

---

---

+ • +

---

---

# THE JOURNAL OF PHILOSOPHY

VOLUME CII, NO. 9, SEPTEMBER 2005

---

---

+ • +

---

---

## THE WRONG KIND OF REASON\*

A good number of people currently thinking and writing about reasons identify a reason as *a consideration that counts in favor of an action or attitude*.<sup>1</sup> I will argue that using this as our fundamental account of what a reason is generates a fairly deep and recalcitrant ambiguity; this account fails to distinguish between two quite different sets of considerations that count in favor of certain attitudes, only one of which is the “proper” or “appropriate” kind of reason for them. This ambiguity has been the topic of recent discussion, under the head “the wrong kind of reasons problem.”<sup>2</sup> I

\* Earlier versions of this material benefited from conversation or commentary from Tyler Burge, Sarah Buss, John Carriero, Sonny Elizondo, Mark Greenberg, Aaron James, Sean Kelsey, Niko Kolodny, Clayton Littlejohn, Matthew McAdam, T.M. Scanlon, and audiences at Princeton University, the University of Michigan, and the Pacific Division Meeting of the American Philosophical Association.

<sup>1</sup> See, for example, Jonathan Dancy, *Practical Reality* (New York: Oxford, 2000); Derek Parfit, “Rationality and Reasons,” in Dan Egonsson et al., eds., *Exploring Practical Philosophy* (Burlington, VT: Ashgate, 2001), pp. 17–39; Joseph Raz, *Practical Reason and Norms* (London: Hutchinson, 1975; reprint, Princeton: University Press, 1990); Raz, *Engaging Reason: On the Theory of Value and Action* (New York: Oxford, 1999); T.M. Scanlon, *What We Owe to Each Other* (Cambridge: Harvard, 1998); J. David Velleman, *The Possibility of Practical Reason* (New York: Oxford, 2000). Obviously, these writers have in mind what are sometimes called “justifying” reasons. What are sometimes called “explanatory” reasons would require different treatment. John Broome takes the explanatory role of reasons as primary, and identifies a justificatory (or “normative”) reason as a fact that explains an ought claim. See Broome, “Reasons,” in R. Jay Wallace et al., eds., *Reason and Value: Themes from the Moral Philosophy of Joseph Raz* (New York: Oxford, 2004), pp. 28–55.

<sup>2</sup> See, for example, Jonas Olson, “Buck-Passing and the Wrong Kind of Reasons,” *The Philosophical Quarterly*, LIV, 215 (2004): 295–300; Justin D’Arms and Daniel Jacobson, “Sentiment and Value,” *Ethics*, cx (2000): 722–48; Wlodek Rabinowicz and Toni Rønnow-Rasmussen, “The Strike of the Demon: On Fitting Pro-Attitudes and Value,” *Ethics*, cxiv (2004): 391–423. The “wrong kind of reasons” problem, as identified by these authors, concerns reasons for “valuing” or for those “pro-attitudes” or “sentiments” that are thought to be appropriate when their object is valuable. So understood, the wrong kind of reasons problem is thought to challenge the viability of a

will suggest that confusion about “the wrong kind of reason” will be dispelled by changing our account of what a reason is. While I agree both that reasons are considerations and that certain considerations count in favor of actions and attitudes and thus provide reasons for them, we should not understand “counting in favor of an action or attitude” as the fundamental relation in which a consideration becomes a reason. Rather, we would do better to think of a reason as *a consideration that bears on a question*. This alternative will allow us to distinguish “kinds” of reasons by distinguishing between kinds of questions on which a consideration can bear, and it will allow us to distinguish the “right” kind of reasons for certain attitudes from the “wrong” kind, by allowing us to consider the relation between the question on which the consideration bears and the attitudes of which it counts in favor.

#### REASONS AS CONSIDERATIONS IN RELATION

Before explaining how the original account generates the ambiguity (and how the alternative avoids it), I will orient discussion by noting some points of commonality between the original formulation and my alternative. Both accounts understand reasons as considerations standing in some sort of relation. In agreeing that reasons are considerations, it may seem I am taking sides in a debate about reasons for action. There has been considerable debate about whether reasons for action are mental states—beliefs and desires, in particular—or rather are “considerations,” that is, the propositions, facts, states of affairs, events, or objects that might serve as the *content* of such mental states, that which the beliefs or desires are about.<sup>3</sup> Which side one

---

certain account of value (the “fitting-attitudes” or “neo-sentimentalist” account). In fact, I believe this discussion involves several distinguishable issues, only one of which I will address: I will consider the question of how to distinguish the “right” reasons for an attitude from the “wrong” kind. I will not focus particularly on “valuing” or on “pro-attitudes,” because, as noted by D’Arms and Jacobson, as well as by Parfit and Christian Pillar (references below), there is a quite general problem about identifying the appropriate reasons for attitudes, a problem that is not restricted to reasons for those attitudes involved in valuing. I hope to resolve this *broader* problem about the “wrong kind of reason.” Even if I am successful, problems about value will remain. I consider these briefly late in the paper.

<sup>3</sup> Donald Davidson is often associated with the view that reasons for actions are belief/desire pairs. See, for example, Davidson, “Actions, Reasons, and Causes,” in *Essays on Actions and Events* (New York: Oxford, 1980), pp. 3–19. It is not at all clear to me, however, that Davidson had the distinction in question in mind. The view the reasons for action are belief/desire pairs has been recently defended by, among others, Michael Smith. See, for example, “The Possibility of Philosophy of Action,” in Jan Bransen and Stefaan E. Cuypers, eds., *Human Action, Deliberation, and Causation* (Boston: Kluwer, 1998), pp. 17–41, and *The Moral Problem* (Cambridge: Blackwell, 1994). The contrasting view has been argued by, among others, Stephen

favors in this debate typically turns on whether one is trying to explain action or rather trying to determine how or whether to act. While it can seem obvious that actions are to be explained by facts about the mental life of the agent, it can seem equally obvious that reasons which settle the question of how or whether to act are facts of a more general sort—including facts about the weather or the time or how much something costs or how far away the movie theater is. For present purposes, I hope it can be agreed on all sides that, once we widen the field to include not only reasons for action but also reasons for beliefs and desires (and even reasons for weather patterns), reasons cannot be restricted to mental states. But “considerations,” broadly construed—facts about the time or the location of the movie theater, or facts about the perceptual system of a believer, the desires of an agent, or the evolution of the species—could, it seems, be reasons of a more general sort. So I will presume that reasons can be understood as considerations, broadly construed.<sup>4</sup>

I also share with the original account the assumption that considerations do not provide reasons in isolation, so to speak. The fact that my eyesight is poor or that my office is on the third floor is not yet a reason. It is just a fact. The same fact could provide many different reasons. The fact that my office is on the third floor might be a reason to take the elevator, the reason it is warm in here, the reason my students do not come to see me, a reason to think I am highly favored, a reason to feel guilty, and so forth. Considerations become reasons only by standing in some sort of relation.<sup>5</sup>

---

Darwall, Thomas Nagel, Raz, Warren Quinn, and Scanlon. A recent, extended, interesting discussion can be found in Dancy, *Practical Reality*. I consider reasons for action in “Reasons for Action” (in progress).

