European University Institute

DEPARTMENT OF LAW

EUI Working Papers

LAW 2009/07
DEPARTMENT OF LAW

A SUFFICIENTIST APPROACH TO REASONABLENESS IN LEGAL DECISION-MAKING AND JUDICIAL REVIEW

Giovanni Sartor

EUROPEAN UNIVERSITY INSTITUTE, FLORENCE
DEPARTMENT OF LAW

A Sufficientist Approach to Reasonableness in Legal Decision-Making and Judicial Review

GIOVANNI SARTOR

EUI Working Paper **LAW** 2009/07

ISSN 1725-6739

Printed in Italy
European University Institute
Badia Fiesolana
I – 50014 San Domenico di Fiesole (FI)
Italy
www.eui.eu
cadmus.eui.eu

Author contact details

Giovanni Sartor
Law Department
European University Institute
Villa Schifanoia
Via Boccaccio, 121
50133 Florence
Italy

Email: giovanni.sartor@eui.eu

An edited version will be published in G. Bongiovanni, G. Sartor, and C. Valentini, eds., Reasonableness and Law, Springer, New York, N. Y., 2009. Supported by the EU projects ONE-LEX (Marie Curie Chair), ESTRELLA (IST-2004-027665), and ALIS (IST-2004-027968). Many thanks to Jerry Postema and Wojciech Sadursky for their comments and suggestions on an earlier draft of this paper.

TABLE OF CONTENTS

Abstract

I shall argue for a sufficientist understanding of reasonableness in legal decision-making: cognitive or moral optimality are not required for reasonableness; what needed is just that a determination—be it epistemic or practical—is sufficiently good (acceptable, or at least not unacceptable). Correspondingly, judicial review on the ground of unreasonableness requires more than mere suboptimality: it requires failure to achieve the reasonableness threshold.

To develop this idea, I shall first analyse the notions of rationality and reasonableness, examining the role they play in cognition. I shall then consider rationality in legal (and in particular legislative) decision-making, focusing on teleological reasoning. I shall consequently develop an idea of sufficientist reasonableness, by combining the idea of bounded rationality with the idea of deference, as required by institutional coordination in the legal process. Finally, I shall consider when a legislative determination can be considered irrational or unreasonable, and how this is related to the violation of constitutional requirements.

Keywords

Proportionality, balancing, rights, constitution

1. Introduction

I shall argue for a sufficientist understanding of reasonableness in legal decision-making: cognitive or moral optimality are not required for reasonableness; what needed is just that a determination—be it epistemic or practical—is sufficiently good (acceptable, or at least not unacceptable). Correspondingly, judicial review on the ground of unreasonableness requires more than mere suboptimality: it requires failure to achieve the reasonableness threshold.

To develop this idea, I shall first analyse the notions of rationality and reasonableness, examining the role they play in cognition. I shall then consider rationality in legal (and in particular legislative) decision-making, focusing on teleological reasoning. I shall consequently develop an idea of sufficientist reasonableness, by combining the idea of bounded rationality with the idea of deference, as required by institutional coordination in the legal process. Finally, I shall consider when a legislative determination can be considered irrational or unreasonable, and how this is related to the violation of constitutional requirements.

2. Reasonableness and Rationality

The concept of reasonableness is often understood as having a larger content (intension) and thus a smaller extension than the concept of rationality understood as cognitive optimality: in order to be qualified as reasonable, a practical determination would need to be both rational and moral. This makes reasonable practical determinations a subset of rational determinations, those qualified by morality (the *differentia specifica* of reasonableness within the genus of rationality), and also a subset of moral determinations, those qualified by rationality (the *differentia specifica* of reasonableness within the genus of morality).

However, a different characterisation is possible, based on sufficiency rather than on optimality: reasonableness pertains to determinations that are good enough though not necessarily optimal; reasonable choices need to "satisfice"; they are not required to maximise (on the notion of *satisficing*, see Simon 1983). This sufficientist understanding of reasonableness, combined with the idea that practical reasonableness requires both morality and rationality, entails that reasonable practical determinations need to be both rational enough and moral enough, as shown in Figure 1.

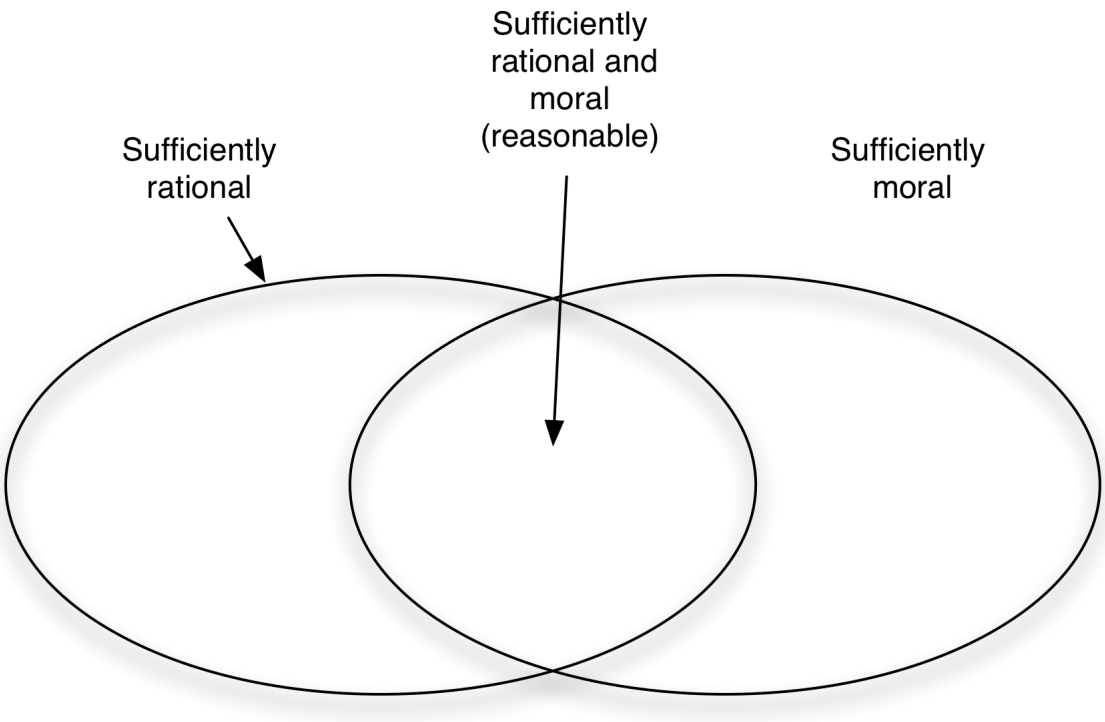

Figure 1: *Sufficientist reasonableness*

Figure 2 illustrates the connection among sufficientist practical reasonableness, rationality, and morality. The oval of the reasonable includes the practical optimum, namely, the set of determinations that are both optimally rational and optimally moral, but is not limited to such a set. It also includes determinations that, while failing to achieve optimal rationality or optimal morality (or both), still reach the sufficiency threshold in both respects. In this contribution, I shall not address moral reasonableness (which is discussed in other chapters of this book), but shall focus instead on cognitive reasonableness. With regard to moral reasonableness, shall just specify that I view morality—understood, in a general sense, as taking fairly into account the interests of others within one's practical reasoning—as a separate aspect of reasonableness. It seems to me that, in legal decision-making, namely, the activity whereby coercible decisions are taken in the name of the collectivity, morality entails the general requirement that everybody (each member of the polity) be treated with equal consideration and respect (on equality, see Sadurski 2008). This is much a stronger requirement than the idea involved in reasonableness in private law, where the moral dimension of reasonableness requires taking into account other people's interests to some extent, but does not require equating them with one's own interests (for a discussion of reasonableness in private law, see Ripstein 2001). In particular, this means, with regard to rights, that the satisfaction of each right-holder's interest in exercising a right should equally be taken into account. However, even with regard to this moral dimension, no more than a sufficientist threshold needs to be respected, a threshold compatible with different understandings of the notion of equal concern and respect.

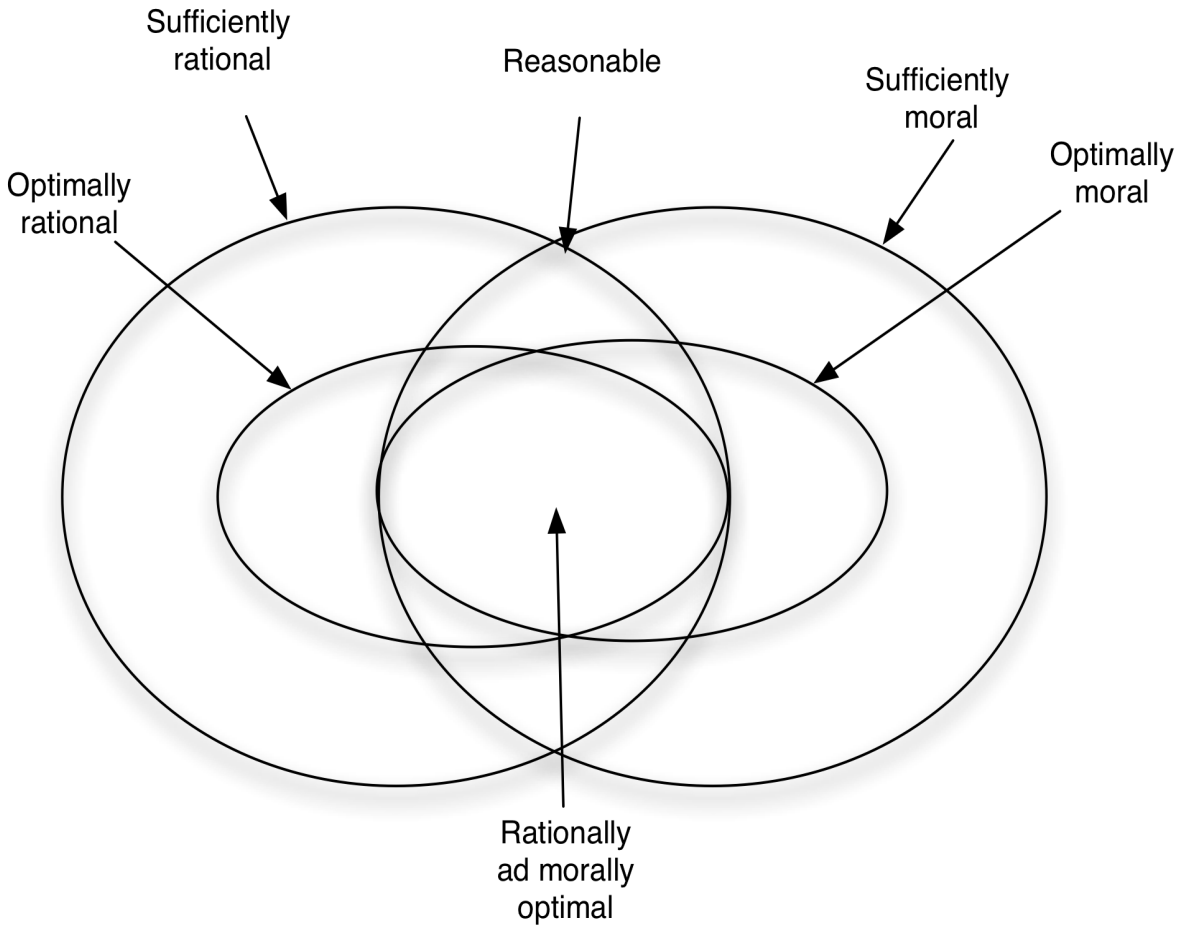

Figure 2: *Optimality and sufficiency in rationality and in morality*

A third aspect (besides rationality and morality) is often included in the notion of practical reasonableness: in order for a determination to be reasonable with regard to a certain context (culture or form of life), it must also be *consonant* (or at least not completely dissonant) with the ideas prevailing in that context (for a development of this idea in the legal domain, see Aarnio 1987), and in particular, with the norms that are practiced in that context. This culturally dependent idea of reasonableness must be distinguished from the trivial assertion that the beliefs of a person about what is reasonable (sufficiently rational and moral) may be influenced by the surrounding culture. The requirement of consonance does not concern what is (possibly mistakenly) believed to be reasonable: it concerns what is reasonable in a certain context. Such a requirement is violated when between a legal determination and general opinion there is a distance that cannot be overcome with the cognitive resources available to people. Consonance with general opinion may entail a certain conservatism, but it corresponds to the idea that legal decisions should be taken in name of the people, namely, of the legal community: though a legal decision-maker may take his decision on the basis of views that are shared only by a part of his community, and may and should rationalise and revise such views when opportune (correcting biases and prejudices, taking into account relevant scientific knowledge, etc.), a certain proximity should be retained between the law and the opinions of the community it is supposed to govern. By adding this further component, the concept of reasonableness represented in Figure 3 is obtained, which results from the intersection of three requirements, sufficient rationality, sufficient morality and sufficient consonance.

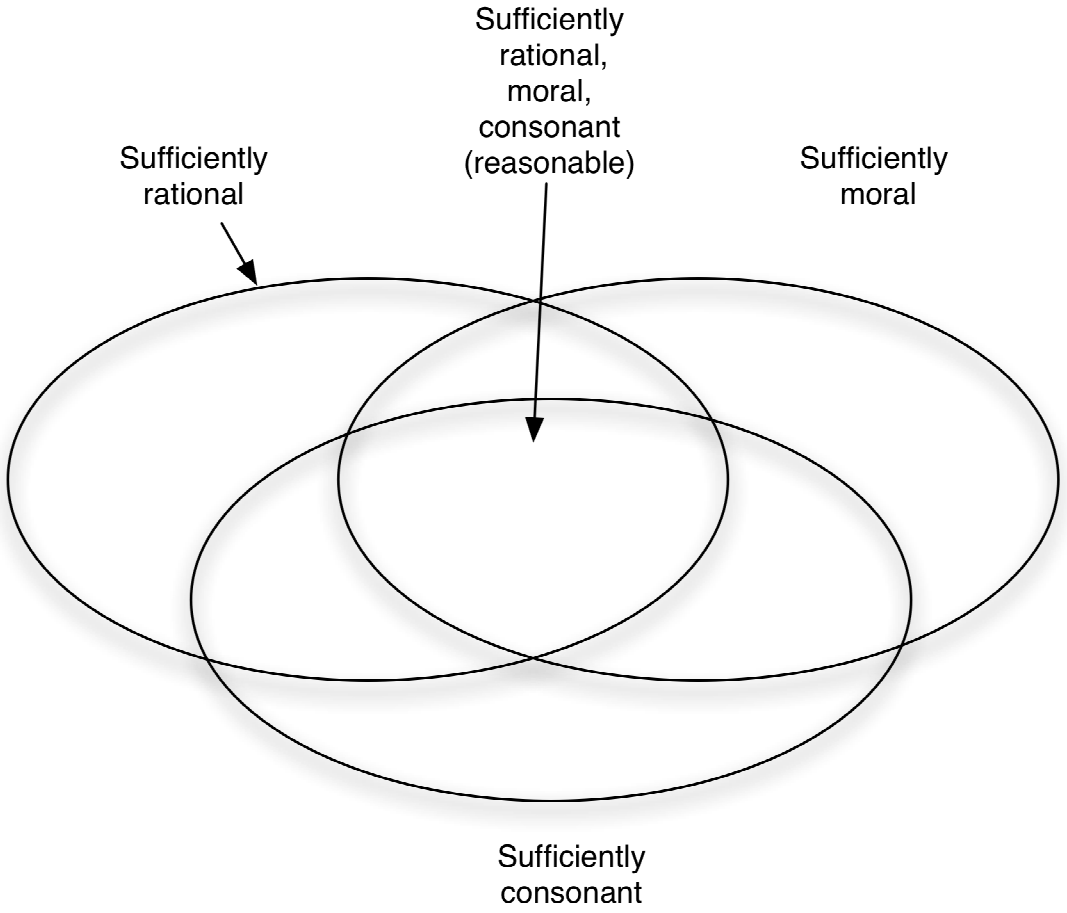

Figure 3: *Reasonableness as sufficient rationality, morality and persuasiveness*

3. The Process of Rational Problem-Solving in Individuals

Though reasonableness is to be distinguished from rationality, it needs to be characterised with reference to the latter. Thus, in the following I shall provide a summary analysis of rationality, based on the account in Pollock (1995), which expands the belief-desire-intention model often adopted in artificial intelligence (see Sartor 2005). This analysis refers to individual action, but its patterns, as I shall argue, also apply to rational decisions of collective bodies, and in particular to legislative decisions.

