

2001

Reasoning with Rules

Joseph Raz
Columbia Law School, jr159@columbia.edu

Follow this and additional works at: https://scholarship.law.columbia.edu/faculty_scholarship



Part of the [Jurisprudence Commons](#), and the [Law and Philosophy Commons](#)

Recommended Citation

Joseph Raz, *Reasoning with Rules*, CURRENT LEGAL PROBLEMS, VOL. 54, P. 1, 2001 (2001).
Available at: https://scholarship.law.columbia.edu/faculty_scholarship/1484

This Working Paper is brought to you for free and open access by the Faculty Publications at Scholarship Archive. It has been accepted for inclusion in Faculty Scholarship by an authorized administrator of Scholarship Archive. For more information, please contact scholarshiparchive@law.columbia.edu.

Reasoning with Rules

J. Raz

What is special about legal reasoning? In what way is it distinctive? How does it differ from reasoning in medicine, or engineering, physics, or everyday life? The answers range from the very ambitious to the modest. The ambitious claim that there is a special and distinctive legal logic, or legal ways of reasoning, modes of reasoning which set the law apart from all other disciplines. Opposing them are the modest, who claim that there is nothing special to legal reasoning, that reason is the same in all domains. According to them, only the contents of the law differentiate it from other areas of inquiry, whereas its mode of reasoning is the one common to all domains of inquiry.

Those of a moderately cynical temperament will not be surprised at the popularity of the ambitious view among lawyers. After all, the more special the law is the more justified are the high fees which make the law inaccessible to all but the rich in so many countries. However, we will not engage in such sociological ruminations. Whoever stands to gain or to lose from the existence or absence of special legal modes of reasoning, the only question to be explored here is whether there are such distinctive modes of reasoning.

Not surprisingly, there is some truth in various rival positions. The most important point to make on the side of modesty is that the core of logic is not domain specific, nor could it be. Numerous arguments establish this point. I will sketch one. Rules of inference are not independent of rules of meaning, and of rules for the attribution of content to concepts and propositions. On the contrary, they are part of the factors that fix the meanings of terms and the content of concepts. The content of concepts is determined in part by the inferential relations that apply to them. That “a is green” entails “a is coloured” is part of what determines the meaning of “coloured” and of “green”. Therefore, if law, morality, physics and medicine, e.g., are each subject to different rules of logic then either they employ distinct terminology or use the same terms in different meanings. In fact while some terminology is special to different domains (for example, “quarks”, “resulting trust”) for the most part we use one and the

same language in all domains, and it would be preposterous to suggest that the same words bear different meanings when used by doctors, lawyers, bus conductors, accountants, etc.

However, not all modes of reasoning belong to the core of logic. Regarding the rest it is more plausible to assume that some differences between domains may exist. Much of what is often called “inductive reasoning” consists in following congeries of rules based on localised experience, or localised probabilities. Perhaps there are domain-specific modes of reasoning consisting in non-deductive rules of warrant. The law applies to all aspects of life. Therefore, legal arguments incorporate modes of reasoning from all domains of thought. But they may add to them. There may be some additional modes of reasoning special to legal thought.

By its nature the law has features which greatly affect the character of legal reasoning. I have in mind three: that the law of every country constitutes **a system of law**, that it consists, if not entirely then at least to a marked degree, of **norms or rules**, and that applying it and following it requires or presupposes **interpretation**. It would be wrong to suggest that these features are unique to law. They are shared by a number of the major religions and by other social organisations. They mark all institutionalised normative systems, and to a lesser degree they can be present in other normative domains as well. But they are central to the law, and they can be claimed to give a special character to legal (and religious, etc.) reasoning. The systematic nature of law, its dependence on rules and on interpretation • these three features of the law are closely inter-related, and being structural features, they can rightly be said to affect modes of reasoning common to legal reasoning. In other words, legal reasoning is just like any other reasoning, but in addition it also manifests features which express the structural, one may say formal, characteristics of the law.

As you see I find what is special in the law in some of its structural-normative features. Many would prefer to single out some of its social-institutional features. The two levels of analysis are not unrelated. The structural-normative aspects of the law affect its social-institutional character, and of course the influence works the other way round as well. There is no reason to be deterministic about the relations between these two aspects of the law, no reason to assume a one-to-one correlation between them. Their

relations are, and have long recognised to be a subject of great interest which we understand very incompletely. This is only one of the reasons for focussing today exclusively on the structural-normative features which I mentioned.

