When and why does haggling occur?

- Some suggestions from a qualitative but computational simulation of negotiation

Bruce Edmonds and David Hales[*]
Centre for Policy Modelling,
Manchester Metropolitan University
`http://cfpm.org`

Abstract. We present a computational simulation which captures aspects of negotiation as the interaction of agents searching for an agreement over their own mental model. Specifically this simulation relates the beliefs of each agent about the action of cause and effect to the resulting negotiation dialogue. The model highlights the difference between negotiating to find any solution and negotiating to obtain the best solution from the point of view of each agent. The later case corresponds most closely to what is commonly called "haggling". This approach also highlights the importance of what each agent thinks is possible in terms of actions causing changes and in what the other agents are able to do in any situation. This simulation greatly extends other simulations of bargaining which usually only focus on the case of haggling over a limited number of numerical indexes. Three detailed examples are considered. The simulation framework is relatively well suited for participatory methods of elicitation since the "nodes and arrows" representation of beliefs is commonly used and thus to be accessible to stakeholders and domain experts.

1. Introduction: About Negotiation

Van Boven and Thompson (2001) (we discovered recently) proposed the basic approach towards negotiation used in this paper.

> *We propose that negotiation is best viewed as a problem solving enterprise in which negotiators use mental models to guide them toward a "solution."*

Where they define "mental models" as …

> *… mental representations of the causal relations within a system that allow people to understand, predict, and solve problems in that system … Mental models are cognitive representations that specify the causal relations within a particular system that can be manipulated, inspected, "read," and "run"…*

According to this picture, negotiation goes far beyond simple haggling over numerical attributes such as price. It is a search for a mutually acceptable solution (albeit at different levels of satisfaction), that is produced as the result of agents with different beliefs about their world interacting via communication until they discover an agreement over action that all parties think will result in a desired state. The motivation behind this is not to discover how to get artificial agents to negotiate nor to determine how a "rational" agent might behave, but to move towards a more descriptive model, which can be meaningfully compared, to human negotiation in order to gain some insights into the processes involved.

With others we distinguish different levels of communication involved in a negotiation process. The most basic is an exchange of offers and requests of actions by the participants (which we call *action haggling*). For example: "If you hold the door open, I will carry the box" or "Can anyone tell me where the nearest shop is, as we have run out of coffee". However in many negotiations such haggling actually takes up a very small amount of the time spent during negotiations. It seems that a lot of time is spent discussing the participants' goals and what the participants believe about the target domain. Thus the next of the two

levels we distinguish are the exchange of opinions of what the target domain is like, in particular how actions of the participants (and others) might change the state of the target domain, we call this *viewpoint exchange*. For example: "If a flood plain was built, this would reduce the severity of any flooding", or "Even if we build higher dykes, this will not prevent all flooding". The third level is the communication and reformulation of goals. This is perhaps the most important but least understood aspect of negotiation, which we will call *goal exchange*. In this paper we do not consider goal *change* or *reformulation* but concentrate on what must be the case concerning the viewpoints and goals for haggling to occur. A fourth level might be meta-communication about the negotiation process itself. The levels are summarised in Table 1.

Table 1. A summary of different levels that can be involved in a negotiation

Level Name	What Communication Concerns	*Example*
Actions	Offers and counter-offers as to possible actions	*I will carry the box if you open the door for me*
Beliefs	What is and is not possible and what states are considered	*Even if we build high flood-defences abnormally high rain could still cause flooding*
Goals	The goals of participants and what states are preferable	*I know you consider this is too expensive but consider how much you will save in the future*
Meta-issues	Suggestions and comments about the negotiation process itself	*We are not getting anywhere, lets go and have lunch*

We *reject* the following assumptions/approaches to modelling negotiation since we have not seen any evidence that they are true of human negotiation:

??   that the participants necessarily have the same view of the world in terms of the causation therein, i.e. they have *different* beliefs as to what actions or events cause what results in the target domain;

??   that the participants necessarily have any sort of "joint utility" or "social rationality" that drives them to seek for agreement or for the common good above their own;

??   that the participants necessarily have any sort of knowledge about others' beliefs or goals, except as revealed or implied by the communications of other;

??   that the participants necessarily agree upon the description of the current world state;

??   that the participants necessarily judge similar resultant states in similar (or even related) ways;

??   that the participants are necessarily able to think out all the possible negotiation possibilities or consequences of their own beliefs.

However we do make (at least) the following assumptions/take the following approaches with regard to modelling negotiation:

??   that the participants have common labels for actions and each other so that they can communicate, since we do observe that people do manage to talk about (and sometimes agree upon) agreements as to what actions will be done;

??   that each possible world state considered by the participants is judged on a limited number of independent properties (e.g. cost, danger, environmental damage) – the dimensions are not necessarily numerical and could well be qualitative characteristics, e.g. colour (this is because it seems computationally implausible that an unlimited number of properties are consciously considered);

??   that the decisions of individuals can be expressed as some sort of program/process based on their beliefs and these judgements, though this may be complex and involve such processes as looking ahead to possible states; making arbitrary choices; can involve qualitative judgements; and even

when alternatives are "weighed" in a way that suggests a numerical function this is not necessarily continuous or convex etc.;

?? and that individual rounds of communication occur off-line, that is no actions are taken during the communication involved in the negotiation of each agreement and that also the world state does not change during a round of negotiation but between rounds (when the actions are performed).

In the model presented we make a few further restrictions purely to make the task feasible:

?? that the utterances that form part of a round of negotiation are accessible and recallable by all;

?? that participants are honest when they make utterances (!);

?? that actions taken are public, that is they are perceivable by all participants;

?? that participant's individual goals or judgements do not change over the course of the model.

All of the above are known to be violated in many negotiations, particularly in multi-party (i.e. more than two participants) where the pattern of alliances and sub-negotiations may involve much cloaking of actions and communication from other participants. How and when alliances and negotiations form is crucial – some early results of Scott Moss seem to indicate that even when one is limiting the scope to multi-dimensional numeric haggling that it is *very* much more difficult to obtain agreement between more than two parties than between only two.

2. Why a qualitative simulation?

One aspect of the approach that is of particular note is that, despite the fact that it is a completely formal computation, it is almost entirely (and can be entirely) a qualitative simulation. That is the design of the simulation (and hence the resulting computation) does not *require* numbers or calculation. This will not be a surprise to logicians or mainline computer scientists, but it has become the overwhelming default to use the calculation and comparison of numbers as an integral part of the modelling of learning and decision making of actors. However we claim that the use of numbers in this way is often simply a result of laziness – we often use numbers as a stand-in for qualitative aspects that we do not know how to program or have not the time to program. I (Bruce Edmonds) have often had the experience of asking a modeller why they used numbers for a particular purpose, only to discover that it was because they could not imagine any other way of doing it. The promotion of utility theory is an example of this, and must take the blame for the widespread unthinking uptake of this approach.

The fact is that it is difficult to satisfactorily use numbers in this way. Numbers are a useful way of representing certain kinds of things: temperature, distance, money or physical time. This is due to the uniformity of what is being represented (which is often due to an underlying averaging or maintenance process). If the nature of what you are representing is such that it is fundamentally not uniform – i.e. it is part of its intrinsic nature that the qualitative differences matter – then its representation by a number ,at best, requires demonstration that this does not distort the results and, at worst, should be simply condemned. We would suggest that properties such as: variety, information, usefulness and relevance are not amenable to being represented as numbers without considerable care and justification in very restricted contexts (unless the nature of these are being redefined such that their names are no longer appropriate). For a more detailed discussion of this issue see (Edmonds 2003).

