

The Teleological Conception of Practical Reasons

ABSTRACT: It is through our actions that we affect the way the world goes. Whenever we face a choice of what to do, we also face a choice of which of various possible worlds to actualize. Moreover, whenever we act *intentionally*, we act with the aim of making the world go a certain way. It is only natural, then, to suppose that an agent's reasons for action are a function of her reasons for preferring some of these possible worlds to others, such that what she has most reason to do is to bring about the possible world, which of all those available to her, is the one that she has most reason to desire. This is what's known as the *teleological conception of practical reasons*. Whether this is the correct conception of practical reasons is important not only in its own right, but also in virtue of its potential implications for what sort of moral theory we should accept. Below, I argue that the teleological conception is indeed the correct conception of practical reasons.

IT IS THROUGH our actions that we affect the way the world goes. Indeed, whenever we face a choice of what to do, we also face a choice of which of various possible worlds to actualize.¹ Moreover, whenever we act *intentionally*, we act with the aim of making the world go a certain way. The aim needn't be anything having to do with the causal consequences of the act. One could, for instance, intend to run even though one's aim is nothing more than to run. So, although an intentional action must always aim at the realization of some end, that end could be nothing more than to bring it about that the act has been performed. The crucial point, though, remains: for every intentional action, there is some end that the agent aims to achieve.

Now, if our actions are the means by which we affect the way the world goes, and if our intentional actions necessarily aim at making the world go a certain way, then it is only natural to suppose that an agent's reasons for action are a function of her reasons for preferring some of these possible worlds to others, such that what she has most reason to do is to bring about the possible world, which of all those available to her, is the one that

¹ I will assume that for each act available to an agent there is some determinate fact as to what the world would be like were she to perform that act. This assumption is sometimes called *counterfactual determinism*—see, e.g., BYKVIST 2003, p. 30. Although this assumption is controversial, nothing that I will say here hangs on it. I make the assumption only for the sake of simplifying the presentation. If counterfactual determinism is false, then instead of correlating each act with a unique possible world, we will need to correlate each act with a probability distribution over the set of possible worlds that might be actualized if S were to perform that act.

she has most reason to desire. This is what's known as the *teleological conception of practical reasons*. Whether this is the correct conception of practical reasons is important not only in its own right, but also in virtue of its potential implications for what sort of moral theory we should accept. At least, it will have such implications if we assume, as many philosophers do, that an agent can be morally required to perform an act only if she has most reason to perform that act.² Below, I argue that the teleological conception is indeed the correct conception of practical reasons.

The paper has the following structure. In section 1, I offer a more precise statement of the teleological conception of reasons (or 'TCR' for short), showing that the view consists in three distinct claims. And I explain why my formulation differs from those of some of its critics in eschewing talk of both value and desirability.³ Then, in section 2, I clear up some common and potential misconceptions about the view. In section 3, I rebut Scanlon's putative counterexamples to TCR, cases where putatively "many of the reasons bearing on an action concern not the desirability of outcomes but rather the eligibility or ineligibility of various other reasons" (SCANLON 1998, p. 84). And finally, in section 4, I provide arguments for each of TCR's three claims and for TCR as a whole.

1. Getting clear on what the view is

Let me start off by stating the view as precisely as I can. Let a_1, a_2, \dots, a_n be the set of mutually exclusive and jointly exhaustive act alternatives available to a subject, S. Let o_1, o_2, \dots, o_n be their corresponding outcomes, where an act's outcome is construed broadly as

² STROUD 1998 calls this the *overridingness thesis*, and she employs it in an argument against impartial moral theories, such as act-utilitarianism, which lack sensitivity to the agent's personal point of view and, consequently, to the reasons stemming from that point of view. The thesis is also sometimes called *supremacy* (DARWALL 2006), *moral rationalism* (VAN ROOJEN 2007), and *the principle of moral categoricity* (SKORUPSKI 1999, 170). For more on the implications of this thesis for moral theorizing, see [redacted].