Of the three: **rules, interpretation and systematic character** the first, **rules**, is the most basic. Because rules play a central role in the law, it has a systematic nature, and interpretation plays a crucial role in legal reasoning. To understand this we need to understand what is so special about rules, and how they determine modes of reasoning.

How do rules figure in practical deliberation? How should they affect action, and the justification of action? It seems that rules are reasons for action. A person may well give the fact that his action is required by a rule as his reason for performing it, and an action may be justified because it conforms to a rule. Yet, rules are unlike most other reasons. Most reasons are facts which show what is good in an action, which render it eligible: It will give pleasure. It will protect one's health, or earn one money, or improve one's understanding. It will relieve poverty in one's country, or bring peace of mind to a troubled friend, and so on. What is the good in conforming to a rule?

This is the question I want to explore today: How can it be that rules are reasons when they do not point to a good in the action for which they are reason? I will call the phenomenon to be explained the opaqueness of rules.¹ I will concentrate on one type of rule, rules which are man-made, and which require conduct unconditionally.² My central case will be rules which are deliberately made as rules. What is true of them is true of other man-made rules, but may not be true of other reason-constituting rules.

Not all rules are reasons. Some "rules" mark regularities, as when we talk of what we do as a rule. Regularities may, but need not be, reasons. In other words, qua

¹ So, a reason is opaque, in the sense intended here, if a complete statement of it fails to show what is good about the action for which it is a reason. But, one may object, does not the fact that the reason is required by a binding (or valid) rule show what is good about it? Not so. As will be made clear below, that only shows that it is required, that we have a reason to perform it. It does not show in what way the action is good. In other words, being the action required by a binding rule is a normative, not an evaluative property of the action.

Admittedly, one may claim that the fact that there is a reason for an action makes it, *pro tanto*, good. But then here the evaluative follows the normative, rather than being its ground. The opaqueness of rules is that a complete statement of a reason does not disclose any good quality in the action which may explain why there is a reason for the action.

² Much of what will be said below applies, with self-explanatory modifications, to other types of rules, especially to other types of legal rules, and need not be discussed here.

regularities they are not reasons.¹ “Rules” are sometimes used to mark any normative proposition, that is any proposition stating what ought to be done², especially those which are natural expressions of common opinion. Such rules are not reasons either. They are statements (true or false) of what we have reason to do, but we have these reasons independently of the rules. Many rules are recipes. They are instructions how to do things: how to bake a cake, assemble furniture, impress an audience, untie knots, find one’s way out of a maze, or win in chess. Such rules state conditional reasons. That is they are not themselves reasons, and the reasons they state are reasons only for those who have some other reason. You have reason to bake this way if you have a reason to bake a cake, for that is how cakes are baked, etc.

Closer to the rules central to our concern are those which are sometimes called “constitutive rules”. The rules of chess, i.e. those which determine which moves are allowed, rather than how to win in chess, are said to be constitutive of the game. It is often said that constitutive rules are not reasons either. I think that they are, but they are (that is their existence constitutes) conditional reasons. They are reasons to behave in this way or that, reasons why you must behave this way and that, if you are to play chess, and therefore, if you have reason to play chess.

Many constitutive rules are themselves man-made. But not all are. Many of them are conditional reasons, whose condition can be avoided. One can avoid playing chess and at least in principle one can avoid a country, or a profession. Some constitutive rules are different. The rules of mathematics and logic are – in the old terminology – rules or laws of thought, constitutive of thinking (or of central kinds of thinking). We do not make them, and we cannot avoid them, at least we cannot avoid observing some of them, so long as we continue to think. But they too may be regarded as setting conditional

¹ Another way of making my point is this: Every fact can be a reason, in the sense of being part of an explanation of why some action, belief, emotion, etc. were appropriate to the occasion. But facts do play different roles in such explanations. Good-making or bad-making properties, and rules, are what I call “operative reasons” (see **Practical Reason and Norms**, (1975, 2nd ed. 1990, reissued by: OUP, Oxford 1999) pages 33-4, and what J. Dancy calls “favouring reasons”).

