

Ecologically Rational Agency

John Debenham, Carles Sierra

► **To cite this version:**

John Debenham, Carles Sierra. Ecologically Rational Agency. Max Bramer. Third IFIP TC12 International Conference on Artificial Intelligence (AI) / Held as Part of World Computer Congress (WCC), Sep 2010, Brisbane, Australia. Springer, IFIP Advances in Information and Communication Technology, AICT-331, pp.3-12, 2010, Artificial Intelligence in Theory and Practice III. <10.1007/978-3-642-15286-3_1>. <hal-01054592>

HAL Id: hal-01054592

<https://hal.inria.fr/hal-01054592>

Submitted on 7 Aug 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Ecologically Rational Agency

John Debenham¹ and Carles Sierra²

¹ Centre for Quantum Computation & Intelligent Systems,
University of Technology, Sydney, Australia

² IIIA, CSIC, Campus UAB, 08193 Bellaterra, Catalonia, Spain

Abstract. The concept of dual rationality in human agents is well understood. It is significant in the work of the economist and Nobel Laureate Friedrich Hayek. In psychology Sigmund Freud described ‘dual information processing’. More recently in cognitive science Daniel Levine has extensively studied the dual mechanisms in the human brain that handle these two forms of rationality. It is perhaps surprising that work in artificial intelligence has focussed heavily on Cartesian constructivist rationality (particularly on game theory, decision theory and logic) as humans deliberate using models based on reason for only a minute proportion of the decisions that they make. This paper aims to open discussion on ecological rationality in artificial, computerised agents, and to examine some of the issues in designing ecologically rational agents.

1 Introduction

This paper is concerned with *bounded rationality* that dates back to David Hume and more recently to the early work of Herbert Simon. Bounded rationality refers to systems that are not founded on Cartesian rationalism; it has been widely addressed in economics [9], and is discussed in all good books on artificial intelligence, e.g. [10]. For over fifty years artificial intelligence research has spawned countless theories and systems that are *not* founded on Cartesian rationalism; one classic contribution being Rodney Brooks’ work reported in his ‘Computers and Thought’ award-winning paper [2]. Despite these advances, work in multiagent systems has been heavily influenced by game theory, decision theory and logic [11].

In economics two forms of rationality are described by the two Nobel Laureates Friedrich Hayek [7] and Vernon Smith [14] as being within ‘two worlds’. Hayek and Smith identify; *constructivist rationality* that underpins rational predictive models of decision making; and, *ecological rationality* founded on the concept of “spontaneous order³” that refers to social institutions and practices that *emerge* from the history of an agent’s interactions and are *not* pre-designed. For computerised, intelligent agents the predominant logical distinction is between *deliberative* and *reactive* logic. Hayek and Smith’s two rationalities relate directly to two distinct forms of deliberation, and have

³The term ‘order’ refers to: traditions, customs, norms, rules and guidelines. An agent may belong to a number of normative systems (or, electronic institutions [1]) whose norms may be shared with, or in conflict with, those of other systems. The ‘extended order’ includes the whole show. If a multiagent system interacts with human society then its norms will respect the rules and laws that apply to society as a whole.

little to do with autonomic reactivity that typically overrides other processes in both the human neuropsychological system and in intelligent agents.

In psychology ‘dual information processing’ is a key concept as Seymour Epstein observed [4]: “There is no dearth of evidence in every day life that people apprehend reality in two fundamentally different ways, one variously labeled intuitive, automatic, natural, nonverbal, narrative, and experiential, and the other analytical, deliberative, verbal, and rational, (p. 710)”. These ideas date back to Sigmund Freud [5]. More recently, the different brain processes that support dual rationality have been extensively studied in cognitive science by Daniel Levine [8]. But, in artificial intelligence the term ‘rationality’ appears to apply exclusively to constructivist rationality.

We assume that the following is true of all real systems and environments: *There are always facts that an agent does not know that may effect the behaviours of other agents or states of the environment.* In particular: *Agent α expects to know less about agent β 's commitments than β knows.* This is *not* about pervasive uncertainty, it is about a limit to that which is known — a necessary ignorance that is distinct from, but related to, Murphy’s law “Anything that can go wrong will go wrong” — we don’t and can’t know the extent of the “anything”.

