

1. Rationality

1.1.A Framework for Rationality

The concept of rationality is applied in many different *contexts*, generally involving an evaluation of how well some set of actions can be said to meet an objective or set of objectives.

We can analyze rationality in a context C defined by the following four attributes:

1. *Actions*. A set (one or more) of moves executed in C .
2. *Constraints*. A set of characteristics that apply in (i.e. that characterize) C , relating inputs to outputs.
3. *Evidence*. Any data available as inputs to the constraints that apply in C .
4. *Objectives*. The desired outcomes of actions executed in C .

EXAMPLES:

Logical inference:

- The evidence is the premises.
- The constraints are the rules of inference relating premises to true conclusions.
- The actions are the inferences made in deductive steps.
- The objective is for the inferences to be true.

Choice of an option:

- The evidence is the data of previous, known experience and data about the decision environment.
- The constraints are information about the available options and their respective outcome consequences.
- The action is the chosen option.
- The objectives are to achieve the most preferred outcomes.

Planning a contingent sequence of moves:

- The evidence is the initial state.
- The constraints are all possible actions, their preconditions, and their effects.
- The actions are the contingent sequence of moves decided upon.
- The objectives are the goals which the plan is trying to achieve.

1.2 Agents and Levels of Rationality

Rationality may be evaluated with respect to an agent, a collection of agents, or a total environment.

An agent, or actor, is an individual that executes actions in a given context. The attributes of the context may be defined with respect to a single agent, a set of agents, or to the context itself, giving rise to the following levels of rationality:

1. *Subjective rationality*. The actions of an agent *A* are subjectively rational in a context *C* to the extent that the evidence available to and the constraints as understood by *A* are consistent with the achievement of *A*'s objectives.
2. *Intersubjective rationality*. The actions of an agent *A* are intersubjectively rational in a context *C* to the extent that the evidence available to and the constraints as understood by the collection *B* of agents in *C* are consistent with the achievement of *A*'s objectives.
3. *Objective rationality*. The actions of an agent *A* are objectively rational in a context *C* to the extent that all of the evidence and constraints applying in *C* are consistent with the achievement of *A*'s objectives.

1.3 Goals for the Evaluation of Rationality

The analysis of actions in a given context may be made with respect to different goals. Two goals for the evaluation of rationality in a context are:

1. *Normative analysis*. How well do possible actions meet objectives in a context *C* at a given level of rationality?
2. *Descriptive analysis*. How well do actually chosen actions meet objectives in a context at some level of rationality?

1.4 Characteristic Evaluations in the Descriptive Analysis of Rationality

When we evaluate the rationality of a set of actions undertaken in a given context, some possible answers regarding the question of whether an agent is (or agents generally are) acting rationally are:

- Yes, usually
- Yes, by definition
- Yes, within the limits of time, information, etc. that they usually have
- Yes, for the problems they evolved to cope with
- Yes, in the aggregate
- No, and we can do something about it
- No, and there isn't much we can do to improve