Giorgio Bongiovanni
Giovanni Sartor
Chiara Valentini
Editors

Law and Philosophy Library 86

Reasonableness and Law

Springer

4 Rationality and Practical Determinations

Epistemic reasoners do not restart from scratch whenever they need to understand or predict their environment: they approach new a situation by using previous epistemic determinations, stored as beliefs in their memory. Similarly, practical reasoners keep memory of their practical determinations to use them at a later time, rather than restarting their practical reasoning from scratch whenever a problem comes up. Once an agent has adopted certain practical determinations, rationality requires the agent to proceed on the basis of these determinations, giving them an appropriate weight in his or her reasoning process, until the agent abandons such determinations. For instance, it would be (procedurally) irrational for me to have an intention to make a €50 donation to charity X and not do so (though it may not be irrational for me to withdraw such an intention in appropriate circumstances, for instance, upon discovering that X is a fake charity, exclusively aimed at enriching its organisers).

Not deriving the practical conclusion supported by our conative states is as irrational as failing to derive the consequences of our beliefs. Failing to proceed in reasoning may be the right thing from the point of view of an external observer who knows that the belief or conative state providing the premise of the inference at issue is wrong (and reliance on it is likely to lead the reasoner to further false beliefs or inappropriate determinations), but it is irrational from the internal perspective of the reasoner, who has nothing else to go on. This does not mean that we should put an absolute trust on our cognitive states (I know that my determinations may be wrong, and my beliefs false, even when I sincerely believe that they are right and true), but the awareness of our fallibillity only justifies continuing the inquiry meant to question such determinations and beliefs (when we have no more-urgent things to do), it cannot justify epistemic and practical paralysis.

In the model of reasoning provided in Section 3, an important role in storing practical determinations is played by intentions: an intention stores the outcome of a teleological deliberation, and it prompts to action. An agent having an intention to do action A under condition B is ready to (unconditionally intends to) perform A when B is met and is committed to perform A, in the sense that it would be irrational for the agent not to perform A, as long as the agent continues to have that intention (and believes that B obtains).

Intentions are not the only practical determinations we store and reuse: by retrieving from memory our preferences and goals, we input them into our present reasoning, and we are indeed justified in doing so, as long as we do not have prevailing reasons to the contrary. However, it seems to me that we often also choose our preferences and goals by making them conscious objects of intentions. This happens in particular when, having questioned a preference or a goal, we come to a determination to adopt it: then we form an intention to have that preference, or to pursue that goal. Subsequently, we retrieve this intention and implement it by adopting the preference or goal contained in the intention. Thus an intention can take different contents:

- a determination to act in a certain way on a particular occasion (I will make a €50 donation to Oxfam tomorrow) or to not act in a certain way (I will not go skiing tomorrow);
- a determination to act in a certain way on all occasions of a certain kind (I will make a €100 donation to Oxfam every year; I shall do physical exercise every evening);
- a determination to qualify in a certain way a certain object or fact under certain conditions (I will consider unjust any act which diminishes human happiness; I will assume that any student making a serious effort satisfies the course requirements, regardless of his or her abilities);
- a determination to aim at a certain goal (I will aim to get fit; I will aim to reduce hunger in the word);
- a determination to view a certain situation as a factor or ground favouring a certain practically relevant conclusion, without committing to the view that such a ground is a necessary or a sufficient condition (I will consider a student's laziness as a factor that supports giving him or her a lower mark; I will consider a person's need as a factor that supports helping him or her).

One may adopt these determinations taking different perspectives: a self-interested perspective (e.g., when deciding what restaurant to go to, an evening when I am on my own), an altruistic view (e.g., when choosing to donate to a charity), or the perspective of a participant in a certain community (e.g., when accepting to follow the rule that professors should reply to student e-mails, or the rule that citizens should respect speed limits). When we are taking the perspective of a participant in a community (which can have different extensions: a private or public organisation, a country, a federation, the whole of humanity, etc.), we adopt determinations assumed to hold for the community itself (or for particular members of it) on the basis of reasons we believe to be appropriate to that community: we each participate in communal reasoning, possibly taking into account our particular role within the community (on participation on collective reasoning and action, see, for instance, Postema 1995; Tuomela 2000). This takes place when we are acting as members of a political body or as a officials participating in the legal process, but also when we willingly practice the norms of our community, coordinating our own behaviour with that of others (I am not considering the case where someone pretends to be acting out of a communal concern while acting out of self-interest).