In the simulation to be described the agents represent the domain they are negotiating about using a network of nodes and arcs representing relevant possible states of the world and actions respectively. These states are judged by the agents when deciding what to do; what offers to make and which to accept. The simulation allows judgement because each node has an associated set of properties. The agent has an algorithm which returns the acceptability of the nodes and allows nodes to be compared with each other. These properties could be numeric (e.g. the price of something), but could be other types (e.g. colour). The goals of the agents are to attain to more preferable states than the one that currently hold and are thus implicit in the results of the judgement algorithm upon the properties of nodes. The structure of the simulation allows for the use of numeric indicators of desirability to be attached to nodes, but does not require them. It also allows for a mixture of numeric and qualitative properties to be used. Suffice to say

that any modelling of cognitive processes which suggests that human actors weigh numeric measures would need to be justified.

3.   An Outline Description of the Simulation

There follows a brief version of the basic structure of the simulation. A fuller specification may be found in Appendix 1. To understand how the simulation works it is probably easier to first look at the examples described in Section 4 and then return to this description, going to Appendix 1 when more details are required. However it is traditional to describe the structure of a simulation before presenting results, so this now follows.

There is a fixed set of agents (the *participants*). They negotiate in a series of negotiation *rounds*, each of which is composed of a subsequence of time instances in which utterances can be made by participants. All utterances are public, that is accessible to all participants. When a set of actions is agreed or no utterances made then the round ceases. Between rounds actions that are agreed are taken – these are public actions, known to all participants. When a round occurs that is identical to the last round and no actions are agreed, the simulation ceases. The simulation output consists of the utterances made and the actions agreed and taken. It is a development of (Hales 2003), this version was roughly aligned with Hale's version (but written in SDML rather than Java) and then extended.

Each participants has the following structures internal to itself and not accessible to others (unless the agent reveals them in an utterance):

   ?? A set of *properties* in which possible world-states are judged;

   ?? A (possibly complex) algorithm which, given the internal information and *current world states,* results in an overall judgement on that set of states;

   ?? A network composed of set of nodes representing what the participant considers possible world states, each node having: a label, a set of properties, a list of possible actions for the agent, and a set of arcs to other nodes (each arc having a condition in terms of actions that have occurred).

The utterances that agents can make to each other are of the following kinds:

   ?? Can someone please do [*these actions*] so that we can read [*these states*];

   ?? If [these actions are done] then [I will do these actions];

   ?? I agree to [I will do these actions] if others agree to [these actions are done].

In addition there are the following reports:

   ?? [*agent name*] has done [*action name*];

   ?? [*agent name*] is in state: [*state name*].

Thus, the input to the simulation is a specification of each agent's beliefs about the world and the output is a *transcript* showing the utterances and actions that occurred.

Basically the simulation works as follows. At the start the agents and their beliefs are initialised from a text file which specifies these (called the *viewpoint file*). Then each agent performs a limited search on its own network for states that it judges are preferable to the current one – if it finds such it makes a conditional offer composed of its own and others' actions necessary to reach that state (if there are none of its own actions this is a request, if no others' actions are needed it simply commits itself to that action without communication). If others make a conditional offer it considers possible combinations of offers to see if there are possible agreements and if there is signals its potential agreement to it. If there is an agreement made by another agent that it judges is acceptable to it, it signals its potential agreement to that. If all the necessary parties have indicated their potential acceptance of an agreement then it becomes binding and all agents commit themselves to doing the actions that is their part of that – the simulation now enters an action phase. During action phases all agents do actions as soon as they become possible – actions may have the effect of changing the current state in agents. When actions have finished then negotiation may begin again etc.

The simulation can be seen as a minimal constraint upon what negotiations could occur, since it would be possible to get almost any output given enough tinkering with the agent's beliefs. In this sense it is akin to an agent programming language. What it does do is relate the agents' beliefs to the viewpoint file that is output in an understandable and vaguely credible manner. The representation of the agents' beliefs with: nodes (states of the world); arcs (the actions between states of the world); and judgements along multiple criteria (the judgement dimensions) are designed to be fairly easy to present as blobs-and-arrows pictures, and thus be amenable to participatory input and criticism. This is in contrast to many agent negotiation set-ups which are couched in pure economic or logical terms.

An interesting aspect of this simulation is that (if the values of the judgements are Boolean or string valued) that it does not *require* the use of numbers at all. This is more of a qualitative simulation than a numerical calculation. Hopefully this will make it more amenable to qualitative interpretation and criticism.

4. Example Runs

We now consider a number of examples using this simulation.

Example 1 – Making Iran Happy

Towards the start of Philip K. Dick's novel, *Do Androids Dream of Electric Sheep?* (aka "Bladerunner"), there is a scene involving a couple, Rick and Iran and their mood-altering device. One can turn a control to make oneself happier or sadder. Rick has turned the control down to make himself less happy and is now so depressed that she does not use the control to make himself happier. Then Rick attempts to rectify the situation by negotiating with Iran.

Let's start by considering just the belief network of Iran. There are two states, happy and sad; happy has the property that it is satisfactory and sad that it is not. Both prefer satisfactory to unsatisfactory states. There are two actions: down to turn the control down so that Iran is sad and up so she is happy. There are (at least) two possible reasons why Iran does not use the control to bring herself into a satisfactory state: (A) that in his depressed state she is not able to make herself turn the control (although she knows this would make her happy); or (B) that he does not think (in his depressed state) that using the control will be of any help – he is quite capable of using the control but he does not think that it will help. The two belief networks representing these are shown in Figure 1. In this the nodes are the states and the arcs are the possible transitions between these states as the result of actions. The properties of the states are in brackets below the node label. Beside the nodes are those actions that are possible for that agent there.
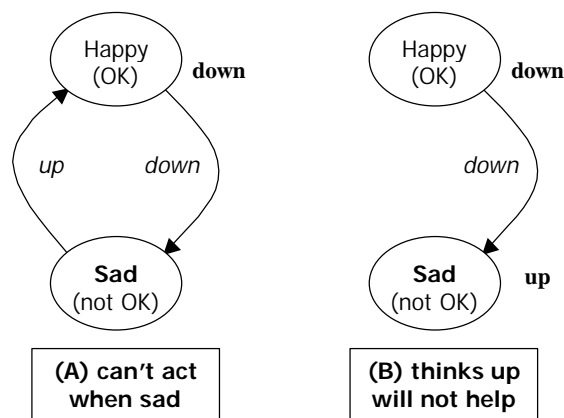

Figure 1. Two belief networks for Iran

Thus, in the first case (A) an up action is there, which reflects that belief of Iran that is up were done she would reach a more desirable state, but up is not one of the actions that is possible for her in the sad state. In the second case up is a possible action from the sad state, but Iran does not do it because she does not think that it will cause her to get to happy. Of course, when Iran is alone it makes no material difference which of these cases holds, but the situation changes when someone else is involved (in this case Rick).

Let us suppose that Rick knows about the control and how it works and can adjust the control, but that she does not necessarily which state Iran is in or what *his* view of the situation is. Rick's belief network is illustrated in Figure 2.
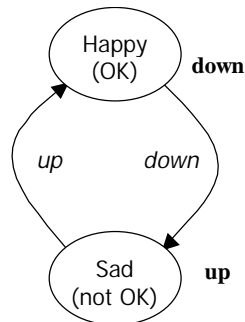

Figure 2. Rick's (more complete) belief network about Iran

Now when Rick interacts with Iran there are four possibilities corresponding to whether Iran has belief network (A) or (B) and independently whether Rick assumes Iran is in state happy or sad. We will label these four runs as HA, HB, SA, and SB (happy with network A, etc.). The template viewpoint file for these and the results are listed in Appendix 1.