Rationality pertains to cognition, namely, to the activity through which we process information in order to come at reasoned determinations. These determinations can be epistemic, i.e., meant to identify the features of the world surrounding us, or practical, i.e., meant to establish the goals to be pursued, the plans of action to be implemented, or the norms to be endorsed. Epistemic cognition consists in the appropriate formation of epistemic states, namely, internal (mental) states meant to model features of the agent's environment. Two types of epistemic states can be distinguished, percepts and beliefs. *Percepts* come about by way of mechanisms of perception, which are activated when inputs are provided to the agent's sensors (on the distinct cognitive function of percepts and beliefs, see Pollock and Cruz 1999, 84ff.). A *belief* consists in the endorsement of a proposition: it consists in adopting a proposition as a premise in one's reasoning and action, as something one is

ready to reason and act upon. An agent endowed with the faculty of epistemic cognition will process external inputs and obtain percepts. The agent will then reason, producing new beliefs on the basis of previously formed percepts and beliefs. Hence, epistemic reasoning is the process through with one builds new epistemic states proceeding from the epistemic states one already possesses.

Practical cognition also consists in forming appropriate cognitive states, but these are conative rather than epistemic. Conative states are not intended to represent one's environment: their function is rather that of guiding one's deliberative process, of playing a role in the process determining behaviour. Thus an agent endowed with the faculty of practical cognition possesses conative states, and has the ability to form new conative states on the basis of his current epistemic and conative states. Practical reasoning is thus the process through with one builds new conative states on the basis of the epistemic and conative states one possesses.

Typically, in epistemic cognition, perception leads to beliefs about the environment. Such beliefs—when combined with other beliefs the agent already possesses, and when they go into a logical process (such as deduction, induction, defeasible reasoning, and probability calculus)—lead to the formation of further beliefs. The set of beliefs one possesses, however, does not consist mostly or even mainly of beliefs one constructs, directly or indirectly, out of one's own perceptions, according to some pattern of logical reasoning. It also includes inputs provided by the built-in cognitive modules of which our mind consists, modules for language, for building three-dimensional representations out of perception, for recognising faces, for making analogies, for inventing hypotheses and theories, for detecting other people's attitudes and the rules they are following, etc. (for different views on the "modularity of mind," see Fodor 1983 and Pinker 1999). These modules do not proceed in a "ratiocinative" way: they do not follow reasoning patterns of which we are aware and which we can monitor and direct. However, they provide sufficiently reliable inputs, inputs our conscious reason needs to accept (or at least to take into serious consideration). Moreover, in a social context, our beliefs mostly include information obtained from others, and through various social processes of cognition (see Goldman 1999; 2006). The rationality of a person's determination to accept the outcomes of a social source is based on the evaluation of the reliability of such a source, an evaluation that should take into account the cognitive limitations of the evaluating person. So it would be irrational for a layperson to refuse to adhere to the outcome of the scientific community, given the impossibility of autonomously forming a judgement on a scientific matter, or given that this judgement would much more likely be wrong than the judgement of the scientific community (unless one has evidence that the formation of the judgement of the scientific community was altered by disturbing factors, such as commercial or political interests, biases, prejudice, etc.).

It is often assumed that in practical cognition, an agent's goals are directly connected with the performance of actions believed by the agent to be necessary in achieving the same goals, thus making it so that actions, rather than mental states, are the conclusions of practical reasoning, as Aristotle says with regard to a syllogism described in *De motu animalium*, (Aristotle 1921, 7). This model, however, is not really adequate for a bounded reasoner, who cannot perform all reasoning at the time of acting, and needs to store the outcomes of his practical reasoning in appropriate ways. Thus, a more articulate view of bounded practical rationality is required, where a key role is played by intentions, which store practical determinations for the purpose of guiding future action. Here I shall adopt the model found in Pollock (1995), who distinguishes three basic conative states.

The first conative attitude consists in having *likings* or *preferences* (I use these expressions as synonymous). By a liking (or a disliking), I mean a generic *pro* (or *con*) attitude with regard to certain situations (present or future) or with regard to certain features of them (Pollock 1995, 12ff.). Consider, for instance, one's dissatisfaction with a job that does not satisfy one's tastes and ambitions, as compared to the (imagined) satisfaction in having a different job.

The second conative attitude consists in having *goals* (the state of having goal is usually called a *desire*).[1] A goal is more specific and focused than a liking. An agent's adoption of a goal has the function of prompting the agent to make plans to achieve that goal. Not every liking gives rise to a goal (for instance, my liking for the idea of being a football player or of winning a Nobel prize will not become my goal, since I know from the start that these outcomes I cannot achieve): only a liking for an objective one views as achievable can give rise to a goal. Consider, for instance, how a dissatisfied worker can adopt the goal of finding a new job.

The third conative attitude consists in having *intentions*. An intention is the state of mind of an agent who has determined that a certain action is to be performed, if need be under certain conditions.[2] When we have formed an intention to perform an action (or a combination of actions), we have made up our minds as to whether we will perform that action, under the appropriate conditions, and are ready to perform the action as soon as these conditions are satisfied. Adopting intentions is our way of storing the plans of action (combinations of instructions) we have adopted, and of staying ready to carry them out: instructions contained in chosen plans provide content to our intentions. Thus, an agent who has formed an intention is *committed* to implementing the corresponding instruction, since this is required by the way in which practical rationality works: we may rationally withdraw our intention, but it would be irrational for us not to implement our intention all the while having that intention. For example, if we have the goal of being hired for a new academic job, we may adopt the intention of applying for a position with certain universities. We may withdraw this intention (if, for instance, we are offered the job we desired, in which case we would interrupt our job search), but it would be irrational for us both to retain the intention and not implement it.[3]

The cognitive states just described are integrated into a reasoning process, which can be called *ratiocination*, where the adoption of certain cognitive states is the reason that leads to, and justifies, the adoption of further cognitive states. In epistemic reasoning:

- having a perception with content *P* is a reason for believing *P*;
- having certain beliefs $P_1,\ldots P_n$ is a reason for believing a certain conclusion *P* linked to the input beliefs according to the various patterns of correct epistemic reasoning (deductive inference, defeasible inference, probabilistic inference, induction, etc.).

In practical reasoning

- liking a state of affairs *P*, and believing that it can be achieved, is a reason for adopting the goal of achieving *P*;
- having goal *P*, and believing that a plan (action) *A* is a teleologically appropriate way to achieve it, is a reason for intending to execute *A*.

By arranging these reasoning schemata into sequences, we obtain downward chains: in epistemic reasoning, from perception to beliefs, and from beliefs to further beliefs; in practical reasoning, from

---

1 A desire may be distinguished from the mere wishing, the latter being a positive attitude toward a state of affairs, regardless of the possibility of achieving it (see Aristotle 1924, III). A desire in the sense here indicated presupposes that the agent does not consider impossible the achievement of the object of desire.

2 I use here the term *intention* in the usual, commonsense meaning (which corresponds to the way in which this term is used in artificial intelligence or practical philosophy: see Bratman 1987). I am not using it to express the "aboutness," or direction toward an object, of a mental state or content of consciousness, as it is sometimes used in the philosophy of mind (on the approach of Brentano 1973).

3 According to Pollock (1995), there is a fourth conative attitude, which he calls a want. A want is an unconditioned impulse toward performing an action, an impulse one feels when one has the intention of performing a certain action given certain conditions (including a certain temporal and spatial framework) and given that one believes that the conditions for the immediate performance of the action obtain. For simplicity's sake, I shall not include the idea of a want in the present model, and shall limit myself to the idea of a unconditioned intention to perform a specific action, without considering the further process leading to action.

likings and beliefs to goals, and from goals and beliefs to intentions (leading to action). These chains may be extended with intermediate steps, where sequences in a chain are recursively activated: the intention to perform an abstract action (e.g., going to a conference) leads to the goal of performing that action (an instrumental goal), and this goal leads to the intention of executing a corresponding plan (registering for the conference, buying a ticket, etc.); the belief that the antecedent of a conditioned intention is satisfied leads to an unconditional intention, etc.

A chain of practical reasonings ends with the agent's intention to unconditionally execute a specific action. If all reasoning steps have functioned properly, we will act in such a way as to be successful, at least in most cases (failure is always possible, in a complex and unpredictable world, even when we have used our cognitive functions as well as we can): by implementing our intentions, we achieve our goals; by achieving our goals, we satisfy our desires; and by satisfying our desires, we adapt the world to our likings. Thus, practical ratiocination enables a reasoner to reach the target of higher-level conative states by achieving the target of lower-level conative states. The fact that a single reasoning step is correct does not mean that the whole sequence, including all previous steps, is correct: a mistake can have taken place above, so that rationally correct inferences can be included in chains of reasoning that, as a whole, are irrational.

It is important to stress that rational problem-solving requires a connection between epistemic and practical reasoning (see Figure 4).

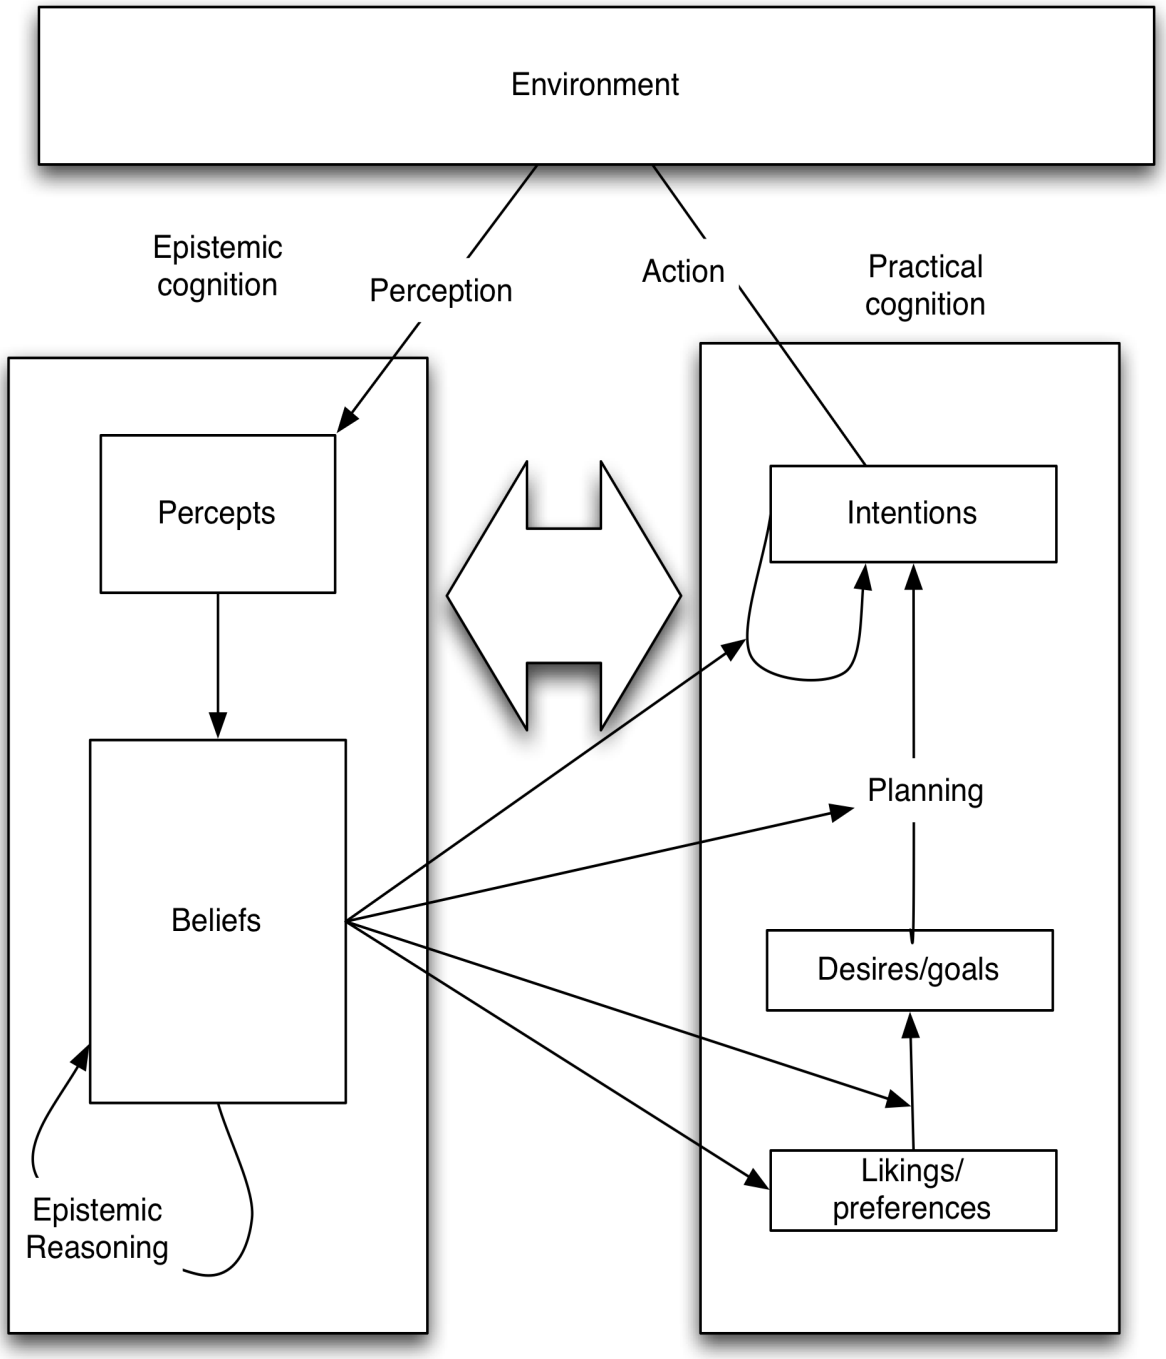

Figure 4: *Rationality = Practical cognition + Epistemic Cognition*

Advances in practical reasoning are dependent upon the results provided by epistemic reasoning: we adopt a goal assuming that it is achievable and we adopt an intention to execute a plan on the basis of the belief that this that plan provides a teleologically appropriate way to achieve its goal (it would be irrational to have a goal to achieve something that we believe to be unachievable, or to have an intention to do something while believing that doing something else would be a better way to achieve

our goals, consistently with our interests and constraints). Thus failures in practical rationality may depend (and do most often depend) on a failure in the epistemic processes providing inputs to practical reasoning. On the other hand, epistemic inquiries can be activated by practical reasoning (they depend on the reasoner's practical interests): when we need new information in order to develop a plan of action (I have the goal of going to a conference, but I do not yet know what means of transportation are available), or in order to establish whether a precondition for an adopted intention holds (I do not know whether I am scheduled for a lecture today, but it is my intention to lecture if I am scheduled), we engage in epistemic-reasoning and auxiliary information-seeking actions aimed at coming to know what we need to know. This does not imply that all epistemic interests are only instrumental to specific "material" goals: humans have an inborn conative disposition for knowing (curiosity), regardless of further specific uses knowledge, and a corresponding liking for the discovery and possession of knowledge, a drive that supports scientific research as well as gossip. As Aristotle observed in *Metaphysics*, 980a, "all men by nature desire to know" (Aristotle 1921).