² Or what must or should, or may or may not be done, etc.

reasons, even though the reason cannot be avoided by anything less than opting out of all rational thought.¹

So, let me return to the rules that I will focus on. These rules are unconditional reasons, and they are man-made. I said that the question to focus on is how can rules be² reasons, given that on their face they do not point to any value in the action for which they are reasons, given that they are opaque? One reply is that the question is based on a false assumption. All normative statements³ (and they are rules, by one use of the term) are opaque. They state what we have to do. It is evaluative statements⁴ which state what is good about doing this or that. There is no puzzle about the opaqueness of normative statements. My puzzle derives from the claim that rules, some rules, are themselves reasons⁵, and not merely statements of what we have reason to do. Its solution is in the denial of the premise. If rules are never reasons, the puzzle disappears.

This last statement has to be acknowledged. However, we should start from the assumption that rules are reasons, for they are commonly treated as such. If, however, the assumption leads to conundrums and paradoxes it cannot be secure until must be laid to rest. In proceeding to explore the puzzle of the opaqueness of man-made rules I will be undertaking to defend the thought that rules are reasons, as well as to explain how they can be reasons in spite of their opaqueness.

The puzzle of the opaqueness of rules is made more acute when added to another: how can it be that people can create reasons just by acting with the intention to do so? This second question is bound to sound familiar. It arises in other cases, case not involving rules. Most prominently, it arises regarding contracts or agreements, but also,

¹ See on the rules of rationality my **Engaging Reason** (OUP, Oxford 2000) chapter 4, where these conclusions are qualified. Notice that the claim is not that we cannot avoid complete compliance with the rules of logic or of thought generally. We all violate them from time to time. The claim is that the reason that conditions their application is not readily avoided, that nearly all have that reason most of the time. These rules are conditional on a reason to live as a rational i.e. thinking being, while one is alive.

² Throughout the rest of this essay I will be using “reasons” to refer to unconditional reasons. Rules which are conditional reasons are opaque only in not disclosing on their face what are the reasons on which they depend. Once these are known the good they serve is made evident.

³ That is statements that we must, or should or ought , etc., do this or that.

⁴ That is statement that this or that have some properties which entail that they are good or of value.

⁵ Strictly speaking it is not rules but their existence which are reasons. I will, however, adopt the shorthand of referring to rules as reasons.

of course regarding promises, and all other voluntary undertakings. Moreover, undertakings and agreements too are opaque. Does the fact that I promised to stay awake tonight show that there is some value, some good in my staying awake tonight? At the very least we can safely say that if it does it is not in the way that the fact that I will be awake by the bedside of an ailing friend does. The similarity between man-made rules, agreements and undertakings helps in answering the second question and through it, it helps in answering the first, which is our main, question.

One analogy between agreements, undertakings and the sort of rules I am focussing on is that of all of them we can ask two different, and relatively independent, questions. The questions will be idiomatically expressed in different ways in different contexts, but they all are versions of the following two: (1) Are they binding, valid, rules (agreements or undertakings)? These questions are equivalent to “ought one to conform to them?” (2) Are they good, wise, justified rules (agreements or undertakings)?

A rule, or a promise, or an agreement can be binding, and it may be wrong to break it, it may be a valid reason for action, and yet it may be a bad rule, which should never have been made, and which should be changed as soon as possible. Rules, agreements and undertakings, I will say, allow for a potential **normative gap**, a gap between the evaluative and the normative, that is between their value, and their normative force.

Contrast this with “ordinary” reasons. That a novel is insightful and subtle is a reason to read it. We cannot here drive a wedge between the evaluative and the normative, between the two questions: “is it good?” and “is it binding or valid?” If being insightful and subtle are good characteristics of novels then they are reasons. There is no gap between being valid reasons and being good or of value, between the normative and the evaluative, as there is in the cases of rules, undertakings and agreements.

Why the difference? It is plausible to think that the explanation has to do with the fact that rules, undertakings and agreements are man-made. Since they are man-made they cannot be reasons unless they pass an appropriate normative test. Not everything which someone intends to create as a reason for himself (as with promises and personal rules), or for others (as with other kinds of rules) is such a reason.

This observation is correct, and may be relevant to an account of the place of rules in practical reason, but it cannot explain the separation between the binding character of rules and their goodness or justifiability. Why not have one test: If rules, agreements and promises are good and wise then they are binding, and if not, they are not? To explain why deficient rules, agreements and undertakings can, nevertheless, be binding we need to rely on something more than the fact that they are man-made.