Table 2. Summary of results of Example 1

	Case A (can't act)	Case B (not worth it)
Case H (thinks Iran is happy)	Iran requests help but Rick does not think this will help	Nothing occurs
Case S (thinks Iran is depressed)	Iran requests help and Rick turns up dial to make Iran happy	Rick turns up dial on own accord to make Iran happy

Here whether Iran requests Rick to help depends on whether it is case (A) or (B) and whether Rick turns up the dial depends on whether he realises that Iran is depressed. It is notable that there is a different outcome in cases SA and SB if Rick has no preference between Iran being happy or sad, for in case SA a request is made of him, so if it is possible and is of no cost he might turn the dial up, whereas in case SB it is solely on his own considerations that he turns the dial up so if he had no preference he might not bother to do so.

Buying a Used Car

This is an example to illustrate the negotiation in a simple transaction of a purchase. In this simple version there is a low price and a high price that could, in theory, be paid in return for the car. In this version one of the (two) properties of nodes could a number, corresponding to the amount of (extra) monetary changes. This is justified because the amount of money is a number. However a fuller model might eliminate this by representing the relevant trade-offs (or opportunities) that that money meant to the actors at the time. Since we only deal with two possible prices (*cheap* and *expensive*). The basic belief networks are shown in Figure 3.
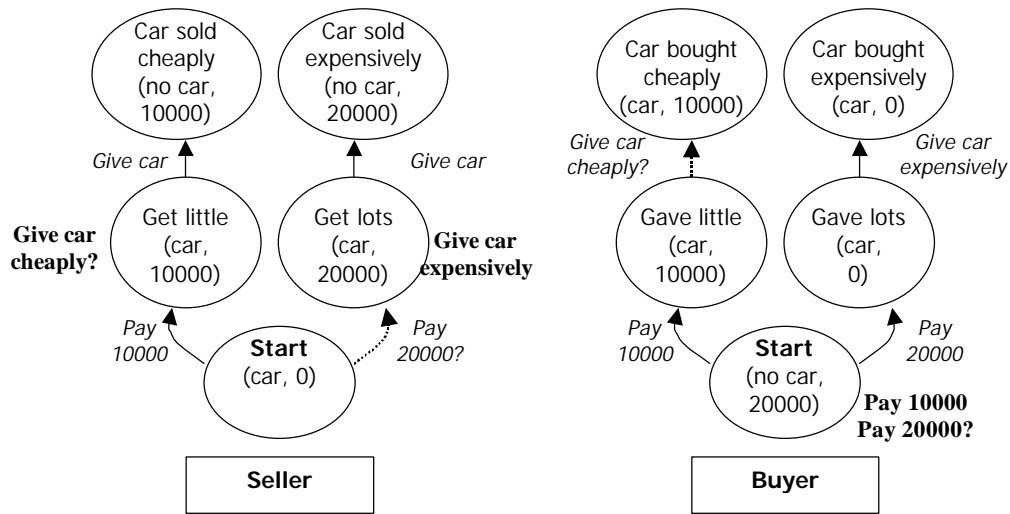
Figure 3. Belief networks of seller and buyer

There are clearly a number of ways of representing the Buyer and Seller's beliefs using this method – we have chosen one. Let us assume that for the Seller the states are ordered thus: Start < Car sold cheap < Car sold expensively < Get little < Get lots; and that for the buyer: Start < Car bought expensively < Car bought cheaply < Get car. There are number of possible variations here: the Seller could mentally rule out the action of Give car cheaply from the state Get little (i.e. only 10000) or not depending on whether this was considered as a possible action; likewise the buyer might or might not consider paying 20000 at the state Get car as possible. Corresponding to these is the existence or absence of arcs in the belief networks of the other agent. So the Seller might or might not have an arc from Start to Get lots depending on whether the Seller thinks that such an action is possible and the Buyer might or might not have an arc from get car to Car bought cheaply for the action Pay 10000 depending on whether the Buyer thinks it will be possible to purchase the car for only 10000.

When this is run there is some initial exploration concerning whether the Seller will give the car for nothing and the Buyer give money for nothing – this is because the agents do not know these would not occur (as we would know).

Given the above there are 2 ? 2 ? 2 ? 2 = 16 possibilities:

?? Seller does (1st u) or does not (1st c) think Buyer would pay 20000 for car;

?? Seller would (2nd u) or would not (2nd c) give the car for 10000;

?? Buyer would (3rd u) or would not (3rd c) pay 20000 for the car;

?? and Buyer does (4th u) or does not (4th c) think Seller would give car for 10000.

Thus the viewpoint file labelled example2-cucu is the viewpoint file where the 1st and 3rd option is commented out (hence the c) and the 2nd and 4th options are left uncommented (hence the u) – this corresponds to the case where: the seller does not think the buyer would pay 20000; the seller would sell the car for 10000; the buyer would not pay 20000; and the buyer does think the seller would sell for 10000. The template for these scripts (with options to comment out the relevant lines) and some example results are listed in Appendix 2. Table 2 below summarises the results of the 16 possibilities.

Table 3. Summary of results from example 2

	Seller does not thinks buyer would pay 20000 *and* would not give car for 10000 (cc--)	Seller does not thinks buyer would pay 20000 *and* would give car for 10000 (cu--)	Seller thinks buyer would pay 20000 *and* would not give car for 10000 (uc--)	Seller thinks buyer would pay 20000 *and* would give car for 10000 (uu--)
Buyer would not pay 20000 *and* thinks seller would not sell for 10000 (--cc)	No agreement	No agreement	No agreement	No agreement
Buyer would not pay 20000 *and* does think seller would sell for 10000 (--cu)	No agreement	*Car Sold Cheaply*	No agreement	*Car Sold Cheaply*
Buyer would pay 20000 *and* thinks seller would not sell for 10000 (--uc)	No agreement	No agreement	Car Sold Expensively	Car Sold Expensively
Buyer would pay 20000 *and* does think seller would sell for 10000 (--uu)	No agreement	*Car Sold Cheaply*	Car Sold Expensively	Car Sold Expensively?

Unsurprisingly, the conditions for the car being sold expensively is that the Buyer would pay 20000 and the seller thinks that the buyer would pay 20000. This is so even if the buyer thinks that the seller would less for less and the seller would be willing to sell for less. This is because of the asymmetry of the belief networks where the payment happens before the handing over of a car (never the other way around); thus the seller explores whether the buyer is willing to pay money without giving the car which delays his more credible offers; this has the effect that the buyer comes down to an expensive offer before the seller makes a cheap offer. The condition for a cheap sale is that seller would sell for 10000 and the buyer knows this, except for the case discussed immediately above. Although this is rather an artificial source of delay in this case, delaying making offers that are less good for oneself is an established negotiation tactic. Most models of bargaining on prices *only* centre on this case (i.e. those represented by the bottom right-hand corner of Table 3.

It is interesting to note that no agreement can result even when the seller would be willing to sell the car for 10000 and the buyer willing to buy the car for 20000 because of their beliefs about what the others will do (e.g. case CUUC). In this example it is clear that the beliefs that each have about the possibilities that exist can make a critical difference to the outcomes.