It seems to me that normative legal reasoning (the reasoning through which one establishes, for oneself and for others, what patterns of behaviour should collectively be enforced) pertains to this kind of reasoning, so that legal norms may be viewed as collective counterparts of individual commitments (intentions). More generally, the process of rationality concerns not only individuals but also collectives and institutions. Also in the latter case, decision-making will proceed stepwise, moving from reasons to conclusions, but reasoning will be performed by different individuals, acting on the basis of the cognitive states (beliefs, goals, intentions) they attribute to the collectivity in which they participate, and in particular on the basis of the

commitments and legal norms adopted by that collectivity through the procedures it recognises.

This requires rationality to be detached from individual self-interest, so that we can each apply it from the different perspectives we may adopt. Thus, we may rationally pursue our self-interest, but we may also rationally act for the benefit of others (regardless of the particular effects on our wellbeing: see Sen 2004b) or in the interest of the policy we represent. In the latter case, we should locate our determinations within the decision-making process of that policy and within its implementation (see Pettit 2002). Indeed, there is nothing mysterious about the fact that we may view ourselves as rationally pursuing the interests of our community, organisation, or institution, both in the public and in the private domain, rather than our individual benefit. Simon (1965, 205) speaks in this regard of *identification*: "a person identifies himself with a group when, in making a decision, he evaluates the several alternatives of choice in terms of their consequences for the specified group." Rational participation in collective decision-making not only requires reference to the interest of the group (collectivity, community, organisation): it also requires that we be aware that our decision is to be located within the group's decision-making process, where we are entrusted with a specific role (legislator, judge, administrator, etc.), involving specific functional requirements and constraints (for further consideration and references, see Sartor 2005, Chapters 9 and 13). Rational participation thus requires that we accept as bases of our own reasoning (as far as our participation in the collective activities is concerned), the relevant practical or even epistemic determinations already adopted by the group in which we are participating, and it requires that we view our choices (made in name and on account of the group) in the framework of the group's reasoning process (namely, as determinations which are guided by earlier determinations of the group and which possibly contribute to guide its successive determinations): thanks to this identification, the group (and in particular a legal community) can be viewed as a subject to which mental states (goals, intentions, commitments) can be rightfully attributed (Pettit 2004), and which may be capable of rational action.

5 The Reflective Dimension of Reasoning

Reflective reasoners not only have certain cognitive states, but they are also able to critically examine, question, and revise their cognitive states. They wonder whether they should have a certain belief, preference, goal, or intention. In case they conclude that they should not have a certain cognitive state they possess, they withdraw it (though this may not be easy or immediate). In case of a conflict between incompatible attitudes (two contradictory beliefs, two intentions leading to incompatible actions, etc.), they consider the comparative strength of such attitudes, withdrawing the less important of them (unless one of these attitudes appears upon reflection to be unacceptable on other grounds).

This reflective attitude leads reasoners to couple their conative states with apparently doxastic states (Pollock 1995): each such state is associated with a proposition,

which can be viewed as a bearer of truth and falsity, and belief in which leads the reasoner to the corresponding practical determination (for instance, the proposition that I should do *A* is a reason for me to adopt an intention to do *A*).

- My intention to act in a certain way on a particular occasion is associated with my belief that I should act in that way on in that occasion (if the weather is nice tomorrow, I should go skiing).
- My intention to act in a certain way on certain or on all occasions of a certain kind is coupled with my belief that I should act in that way on all such occasions (I should make a €100 donation to Oxfam every year).
- My intention to qualify in a certain way a certain object or fact under certain conditions is accompanied by the belief that the corresponding qualities obtain under such conditions (any act diminishing human happiness is unjust; the laws enacted by Parliament are legally valid).
- My having a goal *A* is coupled with my belief that I should pursue *A*, or that *A* deserves to be pursued, i.e., that *A* is a valuable goal, or a value (friendship is a value; privacy is a value, etc.).
- My viewing *F* as a factor favouring a practical conclusion *C* is coupled with my belief that the corresponding connection obtains, namely, that *F* is indeed a factor in concluding C (having committed a crime for the purpose of obtaining certain advantages favours the conclusion that one is not to be entitled to such advantages).