³ Whenever I use the term 'reasons' in an unqualified way, I will be referring to practical reasons—i.e., normative reasons for action. A normative reason for action is some fact that counts in favor of the agent's performing the action. Normative reasons contrast with explanatory reasons (i.e., facts that explain why the agent performed the action). One particularly important subclass of explanatory reasons is the set of motivating reasons, the facts that motivated the agent to perform the action—that is, the facts that the agent took to be her reasons for performing the action. See DARWALL 2006, p. 285.

the possible world that would be actual were the act to be performed.⁴ More precisely, then, the teleological conception of reasons can be stated as follows:

TCR (1) S has more reason to perform a_i than to perform a_j if S has more reason to desire that o_i obtains than to desire that o_j obtains; (2) S has more reason to perform a_i than to perform a_j only if S has more reason to desire that o_i obtains than to desire that o_j obtains; and (3) if S has more reason to perform a_i than to perform a_j , then this is in virtue of the fact that S has more reason to desire that o_i obtains than to desire that o_j obtains.⁵

Having stated the view as precisely as I can, I will now proceed to clarify it, explaining, in the process, how and why it deviates from T. M. Scanlon's formulation of the view, which goes as follows: "the purely teleological conception of reasons [is the view] according to which, since any rational action must aim at some result, reasons that bear on whether to perform an action must appeal to the desirability or undesirability of having that result occur, taking into account also the intrinsic value of the act itself" (SCANLON 1998, p. 84). There are, in fact, five separate points that need clarifying.

§§1.1 *Desirability versus reasons for desiring*: Unlike Scanlon, I have formulated TCR in terms of reasons to desire outcomes as opposed to the desirability of outcomes. To see why, we must first get clear on what the difference is. Let's start with what it means to say that an outcome is desirable. To say that that an outcome, o_i , is desirable is to say that o_i is worthy of being desired—i.e., that it is fitting to desire that o_i obtains. And to say that it is fitting to desire that o_i obtains is just to say that there are sufficiently weighty reasons of

⁴ Again, I'm assuming counterfactual determinism—see note 1 above.

⁵ More concisely, then, TCR is the view that S has more reason to perform a_i than to perform a_j just when, and because, S has more reason to desire that o_i obtains than to desire that o_j obtains. For my purposes, though, it will be useful to keep the three claims separate. And note that, although not stated explicitly here, I take TCR to include the claim that S has a reason to perform a_i just when, and because, S has a reason to desire that o_i .

Michael Smith's (2006) formulation of the view is similar to mine, although he words things a bit differently: "(x)(t)(x at t has all things considered reason to φ in circumstances C iff φ -ing is the unique action of those x can perform at t that brings about what x would desire most happens in C if he had a desire set that satisfied all requirements of...reason."

the right kind to desire that o_i obtains.⁶ What are the right kinds of reasons to desire that o_i obtains? They are all and only those reasons that are relevant to determining whether, and to what extent, o_i is desirable. Call these reasons *fittingness reasons*.

I won't attempt to give a complete account of what sorts of reasons are, and what sorts of reasons are not, fittingness reasons. These are contentious issues, and I have nothing new to add to the growing debate.⁷ Even so, I can plausibly claim that there are some clear cases of reasons that wouldn't count as fittingness reasons. First, if facts such as the fact that an evil demon will torture me unless I desire that o_i obtains constitute genuine reasons for me to desire that o_i obtains and not just reasons for me to want, and to cause myself, to desire that o_i obtains, then these are clearly not fittingness reasons, for such pragmatic "reasons" for desiring that o_i obtains clearly have no bearing as to whether, or to what extent, o_i 's obtaining is desirable.⁸ Now, as a matter of fact, I don't think that such pragmatic "reasons" do constitute genuine reasons for desiring that o_i obtains. I think instead that they constitute only reasons to want, and to cause oneself, to desire that o_i obtains. If, however, I'm wrong about this, I'll need to revise TCR so as to exclude such pragmatic "reasons" for desiring that o_i obtains, for such reasons are no more relevant to whether or not one should act so as to bring it about that o_i obtains than they are to whether or not o_i 's obtaining is desirable.

Second, it seems clear that agent-relative reasons are not fittingness reasons. To see why, consider that, in contrast to some stranger, I might have weightier (agent-relative) reasons to prefer that my child lives. But it would be odd to say that this is because the outcome in which my child goes on living is, other things being equal, better or more

⁶ More precisely, we should first say that, for all states of affairs p and q , it is better (i.e., preferable) that p is the case than that q is the case if and only if, for anyone engaged in the activity of choosing a state of affairs, the set of all the *fittingness reasons* for preferring its being the case that p to its being the case that q is weightier than the set of all the *fittingness reasons* for preferring its being the case that q to its being the case that p . Then we can say that it is good (i.e., desirable) that p is the case if and only if the state of affairs in which it is the case that p is better than (i.e., preferable to) most states of affairs in the contextually supplied comparison class. For more on this, see SCHROEDER 2008a and SCHROEDER 2008b.

⁷ For what I consider to be the best account of how to distinguish fittingness reasons from non-fittingness reasons, see SCHROEDER 2008a.

