

Therefore, I am inclined to conclude that Raz's view of reasons as aspects of things of value can only be sustained if we reject the entailment claim—at least for 'engaging with' reasons (possibly not for reasons of respect). This would, of course, call for a reinterpretation of the normativity of reasons,⁴⁵ which I cannot even attempt to begin in this essay.⁴⁶

⁴⁵ Jimmy Lenman urged me to consider the possibility that a reason statement might not entail an ought statement, but rather a permission to act accordingly. This suggestion is, of course, tempting, because it preserves the link to other deontic notions, such as 'ought'. 'Ought' could then be defined as not being permitted not to do something. Yet, this suggestion seems wrong to me for the following reasons: (1) The reasons we deal with are *pro tanto* reasons only. However, there seems to be no *pro tanto* permission. Permissions are always 'all out': saying that you are permitted to ϕ is tantamount to saying that there is nothing that counts against ϕ -ing. (2) Reasons, it is often said, *count in favour* of actions. Permissions, on the other hand, are only negative: an action is permitted, if nothing counts against it. Hence, there is no contradiction in saying that while you are permitted to ϕ , there is no reason to ϕ ; i.e. while nothing counts against it, nothing counts in its favour either.

⁴⁶ An earlier version of this paper was presented at a conference on 'The Ethics of Joseph Raz' in March 2002 in New York. I am grateful to the participants of the conference for comments and critical discussion. I am also much indebted to Alon Harel, Victoria McGeer, Jimmy Lenman, Joseph Raz, Randall Rose, Jay Wallace, and, especially, Philip Pettit for extremely helpful written comments. At different stages Thérèse Björkholm, Elisabeth Herschbach, and Sherry Orbach helped me with editing the manuscript. I owe a great deal to their suggestions, and I am very thankful for their generous help. Finally, I would like to thank the participants of the Philamore group for a very instructive and insightful discussion.

7

The Truth in Deontology

Philip Pettit and Michael Smith

Let a deontological constraint be a pattern of action that people ought to adopt, even when doing so will not be for the best; more precisely, even when it will not be for the best in promoting conformity to that very course of action on the part of agents overall, or perhaps on the part of that agent over life as a whole. Deontologists argue that such constraints have a widespread and fundamental presence in moral thought. They govern how we should behave across a wide front, whether in our dealings with ourselves, with one another, or with the rest of nature. And the authority they have over us is not derived from the impersonal value of satisfying them, as teleologists or consequentialists would argue; it is in that sense fundamental.

With characteristic sensitivity to the ways in which we reason about what to do in various circumstances, Joseph Raz gives a sympathetic hearing to the claims of deontological constraints.¹ He thinks that deontological patterns of thought are a mark of agents we tend to admire. And, going still further, he sets himself against the idea that consequentialism might be able to explain their authority by pointing us towards neutral goods whose promotion would explain their hold over us. Consequentialists might prescribe the satisfaction of the constraints in normal cases, where satisfying them is for the best, but would take the opposite line in perverse circumstances where it is not for the best. Raz sides with deontology in maintaining that at least some constraints have a fundamental, teleologically underivable status.

¹ Joseph Raz, *The Morality of Freedom* (Oxford: Oxford University Press, 1986), chs. 11–13.

Our aim in this paper is to argue that deontological constraints do indeed have a widespread presence in moral thought, being relevant in an endless range of cases. We try to show that the practice of deliberation is at the core of social life, and that it intrudes deontological constraints to which we must all claim to defer. In doing this, we take a different approach from Raz, but one that he may find complementary; we associate constraints, not with the received conception of admirable agents, but with the framework within which we conduct many of our interactions with others. But towards the end of the paper we break more decisively with his position, arguing that the case made for the widespread presence of deontological constraints does not argue for giving them a fundamental status; it does not rule out consequentialism. There is truth in deontology, but that is not to say that deontology or non-consequentialism is true.

The paper is in five sections. In the first, we show how the simple activity of playing a game can impose deontological restrictions on players; in the second we describe the practice of deliberative exchange, emphasizing the parallel with game-playing activity; in the third we show how deontological restrictions are associated with deliberative exchange; in the fourth we argue that these restrictions—unlike those involved in game-playing—are hard for people generally not to countenance; and then in a brief, concluding section we show why the truth in deontology that is thereby established does not necessarily mean that deontology is true.

1. Deontological Games

One context in which deontological attitudes are at home is that of the game. Think of the game in which different individuals or groups compete for what is deemed to be victory. Think of chess or cricket or football or whatever. In any such game there are rules that people countenance as governing the things that they may and may not do in pursuit of victory. In particular, there are rules such that not to follow them is, by received criteria, to fail to play the game. These are sometimes described as rules that are constitutive of the game, as distinct from rules that regulate how the game is best played, or fairly played, and so on.² They are design

specifications on what has to be done among the parties involved in order for them to count as playing the game. Let those specifications be implemented, and a game of the relevant sort will take place; let them fail to be implemented, and it will fail to take place.

Since these truths about the constitutive rules of a game are going to be available to all, and since game playing is an intentional activity, the success of players in conforming to the rules will manifestly be within their intentional control. They will each conform intentionally to the rules. That being a fact accessible to everyone, they will each expect everyone to conform intentionally. And that in turn being a fact accessible to everyone, they will each expect everyone to expect this, and so on. That they each conform intentionally, in other words, will be a matter of common or mutual or shared awareness between them.³

Other patterns of common awareness will emerge to supplement this common awareness of intentional conformity. In particular, it is almost inevitable that people will come to share in a common awareness that violating a rule is a fault within the context of a game; that each naturally relies on others not to violate rules in that way; that violators will attract a negative reaction from those who relied on them; that this is something that any violator will have been in a position to predict; and so that no one who violates a rule can protest against the negative reaction of others. Violators must be taken, as the vernacular has it, to have licensed or permitted complaint about their violation of the rules; after all, they advertised themselves as game-players who were aware of the rules and aware of the reliance of others.⁴

These claims are all fairly plausible, given familiar facts about human psychology, and given that those facts are generally accessible, that their accessibility is itself generally accessible, and so on. Without trying to document them in detail, we are going to take them as granted in this paper. They entail that the constitutive rules of a game will get established among players as rules that they recognize as binding and rules that can be invoked in licensed or permitted complaint about any breach. The rules will serve within the context of the game as norms that all endorse as guiding ideals and as ideals in the name of which they allow themselves to be judged.