A Public Negotiation Concerning Flood Prevention Methods

This example is loosely derived from reports by members of ICIS at Maastrict about the Maaswerken negotiation process (van Asselt et al. 2001) designed to achieve a consensus about flood prevention measures in the Maas basin. Here there are some flood prevention measures: building dykes; extending flood plains.
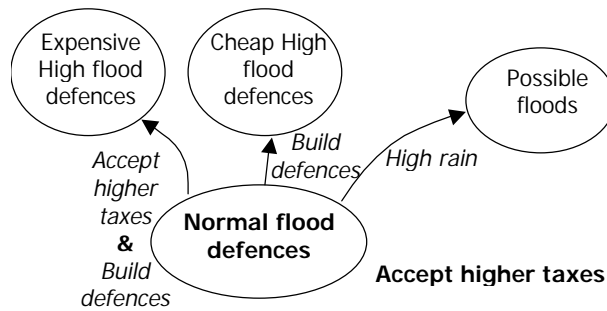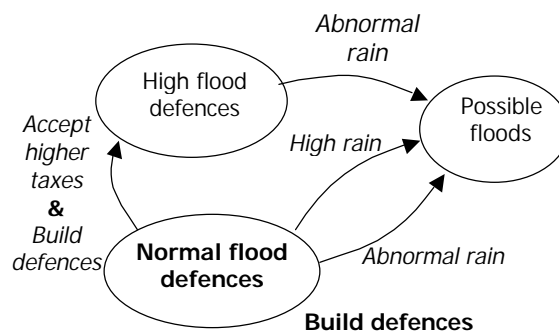
Figure 4. A citizen's view (simple)

Figure 5. The government's view (simple)

For both citizens and government it is overwhelmingly important to prevent getting to the state Possible floods anytime in the future. The citizen thinks it is possible to prevent this by getting to one of the high flood defence states since even High rain will not then cause floods. The government thinks there is a possibility of Abnormal rain which the citizen does not think possible. Hence the government does not think that attaining to the state of High flood defences will prevent the possibility of getting to Possible floods in the future. Other things being equal the citizen prefers not to accept high taxes and the government does not want to build high flood defences.

In this case there is quickly a stalemate (see results in Appendix 3) since in the government's view building high flood defences would not prevent any possibility of flooding because abnormally high rain would overwhelm them. The citizens would prefer high flood defences even at the cost of higher taxes because they think it would prevent the possibility of flooding (since they do not believe in the reality of abnormally high rain).

However, if the view of both parties is expanded to include a new possibility, namely flood plains which are environmentally attractive and will mitigate, but not prevent flooding then the outcome can be very different. These expanded views are shown in Figures 6 and 7. This is the outcome despite the fact that the citizens would prefer high flood defences which (they think) would prevent all flooding. The fact that citizens and government prefer flood plains to the current position means that they can agree upon that. How such "expansions" or changes in beliefs can occur can make the difference between a failed negotiation and one that succeeds. This model does not have mechanisms for such "persuasion" but has the facilities that would make such an extension easy to implement. However, according to my search of the literature very little is known about why, how and when people change their beliefs.
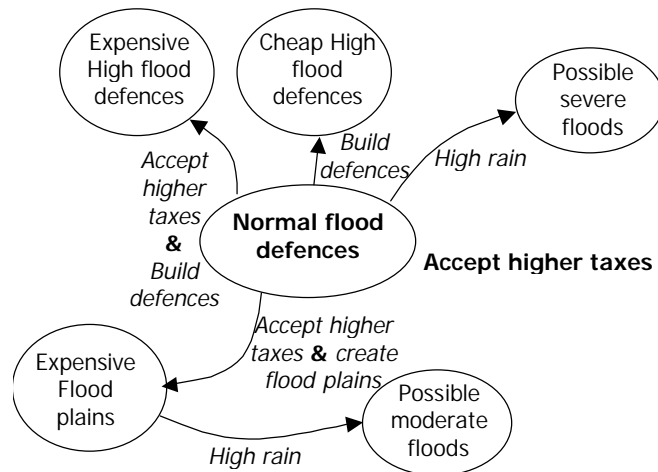
Figure 6. An extended citizen's view
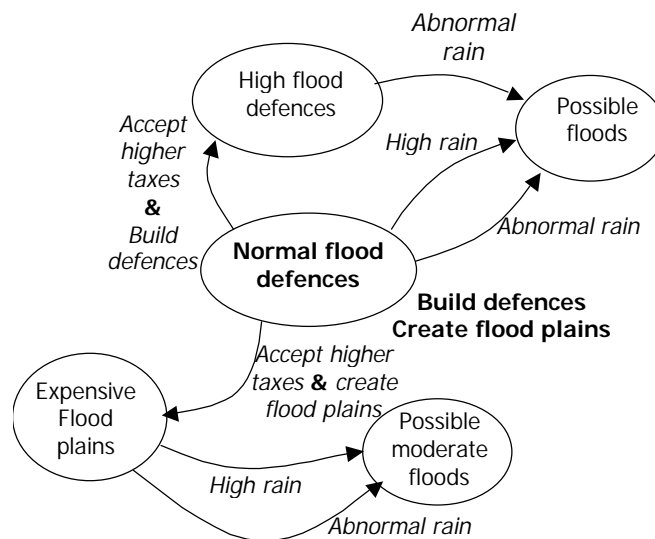
Figure 7. The extended government's view

5. Discussion

Examining the working of this model allows one to form hypotheses about the conditions needed for various outcomes and processes to emerge. These hypotheses can be thoroughly tested in simulation experiments and maybe even proved by examination of the structures and algorithms involved (this is future work!). However the point of extracting them here is that they are candidate hypotheses about real negotiations. Their best use would be to try to validate them against observations of real negotiations. This in turn might suggest a better simulation model and hence better hypotheses, etc. In short this model of simulation is intended as a tool to inform (and hence enrich) our observations of real negotiations. A preliminary scan of the literature on human negotiation that we have made indicates that they are broadly compatible with observations but that most studies are heavily biased by economic frameworks and kinds of model which use drastic assumptions and focus almost entirely on price.

Clearly the participants do not have to have similar beliefs or similar goals for meaningful haggling to occur. However they do have to communicate about something and, minimally, this must include actions. Thus the first condition is this:

Condition 1:   *That they both have (possibly different) understandings of the actions they discuss.*

This condition simply means that communication about actions is possible. It does not mean that the participants actually try to communicate. For this to occur in this model they must need or want some "action" to occur that they can not do themselves (in this model each state change is accompanied by what is called an action so, for example, an agent might want someone else-to-do-an-action rather than self-do-action because it might get to a similar state but without a cost property). Thus the second condition is:

Condition 2:   *That at least one agent cannot get to a state it has considered and prefers using only its own actions.*

Of course, these agents are not necessarily perfect reasoners and so do not necessarily do a complete search of their belief networks. So there may well be states they would prefer involving the actions of others but so distant from the currently holding state that they do not consider it or request actions that might take them there. In this model others only offer actions (possibly conditionally) if they know that these actions might be wanted by others – they don't just go around offering actions in case someone wanted them. Thus participants have to ask for any actions they might want.

Of course, most people would not consider haggling to consist of only one participant requesting actions, but that other parties do so as well. Which would lead to:

Condition 3:   *That at least two agents cannot get to a state they have considered and prefer using only their own actions.*

Thus at least two different requests for actions will be made. Requests for actions might then lead to possible conditional offers of actions. Once these requests have been made there are several possibilities: these requests might be unacceptable to each other because they would not result in states that are preferable for the other; they might not seem to the others that they lead to any states at all; they may *think* that there is not acceptable set of actions that would lead to preferable states; or they may come across a set of possible actions that lead to a state mutually considered preferable and then a deal may be made. For haggling to occur it is not necessary for it to be successful, it may well fail even when there *is* a possible set of actions that would lead to a mutually preferable state (as in the car selling example CUUC).

