32. Deliberation and Decision

PHILIP PETTIT

Sections

The Decision-Theoretic Picture

Does every action originate in a decision to perform that action? It cannot do so if decision is itself an intentional action: if it is a mental act, as some have taken it to be, in which an agent resolves uncertainty about what to do in a given context. For if action were always supposed to originate in decision, and decision were itself an action, then we would face a regress. The regress would be similar to that which Donald Davidson (1980) invoked in criticism of the idea that every intentional action must originate in an act of will, repeating – without apparently being aware of it – a point that Thomas Hobbes (1994: 125) had made in 1640: "a man can no more say he will will, than he will will will, and so make an infinite repetition of the word will."

But decisions do seem to count as actions. Thus, if you claim to have decided something, we naturally assume that you were led to do so on the basis of your desires and beliefs, and we treat the decision in just the way we would treat an action. Then, while we certainly think that some actions begin in distinct decisions, we cannot think that all do so. It seems that we are committed to thinking that there are two types of action: those which originate in decisions and those which do not.

This dichotomy need not cause problems by suggesting a bifurcation in the nature of action. For there is one obvious basis on which actions do fall into two categories of the kind envisaged. Some actions are planned in advance and some are not, and we may equate those that originate in decision with those that are planned in advance, those that do not originate in decision with those that are not diachronically planned. The claim, perfectly plausible in itself, is that some actions originate in prior decisions, others do not.

But does not the tie between decision and action seem to be closer than that? After all, we ask about the decision you made in taking such and such an action even in cases where there is no question of your having had a prior plan. How to build this assumption into the emerging picture?

The way to build it in, I think, is to say that there are prior decisions and there are simultaneous decisions, and to hold that, while only some actions originate in prior acts of decision-making, still all actions are, simultaneously or synchronically, acts of decision-making; all actions involve, incorporate, or display simultaneous decisions. The simultaneous act of decision-making, on this account, is not something distinct from the action, something in which the action might originate; it is that very action, taken under a certain description.

To adopt this picture is to equate deciding with forming an intention, as it is sometimes said. The decision is distinct from the action decided upon when the intention is formed in advance of the behavior that enacts it–when it involves making a plan for that future behavior (Bratman 1987)–and it is not distinct when there is no such temporal lag: when acting thus and so is just acting with this or that intention, not enacting an intention previously formed. Deciding, on this way of thinking, is forming an intention in respect of a certain set of options–if you like, choosing between those options–and it is distinct from enacting the chosen option only in the case where the intention is formed, or the choice made, in respect of options that are not yet available to the agent.

This nexus of claims makes sense of assumptions that are presupposed in decision theory, and I describe it as the decision-theoretic picture (Jeffrey 1983; Joyce 1999). Decision theory is supposed to explicate something that deserves to be described as decision, and yet is meant to apply to every case of intentional action, at least among suitably rational agents. Rational agents will have well-defined probability and utility functions–regimented versions, by most accounts, of systems of belief and desire– and they will act in such a way as to maximize expected utility: they will act so as to maximize the satisfaction of their desires according to their beliefs. Any action of that kind will count as a decision in the sense of the theory. And it will count as a decision, whether it is an action that relates to what is to be done at a future time or an action that relates to what is to be done now: whether or not it is a decision that is distinct from the action decided upon.

To sum up the discussion so far, there are three claims built into the decision-theoretic picture, all of them being well grounded in our common assumptions and ways of speaking:

· First, every decision is an action.

· Second, on pain of regress, while some actions originate in decisions–that is, prior acts of decision-making–not every action does so.

· Third, there is nonetheless no action without decision: actions that do not originate in prior decisions are, simultaneously, decision-making acts.

The Decision-plus-Deliberation Picture

But at this point the notion of deliberation has to be put in the picture. Whenever you make a decision, we are liable to ask about what weighed with you in deciding as you did: what tilted the balance, to use a standard metaphor, in favor of the line you took. We generally think that it makes sense to ask that sort of question, treating you as an authority in seeking the "weighing reasons" that led you one way or another (Broome 2004; but see Dancy 2004). The assumption amounts to an assumption that, wherever

there is a decision, there is deliberation: the agent considers and compares the features of the options available and is led on that basis, as by reasons, to a choice.