Rational thinking and decision-making under complex circumstances are both defeasible and argumentative. Rationality is defeasible in the sense that many of the inferences it validates consist in inferring conclusions on the basis of limited information, conclusions that may need to be revised when additional information is provided. For instance, the perceptual conclusion that an object is pink (having been perceived in this way) should be abandoned when we become aware that the object was viewed under a red light; the conclusion that Tweety is white, being a swan, should be withdrawn when we come to know that Tweety comes from Australia, where swans are black; an explanation of certain aspects of a phenomenon should be abandoned when a more comprehensive explanation is available; etc. Similarly, in the practical domain a goal should be abandoned when new evidence shows that it cannot be achieved; the intention to perform a specific action in keeping with a general intention (e.g., an intention to do exercise today as a consequence of an intention to do exercise every morning) should be abandoned when it conflicts with a prevailing intention (catching an airplane very early in the morning); the belief that a plan is the most appropriate way to achieve a goal should be abandoned when further evidence shows that a better option is available (there is a train that takes me to my destination more cheaply and comfortably than the airplane I had decided to take); etc.

The argumentative nature of rational thinking is related to its defeasibility. If epistemic and practical reasoning is defeasible, then justifying a conclusion will require more than considering the reasons supporting that conclusion: it is also necessary to consider the reasons against it. As Mill (1991, 41) puts it, in "any subject on which difference of opinion is possible, the truth depends on a balance to be struck between two sets of conflicting reasons," so that the use of *positive logic*, which relates a thesis to its supporting grounds, must be supplemented with critical discussion of the opinion to the contrary, that is, by that *negative logic* which "points out weaknesses in theory or errors in practice, without establishing positive truths" (ibid, 109). Recent research on defeasible reasoning has identified for different patterns or schemes of defeasible argument the corresponding "critical questions," pointing to information that (if it were to be accepted) would defeat the argument in question (see, for all of this research, Walton et al. 2008).

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 cannot be rational 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 fallibility 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, chaps. 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 reasoner 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), 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

[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 one adopt one such plan by random choice (rather than being paralysed by an inability to form a preference).

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 Bench-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. 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 or her 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 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

---

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

$EU(\alpha) = \Sigma^{n}_{1} (Pi * Ui)$.

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 **-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 rationality and focusing instead on *systemic rationality*: we should look at how certain forms and styles of decision-making contribute to the functioning of the social systems in which they take place, regardless of how effective they are in achieving the goals pursued by the decision-makers (this is the view famously advanced in Luhmann 1973). Others, such as Habermas (1999, 259), have rejected the idea that rational decision-making involves assessing and comparing impacts on relevant values, affirming that "weighing takes place either arbitrarily or unreflectingly, according to customary standards and hierarchies." Unfortunately, in many cases human reasoners have no alternatives to teleological reasoning. We thus need to overcome such dismissive views, all the while seriously addressing the issue of the limitations of teleological reasoning, especially with regard to the problem of the evaluation of plans.

9. The Evaluation of Outcomes

In evaluating a plan we need to identify its outcome (the results it will produce), distinguish the relevant desirable or undesirable features involved in that outcome, consider to what degree these features are advanced or impaired by the plan, and finally establish what merit is to be accordingly attributed to the plan. For instance, when considering a plan to go to a restaurant r, we may consider to what degree we expect a dinner at r to exemplify the following features: quality of the food, quality of the wines, quality of the service, and price. We would then move from the level of each feature's realisation to the evaluation of their combination. The most delicate step is the last one, namely, moving from single features to their joint evaluation. Analytical reasoning (ratiocination) is in general not very effective at capturing and assessing interconnected sets of features: this is a task we seem to accomplish through a kind of holistic understanding, similar to pattern recognition in perception (as argued in Thagard 2000). However, under certain conditions, analysis may help. The analytical evaluation of multi-featured decisional outcomes is facilitated when all of the following conditions hold:

- there is a numerical measure for the realisation of each relevant feature;
- the worth of realising each feature grows linearly (always in the same proportion or weight) according to the measure of the realisation of this feature; and
- the desired features are all mutually independent.

Under these conditions, we can evaluate an outcome simply by multiplying the measure of the realisation of each feature in that outcome by the weight of the feature, and then adding up the results.[6] This provides an easy way to compute the worth of plans of action, and so an easy way to compare them.

For instance, suppose I assign weights **3** to food, **2** to wines, **1** to service, and **-2** to price. Moreover, I expect that in restaurant r_1 food will score 5, wines 2, service 1, and price 6 (€60), while in restaurant r_2 food will score 3, wines 2, service 1, and price 2 (€20). This allows me to assign points 6 to restaurant r_1 and points 10 to r_2, as Table 1 shows.

[6] The evaluation EV of outcome ω—where ω realises features $f_1, \ldots, f_n$, to the degrees $d_1, \ldots, d_n$, and each feature f_i has weight w_i—is expressed by the simple formula $EU(\omega) = \Sigma_{i=1}^{n}(d_i * w_i)$.

	Food	Wine	Service	Price	Total
r_1	**4***3	**2***2	**2***1	**6***(-2)	6
r_2	**3***3	2*2	1*1	**1***(-2)	10

Table 1: *The expected worth of two restaurant experiences*

Thus, if the input data are correct, and if all assumptions above hold, rationality requires me to go to the cheaper restaurant r_2, even though the food and service it offers are lower-quality than at r_1.

This example shows how this way of evaluating choices gives questionable results, when no precise and objective way is available for quantifying degrees of satisfaction and weights, or when different features interfere. One may rightly challenge my choice for restaurant r_2 (or in any event my procedure for reaching that choice) by pointing out the arbitrariness of assigning degrees of satisfaction and weights—How do I know that food quality at r_1 is 4, and not 5 or 3? How do I know that food quality has weight **3**, and not **4** or **2**?—or by pointing out that these evaluations are contingent on particular individual circumstances (for example, the benefit afforded by the quality of service depends on how much time I have and on how irritable I am on that particular day), or by pointing out their interdependence (having bad service would likely spoil my enjoyment of the food).

However, it may still be possible to compare the degrees of satisfaction of a certain desired feature in different situations. By combining this assumption with the idea of monotonic growth of desires as the degree of the desired feature grows, we obtain some clues on how to develop our preferences.

10. Pareto Superiority

If outcome δ, in comparison to outcome γ, presents at least one feature at a more desirable degree and no feature at a less desirable degree, then it seems that a rational agent should prefer δ to γ (and thus, that it would be irrational for this agent to prefer γ to δ).

Let us state this idea more precisely. Let us write $x \phi y$ to mean that x is strictly preferable to y. Thus, when x and y are degrees of an advantageous feature (a feature that is preferable in a higher degree, like the quality of the food), then $x \phi y$ (x is preferable to y) whenever $x > y$ (x is greater than y). On the contrary, when x and y are degrees of a disadvantageous feature (like the price of the food), namely, a feature that is preferable in a lesser degree, then $x \phi y$ whenever $x < y$ (x is lesser than y). Accordingly, we can say that a rational agent should prefer outcome δ to outcome γ under the following conditions:

- δ and γ share the same desirable features $f_1,\ldots,f_n$, and
- δ presents these features to degrees $d_1,\ldots,d_n$ and γ to degrees $g_1,\ldots,g_n$.
- there is an i such that $d_i \phi g_i$, while for no j, $g_j \phi d_j$.

For instance, if a restaurant r_1 offers better food than a restaurant r_2, and equal or better wines and service, then a rational agent should prefer r_1 to r_2 ($r_1 \phi r_2$).

This is the idea of *Pareto superiority* applied to impacts on different goals or values: when a choice is required between alternative decisions δ and γ, and the two decisions impact on the same values (desirable features), but δ raises some of these values to a higher degree than γ does, the impact on all other values being equal, then δ is to be preferred (for a technical account of multicriteria decision-making, see Keeney and Raiffa 1993). Consider for example, a choice between these two ways of treating videos recorded by cameras located in a street having a high crime rate:

- deleting the data after a week; or
- keeping data for a year.

Let us assume that deleting data after a week enables the value of privacy to be achieved to a higher degree, while there is no difference with respect to the attainment of the value of security (one week being sufficient to check video recordings in connection with serious claims). Under such conditions, a decision to keep the video footage for a year would be irrational according to the idea of Pareto superiority (assuming the only relevant values are privacy and security): an alternative choice would provide a higher achievement of some values without diminishing the level of achievement of any other value.

The idea of Pareto superiority is a useful minimum standard for evaluating decisions. For instance, it seems to subsume some of the standards of *reasonableness* that are used in constitutional and administrative review. For instance, a decision that would undermine certain values without contributing to the realisation of any other value is certainly inefficient, and in this sense it is more than just unreasonable: it is irrational. Similarly, a choice (a) would be irrational on grounds of its Pareto inferiority when it achieves certain values by undermining certain other values, and there is an alternative choice (b) that would realise the same values to the same extent without undermining any other values (or would do so undermining them to a lesser degree): choice (a) would be irrational since it determines a prejudice which is unnecessary to achieve its beneficial outcomes.

The condemnation of Pareto-inferior choices needs to be attenuated with some sufficientist considerations. It would certainly be irrational to choose α rather than β, knowing that α is Pareto-inferior to β. However, we may not know that α is inferior to β, since we may be proceeding on a mistaken appreciation of the impact of the two options on the values at issue. For instance, we may fail to recognise Pareto inferiority owing to our inability to take into account certain complex causal connections (as may happen with regard to choices pertaining to economic policy). When an epistemic mistake remains below the threshold of unreasonableness the resulting choice remains reasonable, though it may rightly appear Pareto-inferior to an observer immune to the mistake.

11. Weighing Alternatives

The idea of Pareto superiority does not help us make choices in those situations where there are alternatives choices α and β, such that α advances certain values more than β, and β advances certain other values more than α. Consider, for example, the problem of making a choice between two restaurants r_1 and r_2, such that r_1 offers better food, r_2 provides better service, and all other relevant features are satisfied to the same degrees. Similarly, consider the problem of choosing whether to delete video footage after one day or after seven days of recording, assuming that deletion after one day yields a higher level of privacy while deletion after seven days provides a higher level of security.

It may be said that such issues can be solved through a comparative analysis that takes into account (a) the degree to which the values are satisfied by different choices and (b) the importance of the values. Choice α is better than choice β if the comparative gains relative to the values that are better promoted by α outweigh the comparative loss relative to the values that are better promoted by β. For instance, rationality requires me to go to restaurant r_1 instead of restaurant r_2 if the comparative advantage in food quality outweighs the disadvantage in service quality. Similarly, it may be said that rationality requires deleting footage after seven days rather than after one day if the gain in security outweighs the loss in privacy.

The difficulty, however, consists in finding a sufficiently precise characterisation of how one should rationally "weigh" such alternatives. It may be said that the weighing judgement depends both on the quantities that are gained or lost and on the importance of what is gained and lost, but this offers very little help to the decision-maker, for whom the problem is exactly that of establishing quantities and relative importance.

Moreover, the importance of a gain or a loss relative to a certain value does not only depends on the "quantity" of the value that is gained or lost (whatever is meant by *quantity*) and on the importance of the value at issue: it also depends on the level from which we measure a gain or loss. With many values, there are two distinct aspects to be considered: the quantitative measure of the value and the benefit (the impact on human wellbeing) of achieving that value up to that quantitative measure. For instance, when buying a flat, a valuable aspect to be considered is the spaciousness of the flat, as measured by the surface area of the flat, but even though the benefit of having a flat of a certain surface increases in proportion as surface does (this benefit being a monotonic function of the surface), this proportion is not fixed: while the additional benefit obtained by moving from 20 to 40 square meters is usually very important, the additional benefit obtained by moving from 200 to 220 is likely to be less significant.

Let us consider now a more significant value—nutrition, the object of the right to food, much discussed nowadays by human-rights scholars—and let us assume that any intake below 2,000 is insufficient to sustain human health. A 1,000-calorie drop below that minimum would thus constitute a very significant failure of nutrition (it would lead to starvation, and probably to death), whereas adding 1,000 calories on top of an already more-than-sufficient intake of 3,000 calories would not bring any additional benefit as far as nutrition is concerned.

A similar analysis would also apply to less-quantifiable values, such as liberty. Having a wider (more inclusive) set of options increases one's liberty. Suppose that $A_1 \subset A_2$, A_1 contains 100 options, while A_2 contains 110, and the additional 10 options in A_2 have the same average significance as the options in A_1. We will then certainly be able to say that having choices A_2, rather than A_1 , is significantly more beneficial as far as liberty is concerned. However, consider two sets of choices A_3 and A_4, such that $A_3 \subset A_4$ and A_3 contains 1,000,000 options while A_2 contains 1,000,010, and the additional 10 options in A_4 have the same average significance as the options in A_3 . In this case, the difference in the benefit provided by having A_4 rather than A_3 would be much less important, probably quite imperceptible. Finally, suppose that one has the possibility of choosing from a range of only 20 options (e.g., kinds of jobs one may aspire to), and that a piece of legislation reduces this range to 10 by eliminating 10 options previously available. This would certainly be a very serious inference in the core of one's freedom of employment.

In Figure 5 you can see the difference between a value providing a benefit that increases linearly (in a fixed proportion) relative to increases in the associated quantitative measure, and a value providing a benefit that increases nonlinearly, bringing a diminishing marginal benefit (as represented by the curve). As Alexy (2002b, 103) observes, the latter pattern characterises not only economic goods (whose marginal utility usually diminishes, i.e., an additional unit of a good G brings less benefit when one has a larger quantity of G) but also legal values.