Some may only consider that haggling is *really* occurring when there is more than one possible set of actions that the participants *might* agree upon (that is they lead to preferable states in both belief networks) but some of these are more preferable to one party and other to the other party. The haggling in this case is not a question of *searching* for a possible agreement but *determining which* of the possible agreements will occur. If one has this picture of haggling then there is a further condition:

Condition 4:   *That there is more than one set of actions which would result in states that are preferable for all parties.*

The primary means in this model for this determination is by dissembling to the other about what is possible so that they accept an agreement that is suboptimal for the other. Thus the car salesman might achieve a better sale through convincing the buyer that he would not sell for 10000 even though he would if there was no other choice. Of course this strategy is risky as the seller might end up with no agreement at all.

Thus in this model there are two sorts of negotiation:

1. Where the parties are *searching* to see if an agreement is possible and
2. Where the parties think more than one agreement is possible and are trying to determine *which* deal.

When a deal is better than no deal then in case (1) it is to the advantage to be honest about what is and is not possible, but in case (2) it can be advantageous to be deceptive about the real possibilities. The case of (2) can be dangerous if the deception means that it *seems* to the parties that no agreement is possible.

Case (2) most closely corresponds to what people commonly refer to by "haggling". This touches on the question of trust and is consistent with (Moore and Oesch 1997) who observed:

> *The good news from this study for negotiators is that there is a real cost to being perceived as untrustworthy. Negotiators who negotiate fairly and earn a reputation for honesty may benefit in the future from negotiating against opponents who trust them more. The bad news that comes from this study is that it is the less trusting and more suspicious party who will tend to claim a larger portion of the spoils.*

It is also interesting to compare this analysis to the observations of negotiations at the Marseille fruit and vegetable market made in (Rouchier and Hales 2003). There, it was observed that there were two kinds of buyer: those who were looking for a long-term relationship with sellers, so as to ensure continuity of supply and those who were searching for the best price. For the former kind, once a relationship had been formed, it is more important for both that they reach an agreement rather than they get the very best price – that is the negotiation, although it may involve some haggling as to the price, was better characterised as a *search* for agreement. This involved a series of 'favours' to each other (the seller giving excess stock to the buyer and the buyer not negotiating about price). The later kind searched the market for the best price, swapping between buyers depending upon the best price of the moment. On the whole, the negotiation with each seller was to find the best price among the possible ones, since if the negotiation failed they can always go to another seller. Thus this is better characterised by the second type – finding the best deal among those possible. This is at the cost of sometimes going without obtaining some product when there is a shortage (since the sellers favour their regular customers in these circumstances).

The structure of the simulation is such that these conditions form a set of hypotheses which it is conceivable that it could be tested using participatory methods and observations of negotiations. A game could be set up where the subjects have to negotiate their actions via a limited computer-moderated script – a web version of the game *Diplomacy* in a form similar to that of the on-line 'Zurich Water Game' might be suitable (Hare et al. 2003c). At suitable stages the subject's views of the game could be elicited in the form of blobs-and-arrows diagrams, possibly using something akin to the hexagon method used in the Zurich water game (Hare et al. 2002a).

Such an investigation might lead to further developments in the simulation model presented above which might, in turn, prompt more investigations as is indicated by (Hare et al. 2002b).

6. Conclusion

We have presented a simulation model which captures aspects of negotiation as the interaction of agents searching for an agreement over their own mental model. The aim of this search for each agent is to achieve states that are preferable to the current one for each participant. Specifically the simulation relates the beliefs about the action of cause and effect in the relevant domain to the resulting negotiation dialogue. The simulation requires that each agent has some means of comparing states to decide which it would prefer, but this does not have to be based on any unrealistic numerical "weighing" of possibilities.

The model highlights the difference between negotiating to *find* any solution and negotiating to *obtain the best* solution from the point of view of each agent. we speculate that the former occurs when it is more important to get an agreement than to risk this trying to get the best agreement, and the later case occurs when there is little risk of no agreement or it is reaching an agreement is less important. The later case corresponds most closely to what is commonly called "haggling".

This approach also highlights the importance of what each agent thinks is possible in terms of actions causing changes and in what the other agents are able to do in any situation. Such views can have a critical effect on the outcomes of negotiations. It seems plausible that this is the reason that the (possibly false) signalling of what is and is not possible is often used to try and obtain a better outcome.

This simulation framework greatly extends other simulations of bargaining which usually only focus on the case of haggling over a limited number of numerical indexes (e.g. price and quantity). The model could be easily extended to include belief extension/change, goal reformulation and even some meta-communication mechanisms. However before this is done there more is needed to be discovered about

how and when this occurs in real negotiations. This model suggests some directions for this research – the simulation framework is relatively well suited for participatory methods of elicitation since the "nodes and arrows" representation of beliefs is commonly used and thus accessible to stakeholders and domain experts.

7.   References

van Asselt, M.B.A. et al. (2001) Development of flood management strategies for the Rhine and Meuse basins in the context of integrated river management. Report of the IRMA-SPONGE project, 3/NL/1/164 / 99 15 183 01, December 2001. http://www.icis.unimaas.nl/publ/downs/01_24.pdf

van Boven, L. and Thompson, L. (2001) A Look Into the Mind of the Negotiator: Mental Models in Negotiation. Kellogg Working Paper 211. (http://www1.kellogg.nwu.edu/wps/SelectDocument.asp?)

Conte, R. & Sichman, J. (1995), DEPNET: How to benefit from social dependence, Journal of Mathematical Sociology, 1995, 20(2-3), 161-177.

Conte, R. and Pedone R. (1998), Finding the best partner: The PART-NET system, MultiAgent Systems and Agent-Based Simulation, Proceedings of MABS98, Gilbert N., Sichman J.S. and Conte R. editors, LNAI1534, Springer Verlag, pages 156-168.

Hales, D. (2003) Neg-o-net – a negotiation simulation testbed. CPM report, CPM, MMU, Manchester, UK. (http://cfpm.org/pubs.html)

Hare, M.P., D. Medugno, J. Heeb & C. Pahl-Wostl (2002a) An applied methodology for participatory model building of agent-based models for urban water management. In (Urban, C) 3rd Workshop on Agent-Based Simulation. SCS Europe Bvba, Ghent. pp 61-66

Hare, M.P.,J. Heeb & C. Pahl-Wostl (2002b) The Symbiotic Relationship between Role Playing Games and Model Development: A case study in participatory model building and social learning for sustainable urban water management. Proceedings of ISEE, 2002, Sousse, Tunisia

Hare, M.P., N. Gilbert, S. Maltby & C. Pahl-Wostl (2002c) An Internet-based Role Playing Game for Developing Stakeholders' Strategies for Sustainable Urban Water Management : Experiences and Comparisons with Face-to-Face Gaming. Proceedings of ISEE 2002, Sousse, Tunisia

Moore, D. and Oesch, J. M. (1997) Trust in Negotiations: The Good News and the Bad News. Kellogg Working Paper 160. (http://www1.kellogg.nwu.edu/wps/SelectDocument.asp?).

Moss, S. (2002) Challenges for Agent-based Social Simulation of Multilateral Negotiation. In Dautenhahn, K., Bond, A., Canamero, D, and Edmonds, B. (Eds.). *Socially Intelligent Agents - creating relationships with computers and robots*. Dordrecht: Kluwer.

Rouchier, J. and Hales, D. (2003) How To Be Loyal, Rich And Have Fun Too: The Fun Is Yet To Come. 1[st] international conference of the European Social Simulation Association (ESSA 2003), Groningen, the Netherlands. September 2003.

