

John Horty

November 12, 2014

(Lightly revised January 15, 2015)

Requirements, Oughts, Intentions

****Comments on Broome for PPR Symposium, first draft****

John Broome's *Rationality Through Reasoning* is a book of tremendous scope, moving from foundational questions about reasons, requirements, and oughts to more general issues concerning the architecture of a rational agent and, especially, the process of reasoning through which such an agent could bring itself to satisfy the requirements of rationality. Although Broome explores a number of these requirements, he is particularly interested in one he calls *enkrasia*, according to which rationality requires an agent to intend to do what that agent believes he or she ought to do.

I agree with most of what Broome says in this book, and I admire all of it, even when I do not agree, for its forthright, constructive approach, and for its clarity. My remarks will be limited to a few areas in which I think ideas from the book might be developed in more detail, or pushed further.

Reasons, requirements, oughts

Broome offers two different definitions of normative reasons. According to the first (Chapter 4), a normative reason is an explanation of why an agent ought to perform some action, or a component of such an explanation. Broome is especially concerned with weighing explanations, which he uses to isolate the concept of a *pro tanto* reason. A weighing explanation is one that explains why an agent ought to perform an action by appeal to the metaphor of

weighing, or balancing, the reasons to perform that action against the reasons not to. Given this concept of a weighing explanation, Broome then offers a functional definition of a *pro tanto* reason as a component of such an explanation that plays a characteristic role—the role of carrying a metaphorical weight, which interacts with the weights of other reasons to determine what the agent ought to do.

I could complain, and have complained elsewhere, about appeal to the metaphor of weighing in this context, and not just because it is a metaphor, but because, even as a metaphor, it suggests too simple a picture of the way in which reasons interact to support oughts. I will refrain from further complaints, however, since *pro tanto* reasons, as Broome defines them, play little role in his larger theory. Instead, he relies almost entirely on his second (Chapters 7 and 8) definition of reasons as *normative requirements*. The general picture within which this definition is developed can be set out in three stages.

First, there are various sources of requirement—morality, rationality, prudence, law, etiquette, fashion, custom, the Catholic church, and so on—each of which sets out various *requirements* to be satisfied by an agent. Broome provides a model of this system of requirements within the framework of possible worlds semantics by stipulating that, if S is a source of requirements, N an agent, and w a particular world, then $R_S(N, w)$ is the set of propositions required of N by S at the world w . To illustrate: if S is the law, N is Jack, and w is the actual world, then if the law requires Jack to pay his taxes, the proposition that Jack pays his taxes will belong to the set of requirements $R_S(N, w)$, along with all the other propositions required of Jack by the law. Note that this treatment allows for conflicting requirements, even from a single source.¹

Second, not all requirements, on Broome's view, need to be taken into account in determining what an agent ought to do. Those that do need to be taken into account are

¹From the standpoint of modal logic, Broome's semantic interpretation of requirements is a version of minimal model, or neighborhood, semantics, which has been used to model conflicting oughts; see, for example, Chapter 7 of Chellas (1980).

the *normative* requirements, or reasons. The requirements of morality and prudence, he says, are like this—these sources of requirement issue reasons. He is unsure, in spite of an extensive discussion (Chapter 11), whether the requirements of rationality are normative or not; and he mentions etiquette and Catholicism as examples of sources whose requirements are not typically normative. Broome allows, however, that certain sources of requirement may have a derivative normativity. If it is prudent to obey the law, for example, then legal requirements may have a kind of normativity derived from prudence; if morality requires kindness, and kindness in some situation can be shown only by meeting the requirements of etiquette, then these requirements too may be normative in a derived way.

Finally, the various normative requirements bearing on an agent are supposed to interact to determine what the agent ought to do. The interactions among normative requirements are complicated by the fact that these requirements may be dyadic, or conditional, and also themselves *pro tanto*, carrying different weights.² To illustrate the first possibility, Broome provides as an example the legal requirement to drive on the left side of the road, under the condition that you find yourself in Britain; to illustrate the second possibility, we can imagine an agent subject to conflicting state and federal regulations who decides to follow the federal regulations on the basis of the legal principle of *lex superiori*, according to which laws emanating from higher authorities carry greater weight.

