

Learning as Collective Belief-Revision: Simulating Reasoning about Disparate Phenomena

T. R. Addis,* D. C. Gooding**

*School of Computer Science and Mathematics,
Mercantile House, University of Portsmouth, Portsmouth PO1 2EG, U.K.

**Science Studies Centre, Department of Psychology,
University of Bath, Bath BA2 7AY, U.K.

*Tom.Addis@port.ac.uk; **hssdcg@bath.ac.uk

Abstract

Computer simulations of learning, discovery and problem-solving generally neglect the social aspect of these processes. EXPLORE v. 4 is a computer simulation that represents learning as a social process of belief revision by a number of interacting agents. Each agent holds a set of beliefs about an aspect of the world that can be investigated empirically. The formation and revision of their beliefs is modeled as a process mediated both by making experiments and by communication between individuals. Agents are defined by attributes including: confidence in one or more of a range of hypotheses, variable sensitivity to the results of their own experiments and to the opinions of other observers, the ability to make decisions about whether to make experiments or to consult other agents and to make further decisions about which experiment(s) to perform or which agent(s) to consult. The paper describes assumptions underlying each of the basic features of the model, its application to the 4-card selection task of Johnson-Laird & Wason, an extension to this task to relate it more firmly to a real discovery process and some initial calibration results generated by the simulation of these experiments. The extension involves the restriction of communication and access to experiments. The results are compared with the related psychological experiments of Gruber. Further simulations including the comparison of the simulation's behaviour for specific historical events are in progress.

1 Introduction

Ten years ago we decided to create a computer model that would combine historical, psychological and computational approaches to scientific discovery well enough to be used to evaluate inferences about the discovery process. To be realistic the model should be developed from historical observations and psychological experiments. This paper describes a simulation of psychology experiments which we treat as simple models of the more complex, historical cases of scientific discovery (Gooding 1990, 1996). This paper is the second of three which describe one of our more comprehensive simulations, EXPLORE v. 4.

Scientists make discoveries in the context of ideas, methods, resources, empirical information and, most important of all – other people. We aim to move beyond the view that discovery is a means of disclosing information about the natural world, in which the interactions of observers plays no role. The rationale for the model described in this paper has been argued in Addis and Gooding (1999) and Gooding and Addis (1999). There we distinguish between models, theories and the simulation. The simulation implements (but is not equivalent to) a theory of scientific investigators working in the

context of a community of investigators. The simulation consists of a number of models. The most important of the models define *agents* as interacting decision-makers capable of learning from the outcomes of their actions; *beliefs* as a disposition to act in certain ways; *decision-making* in a dynamic, changing environment, and *experiments* as sets of results each of which is modalised to a set of contextual factors as well as to a set of hypotheses representing an agents view of the world. The simulation represents features of discovery that historical studies and experimental psychology show to be important, and which therefore belong in an adequate theory of scientific discovery. When fully developed it can be used to evaluate the theory, for example by testing hypotheses adduced by social scientists to explain recorded scientific events.

2 Basic Assumptions

Our model defines a scientist as an agent which holds a range of *beliefs* about the world. The state of the agent depends upon how each of the beliefs is modified by the experience of that individual. These *beliefs* govern the actions of the scientist. A *belief* will be represented as the combination of an hypothesis about the world and a number between zero and one that reflects the subjective

probability that the hypothesis is ‘true’. Where a set of hypotheses are mutually exclusive the sum of the subjective hypotheses will be taken as 1.0.

The model assumes that the agent will benefit in some way if a certain result is achieved. This result is achieved by performing a particular *action* in the context of a correct view of the way the world is. The model of investigation features three types of action:

1. considering a range of hypotheses,
2. making experiments that are expected to bear directly on one or more of the hypotheses, by providing confirmatory or disconfirmatory evidence,
3. consulting an agent whose knowledge will bear directly on one or more of the hypotheses.

Agents receive information either directly through interaction with the world (e.g. experimentation, fieldwork) or indirectly via another agent (e.g. listening to a verbal report or reading an experimental paper). Our model therefore assumes all historically documented hypotheses are available to the agent at all times. However, access to colleagues' views of these hypotheses can be varied, as in real science.

We can thus emulate the generation of new hypotheses through the manipulation of belief. We use *confidence*, a numerical measure of belief, as the primary means of differentiating between hypotheses that are actively considered (or “in play”) and those that are not being considered, either because they have not yet been inferred or because they are no longer regarded as plausible. Similarly, in order to avoid modelling of the generation of new experiments we assume that all historically documented experiments can be known by agents (subject to access restrictions) and that the difference between those in view and those not in view is a question of assessed usefulness (section 4.2).

3 Simulation Mechanisms

In Addis and Gooding (1999) we introduced the notion of a *confidence profile*, a set of values representing an agent’s confidence in each of a range of hypotheses. Agents occupy a changing world of information-bearing experiments and consultations, so the changing probability of an hypothesis being true must be calculated dynamically. Each of these values is modified in the light of the result of *experiments* and *consultations*. We use a modified version of Bayes’ Rule to convert an observed experimental result into a set of modifiers for a confidence profile. Note that we do not use Bayes’ Rule to calculate the accumulation of evidence for an hypothesis (see section 3.1). We first consider the equation from the point of view of an agent who performs experiments in order to determine what hypothesis (model) is most likely.