*Acknowledgements

Mario Paulucci (in suitably smoke-filled rooms) where the ideas were further developed. Some of the ideas concerning goal-exchange were also used in ISTC/CNR's conceptual model, *Pandora*. Later David did the first implementation of *Neg-o-net* (Hales 2003) for a meeting of the FIRMA project (overnight in another smoke-filled room). The name *Neg-o-net* was Rosaria's idea. Useful feedback was provided by the FIRMA partners. Not wanting to be left out of the fun, I reimplemented *Neg-o-net* in SDML. It is an extension of this model that is presented in this paper. There have been many further discussions about negotiation between Juliette, David and I. David has now given up smoking.

Appendix – A specification of the simulation

Each simulation is composed of the negotiation engine (which is the basic algorithm) and a specification of the agents and their beliefs held in a text file. The engine reads in and parses the file to initialise the simulation, the negotiation then proceeds based on this, outputting the result to a transcript and/or text file.

References and Sources

The basic source for the structure of this simulation was based on discussions between Juliette Rouchier, Scott Moss and me in 2000. This simulation is closely based upon (Hales 2003) which was written for a FIRMA meeting in 2002. That simulation was close to the present simulation – the main difference is that in that version the comparison of states was limited to the weighted comparison of numerical indicators whilst this version is fairly flexible about the comparison mechanism. (Boven and Thomson 2001) have independently proposed some of the same sort of structure based upon their observations of negotiations. For a more detailed history see the Acknowledgements above.

Static structure

There is a fixed number of agents, which represent the negotiators. The environment they exist in acts, to a very limited extent as a "chair" of the negotiation.

Each agent has:

- ?? a set of the names of the properties with which the states are labelled (qualitative or numeric);
- ?? a belief network which represents the cause and effects that it thinks will occur in the domain they are negotiating about, this is composed of;
    - o A set of nodes which represent states, each of which has:
        - ✍ A unique label;
        - ✍ A set of actions that can be done by the agent when that state holds;
        - ✍ A set of values for each of the properties;
    - o A set of arcs between nodes, each of which has:
        - ✍ A unique label;
        - ✍ A condition in terms of actions that specifies when the arc can be traversed.
- ?? An algorithm that decides which state of a set of states would be preferable based on the belief network and the properties of the states;
- ?? The label for the state currently considered to hold;
- ?? A memory of the past offers, agreements, commitments and actions made by all agents (including itself);
- ?? A set of *judgement properties* in which world-states are judged;
- ?? A set of strings which are syntactic sugar for nodes and actions – these are substituted for the actions and states in the output transcripts.

Future versions might allow for an "actual" network representing the "real" causality that operates in the world, in this version there is no observation of the outside world only deduction of what would occur as determined by each agent based on their own beliefs.

Temporal structure

Each simulation is divided into a number of negotiation rounds. These occur up to the inbuilt maximum or until it is obvious that no progress is being made when the simulation finishes. Each of these rounds is composed of 1 or more synchronous communication rounds. These communication rounds continue until an agreement is reached or an action is taken. Actions are taken (if at all) at the beginning of negotiation rounds – if this is the case no further communication occurs in that round. This has the effect that the simulation can be interpreted as cycling through two phases:

?? A phase of requests, conditional offers and offers of agreement, which continue until no more offers/requests are made or all the necessary parties agree to the same agreement, then

?? A phase of actions continues where agents do the actions they have committed themselves to by concluded agreements as and when these become possible until all such actions have been done.

Agents decide on present actions and communications independently and in parallel to each other based on what has occurred in past rounds and communication cycles. Thus if one agent makes an offer in one cycle it can only be responded to in the next cycle by another agent.

Dynamic structure

In the version reported here, what changes is:

?? The world states that each agent considers to hold (these are effected by the actions that are done);

?? The offers and agreements made by the agents;

?? And the actions done by the agents.

Future versions might allow for the beliefs of the agents to change, based on what they are told and what they experience, in this version they do not change.

Key parameters and options

The specification of the agent names, their: property names, agents' belief networks, preference algorithm, and initial state are input as a text file (see some of the examples in the following appendices).

Other parameters include:

?? How many arcs of their own belief networks the agents will search for preferable states;

?? The maximum number of negotiation rounds (optional);

?? The maximum number of communication cycles within a round (optional).

The algorithm

The outline algorithm is as follows:

Read in and parse viewpoint file then set up agents with their beliefs and initial world state
The agents plan, decide and act in parallel with each other.
Until the number of rounds reaches the maximum or two consecutive rounds are identical.

Each agent has the following algorithm which it executes in parallel to the other agents:

Repeat negotiation round:
If I have agreed to an action and it is possible then do it
While no actions are done and no agreement finalised do this communication cycle:
If an agreement is suggested
Then

If agreement is satisfactory to me then signal my willingness to agree
            Else
                                    If a conditional offer or request is made
                                    Then
                                                consider possible combinations of offers
                                                if possible agreement exists then suggest an agreement
                                    Else
                                                Search for a preferable state to current one within limit
                                                If offer or request has not already been made this round
                                                Then make appropriate conditional offer or request
            Until either an agreement is finalised or last round is same as this one

Where an "agreement is finalised" means that all parties necessary to an agreement have signalled their agreement to it – the agreement then comes into force and the parties try to do the actions they have agreed to.

A typical sequence that results from this algorithm working in parallel in the agents is cycles of the following phases:

1. A sequence of conditional offers and requests

2. The detection of a possible agreement and it suggestion

3. Agents signal their agreement

4. Agents do the actions they have agreed to and these (possibly) effect the current state in each agent

Of course in some set-ups nothing happens at all (either agents are in a desirable state or no improvement seems possible; some involve only phase 1 (there is no possible agreement); some only phases 1 & 2 (there is a possible agreement but none is found which all necessary parties agree on; and some only phases 1, 2 & 3 (agreement is reached but the actions they have agreed upon are not possible).

Initialization

Unless the algorithm that the agent uses to judge which state is preferable includes a random choice, then the simulation is entirely deterministic. The simulation is initialised by reading in a text file which specifies what agents are involved and what their beliefs and methods of judgement are. Examples of these scripts are shown in the appendices below. Note that the "IndicatorWeights:" line is a hangover from (Hales 2003) and is no longer used but kept for (a sort of) backward compatibility – this has been replaced by the "StateValuationClause:" line which specifies how the agent judges accessible states based on their properties.

Intended interpretation

The simulation is intended to represent the process of two actors seeking agreement as a result of performing a limited exploration of their own "mental model" of the cause and effects they believe to be the case in the domain they are negotiating about. The resulting dialogue is intended to be meaningfully interpreted in this way but not represent the full richness of natural language dialogues. In particular it does not represent any communication between agents about the nature of the domain nor any suggestions as to the reformulation of goals or new suggestions.

Details necessary for implementation but not thought to be critical to the results

Exactly how the preference judgements are implemented is not important as long as the states are judged as relatively preferable in the same cases.

Description of variations

There were no variations in the simulations described except for those caused by having different scripts.

Software Environment

David Hales; original version was implemented in Java – an Object-oriented language developed by Sun (see http://java.sum.com). The version described was implemented in SDML – a declarative forward chaining programming language which has been written specifically for agent-based modelling in the fields of business, management, organization theory and economics (see http://sdml.cfpm.org and Moss et al, 1998).

Code

The code for the SDML version described in this is accessible at http://cfpm.org/~bruce/wawdho. It requires SDML version 4.1 or later.

