

An AI model is as formal as any mathematical model. In its full detail, it is specified so completely that it can be implemented, step-by-step, with a computer program. In this form, it can often be impressive. What should impress us is not the artifact, but the faithful exhibition of a complex range of behaviors through the proper processing of the important data. We have tried to sketch the data and the processing that might be the components of an AI model of negotiation: plans (which might involve probabilities and utilities) and the generation of possibilities, speech acts (which includes at least proposing, querying, and declaring), default knowledge for problem-solving, arguments from precedent, search, and protocol. Details can be found in AI's literature.

In the end, an AI model consists of a collection of fables. Each fable has its own simplicity, and each is used as appropriate in a larger narrative. The aim is not to prove surprising theorems, but to fit the pieces together, so that the output of one could be the input to the next. The authors hope that practicing negotiators and related theoreticians who have been raised on von Neumann's models will be attracted instead to AI's ideas, and that better depiction of negotiation will be the result.

References

Austin, J. *How To Do Things with Words*, Harvard, 1962.

Benham, J. *An Introduction to The Principles of Morals and Legislation*, Hafner, 1948.

Carberry, S. *Plan Recognition in Natural Language*, MIT, 1990.

Fisher, R., W. Ury, and B. Patton. *Getting to Yes*, Houghton, 1981.

Grice, H.P. "Logic and conversation," in *Syntax and Semantics*, Cole and Morgan, eds., Academic, 1975.

Grosz, B., A. Joshi, and S. Weinstein. "Centering: a framework for modeling the local coherence of discourse," *Computational Linguistics 21*, 1995.

Grosz, B. and S. Kraus. "Collaborative plans for complex group action," *Artificial Intelligence 86*, 1996.

Hart, H.L.A. "Ascription of responsibility and rights," *Logic and Language*, Flew, ed., Blackwell, 1951.

Kahneman, D. and A. Tversky. "Rational choice and the framing of decisions," in *Decision making*, Bell, Raiffa, and Tversky, eds., Cambridge, 1988.

Keynes, J.M. *Treatise on Probability*, Macmillan, 1921.

Lewis, D. "Score-keeping in a language game," *Journal of Philosophical Logic 8*, 1979.

Loui, R.P. "Defeat among arguments," *Computational Intelligence 3*, 1987.

Loui, R.P. "How a formal theory can be normative: implementation versus interpretation," *Journal of Philosophy 90*, 1993.

Loui, R.P. "Process and policy," *Computational Intelligence*, 1998 (in press).

Loui, R.P. "Case-based reasoning and analogical reasoning," *MIT*

Encyclopedia of Cognitive Sciences, MIT, 1998 (in press).

Loui, R.P. and J. Norman. "Rationales and argument moves," *Artificial Intelligence and Law 3*, 1995.

McCarthy, J. and P. Hayes. "Some epistemological problems from the standpoint of AI," *Readings in AI*, Webber and Nilsson, eds., 1981.

Pollock, J. *Nomic Probability and The Foundations of Induction*, Oxford, 1990.

Pruitt, D. *Negotiation Behavior*, Academic, 1981.

Quine, W.V.O. *Word and Object*, MIT, 1960.

Raiffa, H. *Decision Analysis*, Addison-Wesley, 1968.

Savage, L. *Foundations of Statistics*, Dover, 1950.

Searle, J. *Speech Acts*, Cambridge, 1969.

Shafer, G. and A. Tversky. "Languages and designs for probability judgement,"*Cognitive Science 9*, 1985.

Simon, H.A. "Rationality as product and process of thought," in *Decision making*, Bell, Raiffa, and Tversky, eds., Cambridge, 1988.

Skalak, D. and E. Rissland. "Arguments and cases," *Artificial Intelligence and Law 1*, 1992.

Sycara, K. "Problem restructuring in negotiation," *Management Science 37*, 1991.

Turing, A. "Computing machinery and intelligence," *Mind 59*, 1950.

von Neumann, J. and O. Morgenstern. *Games and Economic Behavior*, Princeton, 1944.

Wiener, N. *Cybernetics*, Wiley, 1950.

Wittgenstein, L. *Philosophical Investigations*, MacMillan, 1953.