<sup>4</sup> I am at this point overlooking disagreement about whether a reason must be a fact, whether it can be a purported fact, whether it is a proposition, or whether it is the kind of thing that can be the subject matter of a proposition, and so on. I am taking “consideration” loosely enough to avoid commitments on these issues. (From the point of view of the speaker, a reason is always a fact, a true proposition or its subject matter. However, reasoners are often mistaken. When a reasoner is mistaken about a matter of fact, she could mistakenly take a consideration that is not [in fact] a fact to be a reason. In such a case, the reasoner acts on reasons that are not facts. Some will say they are, therefore, not “really” reasons, or not “good” reasons. While I would insist that such nonfacts really are *that person’s* reasons, it seems right to say that they are not good reasons, and, further, that to take something as one’s reason is to take it to be, in this sense, a good reason—that is, that reasons must at least be taken to be facts, from the point of view of the one for whom they are reasons.)

<sup>5</sup> Raz notes that it is natural to think of a reason for action as a relation between “a fact which is the reason...a (class of) person(s) and an action(-type)” (“Agency, Reason, and the Good,” in *Engaging Reason*, p. 22). See also Raz, *Practical Reasons and Norms*. John Skorupski gives a similar, though more complex, account of what he calls “the reasons relation.” He analyzes reasons for action as follows: “the fact that *p* gives [person] *x* reason of degree *d* at time *t* to  $\psi$ , where  $\psi$  ranges over action-types.” He

My alternative differs from the original formulation in the relation in which it takes a reason to stand.

#### COUNTING IN FAVOR OF ATTITUDES

Taking a reason to be a consideration that counts in favor of an action or attitude generates a fairly deep and recalcitrant ambiguity in an important range of cases—it generates the wrong kind of reasons problem.<sup>6</sup>

The problem appears most readily in reasons for believing. A consideration can count in favor of believing in two quite different ways: it can show the content of the belief true or it can show the belief, as an attitude, in some way good (useful, convenient, and so on) to have. The fact that the butler wanted revenge is a reason, of the first sort, for believing the butler is guilty; the fact that believing him guilty is the only way to save your life is a reason, of the second sort, for believing he is.<sup>7</sup>

The ambiguity also appears in reasons for intending. A consideration can count in favor of intending to  $\phi$  either by bearing on whether to  $\phi$  or by showing an intention to  $\phi$  is in some way good to have. Consider the politics of deterrence. You might find yourself convinced that an intention to retaliate is good to have—for its deterrent ef-

provides a parallel account of reasons for belief, “the fact that  $p$  gives [person]  $x$  reason of degree  $d$  at time  $t$  to believe that  $q$ ,” and for what he calls “reasons to feel,” “the fact that  $p$  gives [person]  $x$  reason of degree  $d$  to feel  $\phi$ , where  $\phi$  ranges over types of feelings (that is, emotions, moods, and desires)” (“Reasons and Reason,” in Garrett Cullity and Berys Gaut, eds., *Ethics and Practical Reason* (New York: Oxford, 1997), pp. 345–67, on p. 346). According to Scanlon, the relevant relation is “counting in favor of” some other item. Scanlon takes a reason to be a consideration that counts in favor of something (*What We Owe to Each Other*, pp. 17–18). As Scanlon’s discussion proceeds, one learns that the particular item reasons seem to count in favor of are what he calls “judgment-sensitive attitudes,” attitudes such as belief, desire, intention, fear, resentment, or jealousy. So long as any action involves an intention, Scanlon’s account will cover both reasons for actions. Parfit puts his similar view quite simply: “Facts give us reasons...when they count in favour of our having some belief or desire, or acting in some way” (“Rationality and Reasons,” p. 18). Dancy says, “a reason is a consideration that speaks in favor of action (or belief)” (*Practical Reality*, p. 144). However, Dancy provides some further reflection on what he calls the “ground” for a reason. Whereas, for Scanlon, the notion of “counting in favor of” is primitive (and values are understood through this notion, via the account of reasons), for Dancy reasons are grounded in values, and thereby count in favor of things. My account is closer to Scanlon’s, in that it does not make explicit use of a conception of value.

<sup>6</sup> It also creates difficulty for the already difficult project of explaining action. I consider the explanation of action, and argue for my preferred account of a reason, in “Reasons for Action.”

<sup>7</sup> Why would believing the butler did it save your life? Perhaps the mob has made credible threat to kill you if you do not give convincing testimony of his guilt, and perhaps you are a terrible liar.

fects—without thereby thinking retaliating is to be done.<sup>8</sup> Reasons for wanting are similarly ambiguous: the fact that you must win tonight's game in order to advance to the championships is a reason to want to win; however, the fact that lacking a desire to win would disappoint or offend your teammates is also a reason—of a different sort—for wanting to win.

We could continue to draw out an ambiguity in reasons for fearing, resenting, envying, admiring, forgiving, and trusting, among others.

Both Derek Parfit and Christian Piller have noticed the breadth of this ambiguity. Parfit labels it as a distinction between “object-given” and “state-given” reasons, while Piller uses the labels “content-related” and “attitude-related.”<sup>9</sup> Neither gives a precise account of how to draw the distinction. Both simply suggest that object-given or content-related reasons are provided by or have something to do with the object or content of the attitude, while state-given or attitude-related reasons are provided by or have to do with the state or attitude itself.<sup>10</sup>

As noted by Piller, some might attempt to avoid the problem by denying that attitude-related reasons really are reasons for the attitude. Rather, the supposed attitude-related reasons for believing *p*, say, are really reasons for *wanting* or *trying* to believe *p*. The only reasons that really count in favor of believing *p* are those that show *p* true.<sup>11</sup>

While it is obviously true that only certain reasons are the “real” or “appropriate” kind of reasons for believing—that others are of the

<sup>8</sup> Cf. Gregory Kavka, “The Toxin Puzzle,” *Analysis*, XLIII (1983): 33–36.