This method allows reasoners to detach themselves from their cognitive states and recast practical reasoning as if it were a piece of epistemic reasoning. Consider, for instance, the pattern of practical reasoning according to which someone, having goal *G* and believing that action α would achieve *G*, forms an intention to perform *A*. This can be recast as the apparently epistemic pattern according to which someone, believing that *G* is a value and believing that action α would achieve *G*, forms the belief that he or she should implement α. The reformulation of conative states (goals and intentions) into beliefs seems to transform practical reasoning into epistemic reasoning, but the applicable reasoning patterns remain substantially the same: the conditions under which it would be rational for one to believe the proposition corresponding to the content of a certain conative state are exactly the same as those under which it would be rational to adopt such a state. For instance, the conditions under which it would be rational for me to believe that I should perform action α, given my belief that *G* is a value (a valuable goal), are the same as those under which rationality would lead me to intend to perform action α given that I am pursuing goal *G*. Giving an epistemic form to practical reasoning, however, has some advantages. First of all, by transforming our conative states into beliefs, we can more easily detach our selves from our practical attitudes and submit them to critical analysis. Moreover, we can more easily distinguish what practical propositions (and consequently what determinations) we should adopt from different perspectives (as a fully egoistic-self interested individual, as altruistic parents taking care of our family, as members of a polity acting for its common good, etc.). Finally,

this approach has a distinctive advantage when we get to dialogical argumentation, and views about what one or everyone should do (or should not or may not do) become the focus of moral or legal argument (on the way conative attitudes become quasi-epistemic propositions, see Blackburn 1998).

Reflective reasoners proceed as well in an upward direction: they not only infer certain cognitive states from the cognitive states they already have, but they also move in the opposite direction. When a belief, intention, or goal they have appears to be questionable, they consider whether they may or should accept premises that justify having such an attitude, and they may also question such premises. In case they are unable to find an appropriate justification, they tend to abandon the unjustified attitude, or at least they tend to take it with some reservations.

Finally, a rational reasoner needs to keep coherence (see Thagard 2001) among his beliefs and attitudes. Coherence on the one hand, increases the chance that beliefs are true and, on the other hand, contributes to ensure that preferences and goals are realised over time (which would be impossible if one was randomly adopting contradictory beliefs or was adopting incompatible preferences and goals at different times). The requirement of coherence, however, should not be taken too strictly, and in particular, it should not be understood as requiring full logical consistency. Our persistent cognitive states constitute an argumentation framework, namely, a set of propositions, rules of experience, norms, goals, intentions, and preferences (some of which may be in conflict) including inborn attitudes as well as the stored outcomes of our previous cognitive efforts. This information includes in particular defeasible inference policies, namely, rules that tell us what conclusions to derive from certain preconditions, but only as long as such rules are not contradicted or undercut by prevailing arguments to the contrary. In an argument framework, usability matters more than consistency, that is, the framework must be compact and flexible enough to enable the reasoner to anticipate experience and make appropriate choices.

6 Teleological Reasoning: Using Reason in the Pursuit of Goals or Values

Practical reasoning is broader than teleological reasoning, understood as the procedure though which one constructs plans to achieve goals and becomes committed to implementing these plans (adopts the intention to do so). However, teleological reasoning constitutes the core of practical reasoning (Nozick 1993) and a large part of legal reasoning and problem-solving indeed consists in teleological reasoning.

Teleological reasoning consists in the following: a reasoner that aims to achieve a certain goal constructs and tests possible plans to achieve that goal, and then adopts a plan once he or she is satisfied that it appropriately achieves the goal. The adoption of the plan consists in forming an intention to implement its instructions (in an appropriate sequence). Here is the schema of *teleological inference*:

Reasoning schema: *Teleological inference*

(1) having goal A; AND
(2) believing that plan B is a teleologically appropriate way of achieving goal A

IS A REASON FOR
(3) intending to execute B

By a *teleologically appropriate* way of achieving a goal I mean a way that—though neither necessarily *being* optimal nor necessarily *believed* to be optimal—is better than inactivity, and not worse that any other plan the reasoner has been able to conceive so far through an adequate inquiry.[4] In fact, believing that a better, incompatible plan is available is a sufficient reason for abandoning the previously adopted plan. This is rational since sticking to the old plan would imply a failure to achieve a superior result. Teleological appropriateness thus combines the idea of *satisficing* with the idea of *critical cognition*. According to the first idea (satisficing), we may justifiably act on the basis of a suboptimal plan: even when we know that the plan is suboptimal (we know that a better plan exists, though we cannot identify), it may still be adequate to our needs. However, according to the second idea (critical cognition) teleological reasoning is inherently defeasible: if we come upon a better way to achieve our goals, then we should abandon the inferior one. Suppose, for example, that I have some money I intend to put in a bank. Suppose that the offer of bank b_1 provides the best conditions, among the offers I have collected so far. Suppose, further, that a financial expert, whom I consider to be both competent and sincere, tells me that she knows of a bank b_2 offering better conditions, but she will not tell me the name of bank b_2 (she gives this information only to her clients, and I do not intend to become one of them). Clearly, under such conditions, rationality commands me to choose bank b_1, even though I know that my choice is suboptimal. However, if I succeed, before making the contract, in coming to know which bank b_2 offers better rates, I should retract my intention to put my money in b_1 and should instead go for the more profitable deal. Defeasibility also characterises teleological reasoning in the legal domain (as was observed in particular by research in Artificial intelligence and law, see for all Benuch-Capon and Prakken 2006). Suppose I am a prosecutor, and I am convinced that the man in front of me has murdered a child, but the legal evidence I have only allows me to request his conviction for child pornography. Clearly, under such conditions, I should try to have him convicted for the latter offence. However, if before the end of the trial I come upon evidence supporting his conviction for murder, I should pursue this stronger count.