³ David Lewis, *Convention* (Cambridge, Mass.: Harvard University Press, 1969).

⁴ Margaret Gilbert, *On Social Facts* (Princeton: Princeton University Press, 1989); T. M. Scanlon, *What We Owe to Each Other* (Cambridge, Mass.: Harvard University Press, 1998).

² J. R. Searle, *Speech Acts: An Essay in the Philosophy of Language* (Cambridge: Cambridge University Press, 1969).

The thing to notice now about the constitutive rules that serve in such a normative role is that they identify faults after a deontological pattern. Let us think of a violation of the rules as consequentially productive when it promises to reduce breaches generally, or even to reduce later breaches on the part of the agent. The deontological character of the constitutive rules of a game appears in the fact that a failure to abide by those rules counts as a fault, even if it is consequentially productive in reducing the overall number of breaches. Games would be impossible if players were allowed to duck in and out of compliance with the constitutive rules, depending on whether such violations were consequentially productive in that sense. So the rules which govern game-play serve as deontological norms for the players.

Putting this in other terms, the rules represent binding constraints on how the players are to interact with one another, not attractive constraints whose fulfilment they should try to promote in aggregate.⁵ They determine that certain patterns of behaviour constitute wrongdoing, and they indict wrongdoing even in the case where the wrongdoing allows less wrong to be done overall. Thus they mark a distinction between wrongdoing and allowing the doing of wrong. It is against the rules of the game to violate them; it is not against the rules of the game not to reduce or eliminate violations by others or by oneself at a later time. Players are required to conform to the rules, but beyond the requirements of those rules, everything is permitted within the game.

Given that the context of a game may give rise in the manner sketched to essentially deontological norms, a natural question is whether something of the kind could be the source of those deontological requirements that many see as having an intuitive and intelligible hold on us in day-to-day life. Can we think of any game-like context which is relevant in human interactions generally—not just when we take up chess or cricket or whatever—and which might explain the fact that certain normative prohibitions, such as those on violence and lying, are often assumed to have a deontological character? Can we think of any game-like venture that is not discretionary in the way in which chess and cricket are discretionary and that deontologically shapes the expectations to which we hold one another in ordinary moral talk?

⁵ Robert Nozick, *Anarchy, State, and Utopia* (Oxford: Blackwell, 1974).

We think that there is such a practice that we human beings undertake in common, without paying much attention to the fact, and that it can serve to explain the hold of certain intuitive prohibitions. It underwrites deontological restrictions in the manner of the constitutive rules of a game, but the restrictions underwritten have a much deeper hold on us than game-playing rules. Where we can choose or not choose to play a game like chess or cricket, we can hardly conceive of exiting from this practice. The restrictions associated with it may not be categorical, in the sense of being conditional on nothing; after all, they are conditional on participation in the practice. But they approximate categorical restrictions so far as that on which they are conditional is a mode of interaction that humans cannot easily avoid.

2. Deliberative Exchange

The practice we have in mind is that whereby human beings relate to one another as deliberative or discursive partners, seeking and offering counsel about what one or more of their number should hold or do, consulting about patterns of reciprocation that may appeal to each, or taking counsel together about what they should jointly do or hold. It may involve the assembly that has to make a collective judgement or decision. Or it may involve two or more people working out a pattern of mutually beneficial exchange. Or it may just involve two or three people comparing notes about what one of them should think or do in this or that circumstance.

In any such exchange people claim to direct one another's attention to reasons that are relevant to what they should think or do. Reasons in this sense are either considerations that explain why a decision or judgement is right or considerations that explain why they are right for someone with certain pre-existing commitments.⁶ Deliberative practice purports to facilitate the access that participants have to such reasons, including reasons that may support reciprocal or joint action. It directs people to presumptive, hitherto unregistered facts and values, in particular facts and values that bear on what they ought to believe or do, individually or collectively.⁷

⁶ John Broome, 'Reasons', ch. 2 in this volume.

⁷ Philip Pettit, *A Theory of Freedom: From the Psychology to the Politics of Agency* (Cambridge and New York: Polity and Oxford University Press, 2001); Philip Pettit and Michael Smith, 'Freedom in Belief and Desire', *Journal of Philosophy* 93 (1996), 429–49.

What are the marks of deliberative practice? Deliberative exchange always involves people seeking to influence one another, or at least accepting that what they do may have such an influence. But the influence they seek in deliberation is not that achieved by putting obstacles or penalties or rewards in one another's path. They look for influence in deliberation by trying to affect one another's beliefs, including beliefs about the presence of such factors. But again, they don't look for influence by inducing beliefs in any old manner. They look for influence in the characteristically open fashion of communication, where it is evident that speakers intend to elicit certain beliefs, in particular beliefs that they themselves purportedly endorse; evident that they intend the audience to be aware of this; evident that they intend them to form the beliefs in virtue of that awareness; and so on in the usual Gricean explication of these things. More specifically still, they look for the sort of influence provided by sincere communication, trying to get the audience to believe things that they do actually endorse; to deceive others, after all, is hardly to deliberate with them.