<sup>9</sup> See Parfit, “Rationality and Reasons”; Piller, “Normative Practical Reasoning,” *Proceedings of the Aristotelian Society*, Suppl. Vol. xxv, 1 (2001): 195–216, and “Content-Related and Attitude-Related Reasons for Preferences,” *Philosophy* (forthcoming).

<sup>10</sup> Parfit introduces the distinction when discussing reasons for desires, and says, “Of our reasons to have some desire, some are provided by facts about this desire’s object. These reasons we can call object-given.... Other reasons to want something are provided by facts, not about what we want, but about having this desire. These reasons we can call state-given” (“Rationality and Reasons,” p. 21–22). Piller, who notes that he does not aim to give a precise definition, says, “Take some propositional attitude for which there could be reasons, like wanting *p*, intending that *p*, or believing that *p*. Something is a content-related reason if it refers to some appropriate feature of *p*. For example, believing that *p* is useful would be a content-related reason for wanting that *p*. In other words, *p*’s usefulness is a consideration that speaks in favor of wanting that *p*. Believing that *p* is the best explanation of *q*, where *q* is something believed to be true, is a content-related reason for believing that *p*.... Something is an attitude-related reason for a propositional attitude if it refers to some appropriate feature of the attitudes as a whole and not only to a feature of the content of the attitude. Take wanting that *p*. One of the features of the attitude as a whole would be that *wanting that p is useful*. Believing that *wanting that p is useful* would then be an attitude-related reason for wanting that *p*” (“Normative Practical Reasoning,” pp. 204–05). The account lacks precision because of its reliance on ‘appropriate’.

<sup>11</sup> See Piller, “Normative Practical Reasoning,” p. 206. He argues against this view in “Content-Related and Attitude-Related Reasons for Preferences.”

wrong kind—and while I happily agree that attitude-related reasons for one attitude can serve as content-related reasons for another attitude (namely, a second-order attitude whose content concerns the first), the problem lies in saying why the attitude-related reasons are not “really” reasons for the attitude in question. Reasons, on the original formulation, are considerations that count in favor of an action or attitude. Considerations that show a belief good to have surely “count in favor of” believing, quite independently of whether they show the belief true. Considerations that show something good about believing  $p$  count in favor of believing  $p$  in just the same way that considerations that show something good about attending lecture count in favor of attending lecture. Likewise, the fact that supposing  $p$  would advance the argument counts in favor of supposing  $p$  and is, thereby, a reason for supposing  $p$ . But if the fact that supposing  $p$  would advance the argument can be a reason to suppose  $p$ , the fact that believing  $p$  would advance my career should, for all that has been said, be a reason for believing  $p$ . If a reason is simply a consideration that counts in favor of an action or attitude, attitude-related reasons should qualify as reasons for the attitude of which they count in favor.<sup>12</sup>

Further, for certain propositional attitudes, such as supposing, imagining, and remembering, the “real” reasons seem to be attitude-related, while reasons that bear on the content are reasons for a different attitude. Suppose, for example, that you are asked to imagine there’s no heaven. We can draw a distinction between considerations that bear on the content of the imagining and those that instead simply show imagining in some way good to do. Divine revelation or the human capacity for wishful thinking would bear on the content of the imagining—whether there is a heaven—while the fact that you have been asked to imagine it, or the fact that it’s easy if you try, bear on whether to imagine. Here, the “real” reasons for the attitude seem to be the attitude-related reasons, while reasons that bear on the content seem to be the proper reasons for a different attitude—reasons to believe.

Some might try to show that the content-related reasons are the only “real” reasons for believing by appealing to standards of justification or rationality: although the attitude-related reasons *count in favor* of believing, they are not really reasons for believing, because they do

<sup>12</sup> Rabinowicz and Rønnow-Rasmussen agree with me in this: “Surely, if we are supposed to have reasons for *actions* when the actions have useful effects or are valuable for their own sake, why shouldn’t we have reasons for *attitudes* in comparable circumstances?” (“Strike of the Demon,” p. 413). See also Piller, “Normative Practical Reasoning,” pp. 211–15.

not *justify* or *rationalize* belief. This appeal, simply so put, is unsatisfactory. Once we grant that reasons are considerations that count in favor of an action or attitude, attitude-related reasons will themselves provide corresponding standards of rationality and justification.<sup>13</sup> If you believe that the butler did it because believing this will save your life, then there is a sense in which you have good reason to believe it, and so a sense in which you are justified and even rational in doing so. Of course, this is justification of the wrong kind. But we need some account of why it is the wrong kind. Thus, simple appeal to notions of justification or rationality seems to rely on, rather than elucidate, the distinction we are after. For an appeal to justification or rationality to be satisfactory, we would need an independent account of why only *some* of the reasons that count in favor of an attitude justify or rationalize it—which is just to say, we would need an answer to the wrong kind of reasons problem.

So we are left with the ambiguity. As long as a reason is simply a consideration that counts in favor of an attitude, we are left without an obvious way either to draw a useful distinction between these very different sorts of reasons or to say why one of them seems to be the “real” sort of reasons.

#### AN ALTERNATIVE ACCOUNT

In light of this problem, I suggest we abandon the thought that a reason is a consideration that counts in favor of an action or attitude and look for a different account.<sup>14</sup>

To start reflection, we can note that, most generally, a reason is simply an item in a piece of reasoning. What is a piece of reasoning?

<sup>13</sup> Gibbard notes the connection between reasons and rationality in *Wise Choices, Apt Feelings* (Cambridge: Harvard, 1990). He distinguishes the rationality of “*feeling* or *believing* something” from the rationality of “*wanting* to feel or believe it” at p. 37. D’Arms and Jacobson discuss his treatment. See “Sentiment and Value,” pp. 412–15. Parfit considers the matter of rationality in “Rationality and Reasons” at p. 24. Both treatments are discussed by Rabinowicz and Rønnow-Rasmussen, “Strike of the Demon,” pp. 412–15.

For an argument that pragmatic considerations do not rationalize belief, see Thomas P. Kelly, “The Rationality of Belief and Some Other Propositional Attitudes,” *Philosophical Studies*, cx (2002): 163–96. For an argument that epistemic rationality does not reduce to instrumental rationality, see Kelly, “Epistemic Rationality as Instrumental Rationality: A Critique,” *Philosophy and Phenomenological Research*, LXVI, 3 (2003): 612–40.