The word 'deliberation' derives etymologically from the Latin word *libra* for a weighing scale and it seems appropriate to link it with this sort of process rather than to use it with less demanding connotations, as is sometimes done (Skyrms 1990). I shall assume in what follows that deliberation is a process in which an agent explicitly or implicitly considers the features of the options available in a certain choice and, all going well, is led to the decision which those features recommend as being the best.

Does decision theory in itself make room for deliberation? Decision theory starts from the degrees of utility that an agent attaches to certain states of affairs – that is, roughly, from the intensities with which the agent ranks those states of affairs against one another – and from the probabilities with which the agent links each of an available set of options to those states of affairs, no matter how the probabilities are understood. Then it identifies that option (or sub-set of options) in the relevant set which the agent is rationally required to select, and hence will select, under the assumption of rationality. But does decision theory have to make suppositions about the features of those options that weigh with the agent in anything that might deserve to be called deliberation?

This question may be introduced by analogy with a question raised in the theory of inference – in particular, inference of the sort a robot might display. Suppose that there is a system of artificial intelligence built so as to satisfy, for example, *modus ponens*. Whenever the system counts as endorsing sentences or propositions of the form 'p' and 'if p, then q' – whenever it counts as representing its environment in corresponding ways – then, so long as the conditions satisfy normal specifications, it transitions to endorsing the proposition or sentence 'q' – that is, to forming a q-representation of how things are (Cummins 1983). Should we say that such a system draws inferences in accordance with the *modus ponens*?

In one sense, yes; in another, no. Clearly the system forms new representations on a pattern which conforms to that rule. But, equally clearly, it does not relate to that rule in the way you or I might. It need not have any awareness of that rule as such: not in the way logicians are aware of it. And, more important still, it need not have any awareness of what the rule requires in particular circumstances; it need not form the belief that 'q' on the basis of forming the belief that the truth of 'q' follows from the truth of 'p,' taken together with 'if p, then q.' It need not have the wherewithal to form the belief that 'p' on the basis of forming the belief which most of us would express in the words: 'so q' or 'so therefore, q.' There is nothing about how it performs that would force us to ascribe to it the ability to think in 'so – therefore' terms.

Let us say that a system that draws *modus ponens* inferences without being able to think in 'so – therefore' terms conforms to the rule but does not apply it. The question that confronts us about deliberation and decision theory can be phrased in parallel terms. Does the decision-theoretic agent who conforms to the theory's principles of rational decision-making have to apply such principles, thinking in 'so – therefore' terms? In particular, does the agent have to register properties of the rival options in any choice and let a consideration of those features indicate how it is rational to choose? Or might such an agent merely conform to the principles of rational decision-making, as the robot conforms to the *modus ponens*?

Once this matter is raised, it is pretty obvious that, for all that decision theory says and the decision-theoretic picture supposes, the rational agent is not required to register those features of rival options which might make salient the rationality of choosing this option rather than that; there need be nothing that would correspond to thinking in 'so-therefore' terms. The decision-theoretic picture merely supposes that, when the agent has beliefs and desires which make a certain choice rational, then, other things being equal, the agent will make that choice. It says nothing about how the agent is led to the choice. In particular, it does not suppose that there has to be any process of registering and weighing the features of the rival options or any process which might deserve to be described as 'deliberation.'

This is to say that decision theory as such does not rule in deliberation, making it an obligatory feature of rational choice. But does it actively rule it out?

A theory of rational inference would ideally tell us what it was right to infer on the basis of any given set of beliefs, going beyond what strict logic would deliver (Harman 1986). But it might do this without offering any view on the process whereby the rational agent is moved to draw rational inferences, and without any commitment to there being a process in which the agent is moved by registering what deserves to be described as 'reasons'–a process, for example, in which the agent is moved by the recognition that the truth of *'p'* and of 'if *p*, then *q* ensures the truth of *'q* and provides reasons for believing that *'q.'* The theory of rational inference would not rule in a process of reasoning of this kind. But neither, of course, would it rule out the relevance of such a process. It would just be silent on the matter.

By analogy with a theory of rational inference, a theory of rational decision will tell us what it was right for an agent to do or decide on the basis of a given set of beliefs and desires. But it can do this without any view on the sort of process whereby the rational agent is moved to make rational decisions, and without any commitment to there being a process in which the agent is moved by registering the properties of rival options and is led by them, as by reasons, to make the rational choice. Decision theory will not rule in a process of deliberation, as the theory of rational inference will not rule in a process of reasoning. But neither will it rule out the presence of such a process. It will posit nothing one way or the other.