Is sincere communication enough to ensure deliberation? Whether I seek to communicate sincerely how the world is in some respect or how I am in some respect—say, in respect of belief or desire or intention—I must begin with the sincere communication of what I believe: that the world is thus and so, or that I am thus and so. I must sincerely communicate what I believe—say, that *p*—both in the sense of communicating the content of what I believe and the fact that I believe it; I must put before you both the assumed fact that *p*, and the fact that I believe that *p*, though of course you may not be persuaded of either: communication is not necessarily successful communication. The question which arises now is whether the sincere communication of what we each believe, understood in this way, is always going to count as an instance of deliberating with one another.

The sincere communication of belief is going to count as deliberating with one another, intuitively, only so far as it involves exploring the reasons that we each have for believing, desiring, and doing things. It will involve trying to spell out the available options for belief or intention or action before us, individually or jointly, and then searching out the range and the relevance to choice of those considerations that we countenance as reasons—or, going deeper still, that we can give one another reason for countenancing as reasons. Exploring reasons with one another in this sense is an inherently epistemic project, though it may promise a practical pay-

off in reciprocal or joint action, and may be pursued because of that promise. It contrasts intuitively with the attempt to engineer one another's reasons and manipulate one another's responses.

The engineering of reasons and the manipulation of responses may occur via insincere communication: deception, after all, is a standard means of manipulation. But it turns out that it may also materialize via sincere communication. Thus it is necessary to distinguish between the sincere, epistemic form of communication involved in deliberating with one another and the sincere, manipulative communication that contrasts with it.

The paradigm of sincere, manipulative communication is the threat. We are not thinking of the deceptive threat—not the bluff threat, as it might be called—since it does not count as sincere. And we are not thinking of the predictive threat in which I warn you that if you do *B*, I will do *A*, since your doing *B* would provide me with an independent reason for that response: a reason that would have obtained, even if I hadn't spoken to you about it. Nor, as will appear later, are we thinking of the mesmerizing threat that saps the nerve and the rationality of the audience. We have in mind the threat I might make to penalize your doing *B* by doing *A* in punitive response, where I have no independent reason to respond in that way: no reason independent of the reason I have to declare my intention and push you away from *B*.

This threat is a manipulative attempt to engineer or rig the reasons that are relevant to what you should do, though it may also involve exploring those reasons as they exist after the rigging. Such coercion does not count, intuitively, as an exercise of deliberation. When I threaten you in an attempt to coerce a certain response, I tamper with factors that deliberation would limit me to exploring: the options before you and the considerations available to adjudicate between those alternatives. What I do involves changing the options before you—replacing the *B* option with the option of doing-*B*-and-then-being-penalized—or, equivalently, triggering a new reason in the attempt to influence the choice: creating a penalty that constitutes a reason not to choose *B*. It involves engineering your reasons, as we put it, not exploring them—or at least not just exploring them—with you.

Not only does deliberation require the sincere communication of belief, then; it also has to have an epistemic character, being an attempt to explore your reasons with you. Does this stipulation rule out offers as well

as threats? The deceptive offer—another sort of bluff—will be ruled out as an insincere form of communication. The predictive offer will be ruled in, on the grounds that allowed us to rule in the predictive threat. And the mesmerizing offer will be ruled out on the same grounds as the mesmerizing threat; more on this in a moment. But what about the more regular offer: the offer to reward your doing something, say *D*, by doing *C* in response, where the reason for making the *C* reward available, as well as the reason to let you know about it, is to motivate you to do *D*? Do we have to rule out such regular offers, on the grounds that they are just as manipulative as ordinary threats? Do we have to rule out those offers in which mutually beneficial co-operation may originate, for example, as when you and I make agreements and contracts?

Happily, it turns out that we do not have to take any such drastic line. The offer of a reward that you have no means of refusing—for example, the offer to broadcast your success, should you achieve some result or other—will clearly be a way just of engineering your reasons and will have to fall to the stipulation that rules out strategic threats; it represents a form of manipulation, nothing more. But most offers are not like that, since they are offers of rewards that can be refused by those to whom they are made. And this possibility of refusal makes all the difference.

What I do in making an offer of a reward for your *D*-ing—strictly, a reward that does not mesmerize you, say by connecting with an urgent need—is not, or at least not directly, to engineer the reasons you have in respect of the options before you. I do not interfere with any of the existing factors which deliberation would limit me to exploring: your options or the considerations relevant to the choice of those options. Rather what I do is to expand the realm of available options to include, not just *D*-ing and *E*-ing, but also *D*-ing-and-then-collecting-a-reward. Making such a refusible offer is friendly to deliberation so far as it does not change any of the factors that deliberation would limit me to exploring: your existing options and the considerations relevant to choosing between them.⁸

Making such a refusible offer may be cast, indeed, as itself an instance of deliberating with you. It is designed to provide information, after all, on a new option that you may be very glad to consider: that of *D*-ing-and-then-collecting-a-reward. I might have invited you in clear epistemic vein to

consider, and perhaps enquire into, whether a third party would be willing to reward you for *D*-ing. In the refusible offer I invite you to consider whether I would be willing to do this, and, anticipating your enquiry, I say that yes I would.⁹

This line fits with our intuitive conception of deliberation. If you seek my deliberative advice, inviting me to explore the reasons you have to do this or that, then you may reasonably complain if I directly engineer those reasons by manipulatively putting penalties in place. But you can hardly complain if I respond by expanding the range of options, arranging things so that not only can you do this or that, you can also take one of the existing options and enjoy a reward at the same time; certainly you can hardly complain if the offer is of the non-mesmerizing sort. This is an exercise in which I may seek to influence you, and perhaps enjoy success in the enterprise, but it does not involve engineering your reasons in the manipulative manner of a threat.¹⁰

To sum up the line of thought so far, deliberative exchange involves the sincere, epistemic communication of belief in which people attempt to explore with one other the reasons they have to form various attitudes or perform various actions, whether individually, reciprocally, or jointly. Is this enough, then, by way of a definition? Not quite, though it comes close. For the sincere, epistemic communication of belief may sometimes have the contingent effect of exposing the audience to pressures that rig the reasons available—in the manner, roughly, of coercion—or that reduce the audience's capacity to