<sup>14</sup> As noted, Broome offers a different account, according to which reasons explain, and “normative” reasons explain ought facts. He then argues against the view that counting in favor is the basic normative notion saying, “it cannot be, because it is complex. It incorporates the two elements of normativity and explanation” (“Reasons,” p. 41). It seems to me that this argument will convince only if one is already persuaded that a reason is a fact that explains an ought claim. I hope to have provided reasons to be unsatisfied with the notion of “counting in favor of an action or attitude”

Remaining very vague, we can at least contrast a piece of reasoning with a mere collection of claims, propositions, or concepts, or a mere series of thoughts, contents, or representations, by noting that a piece of reasoning has some direction to it—it is the sort of thing that might have a conclusion, in something like the way an argument has a conclusion.<sup>15</sup> And the items in a piece of reasoning (again unlike the items in a mere collection of claims or propositions or in a mere series of thoughts or contents) stand in certain relations to each other—they bear what we might call “rational relations” to one another—and thereby also bear on the conclusion. Further, reasoning seems called for just in case a conclusion is unknown or called into question. So I suggest, for consideration, the following account of a reason: a reason is a consideration that bears on a question.<sup>16</sup>

This account differs from the original formulation in taking the fundamental relation in which a consideration becomes a reason to be a relation to a question, rather than to an action or attitude. The new formulation will allow us to distinguish between the content- and attitude-related reasons by distinguishing between questions. Perhaps surprisingly, distinguishing content- and attitude-related reasons will not allow us to identify the “right kind” of reasons for an attitude. However, on the new formulation, the relation between attitudes and reasons is mediated by a question. Thus, this formulation requires us to consider the relation between the question on which the consideration bears and the attitude of which it counts in favor. Considering this relation will help to identify the right kind of reason.

#### SORTING REASONS: CONTENT- AND ATTITUDE-RELATED<sup>17</sup>

A consideration can “count in favor” of certain attitudes by bearing on either of two distinct questions. Consider, again, belief. A consideration can count in favor of believing  $p$  by bearing on whether  $p$  or by bearing on whether the belief that  $p$  is in some way good to have. We can draw the distinction between content- and attitude-related reasons for belief accordingly. Reasons that bear on the first question

independently of such an account.

<sup>15</sup> Gilbert Harman has done much to differentiate reasoning from argument. See, for example, *Change in View* (Cambridge: MIT, 1989). I do not mean, here, to assimilate them, nor to choose between an account of reasoning as something concrete agents undertake in time and space and an account of reasoning as a set of abstract relations.

<sup>16</sup> One could say, “a consideration that bears on a conclusion.” I do not think there would be any relevant difference, though I find the idea of answering a question more intuitive for capturing the activities of rational agents.

<sup>17</sup> I also discuss how to draw this distinction in “Controlling Attitudes,” *Pacific Philosophical Quarterly* (forthcoming).



(whether  $p$ ) are content-related reasons for a belief that  $p$ ; reasons that bear on the second question (whether the belief that  $p$  is in some way good to have) are attitude-related reasons for that belief. Likewise for intention: a consideration can count in favor of intending to  $\phi$  by bearing on whether to  $\phi$  or by bearing on whether the intention to  $\phi$  is good to have. Content-related reasons for intending to  $\phi$  bear on the first question (whether to  $\phi$ ); attitude-related reasons bear on second (whether the intention to  $\phi$  is good to have).

We can generalize the method: attitude-related reasons count in favor of an attitude by bearing on the question of whether the attitude is, in some way, good to have. Content-related reasons somehow manage to count in favor of an attitude by bearing on a different question—a question that typically does not mention the attitude itself, but concerns its content.

Notice two features of this distinction, so drawn: first, we have not given a general way to specify the content-related question. In the central cases of belief and intention, the question is clear enough (whether  $p$  or whether to  $\phi$ , respectively). However, when generalized to other attitudes (desiring, resenting, admiring, trusting), the question is not obvious, and we have not given a way to identify it. Rather, we have specified the attitude-related question in a general way (whether the attitude is in some way good to have); we then identified the content-related reasons as those that somehow manage to count in favor of the attitude by bearing on some *other* question, presumably one which has to do with the attitude's content.<sup>18</sup>

Second, note that, even in those cases in which the content-related question is clear enough (the cases of belief and intention), the distinction between content- and attitude-related is not exclusive: the answer to one of these questions can bear on the other question, and thus reasons that provide an answer to one question can, in turn, become reasons of the other sort. For example, quite often (though not always), it is good to have true beliefs. Thus, by showing the content of a belief to be true, content-related reasons might also

<sup>18</sup> This formulation will surely capture a broader range of reasons than the "right kind." For example, a consideration might count in favor of resenting by bearing positively on whether the person resented would be offended by being (instead) forgiven. This reason counts in favor of resenting by bearing on a question other than whether resenting is good to do (though it will, eventually, come to bear on that question), which question has to do with the content of the attitude (the person). One might try to say that the "right kind" of reasons are those that somehow manage to count in favor of an attitude *without* bearing on the question of whether the attitude is in some way good to have. But this will be too restrictive, as shown by the next paragraph.

thereby show the belief is good to have. For beliefs with certain contents (such as the admittedly odd, “This very belief is good to have”) reasons showing the belief good to have may also, thereby, bear on whether its content is true.<sup>19</sup> Likewise, it is almost always good to intend to do what there is reason to do; so content-related reasons for intentions are almost always also attitude-related reasons for that intention (an exception would be the problematic cases in which there is reason to do something only if you do it unintentionally). Finally, reasons that show an intention good to have can readily bear on whether to act. Your complaint that I have no intention of attending your party can be a reason for me to attend.<sup>20</sup> In general, the considerations that answer the content-related question might also, thereby, show the attitude good to have, while considerations that show the attitude good to have might (at least in principle) bear on the content-related question.

Previous attempts to resolve the wrong kind of reasons problem have relied on something like the content-/attitude-related distinction. But, since this distinction is not exclusive, it will not do the job; it cannot distinguish between the “right” kind of reasons for believing or intending and those that are somehow inappropriate.<sup>21</sup> Moreover, because we do not have a general way to specify the content-related

<sup>19</sup> In addition to the peculiar, “This belief is good to have,” consider the following case: Suppose you are interviewing people to run your propaganda campaign, and you find yourself thinking it would be good to believe that a certain person would be best for the job. Moreover, you find yourself thinking the belief would be good to have as a result of the candidate’s own activities. In this case, the fact that it would be good to believe that she will be an effective leader of your campaign might, itself, be a reason that shows that she will. Here, *p* is a proposition whose truth you think is made more likely by the fact that you have a reason for wanting to believe it. I consider this case in “Controlling Attitudes.”

<sup>20</sup> If one doubts that the reason you have given me, here, is really attitude-related, we can change the case. Suppose you have a reliable intention detector. As we sit in your lab, you are interested in whether intentions for different kinds of action produce differing patterns on your instrument, and so ask me to intend to drink a glass of water. I know what to do: I drink the water. My reason for drinking is that I know that, by deciding to drink, I will produce in myself the desired intention. Kavka’s “Toxin Puzzle” provides an interesting case in which reasons that show the intention good to have do not, I would argue, bear on whether to act.