So that is the general picture. Before moving to more detailed remarks, I want to say that I was, personally, surprised to find this picture in Broome. One thing I thought I knew about Broome's work before reading this book was that he wanted to shift our focus, in thinking about normativity, from reasons to requirements.³ But as far as I can tell,

²Broome does not extend his semantic interpretation of requirements to dyadic, or conditional, requirements, but for this he can be forgiven; I have never seen any adequate treatment of conditional requirements within possible worlds semantics. Also, though Broome does not explicitly commit himself to the idea that requirements themselves should be *pro tanto*, at several points (e.g., p. 125) he mentions reasons for thinking so, and I do not see how these reasons can be avoided.

³See, for example, the opening passages of Broome (1999) and Broome (2004).

the picture he offers is structurally identical to the more traditional picture according to which reasons—contributory considerations favoring one side or another—interact to support oughts.⁴ Broome’s requirements seem likewise to be *pro tanto*, or contributory, considerations that interact to support oughts. So I am no longer completely sure how this part of the theory Broome sets out, in terms of requirements, differs from the traditional picture, usually set out in terms of reasons.

Putting this matter of interpretation aside, it is clear that Broome’s theory falls within the general class of those in which oughts are determined by interactions among things that are not oughts. For Broome, of course, these things are requirements, but the idea has been developed in different ways by different people. Bas van Fraassen, for example, sets out a theory according to which oughts are supported by interactions among imperatives; in my own work, I have generalized van Fraassen’s view to one in which oughts are supported by default rules, which I interpret as representing reasons.⁵ Likewise, David Makinson has proposed a theory in which oughts are supported by conditional norms, leading to the development of an extensive literature on input/output logics, where the inputs are the norms that support oughts, and the outputs are the oughts they support.⁶

In any theory of this kind, where oughts are supported by other things—requirements, imperatives, defaults, norms—it is important to distinguish three questions: First, what logical relations hold among the things that support oughts? Second, how do the things that support oughts interact to support the oughts they do? And third, what logical relations hold among the oughts that are supported in this way?

In answer to the first question, Broome argues (Section 7.4) that the logic of requirements is thin, with few relations among requirements. I agree, and in fact would be happy to abandon even the few logical relations that Broome does insist on, such as the idea that

⁴See Lord and Maguire (20xx) for a recent, very useful discussion of this traditional picture.

⁵See van Fraassen (1973), and Horty (2012).

⁶See Makinson (1999), and Makinson and van der Torre (2000).

a requirement that an agent performs some action entails a requirement that the agent performs any logically equivalent action. It is often useful to interpret requirements (as well as imperatives, or defaults) as the actual commands issued by various authorities. Taken this way the set of requirements bearing on some individual might well have no interesting closure properties at all: an authority might issue a command without bothering to issue all logically equivalent commands, for example.

Concerning the second question, how requirements interact to support oughts, Broome writes (Section 7.5) that it may take a substantive deontic theory to describe this process, and that “there is no reason to expect the result to resemble a logic” (p. 128). I am more optimistic, or perhaps simply have a broader conception of what counts as a logic. I agree that an individual faced with conflicting requirements from difference sources—especially when these requirements may carry different, perhaps incomparable, weights, and some may be conditional—will have to engage in a complex process of reasoning in order to determine what he or she ought to do. But it is exactly the task of a logic to help us understand this kind of reasoning. This is what van Fraassen, Makinson, and others try to do, and what I have tried to do in my own work. Reasoning—Broome’s topic, after all—does not just begin with oughts and move in one direction only, forward toward intentions. Instead, it encompasses the process through which oughts themselves are derived from requirements, and often takes place against the background of an existing set of intentions. This is a point I will return to.

Finally, there is the third question, concerning logical properties of oughts themselves, or deontic logic. Broome claims (Section 7.5) that deontic logic is also very thin, like the logic of requirements. At one point he suggests that, if we limit our consideration to a range of actions concerning which the relevant requirements do not conflict, the logic of the oughts they support might simply correspond to the logic of these requirements. I disagree. Suppose, for example, that the law requires you to join the army or perform alternative service, while

morality requires you not to join the army. These requirements, I would argue, support the conclusion that you ought to perform alternative service, even though there is no particular requirement that corresponds to this ought.