We also assume that the following is true of all real systems and environments: *Every fact that an agent knows is uncertain.* This *is* about pervasive uncertainty. When an observation is made it may be momentarily assumed to be true but as time progresses its truth will become more and more uncertain.

The term *rational* means “showing clear thought or reason”, and so to be rational in real environments means: to show clear thought or reason in the presence of necessary ignorance and pervasive uncertainty. Two foundations for rational behaviour are:

Cartesian constructivism. Cartesian constructivism deals with the unknown by assuming a theory that: (i) specifies what facts are required to be known, and (ii) makes assumptions about *some* of the unknowns. See Section 3.1. Game theory is a well known example of Cartesian constructivism, and requires that an agent knows its *actions*, *consequences*, a *consequence function* and a *preference relation*; and assumes that each agent will act so as to maximise its expected utility.

Ecological rationality. This is deliberation that uses past experience and contextual triggers to build action sequences from experiential memory. Past experience is a precursor to ecological rationality. For example, as we have described them previously, trust and honour [12] and reputation [13], are purely ecological concepts [3]. Building action sequences from experiential memory involves more than just retrieval. An agent has: to learn to imitate the actions that it believes that others do, to form expectations of the effect of actions, to select actions from a set of candidates, to adapt actions to suit the current norms and state of the environment, and when things don’t work out to learn to experiment with untested actions.

There may be situations in which neither Cartesian constructivism or ecological rationality may be useful, and there may be situations in which they may both be useful.

As the name suggests, ecological rationality is concerned with a richer form of bounded rationality than simplifying the calculation of a theoretically ‘optimal’ action by: rules for simplifying search, rules for terminating search or heuristic decision rules

to select actions from an incomplete set of options as described in [6]. Ecological rationality is taken in the context of the Hayekian view [7] in which agents evolve themselves together with the norms of the systems they inhabit⁴ whilst their environment changes. This all sounds rather Darwinian, but Hayek is careful to distinguish between genetic evolution and cultural evolution [op. cit. page 23].

In previous work we have described some of the tools that an ecologically rational agent will require [3]. In this paper we describe the essence of an ecologically *intelligent* agent. This paper is organised as follows. Various preliminaries are described in Section 2. Section 2.1 discusses the capabilities that rational agents require. Section 3 discusses dual rationality for computerised agents; Section 3.1 cartesian rationality and Section 3.2 ecological rationality. Section 4 describes a basis for the rational deliberative process, and Section 5 concludes with a discussion on what it means for a rational agent to be “intelligent”.

2 Preliminaries

Agent α is in a *multiagent system* with a finite but unknown number of other agents $\mathcal{B} = \{\beta_1, \beta_2, \dots\}$, and a finite number of *information providing agents* $\Theta = \{\theta_1, \theta_2, \dots\}$ that provide the *context* for all events in the system. Θ^t denotes the state of these agents at time t , and \mathcal{I} is the set of all these possible information states. The information providing agents provide freely available data to anybody who wants to listen. This includes, day of the week, time of day, and so on, as well as news for general distribution. In addition, there is an *information search service*⁵ that proactively seeks information. In so far as the information providing agents are equivalent to an electronic news sheet the search service is equivalent to a search engine, for which, in some cases, a fee may be charged, particularly for confidential information.

The only thing that an agent ‘knows for certain’ is the history of its past communication. Agent α ’s *history* at time t , \mathcal{H}_α^t , is the complete history of: its prior interactions with other agents, observations of information providing agents, and information proactively obtained from the search service. Each *utterance* in the history contains: an illocutionary statement, the sending agent, the receiving agent, the time that the utterance was sent or received. Utterances are organised into dialogues, where a *dialogue* is a finite sequence of related utterances. If a dialogue is not yet complete then it is an *open dialogue*, that is an interaction that an agent is currently “working on”. For convenience we assume that each agent has at most one open dialogue — that is, it participates in at most one “sequence of related utterances” at a time.

An agent acts to satisfy a *need*, v , that are always considered in context, (v, Θ^t) . An agent acts by communicating an utterance, (μ, β) , containing an illocutionary statement, $\mu \in \mathcal{M}$, to another agent, $\beta \in \mathcal{B}$. If an utterance is part of a complete dialogue,

⁴The evolution of individual agents and component systems are not considered in isolation — the whole ensemble evolves in response to itself and to the environment — they are *complex* systems. For example, in Hayek’s extensive writing there is little mention of ethics as it too evolves.