<sup>21</sup> Rabinowicz and Rønnow-Rasmussen initially attempt to solve the wrong kind of reasons problem by appeal to the fact that certain properties both appear in the “content” of the pro-attitude and, at the same time, justify the attitude (“Strike of the Demon,” p. 414). They are relying on the thought that that the right kind of reasons for valuing must appeal to some property of the thing valued. But this runs into counterexamples, because not all properties of a thing are (really) reasons to value it. They imagine an evil demon who requires you to admire him “on account of his determination to punish [you]” if you do not. He seems to have given you a reason to admire him that is of the wrong kind. Yet, he has appealed to a property of the object (himself) as the basis for the admiration.

question, the distinction cannot be made exclusive. We would like to identify those that bear *only* on the attitude-related question as the “wrong kind” of reason, while identifying any consideration that bears on the content-related question (regardless of whether it also bears on the attitude-related question) as the “right kind.” But we cannot do this without specifying the content-related question. Thus, to solve the wrong kind of reasons problem, generally, we must be able to specify the question on which the right kind of reason bears. We can do this, I now suggest, by considering the relation between *settling a question* and *forming or revising an attitude*.

SORTING REASONS: CONSTITUTIVE AND EXTRINSIC<sup>22</sup>

Note that, for certain attitudes, settling a question amounts to forming the attitude. Suppose you take certain considerations to bear positively on whether *p*. In fact, you take those considerations to settle the question of whether *p*. You are convinced by those reasons. If you find the reasons that (you take to) bear positively on whether *p* to be convincing, and so settle for yourself the question of whether *p*, you have thus, *ipso facto*, formed a belief that *p*. I call the reasons that (are taken to) bear on whether *p* the “constitutive” reasons for believing *p*. Finding them convincing and so settling the question on which they bear amounts to believing.<sup>23</sup>

Using Broome’s understanding of what a reason is (that which explains an ought fact), Olsen suggests that the right kind of reasons are those which explain why you ought to have an attitude *without* appealing to “properties of the attitude in question” (“Buck-Passing and the Wrong Kind of Reasons,” p. 299). But surely there could be cases in which a fact that (properly) explains why one ought to adopt an attitude makes reference to the very attitude. Perhaps one ought only to love (romantically) those who respond well to one’s love. Whatever one thinks of this as an account of love, it should not be ruled out by a general formula about what could qualify as the “right kind of reasons” for an attitude.

Both these attempts, then, fail for the same reason. They attempt to specify the “wrong” kind of reason by appeal to whether the reason appeals to facts about the attitude or to facts about its object. But, as I will suggest in the text, finding the “right” kind of reason requires attending instead to agency, that is, asking, not about what the reason appeals to, but about what taking it to be sufficient amounts to.

<sup>22</sup> I introduce and discuss this distinction at some length in “Controlling Attitudes.”

<sup>23</sup> The reasons are not literally constitutive of the belief. You could maintain the belief while forgetting the reasons. Rather, they are the reasons that bear on the question, the settling of which is constitutive of believing.

Note also, the constitutive reasons are those that *bear or are taken to bear* on the question, the settling of which amounts to believing. There is a question, the settling of which amounts to believing *p* (the question of whether *p*). People can disagree about which considerations really do bear on that question, and to what extent. The distinction between the “right” and “wrong kind” of reasons for believing should not, in my view, settle that disagreement. Rather, once we identify the right question,

Consider, now, the *remaining* reasons for believing  $p$ —those which (are taken to) count in favor of believing  $p$  independently of whether  $p$ . Finding these reasons convincing does not amount to believing  $p$ . Suppose that your life depends on your ability to give convincing testimony that the butler did it, but you do not believe he did, and you are a terrible liar. Now, you might think that you have overwhelming reason showing it good to believe the butler did it—it would save your life. Yet, the fact that you take these reasons to be fully convincing does not itself show that you believe the butler did it—because it does not show that you have settled the question of whether he did. Rather, by finding these reasons convincing you have settled the question on which *these* reasons bear, namely, the question of whether it would be good to believe the butler did it. Thus, by finding these reasons convincing, you form a *second-order* belief *about* the belief that the butler did it: you believe it would be good to believe he did. (You might also desire to believe he did, or, if you have the means, you might intend to bring it about that you believe the butler did it.) Since finding these reasons convincing implies nothing about whether you have the belief of which they count in favor, I call them “extrinsic” reasons for that belief. These extrinsic reasons for believing  $p$  are constitutive reasons for a quite different attitude—a second-order attitude about the belief that  $p$ .<sup>24</sup>

This distinction between constitutive and extrinsic reasons for a belief that  $p$  marks the distinction between the “right kind” and the “wrong kind” of reasons for believing  $p$ . The right kind of reasons for believing  $p$  are those that (are taken to) bear on whether  $p$ —that is, those that (are taken to) bear on the question, the settling of which amounts to believing. Extrinsic reasons are not “really” reasons for believing  $p$ , we can say, because they are not the kind of reasons which, simply by finding convincing, one would believe  $p$ . Rather, by

---

we have identified the right kind of reason. There will be further, substantive issues about which among these reasons are good ones. I return to this point in the body below.

<sup>24</sup> The constitutive/extrinsic reason distinction is closely related to the content-/attitude-related distinction. I suggest that we can identify the content-related question(s) as that (those) the settling of which amounts to forming the attitude. From within a single point of view, the class of constitutive reasons for believing will be identical to the class of content-related reasons for believing—and will include those attitude-related reasons which, by showing the belief is good to have, also thereby bear on the truth of its content. Extrinsic reasons are simply those that are not constitutive. As it turns out, the extrinsic reasons will be identical to the *remaining* attitude-related reasons: they count in favor of the belief by bearing on whether the belief is in some way good to have. (I see no other question on which a reason could bear and so count in favor of the belief.)

finding them convincing, one will form a different attitude—a second-order attitude about the belief that *p*. Extrinsic reasons for a belief that *p* are constitutive reasons for a belief about the belief that *p*.

I suggest that we can, in general, distinguish the right kind of reason for an attitude from the wrong kind, not by appeal to whether the reason is content- or attitude-related—not, that is, by appeal to whether the reasons show the attitude good to have, or whether appreciating the reason requires some sort of reflection upon the attitude itself, or whether it appeals to a feature of the content of the attitude—but rather by considering how we form the attitude, when we do so for reasons—that is, by considering our agency over that attitude. The “right kind of reasons” (are taken to) bear on a question, the settling of which amounts to forming the attitude. The wrong kind of reasons do not bear, or are not taken to bear, on that question.