So here is a point in deontic logic on which Broome and I differ. I could list others, but rather than cataloguing differences, I want to consider some ways in which Broome's larger framework, and especially its emphasis on *enkrasia*, provides a setting in which disagreements like this can be situated, and perhaps resolved.

Oughts, intentions, plans

Consider a standard deontic principle: closure of oughts under logical consequence. I have always accepted this principle, on the basis of arguments such as the following: Suppose you have promised to give Jo \$10 tomorrow, so you ought to. Suppose also that you would like to go to a movie, but have only \$20, with no prospect of getting any more, and the movie costs \$12. Closure under consequence then allows you to conclude that you ought not to go to the movie, which seems right. Broome, and many others, reject closure, usually by pointing out that it yields linguistic oddities. A famous example is this: from the fact that you ought to mail the letter, it follows by closure that you ought to mail the letter or burn it. But focusing on linguistic oddities like this leaves the path open for defenders of closure to argue that these oddities are merely pragmatic, and provide no real reason to reject the principle as correct from a semantic standpoint.

Here is where it helps to place the issue in a broader context. According to Broome's principle of *enkrasia*, once you realize that you ought to pay Jo \$10 tomorrow, rationality requires you to intend to do so. How should we understand this intention? Following the work of Michael Bratman, I think it is best to think of intentions as tied up with the human activity of planning, as well as the resulting plans—and I think, further, that it is useful to develop this idea by drawing on some of the work on planning and plan management from

the field of artificial intelligence.⁷

In this literature, plans are usually taken as partially ordered sets of actions aimed at achieving certain goals, with each action having preconditions necessary for its performance as well as specified effects; the set of actions from a plan is subject to complex temporal constraints guaranteeing that the effects of one action do not interrupt the preconditions necessary for another. At any given point, an agent's overall plan will typically be incomplete in various ways: actions may not be fully specified, the preconditions for certain actions may not be guaranteed, and, although constrained, certain actions may not be precisely scheduled. The focus in artificial intelligence is on developing algorithms for constructing complete plans, but it is often best for agents to avoid committing to complete plans prematurely, to allow for future flexibility. Much of what philosophers think of as practical reasoning can be modeled as plan manipulation, or plan management more generally—not only the completion, or elaboration, of plans, but also their modification as new options are presented.

Against this background, when a *basic* ought—one generated directly by a requirement—leads by enkrasia to an intention, it is natural to take that intention as a goal, to be planned for on its own, or more realistically, to be incorporated into the agent's background plan. For example, when you conclude from your promise to pay Jo \$10 tomorrow that you ought to do so, and then move by enkrasia to an intention to pay Jo \$10, it is natural to interpret this intention as a new goal to be planned for; the plan might involve making sure that you have \$10, and then finding a way to get to Jo's house to give her the money.

By contrast, oughts that are *derived* through closure do not necessarily interact with enkrasia in the same way, to yield intentions that can be interpreted as goals. Suppose as before that, although you would like to see a movie tonight, you conclude through closure that you ought not to do so, because paying Jo \$10 tomorrow entails not seeing a movie

⁷The central source for Bratman's view is Bratman (1987), of course; work from artificial intelligence exploring these ideas includes Bratman, Israel, and Pollack (1988), Pollack (1992), Pollack and Horty (1999), and Horty and Pollack (2001).

tonight. Although it may sound reasonable to conclude, in this way, that you ought not to see a movie, the resulting ought is not one that interacts with *enkrasia* to yield an intention that can be interpreted as a goal. You do not have a goal of not seeing a movie—in fact, you would like to see a movie, if only you could. The derived ought can be understood, at best, only as registering the fact that seeing a movie would interfere with other goals to which you are already committed.