To explain the normative gap we should start by noticing its contours and effects. First, the gap is not absolute, nor could it be. When we ask, “What makes rules bind?” the answer will revert to evaluative considerations.¹ The rules of the MASTERGAME chess club may be binding because it is better for the affairs of the club to be governed by its committee than to be organised some other way, or be left in chaos.

It is possible, of course, that that is not so, and that while the rules which give the club committee power to make rules for the governance of the club are binding, they are not good rules. It may be better to leave matters to a general meeting rather than have them decided by the committee, for example.² If so then this is yet another manifestation of a *limited* normative gap. This time the gap exists regarding the constitution of the club, that is those rules which set up the committee, and govern its rule-making activities. To explain why the rules of the constitution of MASTERGAME are binding, that is to explain why they have normative force, we have to rely on evaluative considerations. These may be, e.g., the desirability of not upsetting arrangements which, defective though they are, have governed the running of the club for some time, given that the harm that would be occasioned by a disorderly attempt to overturn them is too great.

¹ Some writers believe that the explanation of what is binding does, at least sometimes, derive from a range of considerations which are altogether independent of evaluative considerations. They are sometimes called deontic considerations. Nothing I say in this essay refutes this supposition. It shows, however, that the very binding nature of rules and the phenomena associated with it do not depend on the supposition being correct.

² I am skirting around a point which must be mentioned however briefly. I assume in the text above that rules can be justified by what I will call below content-independent considerations, relating, e.g., to their mode of origin even independently of the validity of any rules authorising this mode of the generation of rules. This seems to me true in principle. However, most legal system observes doctrines of the rule of law with the consequence that rules are valid in law only if they arise in accord with a legal rule about the proper ways for making law (some original constitutions being the only officially acknowledged exception). My observations in the text are not meant to challenge such rule of law precepts, only to indicate that they are not required by the vary nature of rules.

Were the harm caused by disregarding the rule establishing the committee small, and the advantage of organising matters some other way significant, and the prospect of securing the better way good, then the rules establishing the committee would not have been binding. Normativity is ultimately based on evaluative considerations, but in a way which leaves room for a normative gap.

How can that be? Notice the considerations which justify rules in my example: Take a rule saying that members are entitled to bring no more than three guests to social functions of the club. The considerations which establish that it is binding do not turn on the desirability of members having a small number of guests, nor on the desirability of members having the option to bring guests, but on the desirability of the affairs of the club being organised by the committee which laid down the rule. It is, in other words, an instance of what I call (following Hart) a content-independent justification. It is content-independent in that it does not bear primarily on the desirability of the acts for which the rule is a reason. Here we see clearly how rules differ from other reasons. The insightfulness and subtlety of a novel are reasons for reading it because they show why reading it is good. But the considerations which show why the rule is binding, i.e. why it is a reason for not bringing more than 3 guests, do not show that it is good not to bring more than three guests.¹ They show that it is good to have power given to a committee, and therefore good to abide by decisions of that committee. But that can justify a variety of rules: to have an annual championship competition, to admit new members by a

¹ There are various ways of trying to give a more formal characterisation of content-independent reasons, or justifications. Some of them may well yield somewhat different concepts. The text above suggests breakdown of transitivity of reasons as a mark of content-independent reasons, a point to which I return in the text below. In general, if P is a justification of Q, which is a justification of R, then P is a justification of R. But the justification of a rule is not, in and of itself, a justification for performing the action which the rule requires. It justifies giving the makers of the rule power to make the rule, and not more. Of course, indirectly it justifies the action the rule require, as being an action in accordance with a rule which is thus justified. But, unlike content-dependent justifications, it does not justify the action without these additional mediating premises. Put another way, the justification we consider, like all justifications of rules, is prima facie justification of the acts which fall under the rule. In as much as Roberta's act was to avoid bringing more than tree guests it was justified, because that is what the rule required of her. Overall, of course, matters may be different, for her act has other features as well, some of which may condemn it. All prima facie justifications are description sensitive: The lack of transitivity is that the reasons for the validity of the rule are not in themselves reasons for performing the act required by the rule, as described in the rule. In this rule are the exception to the norm that the reason for a reason for an action is a reason for that action (under the same description). To fully explore the credential of this test goes beyond the ambition of this article. It will require, e.g., a way of distinguishing conditional reasons from content independent reasons.