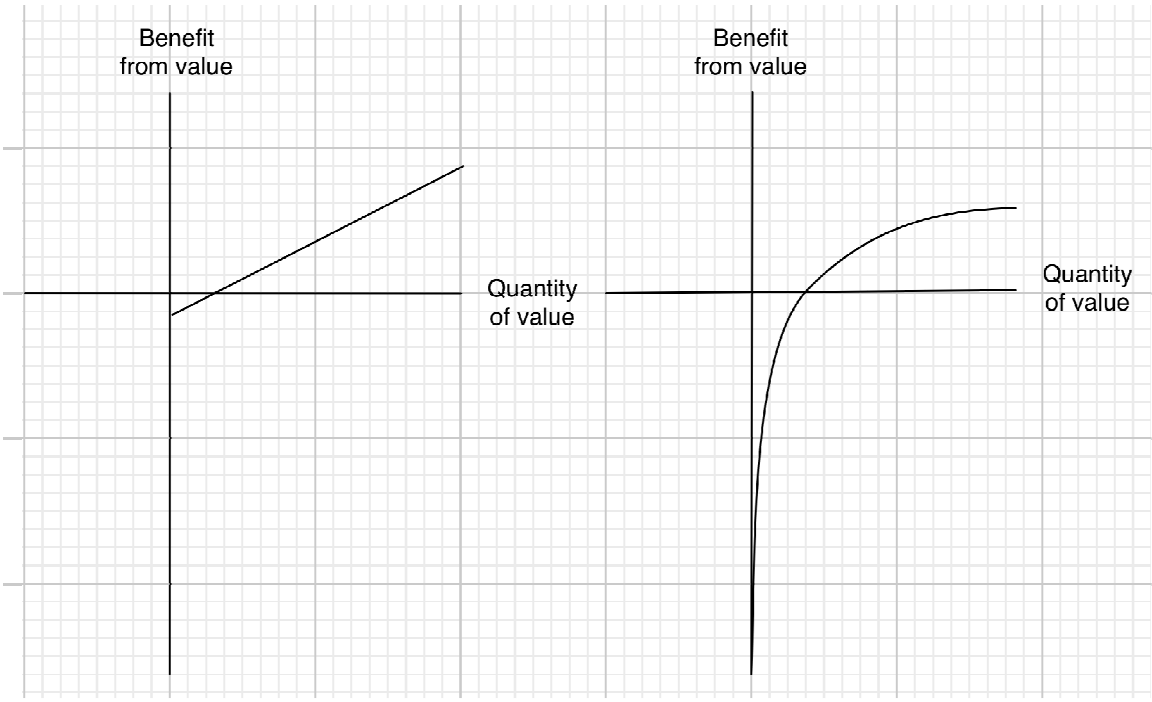

Figure 5: *Quantitative measure and benefit: linear (left) and non linear case (right)*

The difficulties I have just considered should lead us to be wary and critical (maybe even sceptical) of any pretence to "objective" or "scientific" evaluations of decisional alternatives. However, they should not lead us to conclude that the rational comparison of alternatives is impossible, that the tools of decision theory are useless, or that every choice puts us in front of the incommensurable or the "absurd," as existentialists used to say, and that it calls for (or at least presupposes) an arbitrary commitment (cf. Sartre 1993). In fact, we need to approach the problem of weighing and balancing on the basis of our awareness, not only of our failures, but also of our cognitive powers and, in particular, of the power of our implicit cognition. As a matter of fact, we know how to take many decisions impacting on different goals and values, and we can approach this task in a way that, though far from perfect, is sufficiently good for most of our purposes. Humans seem to possess an adequate cognitive faculty to evaluate and compare alternative plans of action under conditions of uncertainty, though we cannot tell precisely how this faculty works, nor can we fully replicate its functioning through explicit reasoning. It is no accident that even decision theorists, when they have to make choices involving multiple aspects in complex domains (choosing a partner or a profession or buying a house or even a car), rely more on their intuitive judgement than on the conceptual tools provided by their discipline. So the fact that our comparative evaluations usually involve an unconscious process does not imply that they are random or absurd. On the contrary, our implicit cognition usually also takes into account data that are explicitly provided, and it processes these data unconsciously, along with information which remains implicit, and which we do not have the ability or the time to express and consciously evaluate. For instance, my choice to buy a certain car may be influenced by the information I find in automobile magazines, or by suggestions by friends, along with various other things I know, though I can only partially articulate this data.

Moreover, we can articulate this implicit knowledge at least in part (and in important cases we *should* articulate it) at the stage where we are critically analysing our choices, trying to rationalise them, as being based on good grounds. For instance, before making a check out to the car dealer, I may try to consider whether my intuitive preference is really based upon relevant grounds, by explicitly listing the pros and cons of a certain particular choice, as compared to the available alternatives. I should stop only at the stage where I have found equilibrium between intuition and reason, that is, when intuitive assessment and explicit reasoning converge on the same result.

The situation is no different in legal decision-making. Consider, for instance, a judge who has to decide, in the absence of a precise rule, whether a certain way of processing data taken through street cameras is permissible, or whether an employer is allowed to have access to e-mails an employee sent and received using the account provided by the company. One can take a stand on such issues only with reference, on the one hand, to the individual and collective legal values at stake (individual security, individual privacy, the efficiency of the economic system, and so forth) and, on the other hand, to the technological and social knowledge concerning the ways in which different arrangements are going to impact on these values. All of this knowledge is brought to bear, usually unconsciously, on one's evaluation as to whether a certain policy or choice unduly undermines the value of privacy as compared to what other possible policies or arrangements might do. This intuitive judgement, however, needs to be rationalised by articulating the grounds for it, and this rationalisation should lead the decision-maker to find an equilibrium between intuition and reasoned assessment.

We can thus say that, in balancing the benefits obtained by achieving different values, explicit reasoning should process and rationalise the outcomes of intuition (implicit cognition) rather than substituting it. Therefore, we can draw two indications. The first is that cognition and rationality can (and should) also govern comparative evaluations. The second is that, in most contexts, the quantitative methods proposed by decision theory should be used to check intuitive choices, analyse their compatibility with similar choices, and fit them into a background theoretical framework rather than to provide a self-sufficient alternative to human intuition (to implicit cognition).

12. Simplifying Evaluations

For some purposes, and under certain conditions, some simplifications may help by providing workable ways of evaluating legal decisions. For example, Alexy (2002a) proposes a method of numerically characterising the impacts on relevant values. He observes that the German Constitutional Court frequently justifies its judgments on the legitimacy of certain laws by examining the impacts these laws have on legal values, and by characterising these impacts as *light*, *medium*, or *serious*. Correspondingly, he recommends that we should qualify legal values according to their low, medium, or high importance, and that we link this qualification to numerical weights (for instance, **1** for values having low importance, **2** for medium importance, and **4** for high importance); we should also qualify gains and losses in the achievement of legal values (gains being positive and losses being negative) as light, moderate, or serious, and should link such qualifications to simple numerical quantities (for instance, **1** for light, **2** for medium, and **4** for serious).

A much rougher simplification than that proposed by Alexy can be had by assuming that that there is a *lexicographic order* of values: values (or sets of them) can be listed in order of importance, and no gain, however big, in a lower-ranked value can outweigh a loss, however small, in a higher-ranked one. For example, it is frequently said that personality rights always prevail upon economic rights. Clearly, such an extreme view cannot be taken literally, or rather, it can only be maintained by sensible persons at the price of hypocrisy and self-deceit (masquerading economic rights as personality rights whenever one feels that they ought to prevail, and vice versa). For instance, it seems undeniable that the modest loss in privacy involved in a vendor's practice of keeping a record of the sale, for a limited time, is outweighed by the possibility of monitoring the performance of the contract.

On the other hand, the idea that personality rights normally (defeasibly) prevail upon economic rights, and that this defeasible presumption can only be defeated when there is a clear and specific reason to the contrary, seems quite sensible.

Lawyers may use further techniques to simplify teleological evaluation. First, we may limit ourselves to considering "the types of decision which should have to be given in hypothetical cases which might occur and which would come within the terms of the ruling," and one may evaluate those consequences by asking about "the acceptability of such consequences" (MacCormick 1978). This means that, rather than examining the social and economic consequences produced by decisions taken in keeping with certain rules, we could simply consider what legal decisions would be taken in certain classes of cases if certain rules were adopted (for a discussion of this idea, see also Luhmann 1974, chap. 4, sec. 5).

A different way of simplifying teleological reasoning is provided by reasoning *per-absurdum*. This consists in focusing on just one negative implication of a certain choice, and on the values that are undermined through such an implication. This very crude way of cutting away at the complexity of teleological reasoning is appropriate when a single consequence of a decision is so detrimental that it will very unlikely be outweighed by any advantageous impacts of the same decision. Often, a vivid impression of the negative impact of a certain choice can be had without even having to specify what values are going to be undermined.[7]

Some criteria of reasonableness used in constitutional or administrative review also give clues for detecting irrationality. For instance, a choice to allocate a certain advantage or burden to certain persons, while not allocating it to others who are in an equal situation (in all relevant respects), does not just violate the principle of equality: is it also an index of irrationality. In fact, where such a choice is concerned, there must usually be a better alternative, which may consist either in allocating the same advantage or burden to these other persons too (if this advantage or burden has a positive impact on the achievement of the values at issue) or in eliminating it completely for everybody (if it has a negative impact). A decision as to whether to extend or eliminate an advantage (or a burden) is a difficult one, but refusing to address it in a reasonable way may lead to very negative consequences. For instance, the Italian Constitutional Court has long been adopting, in the name of the principle of equality, a policy under which every benefit conferred on any person or category (as concerns salary or pension, for example) must also be conferred on every person in the same situation as those already enjoying the benefit. This policy had a very bad impact on public finances, and so the judges had to reconsider it: they now admit that a privilege may have to be eliminated rather than extended, and they usually prefer to stimulate a legislative adjustment rather than act directly.

Similarly, the fact that a choice completely disregards certain values is a strong index of its likely irrationality. Since the loss of benefit brought about by compressing a value increases more than proportionally when the value is achieved to a minor extent, it is very unlikely that the complete non-achievement of a certain value can be outweighed by an increase in the achievement of other values. For instance, an increase in security, though it may justify a reduction of freedom, cannot justify the complete elimination of freedom, and the less freedom we have, the higher the increase in security must be to justify additional restrictions on freedom.

---

[7] Here is how, in *Donoghue v. Stevenson*, Lord MacMillan refuted the thesis that producers owe no duty of care to customers: "Suppose that a baker through carelessness allows a large quantity of arsenic to be mixed with a batch of his bread, with the result that those who subsequently eat it are poisoned. Could he be heard to say that he owed no duty to the consumers of his bread to take care that it was free from poison, and that, as he did not know that any poison had fallen into it, his only liability was for the breach of warranty under his contract of sale to those who actually bought the poisoned bread from him?" (House of Lords [1932] A.C. at 620).

13. The Rationality of Legislative Choices

In what follows, the foregoing analysis of teleological reasoning shall be brought to bear upon legislative choices and their judicial review. As Waldron (1999) has observed, it is often too pessimistic a view that is taken of legislation, a view often coupled with an excessive optimism about judicial decision-making. In fact, jurists often develop normative theories of constitutional adjudication to suggest how judges should reason and act in order to remedy legislative mistakes: legal analyses thus tend to combine a realistic-empirical approach to legislation, focusing on the defective instances of legislation (where legislators are bow to special interests, are moved by prejudice and ideology, pursue particularistic or even sectarian goals, develop ineffective policies, fail to achieve the promised outcomes, etc.), with a normative analysis of adjudication suggesting how judges should remedy such mistakes and thus focusing on the best instances of adjudication (where judges succeeded in protecting fundamental rights or other constitutional values against wrong legislative choices). This approach often entails a reductive view of legislation, failing to take into account that, besides normative models of adjudication, there are also normative models of legislation, models to some extent adopted by the participants in the legislative process, and motivating their behaviour. Effective practice, both in legislation and in adjudication, cannot be reduced to the implementation of a normative model or to the dialectics of opportunistic interests: it instead requires the integration of both aspects, the tension towards a normative standard and the opposite (and often prevailing) pressures to depart from it. So, in analysing legislation, I shall first lay out some aspects of a normative model of legislative rationality, and shall then consider its possible failings.

Legislators (like judges and administrators) should not reason from their private perspective (pursuing their individual interest). When serving as members of a legislative body, they should instead act in the name and for the benefit (the common good) of the polity they are representing, and should make their choices integral to the decisional process of that polity. Thus, when evaluating the teleological rationality/reasonableness of the determinations adopted by a legislature, our reference point should not be the particular private objectives the individual members of a legislative body might pursue, but rather the political goals they adopt according to their vision of the public good, combined with the constitutional values the legislature has to take into account. With this proviso, legislative decision-making can be assimilated to the model of individual rationality I described above, which integrates epistemic and practical rationality (the role for epistemic rationality in legislation is stressed by the idea of evidence-based legislation, on which see, among others, Seidman and Seidman 2001). This assimilation has to be integrated, as I shall argue in what follows, by taking into account the plurality of institutional agents involved in the public decision-making process.

Legislators (supported by their staff, communicating with their constituencies, participating in political debate inside and outside the legislative body) need to first detect a problem-situation, namely, a social arrangement that appears to be unsatisfactory, expressing an unsatisfied social need that they think should be addressed. On the basis of an empirical analysis, they should identify more precisely the issue characterising that problem-situation and the social behaviour from which it emerges. This will enable them to establish what goal (values) should be pursued through legislation in that situation. For instance, let us consider a problem now being discussed by the Italian legislature: a very high number of private telephone communications are wiretapped under police investigations, and the content of such communications often winds up being published in the media, with serious prejudice to its author. A new law designed to deal with this problem-situation should aim to better protect individual privacy, a goal achieving which would in turn also be a way of protecting individual liberty.

Putting such a goal on the legislative agenda would start teleological reasoning, in order to draft a legislative measure protecting privacy with regard to private communications wiretapped under crime investigations. For this purpose, an empirical analysis is required aimed at understanding how possible measures (plans) will impact on the values at stake: not only privacy, but also freedom of speech, freedom of the press, publicity (and the consequent public control) of judicial activities, repression and hence prevention of crimes, and limitation of the costs of judicial inquiries (by reducing wiretapping

costs). For instance, an absolute and unconditional prohibition against wiretapping in crime investigations would increase privacy protection, and would leave freedom of speech untouched, but would seriously limit the possibility of identifying the authors of many crimes, especially those carried out by organised crime rings. It would reduce to 0 the costs of wiretapping, though this may require different kinds of investigations, possibly more expensive ones. By contrast, an unconditional wiretapping authority conferred on every prosecutor in investigations concerning any kind of crime, coupled with an unlimited authority to distribute and publish the wiretapped conversations, would increase the likelihood of preventing crimes (assuming that prosecutors were able to devote their resources to the most effective investigations, on the basis of a correct cost-benefit analysis) and would emphasise freedom of the press.

Such considerations need to be based on empirical analyses that will take into account the complex social connections at issue. It is not sufficient to consider only law in the books; analysis has to extend to law in action. Legislators need to evaluate the probability that legal provisions are not followed, since penalties are nor enforced or fail to deter unwanted behaviour: will a fine imposed on officers and journalists succeed in deterring them from communicating and publishing wiretapped conversations? They must also consider the chance that the legal process is used for deterring legitimate actions: will journalists be deterred from publishing legitimate information concerning legal proceedings of powerful people, fearing the costs and uncertainties of judicial proceedings? Moreover, they need to extend the analysis from the immediate social effects of the intended legislation to its indirect effects: what consequences would the increased impunity, consequent upon the impossibility of using wiretapping in investigations, have on certain kinds of criminality, such as political corruption, extortion and racketeering, or drug trafficking? It may also happen, as when economic policy is involved, that the empirical predictions required to establish the likely outcome of certain measures are very difficult and questionable, being dependent on much-debated theories: will a tax cut boost investment? Will it improve or worsen the condition of the poor?

After considering some alternative measures aimed at solving the problem (not *all* possible measures, since this would exceed human capacities), legislators will need to compare such measures and write into law the measures having the best combined impact on all the values at stake. The analysis of the impacts of a new law on all relevant values can be very difficult. Difficulties may pertain to different aspects: predicting the empirical effects of alternative choices, spelling out the values to be achieved, specifying their content, and establishing their relative importance.

Finally, rational legislators should monitor the outcome of the law, to check whether it achieves the intended objective, or whether it has unwanted consequences, or whether a better solution to the problem can be found, a solution not considered when the legislative choice was made (possibly because certain knowledge became available only later, through advances in the natural sciences or in technology or economics).

Legislative rationality also includes the reflective element I described above, at least to a certain extent: legislators need to represent the interest of the their constituencies, or rather the view that their constituency has of the common good, but they should also subject such views to critical examination, taking into account empirical knowledge, correcting biases, etc.

We can distinguish the substantive and the procedural rationality of legislative procedure, where substantive rationality relates to a decision’s effective capacity to achieve the goals that legislators aimed at, and procedural rationality has to do with following a procedure that reliably tends to provide substantially rational decisions. Such procedural features include the ability to consider different normative and factual opinions, to collect evidence for and against a policy, to carry out empirical inquiries, to stimulate public debate and take its outcome duly into account, etc.