⁵This does not feature in the discussion but it is essential if agents are to proactively seek information and information sources. We just assume it is available.

d , that aimed to satisfy a need then the dialogue is tagged with: the triggering need, v , the prevailing context, Θ^t , and a rating $r \in R$ of how satisfactorily the dialogue satisfied the need. So such a *rated dialogue* will have the form: $d = (d, v, \Theta^t, r) \in \mathcal{H}_\alpha$.

We assume that agents aim to satisfy needs through participating in dialogues that will hopefully achieve a high rating. The *rating* of a dialogue is an *ex post* evaluation of the outcome of the dialogue taking account of the need that triggered it and changes in the world state, or the evaluating agent's state, that occurred during it. A rating acknowledges that the partner agents involved may have taken account of changes in circumstance that occurred during the dialogue, or have even gone "over the odds" and gave more than was expected in some sense. It also includes the possibility that the partner agent may have taken advantage of information asymmetry between the agents, and acted in a way that is contrary to that which α would have desired had he known all the facts. Ratings are on a discrete, finite fuzzy scale.

Ratings are only defined at the *completion* of the dialogue and so it can not play a role *before* or *during* the dialogue — ratings are not 'utility functions' in any sense — they are subjective, *ex post* assessments of outcomes of dialogues that account for the changing context during the entire interaction process. An agent may select an interaction partner and direct the subsequent dialogue in the hope that it will achieve a high rating — in reality that is the best he can do. It will only be possible to guarantee a high rating if the agent knows in advance⁶ all the facts that could effect the rating — in reality *this never happens*.

2.1 Capabilities of Rational Agents

An agent should be able to:

- *observe* and *understand* the behaviour of other agents — understanding involves knowing or guessing the other agents' motivations and *judging* the way in which the behaviours satisfy those motivations. In the special case that the agent receives instruction, it is reasonable to assume that behaviours are presented to the agent with a range of goals or motivations that the behaviours achieve or satisfy. The idea of "teaching or instructing agents" is much simpler than having them observe and understand.
- *imitate* the behaviour of other agents — imitation is more than simple copying what another agent does, it is using another agent's behaviour to satisfy its own motivations so necessarily involves the *assessment* of those behaviours in its own context.
- *remember* what it does and what it observes — the agent's memory should be constructed with efficiency of recall in mind. This means that the agent may *classify* observations based on some value judgement about their potential future use. There is no point in remembering rubbish. It may be useful to *abstract* observations to some degree to facilitate adaptation and reuse, even if this amounts only to simple "typing". In any case each behaviour memorised is *rated* for how well it achieved its

⁶In game theory, utility is in a sense an *a priori* rating. Utility is a rating of outcomes and is required by the postulates of game theory to be known in advance.

goal. In the interests of efficiency it may be desirable for the memory to contain summary measures of trust, honour, reliability, reputation, integrity etc, and for these summaries to be updated each time an observation is made. So in addition to representing the facts of an observation, remembering involves: rating, abstracting, classifying and summarising.

- *recall* relevant past behaviours, trust estimates, reputation estimates etc given a new motivation, need or goal. This may be seen as a complex form of semantic difference.
- *act* rationally — i.e. its acts in a way that it believes will achieve a high *ex post* rating. The agent should be striving to optimise its expected ratings. This does *not* mean that the agent is a game-theoretic agent with a utility function defined by the ratings for the following reasons:
 - ratings are an *ex post* assessment — any prior expectation of their value will be ignorant of circumstances at assessment time and can only be based on wishful thinking and not on any logical model — that is, no method can claim to predict, or to optimise, ratings,
 - as ratings are subjective there is no way that one agent can predict how another agent will rate a dialogue, and
 - an agent may select a non-Cartesian (i.e. ecological) basis for its actions in which case they will not be founded on a model that is capable of forecasting the outcome of performing those actions.

How is an agent to “show clear thought or reason in the presence of necessary ignorance and pervasive uncertainty”? If an agent has not been hard-wired with models of the world or other agents then all that it knows is in its memory. This naturally leads to re-use and adaptation as in case-based reasoning.

3 Dual rationality for artificial agents

In this section we define rationality in both its constructivist and ecological form.