⁹ Many writers, following Habermas, mark a difference between the activity of bargaining with others and debating with them. See Jürgen Habermas, *A Theory of Communicative Action*, vol. 1 (Cambridge: Polity, 1984) and *A Theory of Communicative Action*, vol. 2 (Cambridge: Polity, 1989). And see Jon Elster, 'Arguing and Bargaining in Two Constituent Assemblies', *University of Pennsylvania Journal of Constitutional Law* 2 (2000), 345–421. We believe that there are various ways this distinction could be made in our terms. For example, and perhaps most plausibly, bargaining may be identified with the practice of deliberating with others and making offers in an effort at mutual accommodation, whereas debating is identified with deliberating with others in an attempt to determine what should be jointly done on the basis of jointly accepted premisses.

¹⁰ As a matter of definition, of course, I will have raised the opportunity cost of your existing options: *D*-ing or *E*-ing (and in neither case collecting a reward). Does that involve engineering your reasons, then? We hold not. What explains the fact that you ought, say, to perform *D* and accept the reward, if indeed that is what it turns out you ought to do? What in that sense—the sense signalled in the text—is your reason for taking the option? Not, except more or less tautologically, the fact that it has the lowest opportunity cost. Rather, the fact which explains why that is so: viz., that the considerations that function as reasons for you make it more attractive than the alternatives.

process reasons properly. We now propose to strengthen the characterization of deliberative exchange in a way that restricts it to cases where the interaction is not warped by such extraneous pressures.

It may seem difficult to incorporate the proposed restriction in a tighter definition, since doing so would appear to require, impossibly, a full list of the pressures that might vitiate deliberation. We would define deliberative exchange as the sincere, epistemic communication of belief that is designed to explore the audience's reasons with them, where the communication does not apply extraneous pressures *a* or *b* or *c*... It turns out, however, that no such inventory of extraneous pressures is required. The practice of deliberative exchange involves a heuristic for identifying factors that apply such pressures—we may call them 'vitiators'—and we can invoke those factors, therefore, without having a detailed list to offer. We can define deliberative exchange as involving the sincere, epistemic communication of belief, where the communication is not affected by vitiators of the kind recognized in deliberative practice.

The reason why deliberative exchange involves a heuristic for identifying vitiators goes back to a central aspect of such interaction. The parties to deliberative exchange will not just report judgements and intentions to one another, as if they were matters of incidental, theoretical interest. For reasons parallel to those rehearsed in the discussion of games, such reporting will be transformed by the emergence of certain common beliefs among them. Thus they will report their judgements and intentions in a recognition that others may rely on their holding by the views and the goals in question: that others may make their plans around the assumption that such attitudes are in place, as in commitments to reciprocal or joint action. The parties will elicit this reliance, where it is relevant, as a matter commonly accessible to all, thus depriving themselves—again, in a way that is accessible in common to all—of any basis for protesting against complaints should they not prove reliable; in this sense, they will license or permit such complaints. In order to mark this special way in which the parties to deliberative exchange may report their judgements and intentions, we can say that they avow those attitudes or commit themselves to them. They report the attitudes in a context where reliance is elicited, and exposure to complaint accepted, as a matter of common assumption.¹¹

¹¹ Victoria McGeer, 'Is "Self-knowledge" an Empirical Problem? Renegotiating the Space of Philosophical Explanation', *Journal of Philosophy* 93 (1996), 483–515; Victoria McGeer and Philip Pettit, 'The Self-regulating Mind', *Language and Communication* 22 (2002), 281–99.

But if the practice of deliberative exchange involves avowal and exposure to complaint, it has to involve an adjudicative procedure under which complaints may be validated or invalidated. Deliberators implement such a procedure when they identify cases where a person may not have lived up to an avowal but decide that something about the original occasion of avowal excuses such a failure. Perhaps the person avowed such-and-such a belief, declared such-and-such an intention, made such-and-such a promise. No matter. It is not appropriate to hold the person to that sort of avowal because it was made under this or that extraneous pressure, and, by the criteria implicit in the validating procedure, the pressure provides an excuse—partial or full—for non-performance.

The presence of this validating procedure in deliberative practice means that we can define vitiators as factors that would be taken, under the practice, to excuse someone who made an avowal from acting on it; specifically, they would excuse such non-performance on grounds related to the character of the original exchange. Typical excusers in this category, and so typical vitiators, are the sorts of factors present when we can speak of one person oppressing or exploiting or intimidating another. Suppose I hold out a reward or a penalty—say, a predictive, otherwise acceptable penalty—of a kind that mesmerizes you with an intoxicating or terrifying prospect. Suppose I take advantage of your urgent needs by proposing to sell you a medicine you require but only at an outlandish price. Or suppose I exploit your more general dependence on me, using it—perhaps without either of us being aware of the fact—to get you on side with whatever I say or propose. Under existing deliberative practice, you have a reasonable excuse—perhaps partial, perhaps full—for later refusing to comply with certain commitments you may have made in conditions of those kinds; I cannot expect my complaints about non-compliance to be given significant weight. Those conditions, so the invalidation of the complaints suggests, vitiated the exchange of reasons involved, depriving it of a deliberative, or at least fully deliberative character.¹²

¹² The line taken here for downgrading certain cases where people exchange reasons—in our terms, downgrading them from fully deliberative status—should be distinguished from that adopted, for example, by Jürgen Habermas, *Moral Consciousness and Communicative Action* (Cambridge: Polity, 1990). He argues that in such cases there are always reasons put forward which flout the principle that no consideration should be allowed to carry weight with anyone unless everyone can accept it as a guide for everyone to follow. See too Christopher McMahon, 'Why There is No Issue between Habermas and Rawls', *Journal of Philosophy* 99 (2002),

This leaves us, finally, with a working definition of deliberative exchange or practice. Deliberative exchange involves sincere, epistemic and unvitiated communication on the part of the persons involved, the aim being to explore one another's reasons in the realm of thought or action, including reciprocal and joint action. Or, in full-dress definition, deliberative exchange occurs just so far as

- the parties sincerely communicate their beliefs, openly seeking to get one another to recognize and share them;
- the communication is intended as an epistemic exploration of one another's reasons for believing or doing various things, individually, reciprocally, or collectively;
- there are no vitiators present, so that any avowals of attitude—any commitments—are inescapable under the rules of the practice: only a later change of circumstances can excuse non-compliance.