14. Constitutional Commitments and Legislative Rationality

Rationality requires taking duly into account previously adopted epistemic and practical determinations: it requires that these determinations guide subsequent reasoning until they are withdrawn. While individual reasoners can memorise past determinations as intentions (or duty-beliefs), where a collective agency is concerned, past determinations can be stored in normative sources (official documents, but also shared customs or doctrines) which are publicly accessible and embed norms to be followed and applied by officials, and which are to be modified according to established procedures. There may be different kinds of norms:

- norms establishing general or specific duties or permissions to carry out or not carry out certain acts;
- norms establishing a duty to aim at certain goals (values) or to not prejudice them;
- norms conferring a legal status; and
- norms indicating what factors support certain normative conclusions

All such norms—if they are part of a constitution—should constrain and guide the legislature's deliberative process. It would be irrational for individual reasoners not to act on the basis of a commitment they continue to accept, unless they believe they are in a situation where the commitment is inapplicable or is overridden by a prevailing reason to the contrary (I stand by my commitment to work out every evening, though this evening this commitment is made inapplicable or overridden by my commitment to give a lecture). Similarly, legislative determinations departing from constitutional norms could in a sense be viewed as irrational, i.e., as disregarding some commitments that govern legislative decision-making: for legislators who continue to uphold their commitment to a constitution (as they should when reasoning and acting in the name of the community governed by that constitution, i.e., a community that has undertaken such a commitment), it would be irrational not to respect a constitutional norm, unless they believe they are in a situation where the norm is inapplicable or is overridden by prevailing communal reasons to the contrary. Note that this irrationality only exists when legislators are viewed as members of a legislature acting in name and on account of the community committed to the constitution: violating the constitution to install a permanent dictatorship or to gain immunity from prosecution may be perfectly rational from the perspective of individual self-interest.

How to go about respecting a constitutional norm, however, depends on the content of that norm:

- a constitutional norm establishing duties or permissions to carry out or not carry out certain action is violated when a new legislative determination either directly instantiates a prohibited action or makes permissible what was prohibited or prohibits what was permissible;
- a constitutional norm establishing a duty for the legislature to realise a value (aim at a goal) is violated when the value does not enter in an appropriate way, according to its importance, into the teleological reasoning of legislators, namely, when it is not appropriately taken into account in legislative choice-making;
- a constitutional norm conferring a legal status is violated when a legislative determination denies such a status (similarly a norm denying certain persons a legal status is violated when a legislative determination confers such a status on them);
- a constitutional norm indicating that a factor supports a certain normative conclusion is violated when the factor is not considered in a legal determination where it was relevant.

According to this broad characterisation of the notion of a constitutional norm (which seems to me to tally with the common usage of the term *norm*), it also includes constitutional prescriptions requiring the pursuit of certain values (goods). If legislators have to take into account all constitutional norms, then these norms too will have to direct in legislative decision-making, along with the norms specifying that certain actions be to be taken or omitted.

In order to analyse how a legislature should comply with constitutional commitments, we need to focus specifically on norms establishing rights. On the traditional view that a right protects an individual interest or opportunity (the so-called benefit theory of rights, advanced by authors such as Jeremy Bentham and Rudolf Jhering: see, for a logical analysis and for references to the literature, Sartor 2006), two components are entailed by the statement that "*j* has a right to *A* toward *k*", *where j is the beneficiary of the rights and k* is the counterparty: on the one hand, the situation where *j* enjoys *A* is viewed as valuable and, on the other, it is assumed that there exist guarantees aimed at facilitating this enjoyment, which bear upon *k* (these guarantees can be specified in other norms or may have to be argued from general principles).

Thus a right-conferring norm includes in the first place a value component: the norm stating that "*j* has a right to *A* toward *k*'' entails that the legal system values *j*'s having *A*, or views it as an objective to be pursued through the law. More precisely, a right-conferring norm protects an individualised value, namely, a set of valuable situations pertaining to particular individuals separately considered (my freedom to speak, your freedom to speak, etc.). Consequently, the interest (value, good) protected by a right is essentially non-aggregative: the fact that someone's right is satisfied to an optimal extent does not make up for the fact that someone else enjoys the right to an insufficient extent.

Secondly, there is a guarantee component: *j*'s having a right to *A* toward *k* entails that the law provides some normative guarantees that facilitate *j*'s having *A* and bear upon *k*. For one thing, this right entails that *j* is permitted to have *A* as far as *k* is concerned (i.e., it is not the case that protecting *k*'s interests requires prohibiting *j* from having *A*).[8] The protection provided by a mere (or unprotected) permission to have *A* (see Hart 1982) can be strengthened by what might be called, in Hohfeldian terms, a disability or incapacity, namely, by *k*'s inability to change *j*'s legal standing with regard to *A*, namely, of turning *j*'s permission to have *A* into a prohibition (as would happen if *j*'s right was established under a constitutional norm, one that legislature *k* could not make any exception to). And, for another thing, *j*'s right to *A* toward *k* may include further legal guarantees, consisting in obligations incumbent upon *k* to facilitate *j*'s pursuit of *A*:

- *k*'s goal-duty (an imperfect duty, in Kant's terminology: see Kant 1996, chap. 2; Sen 2004a) to consider in *k*'s deliberative process the goal that *j* should have *A*, recognizing for this goal an appropriate relevance (e.g., the duty to consider freedom of speech when introducing a regulation aimed at protecting privacy);
- *k*'s negative action-duty not to prevent *j* from having *A* (e.g., a duty to not prevent a person—as through imprisonment—from expressing his or her opinion);
- *k*'s positive-action duty to ensure that others do not interfere with *j*'s having *A* (e.g., a duty to protect a person against attempts to prevent him or her from expressing an opinion);
- a positive-action duty to ensure that *j* has the means to enjoy *A* (e.g., a duty to provide access to the media)

Moreover, these duties are often accompanied by the right-holder's power to activate judicial enforcement when some of these obligations are not complied with (a power included in the restrictive notion of a legal right in Kelsen 1967).

In order for a right to exist, it is not necessary that full protection be provided (as would result from the combination of all the duties I have introduced, plus the corresponding powers of enforcement). The protection of certain rights (e.g., some social rights, such as the right to work or to housing) may consist in only a goal-duty, often not judicially enforceable though the right-holder's autonomous action. This would provide a lesser, but not irrelevant, protection of the corresponding individualised values (on how certain rights may consist in only an obligation to take them into account in

[8] For instance, if the legal system *L* prohibits *j* from having an opportunity to express his or her opinion in the interest of the state, we should conclude that *j* has no right under *L* to express his or her opinion about the state.

deliberation, see Sen 2004a). Some rights may operate in different ways with regard to different counterparts (e.g., the right to privacy may be protected by a negative action-duty with regard to administrative authorities, who are prohibited from using personal data unless specifically allowed by the law, but only as goal-duty with regard to the legislature, who can limit the protection of privacy though legislation taking competing interests into account). Certain rights (such as social rights) may be protected only by a goal-duty at a constitutional level, and by action-duties at the legislative level. The view that rights are values also protected (and sometimes only protected) by goal-duties does not mean that all rights are equal. This view is consistent with the assumption that certain individualised values (the enjoyment of civil and political liberties) may carry more weight than other values, and hence have priority over them, and in particular over collective values. This is also consistent with the view that some rights also include protection through defeasible or even indefeasible action-norms. However, outside the domain where an action-norm is to be applied (e.g., the prohibition against torture), goal-norms (norms that deal with values such as individual self-determination and integrity) would still operate.

To understand the distinction between action- and goal-duties, we should go back to our analogy between individual and collective decision-making: just as an individual determination (intention) to perform an action is adopted by a person because he or she considers that action to be an appropriate way to achieve certain goals, so a norm establishing an action-duty is adopted by a certain authority (or collectivity) because that authority (collectivity) considered the norm to be the appropriate way to achieve certain public values. Respecting the authoritative determination that has lead to the adoption of an action-norm requires us to not disregard that norm on the basis of a different comparative assessment of the values considered in that determination. Thus, if a constitution requires that nobody can be detained for more than 48 hours without a judicial warrant, interpreters (legislators and judges in particular) should not disregard this rule on the basis of the value of security (even when they believe that the constitution is wrong, e.g., that it should have established for detention a longer term based on a better balance of the values at issue): the constitution made its evaluation concerning the way to balance security and freedom, and respecting the constitution means respecting this evaluation (this corresponds to Raz's 1978 view of rules as exclusionary reasons).

In other cases, however, the situation is different. A particular constitutional norm obligating the legislature to uphold a certain value, even when the value is individualised and non-aggregative, may only require that the value be taken into account in legislative decision-making according to its constitutional importance. Consequently, this norm does not uniquely determine a legislative decision, which will instead result from a teleological evaluation aimed at achieving not only this value but also the other constitutional values at stake, and to do so in keeping with these values' relative importance. Thus, legislators are obliged to take into account and evaluate all relevant constitutional values: when aiming to guarantee security, for instance, they should also take into account privacy and freedom of speech. Sometimes a constitutional norm will guide such an evaluation by indicating what values should be relevant to this decision (thus excluding that other values may interfere with the outcome, or that they may interfere beyond the limit of evaluation accorded to the decision-maker). Thus, it may be possible to limit freedom of speech only for reasons pertaining to public order and morality, and not, say, for reasons pertaining to scientific progress (which consequently could not be used to justify a ban on advocacy for creationism or homeopathy).

The distinction between action-duties and goal-duties overlaps with another significant distinction, namely, the distinction between a yes/no state of affairs and a scalable state of affairs. A yes/no state of affairs either obtains or does not obtains, while a scalable state of affairs may hold to different extents. For instance, while being a citizen is a yes/no state, being free or unfree is a scalable state of affair (since this is a function of the number and quality of the options within one's reach). When two duties concern the realisation of a yes/no state of affairs, preference should be given to one duty to the exclusion of the other, so that at least one of the two is satisfied (this is the domain of defeasible reasoning). By contrast, when two duties concerning scalable goals have to be satisfied, the best

compromise usually requires that neither of them be completely neglected or completely satisfied. A scalable duty (the duty not to make people suffer when questioned or detained) can become an action-duty with regard to a particular threshold (the duty not to torture people). Action duties concern the realisation of yes/no states of affairs, while goal duties usually concern the realisation of scalable states of affairs.

Given the premises of legislative reasoning—constitutional goals, further legislative goals, preferences for such goals, and constitutional constraints on the pursuit of such goals—legislators should make a teleologically appropriate determination. From the legislators' perspective, this means that after an adequate inquiry, the chosen determination should appear better than inactivity and it should not appear to be worse than any particular alternative determination the legislators have so far identified. The legislative choice would fail to reach teleological appropriateness if the legislators made a choice they believed to be worse than another possible choice they were aware of in achieving the public good (even though the choice may be better suited to advancing the legislators' private interests). Similarly, the legislative choice would fail to reach teleological appropriateness if it were adopted impulsively, without an appropriate inquiry (which would have led to discover a recognisably better option). It would fail as well if it were vitiated by previous epistemic mistakes—in evaluating the evidence, identifying causal connections, examining evidence to the contrary, etc.—when such mistakes would not have been committed through an appropriate cognitive effort.

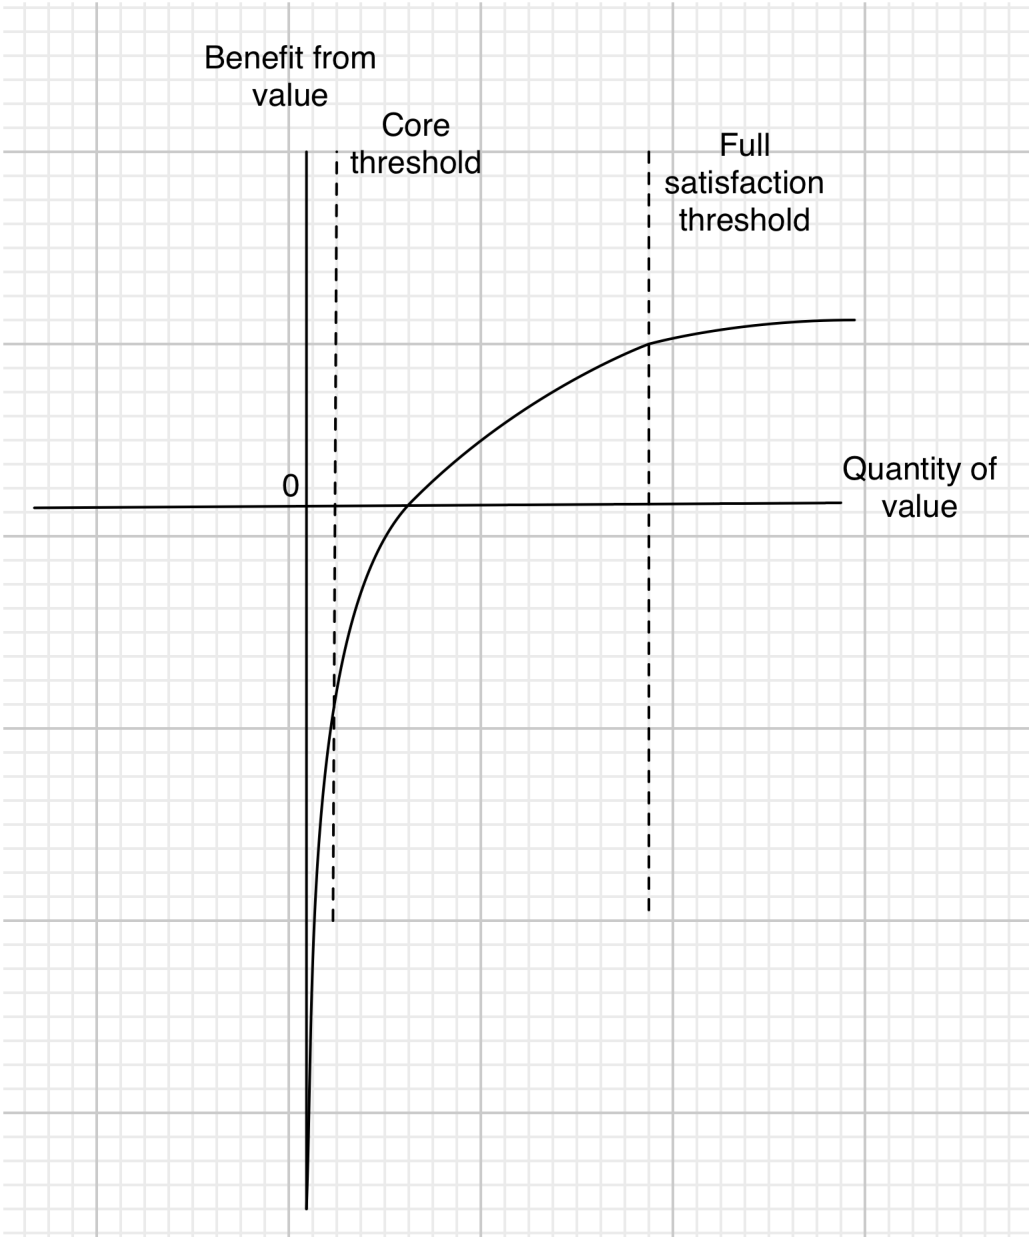

Figure 6: *Core and satisfaction threshold for a value*

Figure 6 shows the connection between the satisfaction of a scalable value and the benefit it provides: a decrease in the satisfaction of the value determines an increasingly significant loss in the benefit deriving from it. We reach a point, the *core threshold*, such that any further decrease in the satisfaction of the value determines a loss of benefit that is unlikely to be compensated by gains in the benefit provided by the greater achievement of other values. The portion of the value line to the left of the core threshold is what may be called the value's core or nucleus. On the other hand, when the level of achievement increases, we come to a point such that any further increase will have little importance. The portion of the value line to the right of this point represents situations where the value is achieved at a fully satisfactory level, so that any further increase, though still positive, may not come within the scope of a legal obligation to advance that value.