Raw Results

Some of the results are shown in the appendices below, the rest are viewable at URL: http://cfpm.org/~bruce/wawdho

Appendix – Scripts and results for "Rick and Iran" example

Viewpoint file Template

```
Agent: Iran : Iran
IndicatorWeights:  happiness 1
StateValuationClause: indicatorValue happiness
InitialNodes: IranIsUnhappy

Node: IranIsHappy : Iran is happy
Indicators: happiness 1
Action: TurnDialDown : make self sad and fatalistic
Link: TurnDialDown => IranIsUnhappy : Iran makes himself sad

Node: IranIsUnhappy : Iran is depressed
Indicators: happiness 0
# Comment next out if it is not possible for Iran to turn dial up when depressed (A)
Action: TurnDialUp : Dial turned up to make Iran Happy
# Comment next out if Iran thinks turning the dial up will not help when depressed (B)
Link: TurnDialUp => IranIsHappy

#-----------------------------------------------------------------------
Agent: Rick : Rick
IndicatorWeights:  happiness 1
StateValuationClause: indicatorValue happiness
# Next line is "InitialNodes: IranIsUnhappy" if Rick thinks that Iran is depressed (S)
InitialNodes: IranIsHappy

Node: IranIsHappy : Iran is happy
Indicators: happiness 1
Action: TurnDialDown : make self sad and fatalistic
Link: TurnDialDown => IranIsUnhappy : Iran makes himself sad

Node: IranIsUnhappy : Iran is depressed
Indicators: happiness 0
Action: TurnDialUp : Dial turned up to make Iran Happy
Link: TurnDialUp => IranIsHappy
```

Results – Run HA

```
Iran: Can someone please TurnDialUp so we can achieve IranIsHappy?
```

Results – Run HB

```
[nothing occurs]
```

Results – Run SA

```
Rick: I will TurnDialUp to achieve IranIsHappy.
Iran: Can someone please TurnDialUp so we can achieve IranIsHappy?
Rick: I will TurnDialUp
Rick has done TurnDialUp.
(State of Rick) is: Iran is happy.
(State of Iran) is: Iran is happy.
```

Results – Run SB

```
Rick: I will TurnDialUp to achieve IranIsHappy.
Rick: I will TurnDialUp
Rick has done TurnDialUp.
(State of Rick) is: Iran is happy.
(State of Iran) is: Iran is happy.
```

Appendix – Scripts and results for car buying example

Viewpoint file template

```
Agent: Seller : The Car Salesman
IndicatorWeights:  car 5000 money 1
StateValuationClause: sum (multiply 5000 (indicatorValue car)) (multiply 1 (indi catorValue
money))
InitialNodes: Start

Node: Start : the start
Indicators: car 1 money 0
Link: Pay10000 => GetLittle : given 10000 by buyer
# Comment out if seller thinks buyer would not pay 20000
# Link: Pay20000 => GetLots : given 20000 by buyer

Node: GetLittle : Seller has 10000 and car
Indicators: car 1 money 10000
# Comment out if seller would not give car for 10000
# Action: GiveCarCheaply : Seller gives car to buyer for only 10000
Link: GiveCarCheaply => CarSoldCheaply

Node: GetLots : Seller has  20000 and car
Indicators: car 1 money 20000
Action: GiveCarExpensively : Seller gives car to buyer
Link: GiveCarExpensively => CarSoldExpensively

Node: CarSoldCheaply : Seller has 10000
Indicators: car 0 money 10000

Node: CarSoldExpensively : Seller has  20000
Indicators: car 0 money 20000

#-------------------------------------------
Agent: Buyer : The Car Purchaser
IndicatorWeights:  car 25000 money 1
StateValuationClause: sum (multiply 25000 (indicatorValue car)) (multiply 1
(indicatorValue money))
InitialNodes: Start

Node: Start : the start
Indicators: car 0 money 20000
Action: Pay10000 : pay 10000
# Comment out if buyer would not pay 20000
# Action: Pay20000 : pay 20000
Link: Pay10000 => GaveLittle : gave 10000
Link: Pay20000 => GaveLots : gave 20 000

Node: GaveLittle : Seller has 10000 and car
Indicators: car 0 money 10000
# Comment out if seller would not give car for 10000
# Link: GiveCarCheaply => CarSoldCheaply : seller gives car for 10000

Node: GaveLots : Seller has 20000 and car
```

```
Indicators: car 0 money 0
Link: GiveCarExpensively => CarSoldExpensively :seller gives car for 20000

Node: CarSoldCheaply : Seller has car and 10000
Indicators: car 1 money 10000

Node: CarSoldExpensively : Seller has car and 0
Indicators: car 1 money 0
```

Results

Due to the length of these I only include a few of the results to give their flavour.

CCCC

Seller does not thinks buyer would pay 20000; seller would not give car for 10000; buyer would not pay 20000; and buyer thinks seller would not sell for 10000.

```
=========================================================================
Buyer: Can someone please Pay20000 and GiveCarExpensively so we can achieve
CarSoldExpensively?
Seller: Can someone please Pay10000 so we can achieve GetLittle?
Buyer: I will Pay10000 if others Pay20000 and GiveCarExpensively.
Seller: Can someone please Pay10000 and GiveCarCheaply so we can achieve CarSoldCheaply?
Buyer: I will Pay10000 if others Pay20000 and GiveCarExpensively.
================================================================ ============
=========================================================================
=========================================================================
(State of Buyer) is: Start.
(State of Seller) is: Start.
```

CUCU

Seller does not thinks buyer would pay 20000; seller would give car for 10000; buyer would not pay 20000; and buyer thinks seller would sell for 10000.

```
=========================================================================
Seller: Can someone please Pay10000 so we can achieve Get Little?
Buyer: I will Pay10000 if others GiveCarCheaply.
Seller: I will GiveCarCheaply if others Pay10000.
Buyer: Can someone please Pay20000 and GiveCarExpensively so we can achieve
CarSoldExpensively?
Buyer: I will Pay10000 if others Pay20000 and GiveCar Expensively.
Seller: I agree to GiveCarCheaply if others Pay10000
Buyer: I agree to Pay10000 if others GiveCarCheaply
Buyer has done Pay10000.
=========================================================================
Seller has done GiveCarCheaply.
=========================================================================
=========================================================================
(State of Seller) is: CarSoldCheaply.
(State of Buyer) is: CarSoldCheaply.
```

CUUC

Seller does not think buyer would pay 20000; seller would give car for 10000; buyer would pay 20000; and buyer does not think seller would sell for 10000.

```
=========================================================================
Seller: Can someone please Pay10000 so we can achieve GetLittle?
Buyer: I will Pay20000 if others GiveCarExpensively.
Seller: I will GiveCarCheaply if others Pay10000.
=========================================================================
================================================================ ========
=========================================================================
(State of Seller) is: Start.
(State of Buyer) is: Start.
```

UCUC

Seller does think buyer would pay 20000; seller would not give car for 10000; buyer would pay 20000; and buyer does not think seller would sell for 10000.

```
=========================================================================
Buyer: I will Pay20000 if others GiveCarExpensively.
Seller: Can someone please Pay20000 so we can achieve GetLots?
```

```
Seller: I will GiveCarExpensively if others Pay20000.
Seller: Can someone please Pay10000 so we can achieve GetLittle?
Buyer: I agree to Pay20000 if others GiveCarExpensively
Seller: I agree to GiveCarExpensively if others Pay20000
Seller: Can someone please Pay10000 and GiveCarCheaply so we can achieve CarSoldCheaply?
Buyer has done Pay20000.
========================================================================
Seller has done GiveCarExpensively.
============================================================= ====
========================================================================
(State of Buyer) is: CarSoldExpensively.
(State of Seller) is: CarSoldExpensively.
```