3.1 Cartesian constructivist rationality

Given a finite set X , let $\Delta(X)$ denote the set of probability distributions over X . A *constructivist rational* agent α with chosen theory T_α and need v , α will first obtain the data $D'_\alpha = T_{\alpha D}(\mathcal{H}'_\alpha, v, \Theta^t)$ that T_α requires given \mathcal{H}'_α , v and Θ^t . Second, α will act using the lottery:

$$C_\alpha : \mathcal{H} \times \mathcal{N} \times \mathcal{I} \times \mathcal{T} \times \mathcal{D} \rightarrow \Delta(\mathcal{M} \times \mathcal{B}) \quad (1)$$

where \mathcal{H} the set of all history states, \mathcal{N} is the set of need states, \mathcal{I} is the set of information states, \mathcal{T} is the set of theory states⁷, and \mathcal{D} is the set of data states as can be extracted from the search service. As above, \mathcal{M} is the set of illocutionary statements and \mathcal{B} is the set of other agents. C_α is a function that encapsulates α 's interpretation of T_α . For example, if α has an open dialogue then Equation 1 determines (non-deterministically) the next utterance that α makes in that dialogue.

⁷The elements of \mathcal{T} are theories that may be chosen.

Constructivist rational agents are non-deterministic. This form of non-determinism is similar to mixed strategies in game theory. The action performed by a *deterministic constructivist rational agent* is determined by: $C_\alpha : \mathcal{H} \times \mathcal{N} \times \mathcal{I} \times \mathcal{T} \times \mathcal{D} \rightarrow \mathcal{M} \times \mathcal{B}$; they are similar to pure strategies in game theory.

3.2 Ecological rationality

An *ecologically rational agent* α with need v in context Θ^t will act using the lottery $E_\alpha(\mathcal{H}_\alpha^t, v, \Theta^t) \in \Delta(\mathcal{M} \times \mathcal{B})$ where:

$$E_\alpha : \mathcal{H} \times \mathcal{N} \times \mathcal{I} \rightarrow \Delta(\mathcal{M} \times \mathcal{B}) \quad (2)$$

where E_α is a function that is *not* founded on an abstraction or theory that models, explains, constrains, describes or prescribes the behaviour of agents or the environment; it encapsulates α 's particular ecological rationality. As above, if α has an open dialogue then Equation 2 determines (non-deterministically) the next utterance that α makes in that dialogue. Ecologically rational agents are non-deterministic. The action performed by a *deterministic ecologically rational agent* is determined by: $E_\alpha : \mathcal{H} \times \mathcal{N} \times \mathcal{I} \rightarrow \mathcal{M} \times \mathcal{B}$. We make two observations:

- E_α , and any other embedded reasoning, should *not* be founded on an abstraction or theory that models, explains, constrains, describes or prescribes the behaviour of agents or the environment. The sum total of the wisdom available to E_α is that which has been observed in \mathcal{H}_α^t .
- Given two ecologically rational agents α and β , suppose $\mathcal{H}_\alpha^t = \mathcal{H}_\beta^t$ there is no expectation that given the same need, v , in the same context, Θ^t , they will act in the same way even if they are deterministic. E_α encapsulates α 's particular ecological rationality, and E_β β 's. This is in contrast to a pair of deterministic constructivist rational agents that would be expected to perform the same actions if they were both expected-utility-optimisers and had both identified a pure dominant strategy⁸ given \mathcal{H}_α^t .

Example. A classical case-based reasoning agent is a deterministic ecologically rational agent.

The shape of the distribution $E_\alpha(\mathcal{H}_\alpha^t, v, \Theta^t)$ enables us to balance exploration and exploitation. The flatter the distribution the greater the degree of exploration.

An agent that operated as in Equation 2 could have been constructivist if the condition that “ E_α is *not* founded on an abstraction or theory that models, explains, constrains, describes or prescribes the behaviour of agents or the environment” had been omitted — the difference between constructivist and ecological lies in the theory T_α . For example, without this condition an agent that operated as in Equation 2 could place expected utility optimising bids in an English auction. It is perhaps more correct to think of E_α as “interpreting the wisdom” in \mathcal{H}_α^t ; just as C_α interprets the wisdom in: T_α , \mathcal{H}_α^t and D_α^t .