We determine the expected hypothesis given a result from an experiment e thus:

$$E_{n-1}(H/R_e) = \frac{E_{n-1}(R_e/H)}{E_{n-1}(R_e)}$$

This generic equation¹ represents a personalised view of the world for a specific agent. Here R_e is the result of an experiment e and H is an hypothesis. The expected result $E_{n-1}(R_e/H)$ for any experiment will depend upon the perceived probability of each hypothesis and an *a priori* understanding of the probability of a result for the hypothesis supposing it were true. The confidence profile is retrospective in that it reflects only past experience. In order to represent an agent’s overall view of the world we also need to calculate a unified value that characterises its expectations about the outcomes of future experiments. In particular we can determine an agent’s view of the expected result for *all* hypotheses as $E_{n-1}(R_e)$.

3.1 Belief Adjustment

Our model of belief-revision is *not* Bayesian. Despite the advantages of a Bayesian approach to the understanding of theory-selection in scientific change, as argued, for example, by Salmon (1990), it is unrealistic to use Bayes Rule for accumulating evidence for an hypothesis. This is because Bayes Rule assumes a constant and unchanging world and that the order of the events is irrelevant.² We retain the assumption of the independence of results, however, a static world cannot be assumed. We therefore take the stance that relative to the number of events that can occur during a change, changes in the world will be gradual. This allows an hypothesis to remain available for consideration, despite a run of apparently falsifying observations. This important virtue is in keeping with scientific practice (see, e.g., Lakatos, 1970).

Hypotheses do not achieve absolute certainty -- or if they do, scientists re-label them as facts, laws or principles. To call something an *hypothesis*, H is to say that there is some support for H given evidence R_e . This support is $E_{n-1}(H/R_e)$ and is a value which lies between 0 and 1 for each hypothesis H . As a new result R_e is gleaned from carrying out an experiment e then the probability which represents the confidence an agent has concerning a particular hypothesis, can be modified to $E_n(H)$ by adapting the above equation to the following one (see Addis 1985, p. 260):

¹ For simplicity these equations are generic, not agent-specific.

² Order is important because changes in belief involve changes in the interpretation of results.

$$E_n(H) = \frac{(N-1).E_{n-1}(H) + E_{n-1}(H/R_e)}{N}$$

We can thus define a concept *flexibility* which ranges from 0 to 1 for each agent such that:

$$\text{flexibility} = 1/N$$

Flexibility reflects an agent's *responsiveness* to evidence.

We assume that when an agent consults another this is equivalent to having access to the consulted agent's complete range of confidence values. This access has the effect of modifying the consulting agent's confidence in each hypothesis in the same way as if it had performed its own experiment. The confidence value of each of the hypotheses that make up the agents belief profile will be modified according to the following equation:

$$E_n(H) = \frac{(M-1).E_{n-1}(H) + E_{n-1}(H_{\text{Consultee}})}{M}$$

M ranges from 1 to infinity. The larger M is, the smaller the effect any evidence has on the change in confidence. We can thus define a concept *receptivity*. This ranges from 0 to 1 for each agent as:

$$\text{receptivity} = 1/M$$

This reflects the influence of any consultee upon the consulting agent.³

3.2 Belief and Indifference

The number of hypotheses actively considered by scientists varies. It is unusual, especially during the exploratory stages of an investigation, to consider just two well-defined alternative hypotheses. Nevertheless, the tendency is always to reduce the number of hypotheses or possible interpretations in play. This means that confidence in a particular hypothesis is affected not only by experiments and consultations but also by the number of hypotheses available for consideration. In order to allow for changing numbers of hypotheses, we introduce a dynamic threshold, the *indifference value*. This defines those hypotheses that are actively being considered (or believed to some extent).

To calculate the indifference value we need a function that changes smoothly between limiting values and is easily calculable from any number of different hypotheses. A quantity that varies in time in this way is the inverse of entropy. Entropy is an expected measure of the log of a certainty.⁴ We use this as a general measure of

³ Bayes' Rule is not applied in consultations because it is needed only to update confidences on the basis of experimental results.

⁴ In information theory entropy (given by $-\sum p \log_2 p$) is the *average degree of certainty* of being able to predict the next bit of information in a stream, e.g. the next character in a string of characters.

an agent's confidence in the world and its ability to act correctly at event n. (Here we will use 'a' and subscript 'a' to denote a particular agent). We call this general confidence measure for an agent the *Model Entropy*, where the term model denotes the set of hypotheses that make up the agent's view of the world. Model Entropy is given by:

$$\text{Entropy}_n(\text{Agent}_a) = -\sum_H E_n(H) * \text{Log}_2(E_n(H))$$

and from this we can obtain an inverse of the entropy which gives an expected value for $E_n(H)$. This will be denoted by $I_n(H_a)$. $I_n(H_a)$ will be called an *Indifference Threshold* for the agent a at event n:

$$\begin{aligned} \text{Indifference Threshold}(a, n) &= \text{Log}_2^{-1}(\text{Entropy}_n(\text{Agent}_a)) \\ &= I_n(H_a) \end{aligned}$$