UUUU

Seller thinks buyer would pay 20000; seller would give car for 10000; buyer would pay 20000; and buyer thinks seller would sell for 10000.

```
========================================================================
Buyer: I will Pay10000 if others GiveCarCheaply.
Seller: Can someone please Pay20000 so we can achieve GetLots?
Buyer: I will Pay20000 if others GiveCarExpensively.
Seller: I will GiveCarExpensively if others Pay20000.
Seller: Can someone please Pay10000 so we can achieve GetLittle?
Buyer: I agree to Pay20000 if others GiveCarExpensively
Seller: I agree to GiveCarExpensively if others Pay20000
Seller: I will GiveCarCheaply if others Pay10000.
Buyer has done Pay20000.
========================================================================
Seller has done GiveCarExpensively.
=========================================================== ===================
========================================================================
(State of Buyer) is: CarSoldExpensively.
(State of Seller) is: CarSoldExpensively.
```

Appendix – Scripts and results for flooding example

Simple view

```
Agent:  Citizen   : The citizens                    # agent name and description
IndicatorWeights:     floodDamage -4                 # weights agent applies to indicators
                      tax -0.5
                      environment 0.5
StateValuationClause: minOf (accessibleStateInNSteps (sumOfAllWeightedIndicatorValues) 1)
InitialNodes: Start

Node:         Start              : no floods, normal flood defences and taxes
Indicators:   floodDamage 0 tax 1 environment 0
Action: accept-higher-taxes              : ambitious internal dykes
Link: accept-higher-taxes and build-defenses => Expensive-High-Flood-defences : build
expensive defences
Link: build-defenses => Cheap-High-Flood-defences : build cheap but effective defences
Link: high-rain => SeriousFloods : high rain causes serious floods

Node:         Cheap-High-Flood-defences                  : high flood defences and low
taxes
Indicators:   floodDamage 0 tax 4 environment -5

Node:         Expensive-High-Flood-defences              : high flood defences and low
taxes
Indicators:   floodDamage 0 tax 10 environment -4

Node:         SeriousFloods            : attractive flood plains up river
Indicators:   floodDamage 10 tax 5 environment -7

#===================================== =======================
Agent:  State           : The government of the citizens
IndicatorWeights:     floodDamage -3                    # weights agent applies to
indicators
                      environment 1
```

```
                        popularity 2
StateValuationClause: minOf (accessibleStateInNSteps (sumOfAllWeightedIndicatorValues) 1)
InitialNodes: Start

Node:          Start                : no floods, normal flood defences and taxes
Indicators:    floodDamage 0 environment 0 popularity 1
Action: build-flood-defences              : ambitious internal dykes
Link: accept-higher-taxes and build-defenses => Expensive-High-Flood-defences : build
expensive defences
Link: high-rain => SeriousFloods : high rain causes serious floods
Link: abnormal-rain => SeriousFloods : abnormal rain causes serious floods

Node:          Expensive-High-Flood-defences              : high flood defences and low
taxes
Indicators:    floodDamage 0 popularity 1.2 environment -0.1
Link: abnormal-rain => SeriousFloods  : abnormal rain means get serious flooding even
having built flood defences

Node:          SeriousFloods            : attractive flood plains up river
Indicators:    floodDamage 10 popularity -2 environment -0.1
```

Result

```
========================================================= ===================
Citizen: Can someone please build-defenses so we can achieve Cheap-High-Flood-defences?
Citizen: I will accept-higher-taxes if others build-defenses.
=======================================================================
=======================================================================
=======================================================================
(State of Citizen) is: Start.
(State of State) is: Start.
```

Extended view

Viewpoint file

```
Agent:  Citizen   : The citizens                # agent name and description
IndicatorWeights:      floodDamage -4           # weights agent applies to indicators
                       tax -0.5
                       environment 0.5
StateValuationClause: minOf (accessibleStateInN Steps (sumOfAllWeightedIndicatorValues) 1)
InitialNodes: Start

Node:          Start                : no floods, normal flood defences and taxes
Indicators:    floodDamage 0 tax 1 environment 0
Action: accept-higher-taxes             : ambitious internal dykes
Link: accept-higher-taxes and build-defenses => Expensive-High-Flood-defences : build
expensive defences
Link: accept-higher-taxes and create-flood-plains => Flood-Plains : create attractive
flood plains
Link: build-defenses => Cheap-High-Flood-defences : build cheap but effective defences
Link: high-rain => SeriousFloods : high rain causes serious floods

Node:          Cheap-High-Flood-defences                 : high flood defences and low
taxes
Indicators:    floodDamage 0 tax 4 environment -5

Node:          Expensive-High-Flood-defences             : high flood defences and low
taxes
Indicators:    floodDamage 0 tax 10 environment -4

Node:          SeriousFloods            : serious disruptive flooding
Indicators:    floodDamage 10 tax 5 environment -7

Node:          Flood-Plains    : attractive flood plains up river
Indicators:    floodDamage 0 tax 8 environment 2
Link: high-rain => ModerateFloods : high rain causes moderate floods

Node:          ModerateFloods : moderate flooding
Indicators:    floodDamage 7 tax 4 environment -4

#===============================================================
Agent:  State               : The government of the citizens
```

```
IndicatorWeights:        floodDamage -3                    # weights agent applies to
indicators
                         environment 1
                   popularity 2
StateValuationClause: minOf (accessibleStateInNSteps (sumOfAllWeightedIndicatorValues) 1)
InitialNodes: Start

Node:          Start                  : no floods, normal flood defences and taxes
Indicators:    floodDamage 0 environment 0 popularity 1
Action: build-flood-defences             : ambitious internal dykes
Action: create-flood-plains             : create flood plains
Link: accept-higher-taxes and build-defenses => Expensive-High-Flood-defences : build
expensive defences
Link: accept-higher-taxes and create-flood-plains => Flood-Plains : create attractive
flood plains
Link: high-rain => SeriousFloods : high rain causes serious floods
Link: abnormal-rain => SeriousFloods : abnormal rain causes serious floods

Node:          Expensive-High-Flood-defences                  : high flood defences and low
taxes
Indicators:    floodDamage 0 popularity 1.2 environment -0.1
Link: abnormal-rain => SeriousFloods   : abnormal rain means get serious flooding even
having built flood defences

Node:          SeriousFloods             : attractive flood plains up river
Indicators:    floodDamage 10 popularity -2 environment -0.1

Node:          Flood-Plains    : attractive flood plains up river
Indicators:    floodDamage 0 popularity 0 environment 2
Link: high-rain => ModerateFloods : high rain causes moderate floods
Link:  abnormal-rain => ModerateFloods : high rain causes moderate floods

Node:           ModerateFloods: moderate flooding
Indicators:    floodDamage 7 popularity -1 environment -0.1
```

Result

```
===========================================================================
State: I will create-flood-plains if others accept-higher-taxes.
Citizen: Can someone please build-defenses so we can achieve Cheap-High-Flood-defences?
State: I will create-flood-plains if others accept-higher-taxes and abnormal-rain.
Citizen: I will accept-higher-taxes if others build-defenses.
State: I will create-flood-plains if others accept-higher-taxes and high-rain.
Citizen: I will accept-higher-taxes if others create-flood-plains.
Citizen: I will accept-higher-taxes if others create-flood-plains and high-rain.
State: I agree to create-flood-plains if others accept-higher-taxes
Citizen: I agree to accept-higher-taxes if others create-flood-plains
State has done create-flood-plains.
Citizen has done accept-higher-taxes.
===========================================================================
=============================================================== =============
===========================================================================
(State of State) is: Flood-Plains.
(State of Citizen) is: Flood-Plains.
```