This definition of deliberative exchange answers, as we have been arguing, to a range of intuitions about when it is appropriate to think that people are deliberating with one another. But it also has another merit that may make it appealing. It represents a dialogical version—a version involving two or more people—of a process that we each conduct in our own case as we try to reason anything through, whether with a view to making a judgement of fact or forming an intention about what to do.

Imagine that I am conducting a bit of theoretical or practical reasoning of this kind. I endorse a consideration, let us say, that I take to be relevant to the conclusion I am to form. Since I will only form that conclusion later, I have to think of myself down that track—perhaps seconds, minutes, or hours down that track—as someone to whom I offer the consideration in the way I offer a consideration in sincere, epistemically oriented communication. And at that later time I have to receive the consideration as my past self's testimony, treating that self in the interlocutor role in which it self-presents. It makes no sense in this epistemically driven enterprise for me at one time to think of forcing my later hand, or deceiving my later self, or putting a coercive intertemporal threat in place. I may

111–29; and Scanlon, *What We Owe to Each Other*. The line taken here promises to downgrade every case that this test would indict. And, more positively, it downgrades certain cases that that test would miss and would be wrong to miss. I may exploit your dependence on me in deliberation despite the fact that every consideration I put forward is reasonable according to Habermas's principle. And yet, intuitively, a case of that kind ought to be downgraded from full normative status.

sometimes try to do that, of course, as when I place little confidence in my later self, and think manipulation is required. But in the process of reasoning, theoretical and practical, I have to relate to myself across time in the manner described.

The definition we have offered of deliberative exchange is supposed to catch an interpersonal analogue of the intra-personal process of reasoning in this way. Just as reasoning with myself requires me to communicate sincerely and epistemically across time, with each self relying on past and future counterparts to live up to certain expectations, so deliberation with others requires analogous moves. It may be some support for the definition offered that what it requires to be in place among partners to deliberation is precisely the sort of thing that must be in place across time whenever people reason with themselves.

3. Deliberative, Deontological Restrictions

The practice of deliberative exchange makes a central place in human life for the recognition of certain deontological restrictions. Such exchange, like the sorts of games described earlier, is governed by constitutive rules. And the existence of these rules means that people who present themselves as aspiring to deliberative exchange, in an understanding of what it involves, must accept that they are going to be expected, as a matter of common knowledge, to endorse those rules and to license others to complain in a deontological way about any failures to live up to them.

Some of the rules that are constitutive of deliberation are rules of conduct—rules for the conduct of deliberation—whereas others are rules of compliance with commitments made within deliberative exchange. By the first clause in the definition of deliberative exchange, to take up the first category, every party must communicate certain attitudes, not seek to impose or inculcate them by other means; and every party must communicate those attitudes sincerely. By the second clause, every party to deliberative exchange must look for an influence on addressees of an epistemic kind. The aim of each must be to draw the attention of others to the options before them in any choice of theory or action and to search out the relevance to their choice of all those reasons that they countenance or can be given reason to countenance. And by the third clause,

a number of further prescriptions and prohibitions follow. I will not be participating in deliberative exchange as such—I will be playing by different rules, playing a different game—if I look for any influence that is supported by deliberatively vitiating pressures such as those that derive from the dependency or the vulnerability or the impressionability of the addressee.

The constitutive rules of deliberative exchange also include rules of compliance. These govern, not how parties should conduct themselves in the course of deliberation, but how they should behave in its wake. They require them to comply with any commitment incurred within the exchange, as in commitments to holding by a certain belief or intention, and acting on it. This sort of commitment may emerge in any deliberative interaction, but it will be typical, in particular, of the sort that leads to co-operative or collective projects. These rules of compliance are constitutive of deliberative exchange in the sense that the practice would cease to exist if participants refused to admit that they could reasonably be challenged with failures to abide by those rules.

The existence and the accessibility of these constitutive rules mean that people who enter into deliberative exchange must disclaim, as a matter of common knowledge, any attitude or action that runs counter to them. Thus no one can aspire to deliberate with someone without being taken to rule out deception, coercion, intimidation, or infidelity to commitments undertaken. And, of course, no one can have such an aspiration without being taken to rule out equally the exercise of force and the like. To set out to deliberate with someone is to announce or present oneself in a certain guise: as a person who can be expected to honour the constraints associated with deliberative exchange.

The constraints that one explicitly or implicitly claims to endorse in assuming the profile of a deliberative partner, like the constraints associated with the games we discussed earlier, are deontological in character. Suppose that I breach a rule of deliberative exchange and try to justify myself on the grounds that this is consequentially productive. The violation involved promises to reduce the number of violations committed in general, I say, whether by me or by others. This justification won't work, if it is meant to show that the violation is to count as a deliberative overture. Deliberative overtures are defined as overtures that are consistent with the rules of deliberative exchange, and any violation of the rules is a departure from that practice. So far as I violate those rules, therefore, I have to

recognize that I have broken with the practice, and that there is no hope of justifying what I do as an instance of what the practice requires.¹³

The rules of deliberative practice, to put this matter another way, dictate that any violation of those rules is deliberatively wrong; they entail that, so far as I violate those rules, I fail to act deliberatively, and can be rebuked by those I let down. But while the rules proscribe such violations, they do not proscribe the act of allowing violations to occur, say on the part of others or on one's own later part. Doing wrong is unjustifiable within the rules of the practice, but allowing wrong to be done is justifiable. And so doing wrong can never be justified as an instance of what deliberative practice requires—it can never be presented as something less than wrongdoing of that kind—on the ground that it reduces such wrongdoing in general. The constraints imposed by deliberative practice are inescapably deontological in character.