⁸In which case we assume that they will have made assumptions about the behaviour of other agents for example that they are expected-utility-optimisers as well.

The “ecological rationality” in E_α is based on the belief that the wisdom in \mathcal{H}_α^t can somehow tell α how to act rationally. This belief is not an abstraction or “theory” (as described above); it is simply a belief that the wisdom embedded in prior observations are a basis for rational action. In a simple form, this belief may be that prior agent behaviour reliably indicates future behaviour⁹. That is, ecological rationality may be founded on *a sense of trust* that α has in agents, i.e. that they will continue to behave with no less integrity than that which they have displayed in the past. Ecological rationality may also be founded on the reputation that another agent has, or on trust in the institution (i.e. a normative multiagent system to which all the agents belong) to ensure that agents behave in some way. In addition, Ecological rationality may be founded on subtle observations mined from \mathcal{H}_α^t . As a simple example, “Whenever John is running a marketing promotion Carles invariable gives excellent service”. In general, ecologically rational behaviour will be derived from \mathcal{H}_α^t using data mining techniques. Ecological rationality, like trust, is an experience-based phenomenon — it can not exist without experience. At the “top end”, ecological rationality embodies all the models that have been described from information-based agents including the important notion of integrity¹⁰.

3.3 Selecting a rational basis for deliberation

We now address the question of when both rationalities may be useful which should be chosen? That is, α will require a mechanism, A_α , to determine which form of deliberation, r , to select in a given circumstance:

$$r = A_\alpha(\mathcal{H}_\alpha^t, \mathbf{v}, \Theta^t) \in \{C_\alpha, E_\alpha\} \quad (3)$$

What really matters is an agent’s *ex post* rating of a dialogue. E.g. an agent may have used game theory to obtain a shatteringly good deal only to hear that the person he was negotiating with (also a good friend) was sacked for agreeing to such good a deal. This may cause the agent not to rate the dialogue highly. One approach is to rate all dialogues and to use the method in Section 4 to select a dialogue and then: if the closest high-rated dialogue in the history to it was generated by a Cartesian method then re-use that method as long as the data is available otherwise use Section 4 to drive ecological deliberation [3].

4 A Basis for Rational Deliberation

An ecologically rational α has only the dialogues in \mathcal{H}_α to work with, and, if α is concerned with ratings, maybe only the rated dialogues. No two dialogues are identical — they will differ at least by the time at which they occurred — at best they will be ‘similar’. This does not necessarily mean that α ’s beliefs will be uncertain about how

⁹Such a belief may incorporate context. For example, “John is most reliable *except* on Mondays”.

¹⁰The extent to which a partner agent will take advantage of his private information when enacting his commitments. E.g. “I haven’t got the strawberries you ordered because they were all rain damaged.”

another agent will behave: “Carles always looks after me — no question about it”. But if the history lacks consistency in a sequence of the dialogue ratings then it may well mean that α will be uncertain about how another agent will behave. One way to approach this uncertainty is to measure the similarity between a current, open dialogue, (d, v, Θ^t) , and a rated dialogue $(d', v', \Theta', r') \in \mathcal{H}_\alpha$. Such a measure, $\text{Sim}((d, v, \Theta^t), (d', v', \Theta', r'))$, may be useful in forming a necessarily uncertain belief about the eventual rating of the open dialogue. If the ratings in the history are totally chaotic then it may mean that the best that an ecologically rational agent can do is to act at random — i.e. E_α will be the uniform, maximum entropy distribution.

Devising a satisfactory definition for $\text{Sim}(\cdot, \cdot)$ is not simple. It would be very convenient if such a definition could be found so that:

$$\mathbb{P}(\text{rating}(d, v, \Theta^t) = r) = \frac{\sum_{d' \in \mathcal{H}_\alpha^t} (\text{Sim}((d, v, \Theta^t), (d', v', \Theta', r')) \mid r' = r)}{\sum_{d' \in \mathcal{H}_\alpha^t} \text{Sim}((d, v, \Theta^t), (d', v', \Theta', r))}$$

for any open dialogue d, v and Θ^t . $\mathbb{P}(\text{rating}(d, v, \Theta^t) = r)$ is the expectation of the rating of the open dialogue (d, v, Θ^t) . The following assumes that $\mathbb{P}(\text{rating}(d, v, \Theta^t) = r)$ may be estimated somehow. \mathcal{H}_α^t can be pruned¹¹.