On this account, the three claims encoded in the earlier picture stand, but they can be supplemented by a fourth claim, as in the following schema, which spells out the elements in the decision–plus–deliberation picture:

- Every decision is an action–that is, an act of decision–making.
- Some actions originate in prior decisions, but not all of them do.
- All actions, however, involve simultaneous acts of decision–making.
- There is no decision, prior or simultaneous, without deliberation.

As already mentioned, we expect people to be able to provide an account of the reasons that weighed with them in any choice, or at least in any choice in the general run. We expect them to be able to speak about their reasons with a certain authority, whatever its source, and not just to speculate about them as they might speculate about hidden influences. Insofar as we expect this, we appear to endorse the decision–plus–deliberation picture. But is that picture really gripping? Are we really as reflective and intellectual as it makes us seem?

It is certainly unlikely that non–human animals live up to the picture. According to this description, the options over which agents deliberate are *abstracta*, not *concreta*: they are ways–things–might–be – ways the agent knows how to make things be – not actual events. Human beings can identify such entities on the basis of how they answer to abstract linguistic specifications: my helping the beggar is that way–things–may–be, under certain contextual constraints, which makes true the sentence "I help the beggar." But it is unclear how creatures without language could ever entertain options as objects of thought and ever deliberate about the features of such options (Pettit 2007).

This restriction does not apply to human beings. But isn't it still implausible to hold that, in the general run of choices, humans invariably deliberate about the options they face: they are invariably in deliberative control?

There are two ways in which deliberation might control a decision and provide support for something like the decision–plus–deliberation picture. It might actively control the decision, as when I go painstakingly through the pros and cons of a particular choice and make up my mind on the basis of this reflection. Or it might control the decision just in a virtual manner. Suppose I go through a certain train of action more or less unthinkingly, as when I drive home or go to the shop to pick up something to eat. There may be no active deliberation whatsoever involved; I may go through the required motions on automatic pilot. But

deliberation will still be in virtual control if the following is true: should things go wrong, as in my taking a wrong turn or beginning to pick up hardware rather than food, then the red lights will go on – I will become aware that something is amiss; deliberation will be activated; and deliberation will put me back on the right track (Pettit 1995).

The decision–plus–deliberation picture is committed, not to the view that deliberation is actively present in all our decisions, but only to the view that deliberation generally has at least this sort of virtual or standby control over what we decide and do. And that observation may help to make the picture more persuasive. The presence of even virtual deliberative control would explain why it generally makes sense to ask a human agent for the reasons which carried weight in a decision. The reasons which carried weight may not have been actively considered. They may be the reasons which would have been invoked, had an otherwise unthinking pattern of action gone amiss. After all, it is true of such reasons that the agent's pattern of behavior was required to satisfy them, however unthinking the agent may have been, on pain of putting on the red lights. They exercised control over the behavior, albeit only virtual control, and in that sense they carried weight.

A Common Mistake

Those who master decision theory may be tempted to think that the properties registered in deliberation are properties that decision theory enables us to identify with precision. Thus they may be tempted to think that the property in virtue of which a given option or a probabilistically associated outcome is found attractive by a rational agent is the property of having such–and–such a utility score: being preferred with such–and–such intensity to alternatives. This line of thought would support the idea that a good way to make people more rational would be to have them use decision theory as a deliberative calculus. They would compute the utility and probability scores of different outcomes, do the usual expectational computations, and in that way identify the option it is rational for them to choose.

This is a mistaken and seriously confused way of thinking (Pettit 2002). In the first place the recommended methodology is not really feasible, since there is no way in which we can tell the degrees with which we desire different scenarios, or even the degrees of probability we associate with them. These may show up in our preferences, in particular our preferences over certain gambles (Ramsey 1990), but they are not available to us at will, as in a process of introspection (Harman 1986).

In the second place, and perhaps even more strikingly, the recommended methodology would force us to revise an established and intuitively attractive way of thinking (Pettit and Smith 1990; Pettit 2006). When I decide to help a beggar on the street, I may do so out of a desire to be rid of a sense of guilt. But in less pathological cases I will do so out of a desire to help a person in obvious need. On the recommended way of thinking, the property of the philanthropic option which should register with me is that it has a high utility count: it promises to satisfy my philanthropic desire. But, were I to go over to that way of thinking, then there would be a serious shift from the way in which I think in the normal, non–pathological case of philanthropy. In the normal case, I think in such a way that, on my view, helping the beggar is the thing to be done in any scenario where he or she remains in need, other things being equal. But were I to go over to the recommended methodology, I would think in such a way that, on my view, helping the beggar is the thing to be done only in scenarios where I continue to have the philanthropic desire.