It is worth noting that, while drawing the constitutive/extrinsic distinction will solve the “wrong kind of reasons” problem, it will not settle any more substantive questions about reasons for an attitude. That is to say, it will not distinguish between (what we might call) good and bad reasons—it will not determine which considerations really do bear on a question, or which considerations really do settle a question. The constitutive reasons *bear or are taken to bear* on a question, the answering of which amounts to having the attitude. These are the right kind of reasons for believing. But the fact that a consideration bears or is taken to bear on a question does not yet mean that it is in fact adequate to settle the question. There is room for various sorts of mistake: mistake of fact, mistake about whether the consideration bears on the question, and mistake about whether a consideration that bears on a question settles the question. Further, there is room to distinguish between “reasonable” and “unreasonable” mistakes of each kind. Thus, there remains room to distinguish good constitutive reasons from bad ones, along each of these dimensions.<sup>25</sup>

#### THE SCOPE OF THE CONSTITUTIVE/EXTRINSIC DISTINCTION

The constitutive/extrinsic distinction can be drawn for any attitude that can be formed or revised simply by settling for oneself a question or set of questions. It seems reasonably clear that belief and intention are such attitudes. Notice that, whenever one has an attitude that can be formed or revised simply by settling for oneself a question or set

<sup>25</sup> Thus it seems to me likely that investigation into the nature of an attitude, while important for drawing the constitutive/extrinsic distinction, will not settle the more pressing question about the justification of that attitude: Which, among the constitutive reasons, would justify it?

of questions (regardless of how the attitude was in fact formed), one is committed to an answer to the relevant question(s). One is committed in the sense that, if one has the attitude, one is answerable to certain questions and criticisms—namely, those questions or criticisms that would be answered by the considerations that bear on the relevant question(s). So, for example, if I believe  $p$ , then I am committed to  $p$  as true, that is, I am answerable to questions and criticisms that would be answered by the considerations that bear on whether  $p$ . Likewise, if I intend to  $\phi$ , then I am committed to  $\phi$ -ing, that is, I am answerable to questions and criticisms that would be answered by considerations that bear on whether to  $\phi$ .<sup>26</sup> I call such attitudes “commitment-constituted.” The reasons that (are taken to) bear on the relevant question (or set of questions) are the “constitutive” reasons for the attitude, not because the reasons *themselves* constitute the attitude (one may, perhaps, form the attitude for no reason, or may maintain the attitude while having long forgot one’s reasons), but because the reasons (are taken to) support the commitment that is constitutive of the attitude.

The constitutive/extrinsic distinction can be drawn for any commitment-constituted attitude. For any such attitude there will be constitutive reasons, reasons which bear on the question, the answering of which amounts to forming the attitude. Further, because we are reflective creatures, capable of thinking about our own attitudes, and because which attitudes we have can matter to us in various ways, we could always, in principle, construct a case (perhaps with the use of science fiction) in which certain considerations show the attitude in some way good to have, *without* supporting the commitment constitutive of the attitude. These will be extrinsic reasons. Thus we can, in principle, draw the constitutive/extrinsic distinction for any commitment-constituted attitude.

However, neither the constitutive/extrinsic distinction nor the content-/attitude-related distinction can be drawn for every attitude or every rational activity. Most obviously, these distinctions cannot be drawn for ordinary intentional actions, such as turning left or preparing dinner. Reasons that count in favor of this or that action bear, unambiguously, on the question of whether so to act. One might try to label these “action-related” reasons, but there is no corresponding content-related question. Likewise, one does not perform an ordinary action simply by settling for oneself the question of whether so to

<sup>26</sup> A point (almost) made by G.E.M. Anscombe, *Intention* (Oxford: Blackwell, 1957).

act—by settling that question, one therein *intends* to act.<sup>27</sup> There are no constitutive reasons for ordinary action.<sup>28</sup>

The distinctions also cannot be drawn for certain propositional attitudes, such as supposing *p* for the sake of argument, imagining there's no heaven, or remembering where you put your keys. There are certainly considerations that bear on whether these are good to do—perhaps it will advance the argument or lead to liberation or allow you to get to the store. One can certainly suppose, imagine, or remember for these reasons. By finding these reasons convincing, one decides or intends or attempts to suppose or imagine or remember. One's attempt can (in principle) fail, as it can for ordinary action. Such propositional attitudes seem to be action-like.<sup>29</sup>

#### CONSTITUTIVE REASONS AND JUSTIFICATION

I hope to have given an independent way to specify the “right kind” of reasons for commitment-constituted attitudes. We can locate this class of reasons by considering the way in which an attitude can be formed. Commitment-constituted attitudes can be formed simply by answering a question or set of questions. Considerations that (are taken to) bear on the relevant question(s) are the “right kind” of reason for these attitudes.

Still, one might want to ask, why are the right kind of reasons for commitment-constituted attitudes those that (are taken to) bear on a question, the answering of which amounts to forming the attitude? While I do not have a complete answer, we can make a start.

Earlier we saw that the “right kind” of reasons correspond to the “right kind” of justification, and that the “wrong kind” of reasons would provide a corresponding, peculiar sense of justification. We can now see how this is, and why the “right kind” might be right.

<sup>27</sup> Preparing dinner is not something one does simply by settling the question of whether to do so. Rather, by settling that question, one therein *intends* to prepare dinner. If all goes well, one will execute that intention in action. But all may not go well, and so the intention and the action remain distinguishable.

<sup>28</sup> An analogous distinction can be drawn in reasons for actions, if the action is picked out according to the reasons for which is performed. Elsewhere I claim that kind, courageous, and spiteful actions, for example, are so discriminated. I explore the analogous distinction, and its importance for moral theory, in “Extrinsic Reasons, Alienation, and Moral Philosophy” (in progress).

<sup>29</sup> Or, they seem to be action-like, in those cases in which they are had for such reasons. In such cases, I am inclined to call them mental acts. I would thus draw the distinction between such “mental acts” and commitment-constituted attitudes by considering whether the attitude can be formed, regardless of its content, simply by settling the question of whether to form it. (This way of drawing the distinction between mental acts and attitudes differs starkly from the common practice of classifying occurrent states as “acts” and standing dispositions as “attitudes.”)

Commitment-constituted attitudes, we have said, leave one open to certain questions and criticisms—that is, they leave one open to certain requests for justification and criticisms for lacking the same. These requests and criticisms would be satisfied by adequate constitutive reasons. Thus we have specified, not only a set of reasons, but also a corresponding sense of justification. This is the ordinary sense. Given typical assumptions about agency and justification, it at least seems appropriate that the reasons that justify an activity, in the ordinary sense, should be those that, by finding convincing, one therein partakes in the activity.