The expression $E_n(H)$ is the expected probability of an hypothesis. Values of $E_n(H)$ above $I_n(H_a)$ are considered to be significant, that is to say, the hypothesis is actively believed by the agent.⁵

A similar value can be calculated for the group of agents, derived from the Group Entropy:

$$\text{Entropy}(\text{Grp}) = \frac{\sum_{\text{Agent}} E_n(H/\text{Agent}) * \{\text{Log}_2(E_n(H/\text{Agent})) - \text{Log}_2(|A|)\}}{|A|}$$

This expression represents the expected result of a single sampling of the confidence in any hypothesis of an agent randomly chosen from the group. The inverse of the group entropy represents the significance threshold for the group of agents in terms of an expected probability. We can also calculate the inverses of entropy for an agent $I(H_a)$, a group of agents $I(H_g)$ and a set of experiments $I(e_a)$ as perceived by a single agent or a group of agents $I(eg)$.⁶

Because the threshold values are independent of the number of hypotheses or experiments we can define the indifference level of a group of actors independently of particular hypotheses that happen to be in play. This is important because it allows us to represent the process whereby an hypothesis changes from being a mere possibility, to being considered plausible, being generally accepted and, finally, coming to have the status of a law or principle. When all the hypotheses have been eliminated except one then both $I(H)$ and $E(H)$ will equal to unity (i.e., certainty). When this happens it can be said

⁵ Since most of the calculations are done by the natural logarithms it is useful to note that:

$$2^x = e^{(x \cdot \log 2)}$$

⁶ An experiment e is defined in terms of the probabilities of results R that are considered possible for a given range of hypotheses H.

of the agents that they are both *indifferent* to the hypothesis and also *certain* of it. The hypothesis is then unquestionably correct.

4 The Choice of an Action

4.1 Assessing Actions

The experimental set-up is represented as a table of real numbers that indicate the *a priori* probability of a result occurring, given that a hypothesis (or model) is ‘True’.⁷ This list of occurrence probabilities *defines each possible experiment*. The list will sum to unity for each hypothesis since at least one of the results will occur if that model were the case.⁸ The choice of experiment is affected by an agent’s bias. In our model the choice of an experiment is derived from its effectiveness in discriminating between hypotheses, as perceived by an agent. This choice is governed by the agent’s confidence $E_{n-1}(H)$. We represent an agent’s view of an experiment by the entropy, represented for each experiment as:

$$\text{Entropy (Expmt } e) = -\sum_R \sum_H E_{n-1}(H/R_e) * \text{Log}_2 E_{n-1}(H/R_e)$$

This equation describes the confidence of an hypotheses given a result as perceived by an agent. The experiment with the *minimum* entropy is the *most likely* to be chosen by the agent in our simulation.⁹ The chosen experiment will give the clearest, i.e. most decisive results for supporting or negating each of the hypotheses in the agent’s belief system. Nevertheless, this will not always be the experiment chosen, as we will show in section 4.2.

The experimental set-up is defined by a table of real numbers (see above). To obtain a vector of certainties we multiply an agent’s *belief profile* by the matrix representing the result probabilities of all experiments that are possible in a given epoch. This generates a *biased perception* for each agent of each of the experiments. Just as in real life, each agent perceives the potential outcome of an experiment differently. It follows from our definition of an experiment that each agent perceives a subtly different experiment being performed. The result of the multiplication is then used to modify the agent’s *a priori* confidence value for each hypothesis. The degree of modification to the belief profile depends upon the agent’s flexibility (see section 3.1). A similar situation occurs when there is a decision to consult another agent rather than do an experiment. The cycle of learning

⁷ We refer to the hypothesis that is actively true within a simulation run as the *objective model*. Simulated experiments will have those outcomes that are most probable in a world in which the objective model is true. For an example, see section 5.3.

⁸ Note that if the set of occurrence probabilities for each experiment is the same for all hypotheses then it is the same experiment, whatever the physical arrangement is in practice.

⁹ ‘Most likely’ because the actions of agents are governed, Monte Carlo fashion, by a probability distribution based upon belief.

through consultation is similar to that for an experiment, except that there is no need to involve Bayes Rule.¹⁰

4.2 Choosing Actions

The pragmatic notion of belief stipulates that the numerical confidence value attached to each hypothesis or model represents the probability of an *action* given a certain state of affairs. We use an agent’s belief to calculate a personalised ‘entropy’ for each experiment (section 4.1). Under certain conditions Game Theory suggests that a ‘mixed strategy’ approach should be taken; where a mix of actions are tried according to some probability distribution.¹¹ The decision mechanism should therefore deploy belief as defined, i.e. as a *probability to act*. We will first consider the easier choice of experimenting or consulting.

Entropy represents the expected probability of any event (in this case an hypothesis). If we assume the expected loss for acting on a wrong belief is one unit and expected gain is zero then we can create a two person zero sum matrix as shown in Table 1.

Table 1. Payoff matrix for a two person zero sum game.