If scalable values have the structure just indicated, then decisions affecting competing values (e.g., privacy and security) take place in a decisional context of the kind represented in Figure 7.

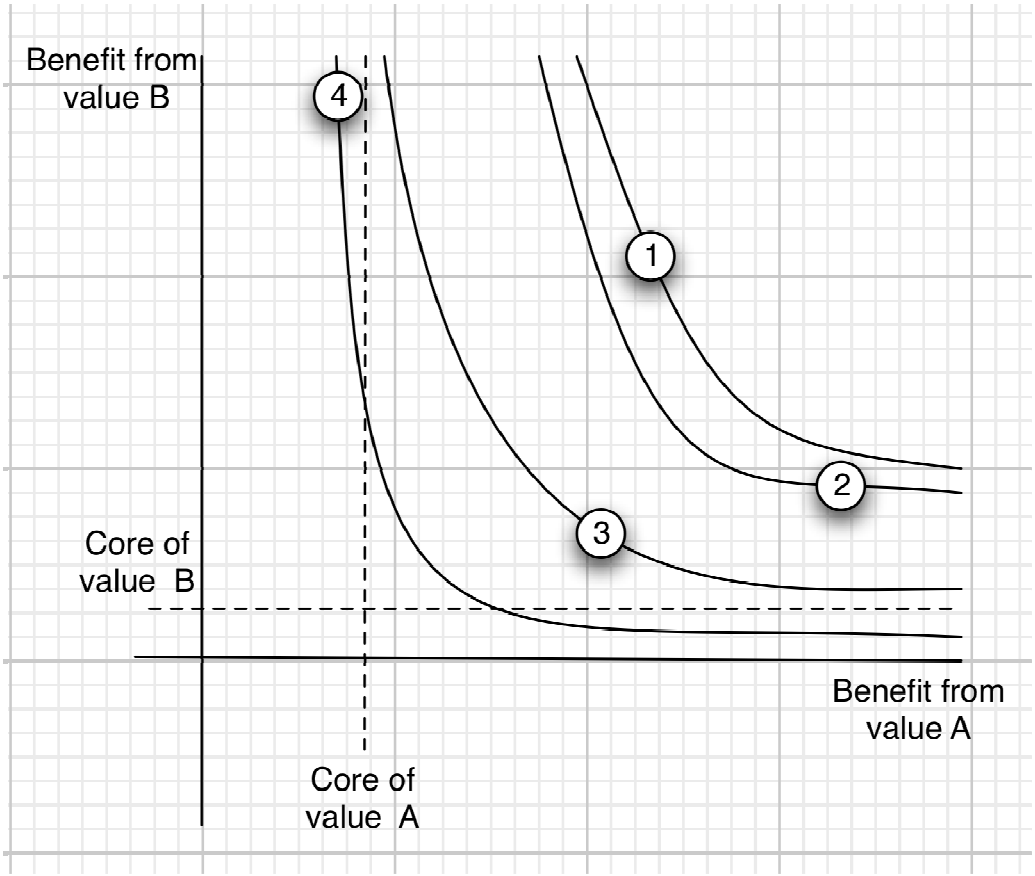

Figure 7: *The balancing of competing values*

The continuous lines indicate indifference curves, namely, combinations of levels of satisfaction of two competing values giving the same compound benefit. For instance, the most external indifference curve shows that achieving level 22 (measured by counting the number of small squares from the origin of the quadrant) with regard to both values *A* and *B* is equivalent to achieving level 40 with regard to *A* and 10 with regard to *B* (both points, [22, 22] and [40, 10], are situated on the same indifference curve). This curve expresses the idea that *B* (e.g., privacy) is less important than *A* (e.g., security): for most curves, a higher quantity of *B* is required to make up for the loss of one unit of *A*. However, when the quantity of *B* decreases, having one additional unit of *A* becomes more and more important, up to the point where any further increase in *B* will no longer make up for a further equal loss in *A*. Let us assume that the decision-maker has choices 1, 2, 3, and 4 available (represented in the figure by way of the numbered circles). Choice 1 is Pareto-superior to choice 3, since it provides not only a higher compound benefit, but also a higher level of satisfaction with regard to both values. Choice 4 is not Pareto-inferior to 1, since it indicates a level of satisfaction for value *B* which is higher than that provided by 1. However, this is obtained at the cost of a very low level of satisfaction for value *A*, a loss which is not made up for by the benefit provided, consisting in an increase in security. Thus choice 4, while ensuring the highest level with respect to *B* obtains the lowest compound score. The conclusion that choice 4 is inferior to 1 thus presupposes a "comparative value judgement," namely, a judgement about the comparative importance of values *A* and *B*. On the basis of this judgement, the loss with regard to *A* in choice 4 is not offset by the corresponding gain in *B*. Even 2 is not Pareto-superior to 1, but this happens for the opposite reason, which is that 2 achieves *A* to a higher degree than 1 but at a cost that is not offset by the loss with regard to *B*.

This quantitative characterisation of the notion of a right's core needs to be integrated by qualitative considerations by taking into account the diversity of the interests protected by a right. A single constitutional right can be analysed into different components, concerning different individual interests, but unified within the same framework (under the same overarching value). For instance, the right to private and family life recognised by the European Convention on Human Rights includes related, but different, components such as protection of the domicile, freedom to establish a family, freedom of sexual orientation, and information privacy (data protection). Interference with each such component takes the nonlinear shape I described above: as the level of satisfaction of a particular component of the right decreases, the negative impact on the corresponding interest becomes more and more important, in an accelerated way. Thus, each right includes a family of cores pertaining to different individual values (interests): for each of the specific constitutional values falling under a single right, there is a point when further losses are unlikely to be matched by gains with regard to other constitutional values pertaining to the same or to other rights. For instance, the fact that a legal system provides full protection of the domicile, along with full data protection, cannot make up for the fact that homosexuality is criminalized: a core of the right to private and family life would still be violated. The same would also happen if freedom in sexual orientation were protected but no data protection were provided. Similarly, a core of the right of freedom of speech would be violated if freedom of speech were fully protected in all respects save for a prohibition against criticising the current government.

15. The Constitutional Evaluation of Legislative Choices: Reasonableness and Deference

When we examined legislative decision-making from the perspective of the decision-makers themselves, we focused on bounded rationality. If sufficientist reasonableness is equated with bounded rationality, then a choice will be reasonable when it remains in the region between bounded rationality and optimal rationality. A reasonable but non-rational choice would be a non-optimal determination, such that no criticism of cognitive ineptitude can be directed at the decision-makers: they appropriately used their cognitive powers (in developing an economic policy, or in designing privacy regulations), only they failed to achieve the best possible result and caused negative outcomes (growing unemployment, citizens' privacy unduly restrained) because of the unfortunate cognitive circumstances in which they were acting (new unexpected social or technological developments, unavailability of good predictive models, etc.). Correspondingly, any departure from bounded rationality (any mistake in acquiring and processing the available legal or factual information) would count as unreasonable.

This does not seem to correspond to the way in which reasonableness (and unreasonableness) is understood in judicial review, with regard to both legislative and administrative choices, where a broader notion of reasonableness is generally preferred, according to which a determination remains reasonable even though it is affected by cognitive faults, according to the reviewer.

With regard to judicial review the analogy we used between individual decision-making and the institutional decisional process of a legal community breaks down: while in case of individual decision-making the same agent is involved in the entire process (agents can consequently review any outcome of their previous reasoning which appears faulty to them), the decisional process of a legal community involves different bodies and institutions, each having its own functions and capacities. It is unlikely that the best integration between a decision-maker and a judicial reviewer will be one where the reviewer can strike down the decision maker's choice whenever the reviewer sees it as failing to achieve complete rationality (this would empower the reviewer to strike down all decisions she views as suboptimal, namely, all decisions she would not have taken had she been in the decision-maker's place, but with the hindsight of someone having all knowledge available the time of the review), even with regard to the achievement of constitutional values (this would empower the

reviewer to strike all decisions she views as failing to maximise the total outcome with regard to all the constitutional values at stake). Nor is the best integration likely to be one in which the reviewer can strike every decision she views as failing to achieve bounded rationality (this would empower the reviewer to replace with her own decisions all the decisions she would have taken differently had she been reasoning with the information the decision-maker had at the time the decision was made).

We must therefore define a different notion of reasonableness, a notion tailored to the institutional role and competence of decision-makers and of their reviewers, and in particular a notion that takes into account the reviewer's deference space, namely, the area within which the reviewer should not attack the measure under review even though she believes a different measure should have been taken (on deference, see Soper 2002). If unreasonableness (where constitutional values are concerned) is understood as providing a sufficient ground for review, then the notion of reasonableness is not independent of deference but is rather delimited by institutionally due deference. In other terms, considerations of institutional deference enable us to identify a sufficientist reasonableness threshold, encompassing not only the decision the reviewer would have taken but also other choices which he or she considers to be faulty but not yet unreasonable (insufficiently faulty to be unreasonable). However, this means that we cannot provide a universal characterisation of deference-based reasonableness, precisely because such a characterisation will depend on institutional deference.

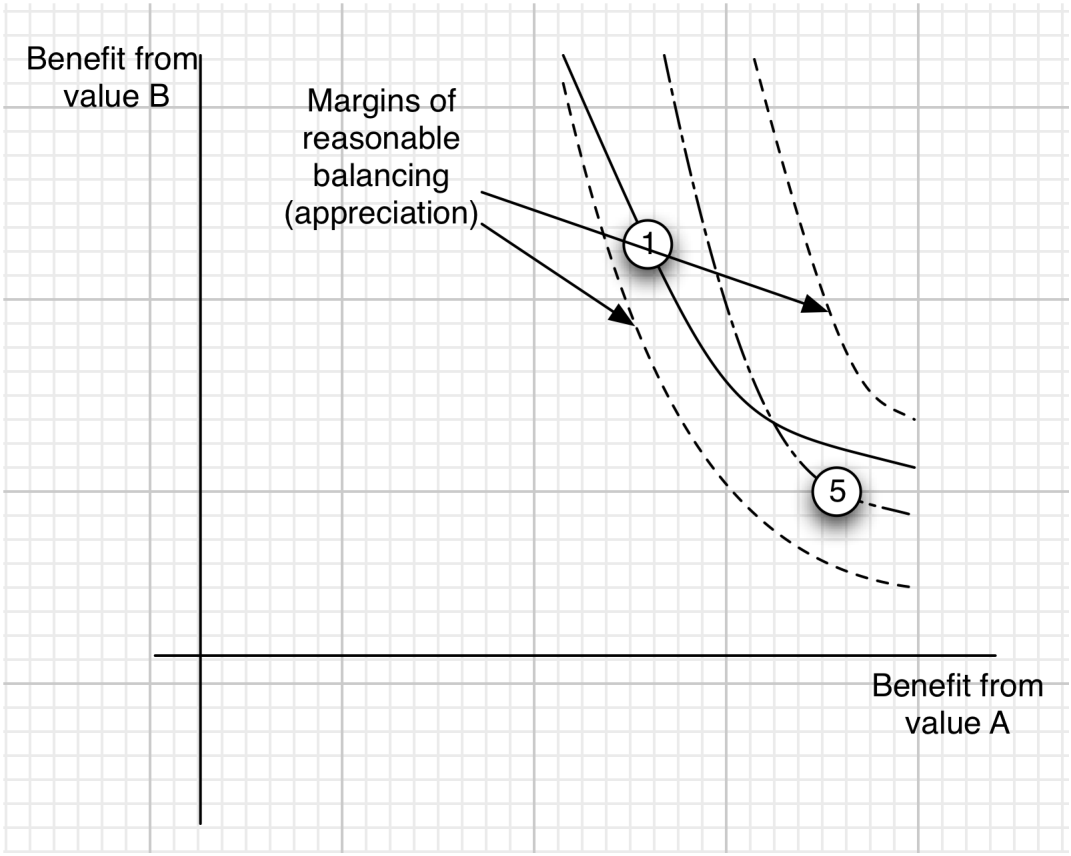

Figure 8: *Reviewer's choice and margin of appreciation*

Figure 8 illustrates how a determination (1) that does not coincide with what the reviewer would choose (5) may still fall within the margins of the decision-maker's appreciation (as indicated by the dotted lines) and may thus escape review: though the reviewer views the decision-maker choice as imperfect (it is based on an indifference curve that in the reviewer's opinion accords too much importance to value *A*), she does not consider it to be attackable, being within the margin of reasonable appreciation.

The idea of a sufficiency threshold applies as well to a legislature's epistemic judgements, which too can determine a failure to appropriately balance the values at stake. For instance, given the factual premise that a terrible terrorist attack is imminent, and the premise that scouring all Internet traffic with data-mining techniques will probably foil the attack, a legislator may be justified in adopting such measures to the detriment of privacy. However, if there are no grounds for accepting either of those premises (no convincing evidence that an attack is underway, and little evidence that unrestricted data-mining will be able to prevent it), then sacrificing privacy may be considered unreasonable. But substituting the court's epistemic assessment for the legislature's seems to require something more than a mere mistake of the latter: it should require a mistake consisting in epistemic unreasonableness, namely, a serious and indisputably ascertainable fault. Thus, this should not be done when the legislature's fault, according to the court, only depends on the adoption of a particular economic or social theory which the court favours (viewing it as more reliable, better supported by the facts), but which other reasonable people reject (as Judge Holmes famously argued in the Locher case).

16. A Sufficientist Understanding of Proportionality

By combining the foregoing analysis of failures in rationality with a deference-based sufficientist idea of reasonableness, we can derive the tests usually available in proportionality analysis (see, among others, Alexy 2003; Stone Sweet and Mattew 2008) as applied to determinations sacrificing constitutional rights. For this purpose, it is useful to stipulate some additional terminology.

I shall say that *α satisfies V more than β does*, in formula $\alpha >^{V} \beta$, to mean that, the level of achievement of value or goal *V* resulting from action *α* is higher than the level resulting from action *β* (and I use similarly, the symbols $<^{V}$, $\leq^{V}$ and $\geq^{V}$). For instance, if a the level of achievement of privacy resulting from legislative determination *α* (keeping the footings taken with street cameras for 1 day) is higher than the level resulting from determination *β* (keeping the footings for 1 year), I shall say that *α* satisfies privacy more than *β* does ($\alpha >^{privacy} \beta$).

By saying that *α sacrifices V,* I shall mean that *α* satisfies value or goal *V* less than the inaction (omitting *α* without replacing it with a different initiative), in formula, $\alpha <^{V} \varnothing$ where ∅ denotes inaction. For instance, if a new regulation *α* allows for the registration of genetic data of all newborn children (which was previously forbidden), I shall say that *α* sacrifices privacy.

Similarly, I shall say that *α advances V*, to mean that *α* satisfies value or goal *V* more than inaction ($\alpha >^{C} \varnothing$). For instance, if a new regulation *α* prohibits processing personal data for commercial purposes without the consent of the person concerned (which was previously permitted) I shall say that *α* advances privacy.