Consider a case: whether I am justified in believing the butler did it depends on the adequacy of the reasons that (I take to) bear on whether he did it. By finding such reasons convincing, I will believe. In contrast, if I find convincing the reasons which I take to bear only on whether it is good to believe the butler did it (extrinsic reasons), then I will, therein, form a second-order attitude about the belief that the butler did it. I will believe that it would be good to believe the butler did it, and I may intend to bring it about that I believe the butler did it. Whether I am justified in these second-order attitudes will be determined by the adequacy of the constitutive reasons for *these* attitudes—considerations which (I take to) bear on whether it is good to believe the butler did it, or on whether to bring it about that I believe the butler did it—which are, of course, the extrinsic reasons for the belief that the butler did it. So it seems that the reasons that justify an activity are those that, by finding convincing, one partakes in the activity.

Earlier we saw that, if believing the butler did it would save your life, and if, for that reason, you somehow manage to believe the butler did it, then there is at least a sense in which you are justified and even rational for doing so. Elsewhere I have argued that you cannot form a belief directly for extrinsic reasons. For such reasons you can only act so as to bring it about that you believe.<sup>30</sup> But of course *this* activity—acting so as to bring it about that you believe—might be justified by the reasons that show the belief good to have. Thus, believing that the butler did it can be, in a sense, “justified” by extrinsic reasons, in the same way that sleeping or stumbling or being angry or falling in love might be, in a sense, “justified” by reasons that show it useful to sleep or stumble or be angry or fall in love. In such cases, you have reason to bring it about that you do something which is not itself the kind of thing that can be done by finding those reasons

<sup>30</sup> See “Controlling Attitudes.”



convincing (or perhaps for reasons at all). Believing the butler did it is, in a sense, “justified” by extrinsic reasons only because (or insofar as) *bringing it about that you believe* the butler did it is justified by those reasons in the ordinary way.<sup>31</sup>

#### THE CURRENT DISCUSSION

Current discussion of the “wrong kind of reasons problem” has focused on reasons for a particular class of attitudes, namely, those “pro-attitudes” such as preferring, desiring, admiring, and even “valuing” which are thought to be appropriate just in case their object is valuable. Current discussion has understood the wrong kind of reasons problem as a threat to “fitting-attitudes” or “neo-sentimentalist” accounts of value.<sup>32</sup>

My ambitions are at once broader and narrower. My ambitions are broader, in that I hope to have solved a broader problem: I hope to have provided a general formula for identifying the “right kind” of reasons for any commitment-constituted attitude (not just for those pro-attitudes involved in valuing). My ambitions are narrower, in that I do not hope to address the metaphysics of value. I will, however, sketch how the account I have provided would apply to the problem under discussion.

The wrong kind of reasons problem is thought to threaten the fitting-attitudes account of value. Such an account attempts to understand value by appeal to whether certain “valuing” attitudes are fitting. (It “passes the buck” from values to attitudes.) Reasons of the wrong kind pose a threat, because they can seem to make an attitude “fitting” without showing its object to be valuable. They thus generate counterexamples.<sup>33</sup> To avoid these counterexamples, we need to identify the reasons that “really” make the attitudes fitting. We need to rule out the wrong kind of reasons.

<sup>31</sup> Phillip Pettit points out some subtlety here. In certain cases, one might want to say that the activity of bringing it about that you believe *p* could not be justified at all—because, in the circumstances, that activity is manifestly impossible. Yet, in such cases, certain reasons may well show the belief good to have. Thus, in such cases, one might want to say that the belief that *p* is “justified,” in some sense, by extrinsic reasons, while denying that the activity of bringing it about is justified, in the ordinary sense. I imagine this case could be handled in a number of ways. I am inclined to think that the extrinsic reasons, in such a case, “justify” the belief insofar as they are associated with a possible justification of bringing about the belief: you would be justified, if it were possible, in bringing it about that you believe *p*.

<sup>32</sup> Scanlon’s “buck-passing” account is characterized in the literature as a “fitting-attitudes” account. I am dubious. I will return to the point in a later footnote.

<sup>33</sup> One such counterexample was noted in an earlier footnote: Rabinowicz and Rønnow-Rasmussen imagine an evil demon who requires you to admire him “on account of his determination to punish [you]” if you do not. The demon seems to have given you a reason to admire him, a reason that makes admiration, in some

I hope to have provided a way to identify the right kind of reasons for attitudes. If I am right, then, in order to determine the right kind of reasons for the vast and interesting range of attitudes with which the current discussion is concerned (attitudes such as admiring, preferring, fearing, envying, valuing, and desiring), one will have to determine the commitment(s) implicit in each. That is, one will have to identify the question or set of questions, the settling of which amounts to having the attitude.<sup>34</sup> The right kind of reasons for these attitudes will be those considerations that (are taken to) bear on those questions.

One might, then, be doubtful of my proposal or of its applicability to these attitudes, because one might doubt that these attitudes are commitment-constituted—one might doubt that they can be formed by settling for oneself a question or set of questions. But I suspect they can, on the basis of the following reflection:

Certain attitudes are subject to direct rational criticism: one can be asked to give one's reasons for them. One can be asked why one believes, intends, resents, supposes, or imagines, where this question does not look for an account of *how it came about that* one believes, and so on, or even why *one brought it about that* one admires, and so on, but rather for the reasons for which one might have formed the attitude. Such requests for justification can, in principle, be satisfied, and it seems that one can satisfy such requests in one of two ways: by providing reasons that show the attitude in some way good to have or by providing reasons that bear on some other question(s), the settling of which amounts to having the attitude. If one can satisfy the request for justification simply by showing the attitude good to have, then it seems the attitude is action-like (like remembering where you left your keys). If you can satisfy the request for justification by

sense, fitting, and yet it is a reason of the wrong sort to show him admirable. See "Strike of the Demon," p. 414.