Action (frequency)	Expected Hypothesis	Expected NOT Hypothesis
(fg) Ask Group Member	Pg . (0)	(1- Pg) . (-1)
(fa) Do Experiment	Pa . (0)	(1- Pa) . (-1)

- Pg is the expected group ‘belief’ that some hypothesis represents a particular state of affairs. It does not matter which hypothesis since this represents some kind of average.
- Expected Gain in accepting the group’s view is $Pg .(0) - (1- Pg) (-1) = - (1-Pg)$

Consulting the group is equivalent to doing an experiment since both affect the agent’s belief profile. What matters are the incremental increases in the probability of an agent’s beliefs being true. Iterating this strategy in each cycle of the simulation makes this application of game theory to science and to learning more appropriate than a single-step decision procedure would be. If fg is the frequency of applying the strategy ‘ask-group-member’ and fa is the frequency of ‘do-experiment’, then we can calculate for each actor the maximum security value (or minimum possible loss) between the agent’s two options, as:

¹⁰ This is because confidence values in the belief profiles of consultants and consultees are already expressed in the same terms, whereas for experiments they are expressed as the probabilities of results, given each hypothesis.

¹¹ Luce and Ruffia, (1957), pp. 67-70.

$$fg = (1 - Pa) / [(1 - Pg) + (1 - Pa)]$$

$$fa = 1 - fg$$

Once this decision has been made there is the further decision as to which experiment to perform or which actor to consult. For this we use the modified *a priori* entropy as expressed through the beliefs of the agent.

The above equations are extended in order to deal with the many cases of what experiment to do or which agent to consult. Thus a consultee is chosen in a similar manner, by extending the calculations to consider the entropy of the average confidence in each hypothesis of every pair of agents, including the agent making the decision. Self-consultation then becomes an option, though in our simulation this rarely occurs. As with experiments, the payoff should invoke a maximum security level from the entropy pairings. So the best bet for an agent is to decide which other agent to be influenced by in order to gain greater overall confidence.

5 Paradigmatic Situations

5.1 The modified four card test

To determine what might be meant by a ‘useful’ experiment consider the four card test (Wason 1960; Johnson-Laird and Wason (1977). This explores the ability of subjects to select an experiment in order to prove or disprove a rule about a specific situation. These well-known tests require subjects to decide how to determine whether a rule about a set of cards is true. Subjects are presented with four cards showing a set of vowels, consonants and even and odd numbers, (such as ‘A’ ‘D’ ‘4’ ‘7’) and are told that every card, of which these are a subset, has a letter on one side and a number on the other side. The rule is that "If a card has a vowel on one side, then it has an even number on the other side.". The task is to say which of the cards should be turned over in order to find out whether the rule is true or false. Wason's results implied that his subjects had difficulty in dealing with tests that involve disproof rather than proof. However, when the task is reformulated into more meaningful terms then a larger proportion of subjects apply a falsificationist strategy (Wason & Shapiro, 1971). This suggests an alternative to Wason's initial conclusion that humans don't reason in a strict, logical way about the empirical consequences of a rule. Instead, the original design shows up the difficulty most people have with problems that are expressed abstractly, without reference to a context of *meaningful actions*.

Our interest in this test is twofold: *first*, the refinement of Wason's 4-card test illustrates the importance of context to the ability to make decisions about actions, and *second*, the 4-card test requires both confirmatory and disconfirmatory experimental strategies. It is there-

fore an appropriate simplification of the scientific situation (see Gooding and Addis, 1999; Tweney 1985.). In Wason type tests psychologists design a world constrained by a particular rule in order to find out how human subjects reason about deciding the truth of the rule. However, natural scientists cannot define the way the world is. To discover constraints they must entertain a wider range of possible hypotheses in order to discover which one(s) may be true. We therefore extend the 4-card experiment to incorporate evaluation of a range of possible hypotheses. Consider the scenario in which subjects are shown a large set of cards, with each card showing only one face. They are told that each card has a numeral on one side and a character on the other and that there is a simple relationship, for all cards, between the number type (odd or even) and the character type (vowel or consonant) on each card. The task is to determine the simple relationship by turning over as few cards as is possible.

We can ask ‘how *confident* is the subject of a particular view before they dare answer?’ Moreover, each turn of a card is an *action* that must count towards increasing this confidence. Our model treats making an experiment and consulting another scientist as *actions* in this sense.

5.2 Modeling the experiment – hypothesis relationship

The specificity of an hypothesis is an important factor in experimental design. Proponents of inductive and deductive models of inference agree that the more specific an hypothesis is, the stronger the relationship between the hypothesis and the data that confirms or falsifies it. For the card situation we are given the known constraints (number <-> character) and the limitation of only four logical connectives (imply, or, and, exclusive-or). From this we can identify $4*4*2 = 32$ possible logical statements of the form exemplified by “even implies vowel”. After removing the logical equivalencies, this number reduces to 10 logically distinct hypotheses that describe all the constraints (see table 2).

Table 2. List of hypotheses for the card test.

Model Number	Constraint
M0	Vowel → Even
M1	Vowel → Odd
M2	Even → Vowel
M3	Odd → Vowel
M4	Vowel ⇌ Even
M5	Vowel ⇌ Odd
M6	Vowel ∧ Even
M7	Vowel ∧ Odd
M8	Even ∧ Consonant
M9	Odd ∧ Consonant

⇌ means XOR

This minimum set can be laid out in the form of ten tables that have the same structure as the example in table 3. This shows the a priori probabilities of a result for each act of turning a card, given one hypothesis, labelled M0.