4. The Significance of these Restrictions

How important is this result? Not very, it may seem, so far as it parallels the sort of thing that is true also in chess and cricket. True, the rules of deliberative exchange may rule out more significantly objectionable activities—force, intimidation, coercion, deception, infidelity—than the rules of any such game. But just as the rules of a game have a hold on us conditionally on our wanting to take part in the exercise, and not otherwise, so the same is true here. Granted I have to disclaim resort to deception and coercion, force and intimidation, so far as I want deliberative exchange. So what? That still leaves me free to embrace those activities, and make it clear that I am doing so. All it means is that I have a choice between playing the deliberative game, thereby licensing complaint about any resort to such behaviour, and parading the fact that I do not mean to conduct myself deliberatively.

This response, however, misses out on an important feature that makes deliberative practice special. Short of explicitly cancelling such a construal, most of the ordinary overtures I make to other people will signal a willingness to deliberate, as a matter of mutual assumption. And making those

¹³ As becomes apparent in the concluding section, there may be hope of justifying an action, period, when there is no hope of justifying it as an instance of what the practice requires.

overtures, therefore, will mean exposing myself to deliberative expectations and, if I fail to perform appropriately, to deliberative complaints. Deliberative practice haunts human interaction, intruding as a default the assumptions on the basis of which people hold one another to deliberative standards. The jurisdiction of deliberative exchange has a resilient hold on human beings, and entrenches routines of assessment and criticism in the name of associated norms.

Consider what happens when I address others with a serious observation or query or proposal; when I speak to them in a way that is not just a bid for phatic stroking or tickling. How are they to take me? Clearly, they will think, I am putting myself forward as having a communicative intention, not just as meaning to pound their ears. And clearly I am putting myself forward as a sincere speaker; to present oneself as deceptive would be self-defeating. I may be threatening and seeking to coerce them, of course. But if I am, my words will make that plain; and short of that being made plain, the default assumption among my addressees must be that this is an overture in deliberative exchange: an invitation to reason about something of concern to one or all of us. Again, I may be doing this in the hope that an extraneous pressure I activate—say, the pressure on a dependant to keep me sweet—will secure the result I want. But in the normal course of addressing others, I will keep that sort of hope in the background, as something unsaid and unsalient, not as something my addressees are invited to see.

Granted that I am going to be taken in normal exchange to have deliberative intent, the question is whether I am likely to acquiesce in being taken that way. Acquiescing would mean acquiescing in its being for me a matter of assumption shared with my interlocutors that we are each disposed to conform to the rules of deliberative exchange; that we are each disposed to rely on commitments made within the exchange; and that we are each making ourselves vulnerable, therefore, to complaints about not proving reliable.

We think that I will have little choice in the normal situation but to acquiesce in this construal of my attitudes. The only alternative will be to renounce deliberation altogether, embracing the coercive or intimidating profile of someone who announces something tantamount to: 'this is an open competition for individual advantage', 'this is a fight for victory'. I face a stark set of options: either to go along with the image of intending a deliberative exchange, whether sincerely or insincerely, with or without

the hope of pressuring others into taking my preferred line; or to renounce that image in favour of a declaration of war. Faced with such alternatives, the option of endorsing the deliberative image will normally prove irresistible. So we surmise; and so everyday evidence suggests.

Acquiescing in something's being a matter of mutual assumption involves not denying that it is a matter of mutual assumption in a context where this, as anyone can see, will otherwise go through by default. And not denying that it is a matter of mutual assumption in such a context amounts, in effect, to putting it forward positively as a matter of mutual assumption. In the case envisaged, then, I manifest a positive claim that the exercise sought or in process is one in which each of us communicates his or her attitudes in a sincere, epistemic, and unvitiated fashion; one in which the norms of deliberation are taken by all to be relevant and are firmly put in adjudicative place.¹⁴

The thesis defended does not mean that deliberative practice itself will have a ubiquitous presence in human life. What will have that sort of presence, rather, is the pretension to deliberative practice, where that pretension may be sincere or insincere, motivated or not motivated by a desire to exert extraneous pressure. To pretend to deliberation is to present oneself as accepting and endorsing the pattern of mutual assumptions on the basis of which deliberative exchange proceeds. It is to put oneself forward as sharing in a mutual belief that everyone involved intends to conform to deliberative norms; that everyone is relying on others to abide by those norms; and that everyone licenses complaints against them, should they fail to conform. Thus it is to cast oneself as operating within the jurisdiction of deliberative practice, inviting others to hold one to the expectation that one will conduct deliberation properly, and that one will comply with any commitments made within deliberative exchange.

¹⁴ The sort of claim envisaged is close to what Jürgen Habermas would describe as an implied claim to validity. See his *Theory of Communicative Action*, vols. 1 and 2. He sometimes represents validity claims as claims about the subjective sincerity of the speaker in making a remark, asking a question, giving an order; about the objective suitability of the world in making that remark true, that question relevant, that order feasible; and about the inter-subjective authority of the speaker in virtue of which the overture cannot rightly be ignored by the addressees. See Habermas, *Moral Consciousness and Communicative Action*. This taxonomy is meant to cover all sorts of speech acts in a single formula, whereas our formula is designed with assertoric overtures primarily in mind. We note, however, that it would seem to ignore the claim implicit in such overtures to be communicating one's attitudes in an epistemic manner and not to be relying on any extraneous pressures for the efficacy of the communication.