<sup>34</sup> The claim that settling a set of questions "amounts to" having the attitude may seem a stronger claim than it is. Settling the question of whether *p* amounts to believing *p* in that, if one settles the question, one is rightly understood as believing, and, if one believes, one is rightly understood to have settled the question—as shown by the fact that one believes just in case one is vulnerable to questions and criticisms which would be satisfied by considerations that bear positively on the question. Likewise, if settling some complex set of questions "amounts to" resenting, say, then, if one settles the set, one is rightly understood as resenting, and if one resents, one is rightly understood to have settled that set of questions—as shown by the fact that one believes just in case one is vulnerable to questions and criticisms that would be satisfied by considerations that bear positively on the set of questions. I consider these connections briefly in "The Will as Reason" (in progress). I am grateful to Mark Johnston, Barry Lam, Victoria McGeer, and Michael Smith for pressing on this claim, though I am confident I have not fully addressed their concerns.

providing reasons that bear on a question, the settling of which amounts to having the attitude, then the attitude is commitment-constituted. Thus, it seems that, insofar as an attitude is itself capable of such direct rational criticism, either it will be action-like or it will be commitment-constituted. The pro-attitudes involved in valuing seem capable of direct rational criticism (I can be asked why I admire or prefer, not merely how it came about or why I have brought it about that I admire or prefer). So it seems that these attitudes are either action-like or commitment-constituted. But if an attitude is action-like, there will be no distinction between constitutive and extrinsic or content- and attitude-related reasons for it. The wrong-kind-of-reason problem, itself, suggests that there is some such distinction. Thus, it seems likely that, insofar as the problem can arise for these attitudes, they are, in fact, commitment-constituted.<sup>35</sup> And thus I am hopeful that, by doing a bit of philosophical excavation, we can come to understand the commitments implicit in them. We would thus isolate the right kind of reasons for them—they are the reasons that support those commitments.<sup>36</sup> (Of course, identifying the commitments implicit in such attitudes will be a difficult and delicate task, one we will likely have to approach piecemeal.<sup>37</sup>)

Thus I am hopeful that we can eventually distinguish the “right” from the “wrong” kind of reasons for those attitudes involved in

<sup>35</sup> Here is another, perhaps stronger, reason to think that the attitudes involved in valuing could not be action-like: anything action-like is voluntary—that is, it is something one can do simply by posing and answering the question of whether to do it. An attitude that is voluntary in this way could not express any conviction that the object is of value.

<sup>36</sup> D’Arms and Jacobson talk about the “norms of fittingness” for different attitudes. See “The Moralistic Fallacy: On the ‘Appropriateness’ of Emotion,” *Philosophy and Phenomenological Research*, LXI (2000): 65–90, and “Sentiment and Value.” They express their doubts about what they call “quasi-judgmentalist” accounts of emotion in Anthony Hatzimoysis, ed., “The Significance of Recalcitrant Emotion (or, Anti-Quasijudgmentalism),” in *Philosophy and the Emotions* (New York: Cambridge, 2003), pp. 127–45.

<sup>37</sup> Elsewhere I have considered whether resentment might be understood as the settling of a complex set of questions—about the genuineness of the wrong, the status of the wrongdoer, one’s own status, and the ongoing significance of the wrong done. Settling these questions would be constitutive of resentment, and reasons that bear on them would be constitutive reasons for resenting. Other reasons, which show resentment useful or costly or inconvenient, would be extrinsic reasons for or against resenting. Forgiveness and excuse, I have suggested, should concern the constitutive reasons for resentment. Excusing involves changing one’s answer to one of the first three questions. In so doing, one undermines resentment. Forgiveness could be understood as overcoming resentment by revising one’s answer to the last of these questions. In contrast, the fact that resentment is a costly and destructive emotion counts in favor of overcoming or eliminating it, but is an extrinsic reason for its elimination; one could not excuse or forgive for such a reason. Rather, one might undertake some sort of drug treatment aimed at eliminating the emotion. See “Articu-

valuing. However, despite my optimism on this score, I doubt that identifying the right kind of reasons for these attitudes will be of much help in providing an account of value—because, as we have seen, the constitutive/extrinsic distinction does not itself isolate “good” constitutive reasons for an attitude from “bad” ones. (Even if we could identify the questions the settling of which amounts to admiring or cherishing something, we would not thereby know whether those questions ought to be answered positively in a given case. Merely identifying the right kind of reason for valuing attitudes will not identify the objects of value.) So, if one hopes to give a general account of value by appeal to the appropriate reasons for an independently identifiable class of attitudes, problems will remain. Thus, even with an answer to the wrong kind of reasons problem, we may not be able to pass the buck from value to attitudes.<sup>38</sup>

#### CONCLUSION

I have offered a suggestion about what a reason is. A reason, I have suggested, is a consideration that bears on a question. The basis for this suggestion is the elementary reflection that a reason is an item in a piece of reasoning. (A good reason is an item in a sound piece of reasoning.) In its support, I have shown one advantage over the thought that a reason is a consideration that counts in favor of an action or attitude:<sup>39</sup> this formulation generates a thoroughgoing ambiguity in reasons for certain attitudes—we cannot distinguish precisely between “content-related” and “attitude-related” reasons. If we instead take a reason to be a consideration that bears on a question, we can distinguish between questions, thus distinguishing these classes. The attitude-related reasons count in favor of the attitude by bearing on

lating an Uncompromising Forgiveness,” *Philosophy and Phenomenological Research*, LXII, 3 (2001): 529–55, and “The Force and Fairness of Blame,” *Philosophical Perspectives*, XVIII, 1 (2004): 115–48.

<sup>38</sup> It is well worth noting that this is not Scanlon’s proposal. Rather, Scanlon proposes “passing the buck” from value to *reasons*. That is, he maintains that *reason* is a more fundamental normative notion than *value*, in that claims about value could, in principle, always be more illuminatingly put as claims about reasons for certain activities—reasons provided by features of the valuable object or state of affairs other than its “value” or “goodness.” It seems to me that maintaining *this* buck-passing thesis does not require that one be able to specify, in any general way, exactly which attitudes are characteristic of valuing, or what the right kind of reasons for them are. The ability to specify which attitudes, in general, are involved in valuing and what the right kind of reasons for them are would be required only if one were hoping to reduce claims about value to claims about certain attitudes—if one were hoping to pass the buck from value to psychology. But that is not Scanlon’s suggestion. See *What We Owe to Each Other*, pp. 95–100.

<sup>39</sup> In “Reasons for Action,” I consider another advantage.

whether the attitude is in some way good to have; the content-related reasons count in favor of the attitude by bearing on some other question. I hope to have shown, however, that this distinction will not isolate the “right kind” of reasons from the wrong kind, because this distinction is not exclusive. The proposed formulation helps to distinguish the right from the wrong kind of reasons by requiring us to consider the relation between the question on which a reason bears and the attitude of which it counts in favor. The constitutive reasons are those considerations that (are taken to) bear on a question or set of questions, the settling of which amounts to having the attitude. The remaining reasons are extrinsic. Drawing this distinction solves “the wrong kind of reasons” problem; it allows us to distinguish the “right kind” of reasons for certain attitudes—the constitutive reasons—from the merely extrinsic ones.

PAMELA HIERONYMI

University of California, Los Angeles