simple majority in a postal vote of all members, to levy a membership fee of £50 a year, etc. Moreover, typically, though not without exception, the very same considerations could justify contradictory rules. They could justify a rule saying that membership will be confined to residents of the district, and a rule that membership will not be limited to the residents of the district, etc. They are, in this sense, content-independent.¹

Do you feel that I am going around in a circle? I stated at the outset that my aim is to explain how rules can be reasons even though they do not show that the acts which they require are desirable or of value, even though they are opaque. And now I have proclaimed as a great discovery the very same fact as if it explains why rules display a normative gap, that is a gap between the normative and the evaluative.

Nevertheless, I believe that we are making progress. First, notice that the thesis about the content independence of the justifications of rules² goes further than opacity, the feature I set out to explain. That was that rules are reasons even though they do not show the value of the actions for which they are reasons. The content-independence thesis makes things worse by adding that even the justification of rules does not bear on the desirability of the actions for which they are reasons.³ Doubling the puzzle makes it easier to solve.

Second, by showing the centrality of the feature to be explained we improve our understanding of rules. We can see now that the opaqueness of rules is a result of their content-independence and their content independence is an aspect of the normative gap

¹ It is important not to confuse content-independence with unlimited jurisdiction. A justification can be, and typically will be, both content-independent and limited. The club committee cannot be authorised to commit or order others to commit murder, etc. It is, if you like, content-sensitive in that it does not allow for any content whatsoever, while being content-independent, in not being specific to one rule. What makes a justification content-independent is not whether it can justify more or less possible rules, but that the considerations which constitute it do not bear on the desirability of having rules with the content of the rules which they can justify. That there are other rules which, because of their content, the justification does not show to be binding is immaterial.

² Strictly speaking not rules, but their justifications are content independent. For the sake of brevity I will, however, refer also to rules as content-independent.

³ Perhaps I should add here, “under that description”. The justification of rules bears on the desirability of actions required by the rules when they are described as “actions required by the rule”, etc. That is not a description of action which can be used in the formulation of the rule (“the rule is that one ought to do whatever this rule requires”, even though true, is not a way rules can be informatively formulated). Put precisely, then, the claim is that the justification of the rule

rules display between the normative and the evaluative. By tying up all these features together, we show them to be robust and central to rules. Of course, we still have to explain them.

But what does it mean that we have to explain them? After all rules are what they are. The task of analysis is to explain their central features. Having isolated three: their content independence, their opaqueness, and the normative gap, and having shown that they are interconnected, we can go on to describe other of their features. But what sort of request is it that we explain them? What more need be said? Of course, analysing rules in these terms does not mean justifying them. We have not shown which rules are binding, nor that there are any rules which bind. But surely, the justification of rules is not our task. The problem is not that it is a normative task. The problem is that it is impossible to justify rules in general. We can consider the justification of this or that rule or group of rules and that cannot be done outside a specific context.

To this we should reply: yes and no. True there is no question of providing a general justification of rules. Some are not justified and are not binding. Possibly, those that are binding are to be justified in arguments of varying patterns, which cannot be exhaustively described in advance. And yet more needs to be said to explain the opacity of rules. At the very least we must show how it is possible for people to believe that rules are binding. For without such an explanation the rest of the analysis is suspect. People do make mistakes and many believe in the validity of many rules which are not valid at all. Yet, unless we can show how it *can reasonably* appear to people that some rules are valid the analysis will be in jeopardy. It is unlikely that so many people, perhaps almost everyone, have normatively similar beliefs, all of which are totally irrational.

Showing how it could be plausible for people to believe that some rules bind is likely to amount to showing, at the very least, that it is possible for rules to bind, and also to pointing to some circumstances under which they do bind.

So we are back with our question: How can rules be reasons when they are opaque? To understand this we need to find a focal point which will open the way from delineating the features of rules to their possible justifications. That focal point is likely

does not bear on the desirability of any action required by the rule, under any description which can be used to formulate informatively the content of the rule.

to be their content independence. The content independence of rules readily explains their opaqueness. It also explains the normative gap. Since the justification of the validity of a rule does not depend on the value of the act the rule is a reason for a normative gap can open. The question then is how can justification be content-independent?