Table 3. An example of a set of a priori probabilities of the outcome of an action, given the constraint M0.

Hypothesis: Vowel → Even The Experiment:	Probability of each Result (displayed on the other side of card)			
	Vowel	Even	Odd	Consonant
TimocraVowel		1.0		
TimocraEven	0.5			0.5
TimocraOdd				1.0
TimocraConsonant		0.5	0.5	

We can also generate a probability of results for each action given each of the ten hypotheses. Such a table suggests a scheme of *generality* for the hypotheses, which we shall develop as part of a model of interaction of hypotheses and experiments. We have observed from these tables that the logical relationships fall into two distinct series and each series forms three groups of decreasing generality. The series and groups are independent of the particular set of statements since membership depends upon the relative form of the truth tables. We can show these relationships as a network of arrows and set names as shown in figure 1.




Figure 1. The Generalisation Groups of two Hypotheses Series

A valuable property of this arrangement is that we have two series of three distinct groups, where group 1 ‘contains’ the results of group 2 and so on. Further, the intersection of group 1 in each series is equivalent to the conjunction of group 3 in the same series.

5.3 Disparate Observations: Gruber's shadow - box experiment

Experiments in cognitive psychology typically focus on the cognitive processes of individuals. Communication, or the sharing of information, is not usually considered relevant. We can use EXPLORE v.4 to vary the Wason experiment by introducing the condition that each agent has access only to a subset of the entire range of experiments. This design is based upon the Gruber shadow-box experiments in which each subject is shown a different two-dimensional shadow-projection of a hidden object (Gruber 1985, 1990). Subjects may be in-

structed to communicate about their perceptual experience or left to decide on this stratagem. Gruber's purpose was to determine whether subjects that exchanged information obtained better hypotheses about the hidden object than subjects that used only the information available to themselves.

The design of this experiment brings us much closer to the scientific case because it requires subjects to exchange information about 2-D images and to reason about these as projections of possible 3-D objects. This introduces ambiguity into the description of observations and of possible 3-D objects that would generate the disparate observations of different observers.




Figure 2. A schematic description of the Gruber shadow-box experiment

There is a natural bias in this experiment in that the subjects are shown only a two dimensional display. The result is that subjects will tend not to agree until an *insight* occurs that what is observed is in fact a three dimensional object with different projections. The significance of this bias will become apparent later (section 6.2).

Our first simulations combine a Gruber-type experimental setup with a Wason card task. We restrict agents' access to each other so as to create two groups, each consisting of a pair of observers. Each pair has access to a different view of the world (i.e. a different set of cards, or a different side of the shadow-box). In this way we introduce important features of real scientific research such as the fact that experiments often display different aspects of a mechanism or process, and the need for communication about such disparate observations. Yet we avoid (for the time being) the ambiguity of verbal descriptions about shadows and of the mental transformations of these into 3-D objects. The two shadow projections are observations that represent two coherent aspects of a single object. Logically these observations are subsets of the total shape and we can say that a three dimensional object is the integration (or sum) of all its projections. Such an integration is a logical equivalent of the existential quantifier (Exists). For two projections such a quantifier reduces to an ‘OR’.

We can, for example, make our *objective model*: "Vowel

∨ Odd" (analogous to placing an object in the box) and limit the experiments that can be done by one team to *vowel* and *even* only. The other team will be limited to the experiments *consonant* and *odd* only. It is therefore not possible for each team to distinguish between the objective hypothesis and their own unique hypothesis,




Figure 3. The ‘card’ equivalent of a cylinder (objective model M5). Agents 1 & 2 might believe in M0, M2 and M6 whereas Agents 3 & 4 might believe in M0, M2 or M9

thus:

- Team A (agents 1 & 2)
Vowel ∧ Even = M6 (cf. circle in figure 1)
- Team B (agents 3 & 4)
Consonant ∧ Odd = M9 (cf. square in figure 1)

6 Simulation Results

Simulation experiments described in the first paper (Addis & Gooding, 1999) showed that bias does not merely delay the discovery of new and better hypotheses; it can prevent their being discovery at all. The important question is, does the behaviour of the simulation suggest that the models that comprise it have captured essential properties of belief revision amongst communicating scientists? We can show that the rational model (section 4) works by comparing the performance of the simulation with a well-designed experimental situation such as the Wason-Shapiro 4-card experiment (the version of the task that relies upon the problem being meaningful to subjects). Where the simulation can reproduce the behaviour of human experimental subjects, we suggest that the rational model of behaviour incorporated into our simulation is coherent with the actual behaviour of human subjects in the Wason-Shapiro experiments..

6.1 Unbiased isolated teams with restricted view

The initial simulation establishes the simple case where there are two teams of two agents. These are isolated from each other. All ten hypotheses are considered equiprobable at the start. Figure 4 shows the initial rise of three believable hypotheses for agent 1: these are M2, M5 and M6 (where "M" denotes a model or hypothesis considered by an agent). Agent 1 no longer believes M2 after 20 cycles nor M6 after 50 cycles. The objective

model M5 is taken to be the most likely by team A. These results can be explained through the generalisation diagrams, figures 1 and 6.