If the rating scale has a linear order it is then a simple matter to estimate:

$$\mathbb{P}(\text{rating}(d, v, \Theta^t) > \text{rating}(d', v', \Theta'))$$

and in particular:

$$\mathbb{P}(\text{rating}(d, v, \Theta^t) > \text{rating}(d', v, \Theta^t))$$

Then in the simple case that we can select an entire dialogue “up front”¹³ given v and Θ^t , one simple strategy is for E_α to select dialogue d with a probability:

$$\mathbb{P}(\forall d' (\text{rating}(d, v, \Theta^t) > \text{rating}(d', v, \Theta^t)))$$

that can be estimated approximately from the above with suitable dependence assumptions. In natural language: “we select a dialogue with a probability equal to the probability that it is the best choice”. The dialogue selected by this method is not necessarily in \mathcal{H}_α^t .

¹¹Pruning \mathcal{H}_α^t . Given a newly triggered need, v , and prevailing context, Θ^t , then let the candidate dialogues for (v, Θ^t) be:

$$\text{Can}_\alpha(v, \Theta^t) = \{d_i \in \mathcal{H}_\alpha^t \mid \text{Sim}((v, \Theta^t), (v_i, \Theta_i^t)) < \sigma(v, \Theta^t)\}$$

where $\text{Sim}(\cdot)$ is a similarity function and $\sigma(v, \Theta^t)$ is a threshold constant such that:

$$\|\text{Can}_\alpha(v, \Theta^t)\| > n \text{ and } \|\{d_i \in \text{Can}_\alpha(v, \Theta^t) \mid r_i \geq \text{'excellent'}\}\| > m$$

where n and m are chosen constants. That is, $\text{Can}_\alpha(v, \Theta^t)$ contains at least n nearest neighbours of (v, Θ^t) and contains at least m with an ‘excellent’ rating¹². The idea then is to use $\text{Can}_\alpha(v, \Theta^t)$ instead of \mathcal{H}_α^t .

¹³This is not unrealistic. The dialogue could be “Please send 1kg tomatoes”, “1kg tomatoes delivered”, “Please send €1”, “Completed”

Now consider a more general situation, given an open dialogue d , α wishes to know what utterance to communicate next, where an utterance, u consists of an illocution, m , sent to an agent, β , $u = (m, \beta)$. Let $d \oplus u$ denote the dialogue consisting of d followed by u . Then following the reasoning above we can estimate:

$$\mathbb{P}(\forall u' (\text{rating}(d \oplus u, v, \Theta^t) > \text{rating}(d \oplus u', v, \Theta^t)))$$

that is the probability for any u that it will be the highest rated choice. Then, as above, we select u with a probability equal to the probability that it is the highest rated choice.

We clarify what it means to “re-use” a dialogue. For Cartesian dialogues we simply reuse the theory and acquire current data to feed into it. For ecological dialogues things are more subtle. No two dialogues are the same, but when we say to our butcher “I need 2kg of prime beef” we are probably reusing an established “dialogue type”, or *pattern*, with that agent (the butcher). In this case the pattern could be something like: “state your need and see what is proposed”. The fact that I may never have purchased 2kg of prime beef previously does not mean that I am not reusing something. An abstraction of dialogues to patterns is required to deal with this. Now given that we understand the sense in which an Cartesian or ecological dialogue is being reused, it is reasonable to suggest that we would be able to estimate the prior distribution $\text{rating}(u, v, \Theta^t)$ for initial utterance u . For example, u could be “I need 2kg of prime beef”. It is then a matter of monitoring how the dialogue, d , is proceeding so that we can update the prior estimate of the distribution to obtain $\text{rating}(d, v, \Theta^t)$. The abstraction process is tricky but quite intuitive.

The above argument “hides a multitude of sins” by assuming that:

$$\mathbb{P}(\text{rating}(d, v, \Theta^t) = r)$$

can be estimated for any d . In particular it enables us to “pluck dialogues out of thin air”. The problem that is neatly hidden is: where does \mathcal{H}_α come from?