For a content-independent justification to be possible there must be reasons for an agent to behave in a certain way other than the value of the behaviour in question. Let us move away from rules (undertakings or agreements) and take an example of content-independent justification for a particular act on a particular occasion. Suppose that you are asked why did you walk to work along Marylebone High St. today, rather than along Baker St.? You may say that there it has more attractive shops and buildings and is less noisy than Baker St. This would be a content-dependent explanation. But suppose you say: because I always do. That reply is content-independent. It shows nothing good about walking along Marylebone High Street. The problem with it is that it is not clear in what way it points to a reason at all. Why should one do what one always does? Depending on the circumstances an explanation may be readily forthcoming. It may be, for example, that you find choosing a route every morning (should it be Baker St? Or Gloucester Place? Or Marylebone High St.? Or Upper Montague St. followed by Montague Sq.?) tiresome. Sticking to a routine is a way of not spending time and energy deciding, when the difference in the merits of the different realistic options does not appear to justify the effort and worry that deciding would involve.

It is obvious that the justification is not entirely content-independent. Had the margin of merit between the different options been greater the reason for choosing the route might have been inadequate. Yet clearly the reason is content-independent in not bearing on the quality of the route chosen. We can therefore examine this case, and consider how some of its features can apply to rules. What enables the content-independent justification to work is the existence of a personal routine of going to work along Marylebone High Street. The very same reason for the desirability of a quick decision, free from detailed consideration of the different options, would have been to no avail but for the existence of the routine (or something to take its place).

As a rule, normative justification, and justification in general, are transitive.¹ If A justifies B and B justifies C then A justifies C. So if there is reason to read the novel because it is a good novel, and if it is a good novel because it is insightful and subtle, then that it is insightful and subtle is reason to read it. And so it goes on. If the novel is insightful because it vividly sheds light on a deep emotional conflict which is usually denied and misunderstood, then that it sheds such light is a reason why it is good, and therefore also a reason for reading it. The opacity and content independence of rules mean that transitivity does not hold. That it is good to uphold the authority of the committee is a reason for the validity of its rules, including the rule that one may not bring more than three guests to social functions of the club. But the desirability of upholding the authority of the committee is not a reason for not bringing more than three guests (not, that is, under this description).

The lack of transitivity in justification seems to me to be among the most important features of rules. They are not, of course, alone. Undertakings and agreements display the same feature. The promise to go to Paris is a reason to go to do so, but the reason the promise is binding (for example, the desirability of people being able to bind themselves) is not a reason to go to Paris (nothing about the value of being in Paris or of travelling there). It justifies going to Paris only indirectly. The same justification could have justified staying away from Paris, had that been what one promised to do.

Some thirty or more years ago much philosophical ink was spilt in debating whether rules make a difference. One party insisted that having rules as a reason for action makes no difference for the guidance and evaluation of action, whereas the other party argued that it does. Put in the terminology I have here developed, those who denied that rules can make a difference relied on the general transitivity of reasoning to argue that rules cannot yield different conclusions from those which would follow without them, for they can only transmit the force of reasons we have anyway. Those who opposed them knew better, but only a few of them realised that the explanation lies in the breakdown in transitivity which is a result of the content-independence of the justification of rules, and of their opaque character. That is why rules, at least man-made

¹ As remarked above, I assume that the justifications concerned are prima facie justifications and therefore that they are description sensitive, in the same way that explanations are description

rules, can make a difference to practical reasoning, and why when valid they are rightly said to be reasons in their own right rather than merely statements of reasons we have independently of them. This then is what we may call the autonomy thesis. It says that rules make a difference. If valid, they constitute reasons which one would not have but for them. While the considerations which justify a rule exist independently of the rule, they do not constitute the same reason for action that the rule constitutes.

You will realise of course that in proceeding with the analysis and identifying further features of rules we are continuously restating the question, or rather stating further and closely related questions. We have not yet answered it. I believe, however, that we have finally found the question which provides the best avenue towards an answer. To explain how can it be that a rule can be a valid reason, even though it is opaque, we need to explain how it is that rules can be autonomous. Once we explain how there can be reasons which are autonomous, in the sense explained, we will have the answer to our question, we will understand why reasons can be opaque.