The contrast between these two kinds of oughts, basic and derived, can be put into sharper relief if we consider, not just planning, but replanning. Suppose your plan for paying Jo \$10 involves driving to her house to give her the money, but you learn that your car has a dead battery. You would then replan to get to Jo’s house, in order to pay her the money, in some other way—perhaps walking, or riding your bike. By contrast, if you discover that you can, in fact, see a movie—perhaps you learn that you had more money than you thought—you would not then replan in order not to be able to see a movie in some other way.

The upshot is this: Broome’s framework provides a setting in which we can see how the principle of closure under consequence can indeed be questioned—contrary to standard deontic logic, and contrary to my own previous opinion. Even when the principle does seem to yield reasonable conclusions, as in the case we have been considering, the oughts that are derivable only through closure play a different role in our practical reasoning, and perhaps carry a different meaning, from the basic oughts generated directly by requirements.⁸ The principle of *enkrasia* applies to basic oughts to yield intentions that can be interpreted as goals by a planning agent, but *enkrasia* does not necessarily apply in the same way to oughts that are derived only through closure.

Consider another deontic principle: consistency, which rules out deontic conflicts, or dilemmas—situations in which an agent ought to perform each of two incompatible actions. I have argued against this principle, since I can imagine, or think I can imagine, situations

⁸See Nair (2014) for a different kind of argument for the claim that ought statements as premises differ in meaning from ought statements as conclusions.

in which an agent ought to perform incompatible actions. Suppose, for example, you have inadvertently promised to have a private dinner with each of two identical and identically situated twins, Jack and Jo, at the same time on the same evening. Then, on the basis of your promises, you are required to have dinner with Jack, and also with Jo. Since these two requirements carry either equal or incomparable weight, neither defeats the other, and so I would say that you ought to have dinner with Jack, and with Jo, though it is impossible to dine with both. Broome, by contrast, endorses deontic consistency, denying the possibility of conflicts like this, largely on linguistic grounds—he suggests that deontic consistency follows from the meaning of the English word “ought.” Many other philosophers have thought so too, though linguists have denied it.⁹ Some allow conflicts between what they call weak deontic modals, such as “ought,” but deny the possibility of conflicts between strong deontic modals, such as “must” or “have to,” though others believe that even strong deontic modals can conflict.

Here, too, I think it helps to shift the discussion from purely linguistic considerations to the broader framework Broome lays out. Let us assume, as I think even those who accept deontic conflicts would allow, that intentions cannot conflict—a rational agent cannot intend to perform each of two incompatible actions. And let us assume also, with Broome, that requirements can conflict. Then we cannot accept both *enkrasia*, so that oughts lead to intentions, as well as the idea that conflicting requirements of equal or incomparable weights generate conflicting oughts. For in that case, your conflicting promises to Jack and Jo, for example, would lead you to the conflicting conclusions that you ought to dine with Jack, and also with Jo, and so by *enkrasia* to conflicting intentions to dine with Jack, and also with Jo.

This is an unacceptable result, I agree, but the question of where to break the chain of reasoning leading to this result is a matter of overall agent architecture, not just linguistic intuitions. It is possible to break this chain right at the start, by stipulating that conflicting

⁹See, for example, von Fintel (2012).

ought are simply never generated, even by conflicting requirements of equal or incomparable weight. But it is also possible to break the chain of reasoning later on by treating *enkrasia* itself as a *pro tanto*, or defeasible, requirement. The idea would then be that, although oughts generally require intentions, there are certain cases in which an agent might realize that he or she ought to perform some action without forming the intention to do so, perhaps because the agent is already committed to a conflicting course of action, or perhaps because, although not yet committed, the agent is forced to choose between conflicting oughts. These two strategies—ruling out conflicting oughts vs. weakening *enkrasia*—need to be compared from the standpoint of agent architecture. My own feeling is that the strategy of weakening *enkrasia* is very attractive.

Having explored ways in which Broome’s framework allows us to see issues from deontic logic in a new light, I want to suggest one way in which the framework itself can be broadened.