For human agents, \mathcal{H}_α^t is built up under guidance from parents, school teachers and from their own experience as they as young adults gradually interact with the world in richer and richer ways. Computerised agents are orphans, and can only initially rely on the ideas of their creator; these may well include theories such as game theory. Cartesian deliberation generates dialogues that may be rated and will be represented in \mathcal{H}_α^t , but they may not provide suitable fodder for ecological deliberation because the rationality in former Cartesian dialogues will be grounded in the particular theory chosen, T_α , and not in the agent’s experiential history.

For computerised agents the challenge for (ecological) deliberation is to strategically shape, develop and maintain its experience base so that it contains sufficient wisdom for the (unknown) future. This is what smart business men do. It is what smart social operators do. Computerised agents need to be able to “observe, understand, evaluate, imitate, remember and recall”.

5 Discussion: What is an *intelligent* rational agent?

It seems appropriate to end this discussion on rationality by considering the question “What is an *intelligent* rational agent?”. Intelligence is a comparative and not an absolute concept. The question “Is α intelligent?” is meaningless, whereas “Is α more

intelligent than β ?” may make sense. Intelligence has nothing to do with knowing scientific theorems or methods that are reported in the public domain. If a software agent was endowed with every known scientific theorem then that in itself would not make the agent intelligent. *Intelligence* is all about the following four questions:

- information about the world and about other agents — how do you source it, how do you rate its integrity in time?
- how to build trusted relationships with other agents that enable you to benefit from the present and, equally important, to cope with the unknown future?
- how to design and manage an agent’s memory to benefit from the present and, equally important, to cope with the unknown future?
- how to select a basis for rational action from all the stuff an agent knows — in plain language, how to select a method to achieve a particular goal?

What makes one agent more intelligent than another is the way it does these four things.

References

1. Josep Lluís Arcos, Marc Esteva, Pablo Noriega, Juan Antonio Rodríguez, and Carles Sierra, *Environment engineering for multiagent systems*, Journal on Engineering Applications of Artificial Intelligence **18** (2005).
2. Rodney A. Brooks, *Intelligence without reason*, Proceedings of the 12th International Joint Conference on Artificial Intelligence (Sydney, Australia) (R. Myopoulos and J. Reiter, eds.), Morgan Kaufmann, August 1991, pp. 569–595.
3. John Debenham and Carles Sierra, *Dual rationality and deliberative agents*, Research and Development in Intelligent Systems XXVI (Max Bramer, Richard Ellis, and Miltos Petridis, eds.), Springer Verlag, 2009, pp. 79–92.
4. Seymour Epstein, *Integration of the cognitive and the psychodynamic unconscious*, American Psychologist **49** (1994), no. 8, 709–724.
5. Sigmund Freud, *Formulation of the two principles of mental functioning*, Hogarth, London (originally published in 1911), 1958.
6. Gerd Gigerenzer and Reinhard Selten (eds.), *Bounded rationality — the adaptive toolbox*, MIT Press, Cambridge, MA, 2002.
7. F. A. Hayek, *The fatal conceit: The errors of socialism*, University Of Chicago Press, 1991.
8. Daniel S. Levine, *Brain mechanisms for making, breaking, and changing rules*, Advanced Intelligent Computing Theories and Applications. With Aspects of Contemporary Intelligent Computing Techniques (De-Shuang Huang, Donald C. Wunsch II, Daniel S. Levine, and Kang-Hyun Jo, eds.), Springer Berlin Heidelberg, 2008, pp. 345–355.
9. Ariel Rubinstein, *Modeling bounded rationality*, MIT Press, Cambridge, MA, 1998.
10. S. Russell and P. Norvig, *Artificial intelligence: A modern approach*, 2nd ed., Prentice Hall, 2002.
11. S.J. Russell, *Rationality and intelligence*, Artificial Intelligence **94** (1997), no. 1-2, 57–77.
12. C. Sierra and J. Debenham, *Trust and honour in information-based agency*, Proceedings Fifth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2006 (Hakodate, Japan) (P. Stone and G. Weiss, eds.), ACM Press, New York, May 2006, pp. 1225 – 1232.
13. ———, *Information-based reputation*, First International Conference on Reputation: Theory and Technology (ICORE’09) (Garganza, Italy) (Mario Paolucci, ed.), 2009, pp. 5–19.
14. Vernon L. Smith, *Rationality in economics: Constructivist and ecological forms*, Cambridge University Press, 2007.