There are some desires with which it makes sense to endorse the options they motivate only so long as the desires remain in place. I will make plans to have a smoke after a meeting, or to have a snack later, only on the assumption that I continue to have a yen for a smoke or a snack. But with most desires this is not the case (Parfit 1984). I can sensibly make plans to fulfill such a desire even when it is possible, as I recognize, that I will have ceased to have the desire at the time I am required to act. In such a case, the property of the option which prompts me to desire that very option is not its promising to satisfy my desire – to relieve me of it, as I might be relieved of an itch – but rather an independent property: say, its being just, or generous, or spontaneous, or likely to improve the existing state of things in some way.

Were I to adopt the recommended methodology, the effect would be that I would treat all desires in the manner in which it is appropriate to treat only a yen such as for a smoke or a snack. I would help the

beggar, not on the grounds of his or her need, but rather on the grounds that doing so provides me with the best available opportunity for desire–satisfaction. Decision theory may be consistent with deliberation, as on the decision–plus–deliberation picture, but it is not a good guide as to the form which deliberation ought to take.

See also: collective action (9); motivating reasons (19); humeanism about motivation (20); practical reasoning (31); akrasia and irrationality (35); rationality (36); motivational internalism and externalism (37).

References

Bratman, M. (1987). *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.

Broome, J. (2004). Reasons. In J. Wallace, M. Smith, S. Scheffler, and P. Pettit (eds), *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*. Oxford: Oxford University Press, 28–55.

Cummins, R. (1983). *The Nature of Psychological Explanation*. Cambridge, MA: MIT Press.

Dancy, J. (2004). *Ethics Without Principles*. Oxford: Oxford University Press.

Davidson, D. (1980). *Essays on Actions and Events*. Oxford: Oxford University Press.

Harman, G. (1986). *Change in View*. Cambridge, MA: MIT Press.

Hobbes, T. (1994). *Human Nature and De Corpore Politico: The Elements of Law, Natural and Politic*. Oxford: Oxford University Press.

Jeffrey, R. C. (1983). *The Logic of Decision*, 2nd edn. Chicago: University of Chicago Press.

Joyce, J. M. (1999). *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.

Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.

Pettit, P. (2002). Decision theory and folk psychology [1991]. In P. Pettit (2002), *Rules, Reasons, and Norms*, Oxford: Oxford University Press, 192–222.

Pettit, P. (1995). The virtual reality of *Homo economicus*. *Monist* (78) : 308–329. [Expanded version in U. Maki (ed.), *The World of Economics*, Cambridge: Cambridge University Press, 2000; reprinted in P.Pettit (2002) *Rules, Reasons, and Norms*, Oxford: Oxford University Press.].

Pettit, P. (2006). Preference, deliberation and satisfaction. In S. Olsaretti (ed.), *Preferences and Well–Being*. Cambridge: Cambridge University Press, 131–153.

Pettit, P. (2007). Rationality, reasoning and group agency. *Dialectica* (61) : 495–519.

Pettit, P., and Smith, M. (1990). Backgrounding desire. *Philosophical Review* (99) : 565–592. [Reprinted in F. Jackson, P. Pettit, and M. Smith (2004), *Mind, Morality and Explanation*, Oxford: Oxford University Press.].

Ramsey, F. P. (1990). *Philosophical Papers*. Cambridge: Cambridge University Press.

Skyrms, B. (1990). *The Dynamics of Rational Deliberation*. Cambridge, MA: Harvard University Press.

Cite this article

PETTIT, PHILIP. "Deliberation and Decision." *A Companion to the Philosophy of Action*. O'Connor, Timothy and Constantine Sandis (eds). Blackwell Publishing, 2010. Blackwell Reference Online. 28 January 2013 <http://www.blackwellreference.com/subscriber/tocnode.html?id=g9781405187350_chunk_g978140518735034>

Copyright

Bibliographic Details

A Companion to the Philosophy of Action

Edited by: Timothy O'Connor and Constantine Sandis
eISBN: 9781405187350
Print publication date: 2010