As I observed in the previous section, practical reasoning could be given a doxastic form, where conative states are substituted by corresponding beliefs.

[4] I say "not worse" to cover "Buridan ass" cases, namely, cases where the reasoner views two alternative plans as equally good (and both better than inactivity): in such a case, rationality requires that the reasoner adopt one such plan by random choice (rather than being paralysed by an inability to form a preference).

Correspondingly, teleological inference could be rephrased in a pattern where goals are substituted by values and intentions by duties ("shoulds"):
Reasoning schema: *Teleological inference*

(1) believing that *A* is a value; AND
(2) believing that plan *B* is a teleologically appropriate way of realising value *A*

IS A REASON FOR
(3) believing that *B* should be executed

To prevent possible misunderstanding, I should specify what is meant here by a *value*. In the *New Oxford American Dictionary*, this term is said to have the following meanings (among others): "1. the regard that something is held to deserve; the importance or preciousness of something; 2. (values) a person's principles or standards of behaviour; one's judgement of what is important in life." Here I shall not use the term *value* in either of these two senses: for the first of them I shall instead use the term *worth*, and for the second the term *norm*. By *a value*, as is often done in legal and constitutional debates, I shall instead mean a property (a feature or pattern of states of affairs) that deserves to be pursued (to constitute a goal), since states of affairs instantiating this property are better (more valuable) than those not instantiating it. For instance, when I say that freedom is an individual legal value, I mean that, according to the law, each person's freedom deserves to be pursued (this is a goal for the legal system), since the law prefers one's freedom to one's unfreedom. Similarly, when I say that science is a communal legal value, I mean that, according to the law, widespread scientific and technological competence and the ability to produce scientific advances deserves to be pursued, since the law prefers, with regard to science, knowledge and competence to ignorance and incompetence.

A value may be individualised, i.e., its realisation may consist in the fact that the relevant property is satisfied with regard to each single individual, or it may be collective, i.e., it may consist in the fact that the relevant property is satisfied by the community as a whole. For instance, freedom of speech is an individual value, since it is satisfied when each individual enjoys the opportunity to express him- or herself: the situation where a small minority is completely deprived of its freedom of speech, while all other people enjoy it at the maximum level, would still entail a serious subversion of this value. Science, by contrast, is a collective value (e.g., according to the Italian Constitution, which says that the Italian Republic is to promote science), since the value is achieved when there is widespread scientific knowledge and competence, even if some people do not have any scientific knowledge (though an individual's right to participate in science may be violated if ignorance has been imposed upon that individual, rather than depending on his choice or incapacity). On this understanding of the notion of a value, value-based practical reasoning cannot be separated from norm-based practical reasoning. Certain norms directly require value-based reasoning, i.e., they prescribe the pursuit of certain collective or individual values (e.g., culture, privacy, freedom of speech), and complying with these norms requires engaging in the pursuit of these values through teleological reasoning. Other norms do not directly require the pursuit of values, but

rather prescribe particular positive or negative actions (consider for instance the prohibition to prohibition to process sensitive data without the express consent of the data subject), but these norms are justified by the fact that compliance with them would contribute to the achievement of the values at issue (privacy).

7 The Evaluation of Plans

When we have constructed a plan to achieve a certain goal (value), we need to evaluate the plan and decide whether to adopt it. This decision may require a comparison with alternative plans to achieve the same goal. The most abstract model for evaluating and comparing decisional alternatives is provided by *decision theory* (see, for all, Jeffrey 1983). Decision theorists usually assume that the value of an outcome consists in a numeric measure, which is called the *expected utility* of that outcome. Rationality (as it is understood in decision theory) recommends choosing the plan that provides the highest utility, which is a very difficult task.