Some people think that the very idea of autonomous reasons is incoherent, or at least that it is logically impossible for such reasons to be valid. After all, the autonomy of rules is an expression of the breakdown of transitivity and that means that the force of some reasons does not carry through. It is thwarted by other considerations which in themselves are neither reason for the action concerned, nor against it. Think of the chess club example again. There may be a reason for a member of the club to bring more than three guests to a particular function. He cannot do so because the rule forbids him. The rule is a reason against the action, so everything looks ordinary. But the reason for the validity of the rule is the good of having the committee regulate the club. And that is not a reason against bringing four guests to the function. Small wonder that rules are opaque. They do not show what is good about the action they require for they do not in fact rely for their validity, for their force as a reason, on anything which makes the action they require good. (Notice that it is the same with agreements and undertakings: the reason is that you are committed, rather than that what you are committed to has any value). So we have in the rule a putative reason against four guests which does not depend on there being anything wrong with four guests but which stops one from acting on reasons for

having four guests. Is not that irrational? Does it not follow that there cannot be binding rules?

As you know, some have thought that to dissolve this apparent paradox one needs to invoke considerations of a different order altogether, considerations which are not subsumable under the category of the good. They are sometimes identified as deontological considerations. I will say nothing to refute the thought that there are normative considerations which cannot be subsumed under the good. I do not believe, however, that they need to be invoked to explain the autonomy or any other feature of rules, or to account for the possibility of plausibly believing in the validity of some rules. So far as the considerations we consider here are concerned all can be explained on the assumption that the normative derives entirely from the evaluative, that reasons depend exclusively on values.

Long ago¹ I suggested that rules, rules of the kind we are considering, are not simple reasons but a structure of interrelated reasons. They are, first, reasons for the acts they prescribe, but they are also, second, reasons not to act for some competing reasons. The rule that no more than three guests may be invited by a club member is, first, a reason for members not to invite more than three guests, and also, second, a reason not to act on some reasons for inviting a fourth guest. To introduce some more terminology, I call the second kind of reason rules are exclusionary reasons (for they exclude action for some reasons) and I call the rule itself a protected reason, for the reason for the action it prescribed is protected by these exclusionary reasons.

Any complexity of this kind is unwelcome. Why do I think that it helps to explain how rules function? Take again our humble, and long suffering, example. I repeatedly said that the reason for the validity of the rule is that it is best if club affairs are regulated by the committee which made the rule. That is presumably because, on the whole, if members follow the judgement of the committee their actions will track reason better than if they act on their best judgement without taking account of the judgement of the committee. Usually when this is the case it is so through a combination of two factors. First, the good judgement of the committee. And, second, the fact that it can secure desirable co-ordination among people, which, left to their own devices, the members are

less likely to secure. These factors are not a reason against inviting a fourth guest. But they are a reason for not second guessing the decision of the committee. So if the committee, having had the opportunity to weigh the pros and cons of imposing the no-more-than-three-guests rule, has approved it then all members have reason not to challenge that judgement, and that means that they have reason not to act on the reasons for or against bringing a fourth guest. Rather, they should regard the rule as displacing the reasons which the committee was meant to consider in issuing the rule. That is what I mean when I say that the rule is an exclusionary reason.

Obviously it is also a reason for the action required by the rule. In all it is a protected reason for that action. If this example can be generalised, and I believe that it can, then we have here an explanation why rules are opaque, content-independent, autonomous reasons for action, and how they can be rational even though they violate the transitivity of reasons. I finally call this an explanation, for it includes an account of how it is possible for rules to be valid. More than that, it makes it clear that often rules are valid protected reasons. Q.E.D.

Armed with this skeletal analysis we can make its meaning and implication clearer by considering a few of the contexts in which we may find some valid rules. An obvious and very important context consists of all those cases where there is a good case for enabling people, organisations, or other agents, to pre-commit themselves. To achieve its purpose pre-commitment must achieve closure: once the commitment is made it is to be adhered to. Absolute closure would mean that the commitment is to be adhered to however much circumstances change, whether or not one changes one's mind, and whether or not one realises that one made a mistake in making the commitment. Arguably, no valid commitment can be absolute. There are no circumstances which would warrant absolute commitments. But a pre-commitment need not be absolute. It may be designed to achieve closure so long as the situation does not change radically, or closure from change of mind motivated by considerations of one kind or another, and so on. But how can any pre-commitment be rational? Does not reason require reassessment of the proposed action just before it is undertaken? My suggestion is that it works because when there is a case to enable an agent to pre-commit himself, the pre-