I shall also say *α satisfies a set* {*V*1, …,*Vn*} *more than β does* ($\alpha >^{\{V1, \ldots, Vn\}} \beta$) to mean that the *α's* compound impact on all values in the set is better (higher, more valuable) than *β*'s. In assessing this compound impact the relative sacrifice in one value can be outweighed by the relative advance in another (in case numerical indicators can be given, the satisfaction of a set of values would be the weighted sum of the satisfaction of each single values in the set, as we observed in section 4). Consequently, it is possible that *α* satisfies the value set {*V*1, …,*Vn*} more than *β* does even though *β* satisfies certain values the set more than *α* does. This happens when *α* satisfies certain other values in the set more than *β* does, and the comparative advantage provided by *α* with regard to the latter values outweighs the advantage provided by *β* with regard to other values. Assume, for instance (see Section 4), that *β* consists in prohibiting street cameras, and *α* consist in allowing cameras and requiring the destruction of the footage after 1 day, and that the comparative advantage in security provided by *α* outweighs the comparative disadvantage in security provided by *β*. Then it can be concluded that *α* satisfies the set {*privacy, security*} more than *β* ($\alpha >^{\{privacy,\ security\}} \beta$), even though *β* satisfies privacy to a higher extent than *α* does ($\beta >^{privacy} \alpha$).

As we observed above, a more rigid kind of comparison, excluding tradeoffs between different values, is provided by Pareto superiority. For α to be Pareto superior to β (and β Pareto inferior to α) with regard to a set of values $\{V1, \ldots, Vn\}$ it is required that α satisfies at least one value in the set more than β and that α satisfies every other value no less than β (there exists no value in the set such that β satisfies it to a higher extent than α). This condition would not be satisfied in the latter example, where β satisfies privacy to a higher extent than α. Thus in this example, even though α satisfies the set {*privacy, security*} more than β does, α is not Pareto-superior to β. The advantage of the criterion of Pareto-superiority is that it allows us to conclude for the preferability (rationality) of a choice without engaging in balancing, understood as the (usually controversial) assessment of tradeoffs between impacts on different values. However, when we have to compare choices α and β that impact different values, α satisfying more certain values and β certain others, then Pareto-superiority is not applicable, and balancing is required. Let us consider the four usual proportionality tests.

Test 1: Legitimate aim 1 (permissible intended purpose). A legislative determination sacrificing a constitutionally protected value must aim at advancing a constitutionally permissible goal. This requirement is violated when the legislature adopts a determination sacrificing a constitutional value, and does so in order to pursue a goal that it is constitutionally impermissible. Such a decision would fail to be rational (or, for that matter, reasonable) with regard to the decisional context as emended by removing the impermissible goal. In other words, when the impermissible goal is eliminated, the only relevant intended impact of the decision is its negative impact on the constitutional right/goal, which makes that choice internally irrational (Pareto-inferior to the choice that consists in inactivity).

Test 2: Legitimate aim 2 (legitimate outcome). A legislative determination sacrificing a constitutional value must effectively advance a constitutionally permissible goal. In other words, if a constitutional value *C* is sacrificed by α, there must exist a legislative goal *G* that is advanced by α, where the pursuit of *G* is constitutionally permitted. Assuming that requirement 1 is met, i.e., that a permissible goal *G* is pursued, the *legitimate-outcome* requirement is violated when the legislature adopts a determination α that sacrifices a constitutional value without advancing the pursued legislative goal (e.g., a useful medical therapy is prohibited assuming, on the basis of wrong medical information about the effects of this therapy, that it damages patients' health). This determination will be unreasonable when the legislature's epistemic mistake is very serious and indisputable (unreasonableness can also involve retaining the new legislation when new indisputable evidence is made publicly available showing the previous choice to be mistaken). It seems to me that the reviewer's consonance and the legislature's dissonance with regard to the scientific community's evaluations should often play a decisive role in justifying the reviewer's intervention.

Test 3: Pareto-necessity. A legislative determination that advances a certain goal while sacrificing a constitutional value, must not be Pareto inferior (where constitutional values are concerned) to any determination that equally advances the goal with a smaller sacrifice of the constitutional value. This requirement is violated when the determination α—while respecting requirements 1 and 2, i.e., pursuing a permissible goal *G* through effective means—sacrifices a constitutional value *C*, and there exists an alternative determination β that advances *G* at least to the same extent, sacrifices *C* to a lesser extent, *a*nd imposes no additional burden on any another constitutional value. Under such conditions, we may that α is not necessary to advance *G*, since β rather than α could have been taken (avoiding the unnecessary sacrifice at no cost). Thus, when Pareto-necessity is violated, the following conditions hold: α sacrifices *C*, and there exists an alternative choice β such that β satisfies *C* more than α *does*, and β is Pareto-superior to α with regard to a value set $\{C,C1,\ldots Cn,G\}$ comprising all constitutional values $C,C1,\ldots,Cn$ plus the pursued legislative goal *G*. Note that for β being Pareto-superior to α it is necessary that β satisfies each of *C1,…Cn, and G* no less than α. Consequently, if α satisfies any of the $C1,\ldots,Cn,G$ more than β does then β fails to be Pareto-superior to α, which means that α satisfies Pareto-necessity (as far as β is concerned). For instance, a law α requiring DNA samples being taken of all citizens (for their use in future police investigation), would stand the test of Pareto-necessity as compared with the choice β of getting samples only from suspects for particular crimes: even though α

sacrifices privacy to a very high extent, if satisfies the value of security more than β does, and this is sufficient to exclude that β is Pareto-superior to α. Failure to satisfy Pareto-necessity may depend on an epistemic mistake in assessing the impacts of α and β (e.g., the mistaken belief that α sacrifices *C* less than does β), or it may depend on a heuristic failure to identify determination β. When such mistakes overstep the unreasonableness threshold, review is justified.

Test 4: Comparative balancing. A legislative determination sacrificing a constitutional value must satisfy the set of all constitutional values no less than inaction ($\varnothing$) would do. Let us assume that a determination α sacrificing constitutional value *C* meets all requirements 1 through 3, i.e., that α effectively pursues a permissible goal *G* and α is not Pareto-inferior to any other possible determination β. Under these assumptions, α complies with the requirement of Test 4 if α satisfies no less than $\varnothing$ the value set {*C, C1,…,Cn*}, including all constitutional values. Given that α sacrifices value *C*, this requires that α advances other values, to an extent offsetting *C*'s sacrifice. Conversely, failure to satisfy Test 4 would mean that *C*'s sacrifice is not offset by the advancement with regard to other values. For instance, the conclusion that a law requiring DNA samples being taken from all citizens fails to meet the requirement of comparative balancing (as compared to the current state of affairs when samples are only taken on suspects) requires that the advance this law provides with regard security is outweighed by the sacrifice it imposes on privacy. Failure to satisfy Test 4 may derive from the legislators' practical mistake in assessing the relative importance of two or more values at issue (giving too much importance to some values advanced by α or too little importance to some value sacrificed by it), or from their epistemic mistake in assessing the impact of their decision on those values. In either case, in order for α to be unreasonable, these practical or epistemic mistakes must also have been unreasonable.

Special care is required in Test 4, especially when unreasonableness depends on a mistaken evaluation of the comparative importance of the constitutional values (goals) at stake. Here the court's interference with the legislature's political autonomy would be highly controversial (given that disagreements over matters of fact tend to more easily be resolved than disagreements over matters of value). Such an interference can be more easily justified (considering the democratic derivation of the legislature's power) when the judges' evaluation appears to comport with people's assessment of the relative importance of the values (or with the assessment of the people who are interested in the matter and have been discussing it), namely, when it appears that the legislature has failed to take duly into account the idea of reasonableness as agreement or consonance with general opinion (see Section 2).

The way in which balancing in a strict sense has been characterised in Test 4 above does not exhaust all possibilities. Let us consider two variations, the first one making judicial balancing *less* intrusive and the second one *more* so.

The first variation consists in admitting the relevance of non-constitutional goals pursued by the legislature. This is excluded by the Test 4 as characterised above, which only contemplates constitutional values: a diminution (meeting the required seriousness threshold) in the combined satisfaction of constitutional values is sufficient to strike down legislation regardless of any increase in the satisfaction of non-constitutional values. This means that constitutional values are viewed as lexically superior to non-constitutional values (no achievement of any level of a non-constitutional value can justify a diminution with regard to constitutional values, that is, non-constitutional values are irrelevant to constitutional balancing).

Let us say that a choice α *outbalances* β with regard to values *V1,…,Vn*, to indicate that α satisfies the values set {*V1,…,Vn*} more than β, to an extent sufficient to make it unreasonable to prefer β over α. On the approach under Test 4, if *C,C1,…,Cn* are all the constitutional values involved in the case, and inaction ($\varnothing$) outbalances α with regard to those values, α fails the test, regardless of α's impact on non-constitutional values. For instance, let us assume that privacy and security are constitutional values, while cutting down on public spending is not. Let us also assume that choice α (e.g., introducing body-scanning in airports) impacts with a negative combined outcome on both privacy

and security (and on no other constitutional value besides), but that implementing α is much cheaper than implementing the previous policy (which entails costly inspections by airport customs). It then follows, on the approach under Test 4 above, that if a judge believes that inactivity (maintaining the previous policy) outbalances α with regard to the set {*privacy, security*}, this judge should disregard the older policy's higher costs (assuming that cost-effectiveness is not a constitutional value) and should accordingly strike down α. The same judge could reach a different conclusion is non-constitutional values were also relevant in the comparison. In the latter case, then α could stand the review if the judge believed that α's gain with regard to the non-constitutional legislative goal (cost-cutting) would balance α's loss with regard to the constitutional values (or at least would keep the imbalance within the unreasonableness threshold). This leads to the following weaker version of Test 4.

Test 4.1: *Comparative balancing* (weaker version). A legislative determination α to achieve a certain goal *G* sacrifices a constitutional value *C*, it must not be outbalanced by $\varnothing$ as concerns both the impact on constitutional values {*C,C1,…,Cn*} and the impact on the larger set {*C,C1,…,Cn,G,L1,…,Lm*}, comprising the goal pursued plus any other objectives *L*1,…,*Lm* valued by the legislature. This requirement can be considered unfulfilled only when both $\varnothing$ outbalances α with regard to both sets.

For instance, in the example above, it would be possible for Test 4.1 to be satisfied even when $\varnothing$ outbalances α with regard to {*privacy, security*}, i.e., when Test 4 fails to be satisfied. This would happen when the outcome of the comparison would change by also including the goal of *cost-cutting*. In other terms, while under Test 4, the only condition that needs to met in order for a legislative determination α to be struck down is that $\varnothing$ outbalances α with regard to {*privacy, security*}, Test 4.1 also requires that $\varnothing$ outbalances α with regard to {*privacy, security, cost-cutting*}.

A different, more intrusive variation of Test 4 would consist in substituting the requirement that inaction ($\varnothing$) not outbalance α with the requirement that there be no alternative determination β outbalancing α: the reviewer would strike down α not only when α worsens the preexisting combined achievement of constitutional values but also when α, while resulting in a combined achievement of constitutional values higher than (or equal to) the outcome afforded by $\varnothing$, falls short of a maximal combined achievement of the constitutional values at stake, since there is an alternative determination β that would provide an even better outcome. Let us consider the following hypothetical example (which follows to some extent the case *S. and Marper v. the United Kingdom*, decided by the European Court of Human Rights on 4 December 2008). Suppose that non-voluntary storing of genetic data was prohibited under a preexisting legal framework r_1, and that a new rule is issued making it possible to store genetic data collected in the course of a criminal investigation: this would yield a new regulatory framework r_2. Suppose now that a reviewer agrees that framework r_2, while sacrificing privacy to security, yields an outcome (a combined satisfaction of the two values of privacy and security) better than the outcome yielded by r_1 (the loss of privacy being outweighed by the advance in security). And suppose, finally, that that the reviewer also believes that a different framework r_3 (e.g., making it possible to store data but requiring deletion in case charges are dropped) would more suitably balance privacy and security. The reviewer will thus conclude that that r_3 outbalances r_2 , and that r_2 outbalances r_1 ($r_3 >> r_2 >> r_1$). However, r_3 is not Pareto-superior to r_2, since r_2 satisfies security more than r_3 does, though, in the reviewer's judgment, r_3 provides an advantage in privacy that outweighs the advantage in security provided by r_2. In such a case, then, the reviewer's decision to strike down r_2 would not pass any of the previously listed proportionality tests (it will fail Test 3, since r_3 is not Pareto-superior to r_2, and it will fail Test 4, since r_2 is not comparatively inferior to r_1, which would have resulted from inaction. If we are to justify striking down r_2, we have to introduce the following additional test.

Test 5: Comparative counterfactual balancing. A legislative determination sacrificing a constitutional value must not be outbalanced, with regard to the set of all constitutional values by any alternative determination the legislature could have taken. A legislative determination α violates this requirement when α sacrifices a constitutional value *C* and there exists an alternative determination β that would

produce a lesser sacrifice of *C* and would outbalance β with regard to the set of all constitutional values. This would be the case of the last example, where r_2 sacrifices privacy advancing security, and r_3 would provide a lesser sacrifice in privacy and a better compound achievement of all constitutional values at stake (including both privacy and security).

The application of this last test seem to me particularly problematic in the cases where Test 4 is not satisfied, i.e., in the cases where the contested decision α satisfies the set of all constitutional values more than the preexisting state of affairs (the outcome of inaction), though less that the alternative measure β being considered. If the reviewing court can only act as a negative legislator, that is, its only power is to remove α from the legal system on the basis of Test 5, the court would bring back the preexisting state of affair, worsening the compound achievement of the constitutional values (as would happen in the previous example, if r_2 was cancelled, reinstating r_1). This paradoxical effect can be avoided if the court can strike down α without cancelling all of its effects (e.g., it can impose β as a particular interpretation of the language of text α, as has happened with certain interpretive decisions of the Italian Constitutional Court, or it can rule that the government pay out compensation for the injury owed to by α, all the while preserving α's legal validity, as has happened with decisions of the European Court of Human Rights). However, it seems that deference to legislative authority requires that such interventions only take place in extreme cases: β's advantage over α must be truly uncontroversial (or at least there must be broad support for this view in public opinion), and β must be obtained through a modification of α (by introducing an exception or extension to it), reframing α's content rather than going to the extreme measure of striking down α. A variant of Test 5 could be obtained by also including non-constitutional values, as in Test 4.1.

Let us now go back to Test 3, namely, Pareto-*necessity*. As we have observed above, the idea of Pareto necessity does not involve an assessment of tradeoffs between impacts on different values (as included in Test 4). This matches the use of the term "necessity" in some legal context, but does not correspond to the way in which this term is used in other contexts. For clarifying this issue, I shall identify, besides the Pareto necessity a different notion of necessity, which I shall call *balanced-necessity*. We must firstly consider that for the purpose of proportionality – in assessing whether a measure α, sacrificing a constitutional value *C*, is necessary to achieve a goal *G* – we cannot use the notion of necessity in its common sense, namely, as meaning that α is a *conditio sine qua non* of *G*, i.e., that unless α is adopted, then *G* will not be achieved (at a level equal or higher to that which it would be achieved though α), i.e., that there is no alternative measure β such that, if β had been taken instead of α than G would have been achieved. This understanding of necessity would be too restrictive for the purpose of proportionality: it would lead us to conclude that α fails to be necessary to achieve G whenever there exists another measure β that achieves G, even though β sacrifices C more than α does. It could then could be argued, for instance, that body scanning is not necessary for security in airports since there is an alternative measure—namely, having people stripped of all their clothes—which would achieve the same level of security (even though the latter measure would sacrifice privacy more than body-scanning). To avoid such absurdities, when evaluating whether α, involving C's sacrifice, is necessary in order to achieve G, we must only consider those alternative measures that, while achieving G entail a lower sacrifice of C. This is done with the notion of Pareto-necessity ad characterised in Test 3 above, which extends considerably the scope of necessity. When evaluating the Pareto-necessity of a measure α sacrificing value C, the alternative measures excluding necessity are restricted to those that would not entail a lesser satisfaction nor of the goal been pursued nor in any other constitutional value: α fails to be Pareto-necessary only if there exists any alternative measure β that besides sacrificing *C* less than α, also satisfies no less than α does the pursued goal as well as all other constitutional values.