Broome writes—as I mentioned earlier—as if practical reasoning moves in one direction only, forward from oughts to intentions. But reasoning moves in the other direction as well. An agent’s intentions and plans are not merely responsive to the agent’s conclusions about what he or she ought to do; they also helps to shape the reasoning that leads to these conclusions. Suppose, for example, you would like a new shirt. Perhaps you think aesthetic considerations require you to buy one. But then you realize that buying a shirt involves a trip to the mall. You know how depressing the mall is—the waste, the consumer culture, the bad food—and so you decide, all things considered, that you ought not to buy a new shirt. Here, then, is one simple way in which reasoning about intentions and plans affects reasoning about oughts. Constructing a plan to achieve a goal brings you to understand that the goal is not worth the cost of achieving it, so that what you thought you ought to do is now something you think you ought not to do.

But suppose you remember that it is your son’s birthday next weekend, and that you promised to buy him a train set for his birthday, leading you to conclude that you ought to

do so. Through enkrasia, this ought leads you to an intention, or goal, to buy a train set. You realize that the only place to buy a train set is the mall, and so you adjust your overall plan to include a trip to the mall. At this point, realizing that you will be going to the mall after all, it is natural to revise your earlier decision and conclude that, while you are there, you ought to buy yourself a shirt. So here is another way in which reasoning about plans affects reasoning about oughts. What you ought to do can vary depending on the plan to which you are already committed. In the context of a plan that does not already include a trip to the mall, it is reasonable to conclude that you ought not to buy a shirt, but in the context of a plan that does include a trip to the mall, it is reasonable to conclude that you ought to.

My suggestion is the enkrasia requirement should be generalized to reflect the fact that oughts might be conditional on an agent's current plans, and also the fact that plans themselves might be more complex than simple intentions. Setting aside the various restrictions on enkrasia explored in Broome's book, and the further restrictions suggested here, the new principle would read: Rationality requires that, if an agent believes that, in the context of a plan to which the agent is committed, the agent ought to perform some action, then the agent's plan includes that action as a goal. This generalization would collapse into Broome's original version of enkrasia in the special case in which the agent has no current plan, but would require of you, in case your current plan involves going to the mall, that your plan should include buying a shirt.

References

- Bratman, Michael (1987). *Intention, Plans, and Practical Reason*. Harvard University Press.
- Bratman, Michael, Israel, David, and Pollack, Martha (1988). Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4:349–355.
- Broome, John (1999). Normative requirements. *Ratio*, 12:398–419.

- Broome, John (2004). Reasons. In Wallace, R. Jay, Smith, Michael, Scheffler, Samuel, and Pettit, Philip, editors, *Reason and Value: Essays on the Moral Philosophy of Joseph Raz*, pages 28–55. Oxford University Press.
- Chellas, Brian (1980). *Modal Logic: An Introduction*. Cambridge University Press.
- Horty, John (2012). *Reasons as Defaults*. Oxford University Press.
- Horty, John and Pollack, Martha (2001). Evaluating new options in the context of existing plans. *Artificial Intelligence*, 127:199–220.
- Lord, Errol and Maguire, Barry (20xx). An opinionated guide to the weight of reasons. In Lord, Errol and Maguire, Barry, editors, *Weighing Reasons*. Oxford University Press.
- Makinson, David (1999). On a fundamental problem of deontic logic. In McNamara, Paul and Prakken, Henry, editors, *Norms, Logics, and Information System: New Studies in Deontic Logic and Computer Science*, pages 29–53. IOS Press.
- Makinson, David and van der Torre, Leendert (2000). Input/output logics. *The Journal of Philosophical Logic*, 29:383–408.
- Nair, Shyam (2014). Consequences of reasoning with conflicting obligations. *Mind*, 123:753–790.
- Pollack, Martha (1992). The uses of plans. *Artificial Intelligence*, 57:43–68.
- Pollack, Martha and Horty, John (1999). There’s more to life than making plans: plan management in dynamic, multi-agent environments. *AI Magazine*, 20:63–72.
- van Fraassen, Bas (1973). Values and the heart’s command. *The Journal of Philosophy*, 70:5–19.

von Fintel, Kai (2012). The best we can (expect) to get? Challenges to the classic semantics for deontic modals. Paper presented at the Central APA, February 17, 2012; available at <http://web.mit.edu/fintel/fintel-2012-apa-ought.pdf>.