The pretension to deliberative exchange, understood in this way, may be present without such exchange actually materializing. Recognizing that I will inevitably trigger the mutual assumption that we are involved in deliberation—and that I will inevitably expose myself, therefore, to adjudication in terms of deliberative norms—I may decide to make the best of this in self-serving terms. I may be prepared, where doing so is advantageous, to tell lies or break promises or rely on unspoken pressures that I can intrude in my dealings with others. And what I may do, of course, may be done by many others. Deliberative pretension does not ensure the realization of deliberative practice, then. All it means is that the norms associated with deliberative practice will be assured of relevance; what I and others do may not live up to those deliberative norms, but it will be subject to their jurisdiction.

At the beginning of this discussion we asked whether the restrictions associated with deliberation had any deeper bite than those associated with games. The suspicion raised was that they don't. One may apparently choose not to play the deliberative game in the way in which one may choose not to play chess or cricket. But we now see that this thought misses out on the fact that deliberative pretension has a resilient presence in social life. In routine overtures with others, one will naturally seem to be playing the deliberative game, whether that is something welcome or not. And in such overtures it will be difficult not to acquiesce in the assumption. Once one acquiesces in that assumption, of course, one exposes oneself to the jurisdiction of deliberative practice and to adjudication in the name of deliberative rules. Those rules will attain relevance in the assessment of how one behaves.

For all that this says, as we mentioned, it remains possible that people will pay little or no attention to the norms of deliberative exchange, flouting them with abandon. But the line of thought we have been running also gives us reason to think that people will tend to take the norms quite seriously. Suppose that someone routinely breaches deliberative norms and, as is likely to happen, that this becomes widely known. That will mean that the person will not be able to pretend effectively to deliberation with acquaintances, since they will take his or her past behaviour to show that in all likelihood the pretension is not sincere. But if the person is unable to pretend effectively to such deliberation, then that will ensure a sort of ostracism from their community. It will mean that the person cannot hope to enjoy the fruits of deliberative exchange with acquaint-

ances—or of the simulation of deliberative exchange—in any area where such exchange makes others vulnerable.

The sort of ostracism envisaged will have quite enormous costs. First of all, it will severely reduce the ways in which those ostracized may hope to exercise even unwelcome influence over others. They may set out to force others to do things. Or they may shamelessly threaten and try to coerce them. Or they may seek to rig things so that others are unwittingly led to act as they want. But that is about all. The reserves of influence associated with deception and fraud, insinuation and innuendo, will remain closed. They are only there to be mined, after all, by the person who can effectively pretend to deliberation.

Not only is the ability to pretend to deliberation indispensable as a tool for the abuse of others in these ways, it is also the only instrumentality available whereby people can achieve certain mutually beneficial kinds of influence. We can depend on each other to be able to expand our own reserves of evidence and understanding, and so to improve the quality of our individual judgements and decisions, only so far as we can establish deliberative relations with one another, proving ourselves to be mutually reliable over the long haul of trial and test. And we can unite with others to achieve initiatives of joint action only so far as we can establish among ourselves that we are of a common, deliberatively shaped mind. Those who rely on force or coercion to get others to act with them will only be able to have faith in those others when they are under effective scrutiny. And those who seek to expand their own evidence and understanding by interrogation of those who fear them will be exposed to the paradox of the counsellor: the paradox that only the king who abjures his kingly power can hope to have honest counsellors.

Where the resilience of deliberative pretension facilitates entry to the jurisdiction of deliberative norms—it makes it almost impossible not to enter—the feature we have just been reviewing explains why that jurisdiction may be expected to have a substantive role in determining how people in the general run behave. Let people conduct themselves with relative indifference to the norms of deliberation, and they will find themselves deprived of the ability to pretend to deliberate, and so deprived of access to those overtures that involve deliberative pretension. They will find themselves cast out from the society of their immediate fellows. The only life possible for them will be a life on the hoof, in which they move on at every point where their deliberative recalcitrance begins to register with associates.

The considerations canvassed should be sufficient to establish that the practice of deliberative exchange has a compelling character that marks it off from the games discussed earlier. The fact that that practice imposes deontological constraints on us, therefore, means that such constraints have a central place in our dealings with one another; they assume an inescapable and imposing authority in our lives. People can parade an indifference to deliberative constraints only in those cases where they are dealing with others whom they cast unthinkingly beyond the pale of ordinary conversation and commerce. And such cases are likely to be infrequent. They may arise in the relations between despots and their underlings, or between two groups which are sworn to mutual enmity, but they are hardly robust phenomena on the settled, social scene.

The centrality of deliberative norms in social life shows up in the fact that we can usefully treat them as explicating the demands of respect, at least on one interpretation.¹⁵ On this interpretation, to respect other people—to treat each as a person—is to deal with them in a deliberative fashion or, at the least, to deal with them in a way that leaves open the possibility of deliberation. Negatively, this requires not ostracizing them from one's world and not treating them with aggression or coercion or the like. Positively, it requires, first, addressing them as minds with which one can communicate in a sincere, epistemic, and unvitiated manner; and second, giving them a hearing as voices that can reach one's own mind in reciprocal fashion. When one addresses others in this way, the epistemic character of the deliberation projected means that one never treats them as means only, but always also as ends.

This interpretation of respect is a distinctively social one. It explains why one cannot fully sustain respect in relation to people who do not reciprocate in kind. And it explains why there are different ways in which social arrangements may jeopardize the possibility of respect without anyone's being at fault. Respect will be put at risk by the sort of asymmetry which means that one party has to condescend to others, as if in gratuitous benevolence, in addressing them or giving them a hearing. That is why Kant, following a long republican tradition, insists that it is in the nature of respect that it is not given as a gift.¹⁶ One cannot receive respect as a gift

conferred by others; one can only have the respect of others when one is sufficiently on a par to be able to command it.