¹ In **Practical Reason and Norms**, and in subsequent writings.

commitment constitutes an exclusionary reason for not acting on those considerations that it, the pre-commitment, was meant to exclude (subject to emergencies, or whatever other exceptions reason imposes on the power to pre-commit).¹

Pre-commitment is useful, sometimes even necessary, in a large variety of contexts and for a number of reasons. It can facilitate forward planning, it can enable co-ordination, it extends people's abilities to form ties with others, and also their ability to enter into mutually profitable arrangements with others. Finally, it is a pre-requisite of many arrangement based on a division of labour between different agents. Pre-commitments take many forms: promises, vows, and agreements are familiar cases. So are personal rules (to have only two cigarettes a day, or to go jogging daily) and decisions.

Some people regard rules made by authorities as pre-commitments. They regard communities as agents and governments as their agents. Laws passed by the government are seen as commitments of the community. This may be the right way to think of some communities, but for reasons we need not go into here, it seems no more than a fiction when considering most states today. This does not mean that legal rules cannot be binding. Governments have useful, and in the conditions of our life, essential functions in securing co-ordination, overcoming collective action problems, and utilising hard to master information for the benefit of their communities and beyond. There is much more to be said about rules of political communities, such as legal rules. I will leave the subject with a couple of comments on rules and disagreements, and rules and division of power among legal institutions..

There are many sources of disagreement. Hobbesian and market-oriented theories tend to regard all disagreements as expressions of conflicts of interests. This is an exaggeration, but no doubt many disagreements are a result of conflict of interests. Many political theorists of left-liberal persuasion tend to emphasise the prevalence of disagreements about morality and about values more generally. The disputes about abortion, surrogate motherhood, the rights of gays and lesbians, and many others are predominantly, though perhaps not exclusively, such disagreements. But there are other

¹ There are alternative explanations, but I believe that they either are equivalent to mine, or fail to account for all the aspects of pre-commitments.

sources of societal disagreement. Disagreements can arise between people who share the same values and whose interests do not conflict. They can arise, of course, because of factual disagreement: economics, e.g., is far from a secure science, and disagreements about the likely effects of various social or technological changes lead to disagreements about governmental policies. Finally, disagreements are liable to arise where reason suggests that people should co-ordinate their conduct but it allows for various schemes of co-ordination, without judging between them.

The actual situation is much more complicated than this thumbnail sketch suggests. Not only are the causes of disagreement often mixed, but in addition, it is often far from clear what they are. People faced with a problem of co-ordination where reason under-determines the solution may believe that the problem is of disagreement over values, etc.

Where the law is concerned another complication looms large: the co-ordination the law aims to achieve is multi-layered. The co-ordinates not only the conduct of individuals but that of legal institutions as well. These institutions themselves are inter-related in complex ways: Electorates, legislatures, executive and administrative bodies, central and provincial bodies, as well as complex hierarchies of courts and tribunals, all have to function in an orderly and co-operative fashion.

These brief reminders of what we all know connect with the analysis of rules, as autonomous, opaque reasons. To some degree the sensible reaction to disagreement is to avoid common policy and common action which relies on the disputed beliefs. Often that is neither possible nor desirable. Yet common action requires some measure of agreement, at least on the part of the officials who may be involved in implementing the disputed measures. Moreover, in general it is desirable that common actions shall command agreement. Rules allow agreement in the face of disagreement. They do so by allowing for agreement on the decision procedure in spite of disagreements about the measures it should yield, or because of agreement on measures, in the face of disagreement about their justification.¹ Again we can see how rules are the inevitable backbone of any structure of authority, of which the law is a paradigm example. They

¹ A point well highlighted by Cass Sunstein in various publications. See, for example, Sunstein, C. R.. **Legal reasoning and political conflict**. (New York: University Press 1996)

also begin to indicate how the centrality of rules, and of the factors which justify it, make interpretation crucial to much legal reasoning, and make much interpretation concerned with giving effect to the systematic nature of law. These too are concomitants of the fact that the law is a structure of authority. But that is a matter for another occasion.