The notion of Pareto-necessity, however, is too permissive to match the way in which the notion of necessity is used in certain legal contexts. Consider for instance Art. 8 of the European convention on human rights, stating that limitations of the right to privacy must be necessary "for the prevention of disorder or crime, for the protection of health or morals, or for the protection of the rights and

freedoms of others". In the application of this norm, necessity has been denied even with regard to legislative choices that would satisfy Pareto-necessity (since the alternative choices available to the legislator where indeed inferior, under certain regards, to the legislative choice considered non-necessary by the judges). For a certain measure α to be necessary for achieving a goal *G*, in the sense assumed by Art. 8, it seem indeed that there must be no alternative measure β that, while achieving *G* to a satisfactory extent (though possibly less than α would do), sacrifices *C* to a lesser extent, and provides better combined impact on all relevant values at stake (including *C* and *G*). In other words, this decision β must be such that the advantage it gives with regard to the sacrificed value *C* is not offset by α's advantage with regard to *G* and other relevant values.

Test 3.1 Balanced necessity. A legislative determination α sacrificing a constitutional value *C* to achieve a goal *G*, must be such that there exists no alternative determination β, such that β satisfies *G* to a satisfactory extent, sacrifices *C* less than α does, and outbalances α with regard to the set including all constitutional values and goal *G*.

So, the balanced-necessity test seems to be an instance of Test 5 (comparative balancing), qualified by the additional requirement that the alternative determination β still satisfies the legislative goal up to a satisfactory level, while sacrificing less the value diminished by α. This seems the idea of necessity adopted by the European Court of Human Rights adopted in case Marper v. the United Kingdom. In fact the attacked measure (preserving all DNA samples taken from suspects, even after are acquittal), provides a higher level of security (maintaining more data potentially useful for future investigations) as compared with the alternative measure considered by the Court (preserving only the samples taken from convicted persons), and thus appear to meet Pareto-necessity. By excluding that the preservation of all samples is necessary for security, the Court assumes that necessity may fail even when the less interfering measure satisfies the legislative goal (security) less than the contested legislative measure does. Obviously, this understanding of necessity allows for a more intrusive judicial control than mere Pareto-necessity, a control that is based on assessing tradeoffs between different values. Thus review based on violation of the standard of balanced-necessity needs to be constrained appropriately though sufficientist considerations (by requiring that the legislative goal to which the legislative measure α was directed really is satisfactorily achieved also by the alternative measure β, that the advantage provided by β with regard to other values at stake strongly outweighs β's lower achievement of the legislative goal, that β is obtained by conveniently extending or restricting α*'s* content).

17. Constitutional Teleology and Implied Constitutional Norms

Constitutional rights, as inputs of legal reasoning, can be viewed in two complementary ways. The first view understands them as entitlements directly deriving from the unqualified recognition of the corresponding (individualised) values. Accordingly, a constitutional right primarily operates, with respect to the legislature, as a guide for the legislature's teleological reasoning: the legislature has the goal-duty to take into advance the corresponding value (freedom of speech, privacy, participation in science and culture, etc.) taking it into account in legislative determinations. And so, where judicial review is concerned, rights operate as criteria (values) for assessing the reasonableness (proportionality) of legislative choices.

The second view understands constitutional rights as articulating a set of normative positions established by more-specific norms, which are not expressly set forth in the constitution, but whose implied existence can be teleologically argued for with reference the rights (values) explicitly so recognised: given the constitutional values, their relative importance, and the social and institutional conditions for their implementation, more detailed right-conferring norms are extracted through interpretative arguments affirming that the application of these norms is likely to provide an appropriate balance of the values at stake. In particular, norms can be devised that unconditionally prohibit certain interferences with a right (these being inferences encroaching upon the core of the right, inferences that consequently cannot normally be balanced by the need to satisfy other values), or

norms can be devised stating that a right may not be limited by certain values (whose increased satisfaction is unlikely to justify a limitation of that right). In this way, a right's normative content can be characterized in a more casuistic way (but still in general terms, concerning classes of cases) by introducing norms to be used as "trumps" against competing values, it being assumed that the protected right outbalances the competing values in the cases identified in these norms themselves (e.g., "The right to private and family life includes self-determination as concerns sexual orientation and reproductive choices, within the following limitations: … This right includes in particular … and can only be constrained with respect to the following cases: … It also includes the right to exercise control over one's personal data, specifically as it pertains to … under the following limitations: … This right includes the right to access pornographic materials, with the exclusion of child pornography," and so on).

Different legal systems may place a different emphasis on the two perspectives I have described. Some systems may focus directly on constitutional rights in the abstract, thus entrusting the judges with the task of directly evaluating particular legislative decisions in view of their impact on those constitutional values. Other systems may rely instead on the judicial and doctrinal definition of lower-level rules constraining legislative decision-making, thus entrusting the judges with the task of framing and applying such rules, and then evaluating the legislative decisions on that basis (see Nimmer 1968, who introduced the term "definitional balancing" to describe this idea; for some criticism, see Aleinikoff 1987, 979). The first approach seems to correspond to some extent to the practice of European constitutional courts, while the second is more often used by the U.S. Supreme Court.

Which strategy (or which combination of them) is most appropriate depends on different institutional structures and legal traditions, but these two strategies may be considered to some extent as functionally equivalent. In common-law jurisdictions, based on the idea that judicial decisions produce binding rationes decidendi, constitutional judges may feel conformable with explicit rule-making. Consequently, rather than attacking a legislative determination for its failure to appropriately balance constitutional values, common law judges may prefer to extract from the constitutional recognition of such values (e.g., the right to free speech) some general rules (e.g., the rule that no content-based restrictions on speech are admissible, unless conditions of strict scrutiny are meet; or the rule that child pornography is not covered by freedom of speech) whose application will likely lead them to strike down the legislative determinations failing to effect an appropriate balance. They will then be able to decide cases (e.g., striking down a law that establishes a content-based limitation on free speech, or not striking down a law that makes child pornography illegal) by evaluating legislation in light of implied constitutional rules rather than in light of the underlying values to be balanced.

But it will still be necessary to rely on the underlying values in justifying and interpreting the implied constitutional rules or in working out conflicts between them (by giving priority to the rule whose application, in the case at hand, leads to a higher combined satisfaction of the values at stake: for instance, in cases involving hate speech, freedom of speech can prevail on dignity and non-discrimination, or vice versa). Moreover, when application of the implied rules fails to provide an appropriate outcome (it would lead to striking down a legislative norm providing an appropriate balance, or to preserve an unbearably unbalanced one), the judges would need to reformulate such rules or to supplement them with exceptions.

I cannot consider here advantages and disadvantages of the two approaches (greater contextual flexibility as against greater predictability, a clearer perception of the interests at issue as against an incremental refinement of precedent-based choices), for this would in turn have us compare rule-based decision-making with a more casuistic style of decision and weigh the pros and cons (see Schauer 1991; and for a discussion of some problems involved in case by case balancing, see Kumm 2007). We should bear in mind, finally, that both perspectives recognise the important role that teleological reasoning plays where constitutional values are concerned, and that this role is framed in different ways in the two approaches, which in this respect can be considered complementary (the constitutional judge/interpreter can to go back to goal-norms when implied rules are not applicable, or can revise rules when they fail to appropriately balance the constitutional goals).

18. Conclusion

This contribution has presented two main theses. The first concerns the correspondence between the general notion of rationality in decision-making and rationality in legislation: the argument here was that legislative decision-making is guided at its core by teleological reasoning, and that such reasoning can be analysed and evaluated according to general patterns of rationality. Moreover, a correspondence has been established between the way in which individual rationality is guided and constrained by commitments (intentions) and the way in which legislation is guided by constitutional norms. Combining the ideas of teleology and of norms made it possible to argue that goal-norms play a key role in legislative decision-making, and so we looked at their particular structure and function.

The second thesis concerns the development of a sufficientist understanding of reasonableness conceived as a standard for the constitutional evaluation of legislation: in order for reasonableness to be achieved, a sufficient level of rationality and morality is required, and this is a lower level than that of cognitive and moral optimality. This idea, introduced in Section 2, was applied to the complexity of legislative decision-making and modelled after the proportionality test: it says that constitutional review must leave the legislature a margin of epistemic and practical appreciation, even when constitutional values are at issue.

In conclusion, the achievement of constitutional values through legislation is a difficult, uncertain, and complex task, open to reasonable disagreement, and this appreciation requires that constitutional review of legislation be based on a modest (sufficientist) understanding of reasonableness.

References

Aarnio, A. 1987. *The Rational as Reasonable*. Dordrecht: Reidel.

Aleinikoff, T. 1987. Constitutional Law in the Age of Balancing. *Yale Law Journal* 96: 943–1005.

Alexy, R. 2002a. On Balancing and Subsumption: A Structural Comparison. *Ratio Juris* 16: 33–49.

– – –. 2002b. *A Theory of Constitutional Rights*. Oxford: Oxford University Press.

– – –. 2003. Constitutional Rights, Balancing, and Rationality. *Ratio Juris 16*, 131–40.

Aristotle. 1912. *On the Motion of Animals (De motu animalium)*. Oxford: Clarendon.

– – –. 1924. *Metaphysics*. Oxford: Oxford University Press.

Bench-Capon, T. J. M. and H. Prakken (2006). Justifying Actions by Accruing Arguments. In *Computational Models of Argument.* Eds. P. E. Dunne and T. J. M., Bench-Capon, 247–258. Amsterdam: IOS Press.

Blackburn, S. 1998. *Ruling Passions: A Theory of Practical Reasoning*. Oxford: Clarendon.

Bratman, M. 1987. *Intentions, Plans and Practical Reasoning*. Cambridge, Mass.: Harvard University Press.

Brentano, F. 1973. *Psychology from an Empirical Standpoint*. London: McAlister.

Fodor, J. 1983. *The Modularity of Mind*. Cambridge, Mass.: MIT Press.

Goldman, A. 1999. *Knowledge in the Social Word*. Oxford: Oxford University Press.

– – –. 2006. Social Epistemology. In *Stanford Encyclopedia of Philosophy*. Stanford University.

Habermas, J. 1999. *Between Facts and Norms*. Cambridge, Mass.: MIT Press. (1st ed. in German 1992.)

Hart, H. L. A. 1982. *Essays on Bentham*. Oxford: Clarendon.

Jeffrey, R. 1983. *The Logic of Decision*. Chicago, Ill.: University of Chicago Press.

Kant, I. 1996. Groundwork of the Metaphisics of Morals. In *Practical Philosophy*. Ed. M. J. Gregor, 37–107. Cambridge: Cambridge University Press.

Keeney, R. and H. Raiffa. 1993. *Decisions with Multiple Objectives: Preferences and Trade Offs*. Cambridge, Mass.: Cambridge University Press.

Kelsen, H. 1967. *The Pure Theory of Law*. Berkeley, Cal.: University of California Press. (1st ed. in German 1960.)

Kumm, M. 2007. Political Liberalism and the Structure of Rights. In *Law, Rights and Discourse. Themes from the Legal Philosophy of Robert Alexy*. Ed. G. Pavlakos, 131–66. Oxford: Hart.

Luhmann, N. 1973. *Zweckbegriff und Systemrationalität*. Frankfurt am Main: Suhrkamp.

– – –. 1974. *Rechtssystem und Rechtsdogmatik*. Stuttgard: Kohlhammer.

MacCormick, N. 1978. *Legal Reasoning and Legal Theory*. Oxford: Clarendon.

Mill, J. S. 1991. On Liberty. In *On Liberty and Other Essays*, Ed. J. Gray, 5–130. Oxford: Oxford University Press. (1st ed. 1859.)

Nimmer, M. 1968. The Right to Speak from Times to Time: First Amendment Theory Applied to Libel and Misapplied to Privacy. *California Law Review* 56: 935–67.

Nozick, R. 1993. *The Nature of Rationality*. Princeton, N. J.: Princeton University Press.

Pettit, P. 2002. Collective Persons and Power. *Legal Theory 8*, 443–70.

– – –. 2004. Groups with Minds of Their Own. In *Socializing Metaphysics*. Ed. F. Schmitt, 167–93. New York, N.Y.: Rowman and Littlefield.

Pinker, S. 1999. *How the Mind Works*. London: Penguin. (1st ed. 1997.)

Pollock, J. L. 1995. *Cognitive Carpentry: A Blueprint for How to Build a Person*. New York, N.Y.: MIT Press.

Pollock, J. L. and J. Cruz. 1999. *Contemporary Theories of Knowledge*. Totowa, N.Y.: Rowman and Littlefield. (1st ed. 1986.)

Postema, G. 1995. Morality in the First Person Plural. *Law and Philosophy* 14: 35–64.

Raz, J. 1978. Reasons for Action, Decisions and Norms. In J. Raz, ed., *Practical Reasoning*, pp. 128–43. Oxford: Oxford University Press.

Ripstein, A. 2001. *Equality, Responsibility and the Law*. Cambridge, Mass.: Cambridge University Press.

Sadurski, W. 2008. *Equality and Legitimacy*. Oxford: Oxford University Press.

Sartor, G. 2005. *Legal Reasoning: A Cognitive Approach to the Law*, Volume 5 of *Treatise on Legal Philosophy and General Jurisprudence*. Berlin: Springer.

– – –. 2006. Fundamental Legal Concepts: A Formal and Teleological Characterisation. *Artificial Intelligence and Law* 21: 101–42.

Sartre, J. P. 1993. *Being and Nothingness*. New York, N. Y.: Washington Square Press. (1st ed. in French 1943.)

Schauer, F. 1991. *Playing by the Rules*. Oxford: Clarendon.

Seidman, A. and R. Seidman. 2001. *Legislative Drafting for Democratic Social Change: A Manual for Drafters*. London: Kluwer.

Sen, A. 2004a. Elements of a Theory of Human Rights. *Philosophy and Public Affairs* 32: 315–56.

– – –. 2004b. *Rationality and Freedom*. Cambridge, Mass.: Belknap.

Simon, H. A. 1965. *Administrative Behaviour: A Study of Decision-Making Processes in Administrative Organisations.* 2nd ed. New York, N.Y.: MacMillan.

– – –. 1983. *Reason in Human Affairs*. Stanford, Cal.: Stanford University Press.

Soper, P. 2002. *The Ethics of Deference*. Cambridge: Cambridge University Press.

Stone Sweet, A. and J. Mattew. 2008. Proportionality Balancing and Global Constitutionalism. *Columbia Journal of Transnational Law* 47.

Thagard, P. 2000. *How to Make Decisions: Coherence, Emotion, and Practical Inference*. Cambridge, Mass.: MIT Press.

– – –. 2001. *Coherence in Thought and Action*. Cambridge, Mass.: MIT Press.

Tuomela, R. 2000. *Cooperation: A Philosophical Study*. Dordrecht: Kluwer.

Waldron, J. 1999. *Law and Disagreement*. Oxford: Oxford University Press.

Walton, D. N., C. Reed, and F. Macagno. 2008. *Argumentation Schemes*. Cambridge, Mass.: Cambridge University Press.