5. Deontology versus Teleology: A Comment

Where, finally, does this leave the more familiar debate between deontology and teleology, non-consequentialism and consequentialism, in moral theory? We have argued that certain deontological constraints have a widespread presence in human life, coming into play with deliberative pretension and deliberative practice. The question, however, is whether this means that they have a fundamental status as norms. Can they be derived as norms we generally ought to countenance, given the neutral good that doing so promotes, as consequentialists say? Or do they resist any such derivation, as deontologists would claim?¹⁷

There are two reasons why the truth that we have found in deontology—the truth that deontological constraints are so widespread—may seem to establish the truth of deontology. The first is that the deontological practice of deliberative exchange is absolutely central to human interaction, giving content to the very notion of respect. The second is that deliberative practice offers the only site at which the question between deontology and teleology can be debated and, being itself subject to deontological rules, makes a judgement in favour of teleology difficult or even incoherent. Neither consideration, however, is decisive.

The second consideration, to take them in reverse order, starts from the observation that the question between deontologists and teleologists arises within deliberative exchange, so it has to be debated under the deontological rules of that exchange. I give my reasons for why I think that whether something is the right way to behave is determined by the neutral consequences associated with it. You counter with your reasons for denying

¹⁷ For the record, though one of us is committed to the truth of consequentialism, the other isn't. See Philip Pettit, 'A Consequentialist Perspective on Ethics', in M. Baron, M. Slote, and P. Pettit (eds.), *Three Methods of Ethics: A Debate* (Oxford: Blackwell, 1997), 92–174; Philip Pettit and T. M. Scanlon, 'Consequentialism and Contractualism', *Theoria* 66 (2000), 228–45; and Michael Smith, 'Immodest Consequentialism and Character', *Utilitas* 13 (2001), 173–94. For an overview of issues that arise for a consequentialism that evaluates practices and other items, as well as acts, see Philip Pettit and Michael Smith, 'Global Consequentialism', in Brad Hooker, Elinor Mason and Dale E. Miller (eds.), *Morality, Rules and Consequences* (Edinburgh: Edinburgh University Press, 2000), 121–33.

¹⁵ Stephen Darwall, 'Two Kinds of Respect', *Ethics* 88 (1977), 36–49.

¹⁶ On this theme in republicanism see Philip Pettit, *Republicanism: A Theory of Freedom and Government* (Oxford: Oxford University Press, 1997), chs. 1–2. On Kant see J. B. Schneewind, 'Autonomy, Obligation, and Virtue: An Overview of Kant's Moral Philosophy', in P. Guyer (ed.), *The Cambridge Companion to Kant* (Cambridge: Cambridge University Press, 1992), 309–41.

this. And both of us have to conduct that debate—where else?—in the context of deliberative, deontological rules.

But the fact that the issue has to be decided under deontological rules does not guarantee a victory for deontology. Victory would be guaranteed only if it was impossible within deliberative exchange to talk about why the practice of such exchange is valuable, and to ask after whether it may not occasionally be better—however regrettable—to eschew deliberation in favour of initiatives like coercion or outright force. But of course this is not impossible.

No impossibility blocks me from trying to convince you in deliberative exchange that what makes such exchange valuable, and in general better than alternatives, is the fact that it produces various attractive results. And no impossibility blocks me from arguing that whenever circumstances are perverse enough to make the prospective results of exiting the practice and resorting to force or coercion better than those of sticking with it, then that is what I should do. If I can argue this line successfully, I will be able to vindicate consequentialism; if not, not. And for all we have seen, both possibilities remain open.

The second reason why the truth in deontology may seem to ensure the truth of deontology derives from the centrality in human life of deliberative practice. The considerations rehearsed in the paper direct us to a canonical mode of interaction between human beings. And they show that this mode of interaction is conducted, in the nature of the case, under deontological restrictions that are admitted by all. One cannot deliberate with others and not disavow consequentialist calculations about what one should say or do, just as one cannot assume the profile of a friend in dealing with others and not eschew such thoughts. But deliberating with others is the very paradigm of interpersonal relations, and its non-consequentialist character may make deontology seem like an ineluctable doctrine.

Consistently with showing that deliberative practice has such a character, however, the considerations rehearsed leave it an open question whether rightness is determined in the last analysis by the neutral value of associated consequences. Deontologists are certainly on the right track in holding that the code implicit in the practice of deliberative exchange—the practice, as we take it, of interpersonal respect—is inherently deontological; and some may think that this is all they ever really wanted to maintain. But they are not on the right track if they go beyond this and hold that this in itself establishes the truth of deontology.

No matter how central deliberative practice is in human life, it is still possible that fidelity to this practice is required, as consequentialists maintain, only so far as it makes for the best overall. Who is to say that if fidelity to the practice is going to have disastrous results, for example—say, because it will hinder the prospect of saving innocent lives—I should still remain faithful to it, abjuring the force or coercion or whatever that would save those lives?¹⁸

It is common for consequentialists to argue that while a practice like friendship involves eschewing consequentialist thoughts in dealing with friends, still the practice should be espoused so far, and only so far, as it promises to be for the best overall. Its centrality notwithstanding, there is nothing to block consequentialists from taking a similar line on deliberative practice. The deontological character of that practice shows that there is truth in deontology, and perhaps more truth than consequentialists have traditionally allowed. But it does not show that deontology is true; it leaves the consequentialist problematic intact.¹⁹

¹⁸ Commitment to even the interpersonal analogue of deliberative practice described above looks like it too will be vulnerable to the value of that commitment. See Michael Smith, 'Search for the Source', *Philosophical Quarterly* 49 (1999), 384–94.

¹⁹ An earlier version of this paper was presented at a colloquium at the Australian National University. We were greatly aided by comments from participants at that colloquium. We are also especially grateful for comments received from Jay Wallace.