The change in belief profiles for team B is shown in Figure 5. This has a similar pattern except it is M9 rather than M6 that is a major contender for the most believed hypothesis. From an experimental view there is no difference between M6 and M5 for team A and M9 and M5 for team B. Because of the indeterminacy between the two hypotheses for each team the final belief in M5 will tend to 0.6 rather than 1.0.

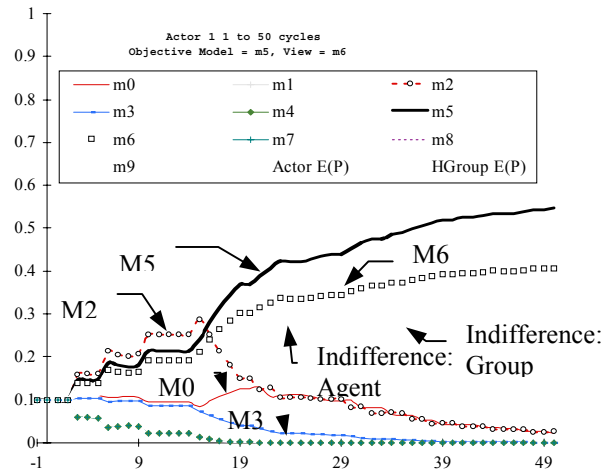


Figure 4 The change of beliefs as Agent 1 in team A acts. No initial bias. Objective model is M5

Suppose evidence is found only for M6. This evidence is consistent with the more general cases M5, M0, M2 and M3. However, because of the restrictions on what experiments can be done by each team of agents then for team A there is always the possibility of observing further positive evidence to support M0, M2 and M3. For example the turning of a vowel card could provide two possible results for M2 and M3. Turning an even card could show two possible results for M0 and M3. However, this is not the case with M5 which remains coherent for all experiments for both teams. This is summarised in table 4. A similar result can be constructed for team B as shown in table 5.

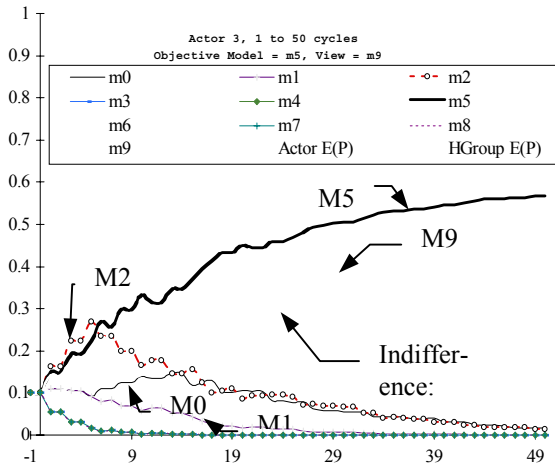


Figure 5 The change of beliefs as Agent 3 in team B acts. No initial bias




Figure 6. Explaining the results through a generalisation graph

Table 4. Team A and the value of results

Experiment M6	M0	M2	M3	M5	M6
Vowel	Even	Both	Both	Even	Even
Even	Both	Vowel	Both	Vowel	Vowel

Table 5. Team B and the value of results

Experiment M9	M0	M1	M2	M5	M9
Odd	Consonant	Both	Both	Consonant	Consonant
Consonant	Both	Both	Odd	Odd	Odd

Given a restricted view and a starting position where all hypotheses are of equal standing, *the general rule seems to be that agents prefer the most general hypothesis, where the results are coherent with a priori expectations.* In both cases this hypothesis will be M5. The results of M0 and M2 are consistent but not coherent with observations because the 'alternative' result never occurs. The incoherence is even more marked in the special cases of M1 and M3. M2 tends to be preferred to M0 because the a priori expectations are different in their relationship to the M5 projections.

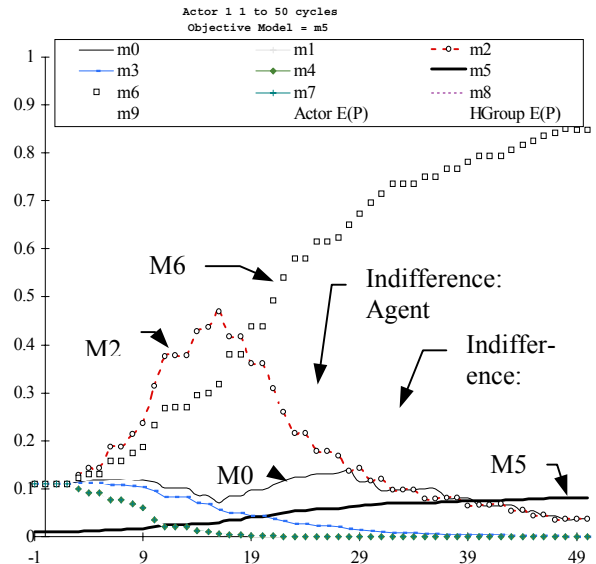


Figure 7. The change of beliefs as Agent 1 in team A acts. An initial bias against M5.