Let us consider the simplest case first, a fully *deterministic* plan: the plan has just one possible outcome, about which the planner is absolutely certain. The merit of the plan is then to be determined by the expected utility of this outcome alone. Consider the following example: a judge is tasked with deciding whether a convicted criminal should be granted parole, and the judge is absolutely certain that the convict has now essentially changed and will no longer commit any serious crime. The judge thus believes that the only relevant outcome of his decision will be a very positive one: the convicted person will be free again and will probably find a job and contribute to supporting the family. The alternative decision (keeping the convict in prison) will achieve, with equal certainty, a negative outcome: the convict is likely to slide into drug abuse and to be introduced into serious forms of criminality, and that will diminish for this person all chances of finding a job and providing for the family. When one is so lucky as to find so simple a decisional context, the decision is easy, even when one cannot assign numerical utilities: one knows for certain that one decisional alternative is better than all the others.

The situation is more complex when the plan is non-deterministic, that is, when the plan may have different outcomes having different utilities. Consider, for example, the situation of another judge: she has to decide whether to release on parole someone convicted of paedophilia. Suppose the judge believes there is a good chance that the paedophile will now be able to control his impulses, but she is also aware that there is some chance he will relapse and repeat the same crime. According to decision theory, she needs to evaluate each action she may take by considering the utility of each possible outcome of that action, multiplying this value by the chance of that outcome, and adding up all the results she obtains for the different possible outcomes.[5] For example, suppose that the judge makes the

[5] In general, when plan a may lead to n mutually exclusive outcomes $O1, \ldots, On$—each outcome O_i having probabilities Pi and utility U_i—then α's expected utility, $EU(a)$, is expressed by the formula $EU(\alpha) = \Sigma_1^n(P_i^* U_i)$.

following utility assignments: utility 1 to the situation where the paedophile will not relapse and utility −6 to the situation where he will. If there is only a 10% chance that the convict relapses, then a decision to let him free will have a positive score, with an expected utility of 0.3, according to the following calculation:

$$(1^{*}0.90) + (-6^{*}0.10) = 0.30$$

Even if, for the sake of simplicity, we discount the problem that a plan may have multiple possible outcomes, depending on unknown circumstances, we still face a very tough problem when applying such calculations in practical cases: it is very difficult, in many practical domains, to assign a numerical utility to the outcomes of possible plans in such a way as to make it possible to establish their comparative merit. In fact, it seems that in order to "rationally" compare alternative plans (choices, decisions), we need to first analyse the expected outcomes of each one of those plans, by identifying the desired (valuable) features that characterise each outcome and establishing to what degree each feature will be satisfied (promoted) by that outcome. We will then need to assess the total worth of each plan, by considering the plan's combined impact on all those features. Finally, on the basis of such an evaluation, we will need to compare the different alternative plans that we have been able to devise.

For example, when considering a plan to go to a restaurant r, I may consider to what degree I expect that a dinner at r would exemplify the features I desire relative to the food, wine, service, price, and so on. I would then need to compute the overall expected worth of the experience of going to restaurant r as being characterised by the fact that the desired features are satisfied to such a degree. Having done that, I would be able to compare plans to go to different restaurants, each of which offers a different combination of levels as to the quality of the food, wine, and service, and as to price.

Similarly, a judge, when considering different alternative ways of deciding a case, may examine how each possible choice may impact on legal values. For example, a decision that permits putting video cameras in public spaces and storing the footage for a year would impact on two individual legal values, privacy and security. To evaluate this decision, and compare it with possible decisional alternatives (prohibiting cameras altogether, or allowing them only if footage is deleted after a very short time), we need to assess how the decision impacts on each of these values, and need to provide, on the basis of such an assessment, a comprehensive evaluation.

According to the procedure that is usually suggested by decision theory, making the evaluations we have just described requires a mathematical characterisation of both

- the information on the basis of which a plan is to be evaluated; and
- the procedure that computes, on the basis of that information, the merit of the entire plan.

In the simplest case, this is done by

1. assigning a (positive or negative) weight to every relevant feature of the outcome;
2. quantifying the degree to which every feature will be satisfied by the expected outcome of a certain choice;
3. multiplying the degree of satisfaction of each relevant feature by its weight; and
4. adding the results that are obtained in step (3).