6.2 Biased, isolated teams with a restricted view

If we bias the starting point so that the objective hypothesis M5 is not favoured then we get a very different result: each group favours the model to which their experiments are limited (see figures 7 and 8). Confidence in M5 in each case levels off. Note the initial tendency is to believe the more general hypothesis. Although M5 is unlikely to win out its existence prevents the favoured hypothesis from achieving so high a level of confidence as to become a *fact*. These results are in accord with Gruber's finding, that where subjects fail to share information they achieve fewer solutions, and less good solutions than when they communicate about the implications of their disparate experiences.

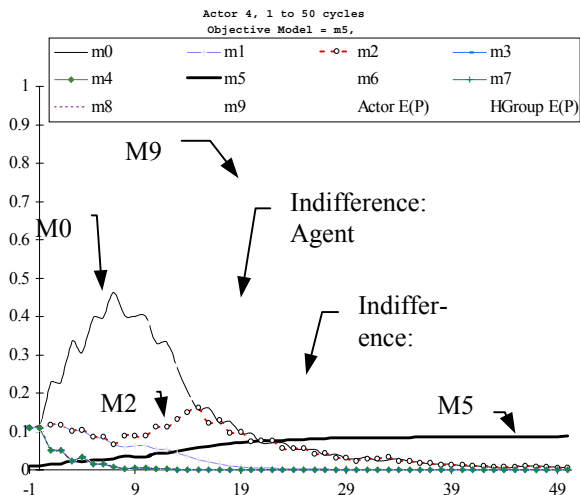


Figure 8. The change of beliefs as Agent 4 in team B acts. An initial bias against M5.

6.3 Biased teams with restricted view and communication

When we introduce communication we get a dramatic difference, but to show this completely we have to go to 100 event cycles. (Arrows in figure 9 indicate the relation of consultor to consultee as the direction of information flow, thus Agents 1 and 2 consult each other; agent 3 consults 1 and 2 and agent 4 consults 3 only). In the first instance there is very little difference between this limited communication and the initial results for isolated agents. We see in figures 10 that agent 3 starts off convinced that M2 is the right model. This accords with the initial view of agents 1 and 2.




Figure 9. The communication channels for four agents in the second experiment. Arrows indicate direction of information flow.

However, leaving agent 3 behind, agents 1 and 2 change their view and adopt hypothesis M6 just as they did before. This influences agent 3 through the insight that it is M5 that is a good possibility. This evokes more experimentation, but some doubt arises after consulting agents 1 and 2. This doubt seems to provoke further contact. The contact produces even more doubt. However, confidence in M6 rises for a short while instead of M5 as agent 3 consults agents 1 and 2. This provokes agent 3 to experiment and confidence in M5 takes over once again. This produces a rising cycle as shown in figure 10.

Agent 4, who only consults agent 3, does not go through these sequences of fluctuating confidence. The change over from M9 to M5 is smooth, unhurried and rises to a higher level of confidence than agent 3 (figure 10). Given the opportunity to 'talk' (either directly or indirectly) to agents in possession of alternative experiments both agents 3 and 4 eventually find the objective model M5. However the isolated agents 1 and 2 stick with their own views as before.

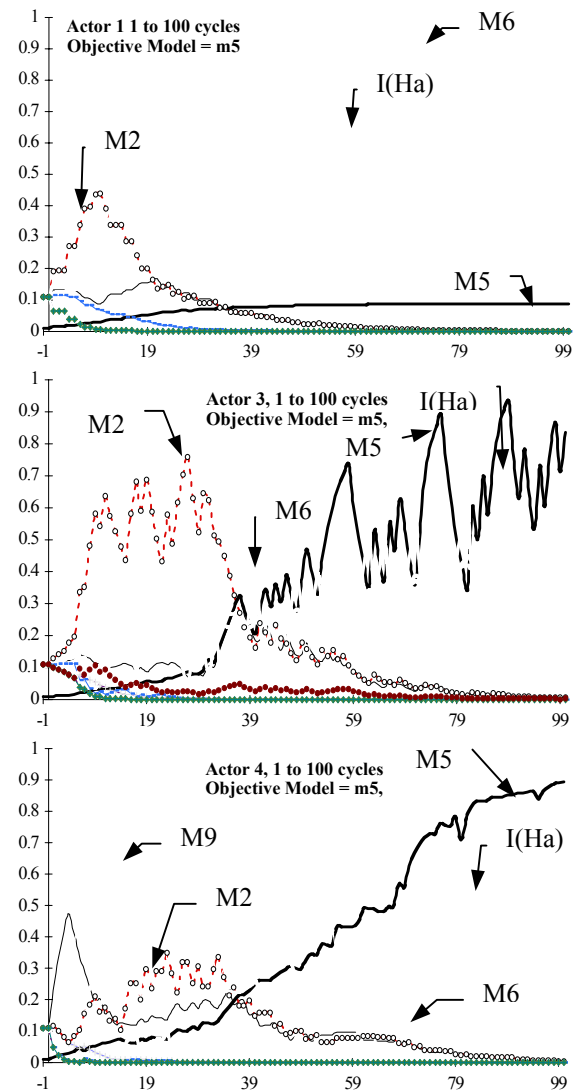


Figure 10. The effect of communication on belief

7 Conclusions

We have taken the pragmatic view of C. S. Peirce that the "essence of *belief* is the establishment of a habit, and different beliefs are distinguished by different modes of *action* to which they give rise" (Peirce 1878, quoted in Weiner, 1966, p.121.). The measure of belief reflects an agent's potential to action and is thus an objective measure. It is not intended to represent a numerical value of the subjective intensity of belief nor carry with it many of the concepts that make the word rich within its general usage. We limit ourselves to a mere action probability and as such the indifference threshold that delineates belief from disbelief changes as it reflects the number of hypotheses under consideration by an agent.

The simulation runs described here suggest that when there is no bias and no observable distinction between two hypotheses drawn from several, then the most general hypothesis will tend to be favoured as a basis for

action. However, where there is a bias then preference for general hypotheses persists, until new experiments provide evidence that makes a distinction or consultation provides access to opinions based on evidence available to other agents. The 'go between' (agent 3) has the most traumatic role to play in communication; this agent's loyalties to hypotheses fluctuate wildly depending upon which group is being consulted. The 'scholar' (agent 4) will move more confidently and steadily towards the 'truth' (the objective model) than others. However, as in real science and as Gruber's shadow-box experiments show, the vital ingredient for reaching the best hypothesis is that communication occurs.

There is still much to be done with this simulation, both to explore the paradigmatic situation described here and to apply it to more detailed, historical cases. Once the range of its behaviour is calibrated against experimental studies we will continue to the next stage, which is to interpolate between documented historical events and thus throw light on some of the hidden details of history. The work described here shows the potential of the EXPLORE v.4 simulation.

Acknowledgements

Much of the work on CLARITY has been supported through the projects *Knowledge Design from Natural Science Experiments*. MRC-SERC-ESRC (D. C. Gooding & T. R. Addis, Project No. SPG 9107137) and *Graphic Programming Methods for Modelling Cognitive & Social Process of Science*, ESRC (D. C. Gooding & T. R. Addis, Project No. R000235286) This work would not have been possible without the support and development provided by Jan J. Townsend-Addis who developed the CLARITY programming environment with which we designed the simulation. See ADDIS T. R. & TOWNSEND ADDIS J. J. (1998).

References

- ADDIS, T. R. & GOODING, D. C. (1999). Interacting Agents: Simulating a Community of Scientists, *International Journal of Human Computer Interaction* (submitted).
- ADDIS T. R. & TOWNSEND ADDIS J. J. (1998) 'A Functional Schematic Interpreter: an environment for Model Design', *International Journal of Systems Research and Information Science*, 7: 263-291.
- ADDIS, T. R. (1990). Knowledge for Design. *Knowledge Acquisition*, 2: 95-105.
- ADDIS, T. R. (1985). *Designing Knowledge-Based Systems*. London: Kogan Page / New York: Prentice-Hall.
- GOODING, D. C. (1990). *Experiment and the Making of Meaning*. Dordrecht and Boston: Kluwer Academic.
- GOODING, D. C. (1996a). Discovery as creative exploration, *Creativity Research Journal*, 9: 189-205.
- GOODING, D. C. & ADDIS, T. R. (1999), A Simulation of model-based reasoning about disparate phenomena, *forthcoming in* Magnani, L., Nersessian, N. and Thagard, P., eds., *Model Based Reasoning in Scientific Discovery*, London and New York: Kluwer/Plenum.
- GRUBER, H. (1985). From epistemic subject to unique creative person at work', *Archives de Psychologie*, 53: 167-185.
- GRUBER, H. (1990). The Cooperative Synthesis of Disparate Points of View, in Rock, I., ed., *The Legacy of Solomon Asch: Essays in Cognition and Social Psychology*, pp. 143-158.
- JOHNSON-LAIRD, P. N. & WASON, P. C. (1977). A theoretical analysis of insight into a reasoning task, in P. N. JOHNSON-LAIRD & P. N. WASON, Eds., *Thinking: Readings in Cognitive Science*, Cambridge: Cambridge University Press, pp.143-157.
- LAKATOS, I. (1970). Falsification and the methodology of scientific research programmes. In I. LAKATOS & A. MUSGRAVE, Eds., *Criticism and the Growth of Knowledge*, Cambridge: Cambridge University Press, pp. 91-196.
- LUCE, R.D. & RAIFFA, H. (1957). *Games and Decisions*. New York: John Wiley.
- SALMON, W. (1990). Rationality and objectivity in science, OR, Tom Kuhn meets Tom Bayes, in C. W. SAVAGE, Ed., *Scientific Theories*, Minneapolis: University of Minnesota Press, pp. 175-204.
- TWENEY, R. D. (1985). Faraday's discovery of induction: a cognitive approach, in D. GOODING & F. JAMES, Eds., *Faraday Rediscovered*, London: Macmillan / New York: APA Press, pp. 189-209.
- WASON, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology*, 12: 129-140.
- WASON, P. C. & SHAPIRO, D.A. (1971). Natural and contrived experience in a reasoning problem. *Quarterly Journal of Experimental Psychology*, 23: 63-71.
- WEINER, P. P. (1966). *Charles S. Peirce: Selected Writings*. New York: Dover.