Note that weights are negative for those features which impact negatively on the outcome: the higher the quantity of these features, the worse their outcome (all the rest being equal). So, for example, suppose I assign weight **4** to food, **2** to wines, **1** to service, and $-\mathbf{3}$ to price, and that I expect that restaurant r will score 3 on food, 2 on wine, 1 on service, and 2 on price (0 indicates average, so that 2 describes a fairly high price). The expected worth of a choice to go to r will then be

$$(3^*\mathbf{4}) + (2^*\mathbf{2}) + (1^*\mathbf{1}) + (2^*\mathbf{-3})) = 11$$

Similarly, suppose that a judge gives weights **3** to security and **2** to privacy, and expects that a situation where footage is recorded and stored for a year will satisfy security to level 3 and privacy to level -2. The expected worth of such a choice would then be

$$(3^*\mathbf{3}) + ((-2)^*\mathbf{2}) = 5$$

Such a numerical procedure appears intuitively correct, and even upon reflection it appear to be free from apparent flaws. The issue we need to address, then, is why humans rarely perform such numerical calculations, especially when taking important decisions: few people use arithmetic when choosing their partner, their house, or even their new car. We may conjecture that the reason for this apparent "irrationality" is that our natural (implicit) cognitive capacities include more-powerful unconscious mechanisms for evaluating plans. It is not clear at all how such mechanisms may work, but they are certainly there. This does not imply that explicit plan evaluation (and even the assignment of weights and numbers) is always useless, since our unconscious decision-making processes, though far better than any approach decision theorists have yet been able to provide, are far from infallible. We should rather say that explicit evaluation of plans (according to the indications of decision theory) should be used to check the intuitive results that are provided by our implicit cognition. It would be improper, in most cases, to use it on its own, as an independent procedure for decision-making.

8 Bounded Rationality and Teleology

Teleological reasoning represents the core of human problem-solving and provides the pivotal link between epistemic and practical reasoning: (1) practical reasoning provides epistemic reasoning with goals, (2) epistemic reasoning constructs and evaluates plans according to one's likings and beliefs, and (3) practical reasoning

endorses an intention to implement a sufficiently good plan (which must be one of the best among the constructed ones). However, due consideration should be given to the practicability of teleological reasoning: teleological reasoning requires an enormous amount of knowledge, which often is not available. Such knowledge is required not only to address the formidable problem of planning (constructing plans) but also to compare and evaluate the constructed plans. Optimal planning thus seems to exceed human cognitive powers in many contexts.

In fact, in order for there to be a guarantee that a decision-maker will choose the optimal plan, the decision-maker must succeed in both (a) constructing a set of candidate plans that includes the best possible one and (b) making the right choice among the constructed plans. In both regards, optimisation is often out of reach for a bounded decision-maker. Firstly, we cannot consider all possible strategies for achieving certain objectives, and so we may fail to construct the best strategy. For example, in planning an out-of-town dinner, I may fail to detect the restaurant that is better suited to my tastes, since I am not aware of its existence. Similarly, consider how a legislature may fail to see what the most effective solution to economic growth is, and so may adopt a wrong decision (for example, cutting taxes may trigger a recession and a huge deficit rather than favouring economic development, as expected), or how judges or legal scholars may fail to discover optimal solutions to the problems they are considering (for example, punishing certain crimes too harshly may impede rehabilitation and lead convicts to commit further, more serious, crimes).

Secondly, even when we have constructed the best plan (together with other candidate plans), we may not be able to realise that it is the best one, and so may choose an inferior option. Failure to rank the available options according to their merit may depend in particular on the following:

- we have very little knowledge of the factual consequences of many of our choices;
- we have very confused ideas about what ends should inspire our evaluations, and about their relative importance, in various possible situations.

This problem concerns individual psychology, but also the functioning of organisations. It frequently happens that "the connection between organisation activities and ultimate objectives is obscure, and these ultimate objectives are incompletely formulated, or there are internal conflicts and contradictions among the ultimate objectives, or the means selected to achieve them" (Simon 1965, 64). Obviously, such problems are particularly serious in political and legal decision-making, which should ideally take into consideration all valuable goals, namely, all values that are relevant to a community. This makes it very difficult to assess the rationality of decisions impacting on different values by way of a combined assessment of resulting gains or losses with regard to all relevant ends. Consider for example, how difficult it is to assess the rationality of decisions in issues of Internet law, where one has to balance such diverse values as privacy, freedom of information, individual liberty, democracy, economic growth, and technological and scientific development.

Various views have been expressed in this regard. Some authors seem to believe that we can understand and justify decision-making by moving beyond teleological