⁸ This is known as the *wrong-kind-of-reasons problem* for the fitting-attitude or buck-passing account of value (or desirability). For more on this problem and for some potential solutions to it, see, for instance, RABINOWICZ AND RØNNOW-RASMUSSEN 2004, PARFIT 2008, and SCHROEDER 2008a.

desirable than the outcome in which her child goes on living. Other things being equal, the outcome in which my child lives is neither more nor less desirable than the outcome in which her child lives. So although one can have agent-relative reasons for preferring one outcome to another, this does not entail that the one is better than, or preferable to, the other.⁹ And, thus, agent-relative reasons for preferring one outcome to another are not fittingness reasons for preferring the one to the other.¹⁰

It is this last exclusion that makes trouble for formulating TCR in terms of the value or desirability of outcomes. If TCR is to be formulated in terms of value/desirability and value/desirability is to be understood exclusively in terms of agent-neutral reasons for desiring, then TCR will automatically disallow agent-relative reasons for action, such as the agent-relative reason that I have to save my own child. Yet there is no reason why the teleologist should exclude the possibility of agent-relative reasons for action, as even critics of TCR admit. Scanlon, for instance, says, “The teleological structure I have described is often taken to characterize not only ‘the good’ impartially understood, but also the good from a particular individual’s point of view (the way she has reason to want things to go)” (1998, 81). So if we are to allow that one agent might have a reason to bring about a state of affairs that another has no reason to bring about and that one agent might

⁹ This is known as the *partiality challenge* to the fitting-attitude or buck-passing account of value. See OLSON 2008, SUKKANEN 2008, and ZIMMERMAN 2008 for some potential solutions to this particular problem.

¹⁰ There are also time-relative reasons for preferring one outcome to another (and these too are not fittingness reasons). Sometimes, for instance, the preference that an agent should have before choosing to perform some action is the opposite of what it should be after performing the action. Suppose, for instance, that Ana decided not to have an abortion even after learning that her fetus might have Down syndrome. Consequently, she gave birth to a boy, named Bill, with Down syndrome, who is now eight. She ought, at this point, be glad that she didn’t have an abortion. After all, she does, at this point, have a very special bond with her son Bill, whom she has come to love for exactly who he is, which includes the fact that he has Down syndrome. If she could somehow take it all back and make the decision over, she should not prefer that she has the abortion and then later gives birth to a child without Down syndrome. Interestingly, though, this is compatible with our thinking that at the time of her initial decision, a time before which she had formed any special relationship with the child that she was carrying in her womb, she should have preferred the outcome in which she has the abortion and then gives birth to a child without Down syndrome. That would have been the better outcome, and, at the time of her initial decision, she had no reason to prefer the worse outcome in which she has a child with Down syndrome. Given that what it is reasonable to prefer can change over time, TCR should actually be time-indexed. Consequently, TCR should be formulated as follows: S has more reason to perform a_i at t_i than to perform a_j at t_i just when, and because, S has more reason, at t_i , to desire that o_i obtains than to desire that o_j obtains. For more on the issue of how it can be reasonable to prefer that a loved one exists even though one recognizes that there is a state of affairs in which that loved one doesn’t exist that’s preferable to the one in which that loved one does exist, see HARMAN 2005.

have more reason to bring about some state of affairs than another agent does, we should eschew talk of the value/desirability of states of affairs and talk instead of the reasons that various agents have to desire those states of affairs, where some of these reasons will be agent-relative reasons.¹¹

Interestingly, Scanlon is aware that agent-relative reasons for valuing/desiring are not fittingness reasons. He says, “To claim that something is *valuable* (or that it is ‘of value’) is to claim that others also have reason to value it, as you do” (1998, p. 95). Furthermore, he claims that “[w]e can, quite properly, value some things more than others without claiming that they are more valuable” (1998, p. 95). Likewise, we can properly desire some outcomes more than others without claiming that they are preferable to the others. Why, then, does Scanlon formulate the teleological conception of reasons in terms of value/desirability when he clearly wants to allow that the teleologist can accommodate agent-relative reasons for valuing/desiring and, consequently, agent-relative reasons for action? The answer is that when Scanlon talks about value/desirability, he means for this to include agent-relative value (i.e., “the good from a particular individual’s point of view”). Indeed, he brings up the teleological conception of reasons to explain why some think that we must assign agent-relative disvalue to an agent’s killing in order to make sense of agent-centered restrictions against killing (1998, pp. 94-95). But although Scanlon seems content to talk as if agent-relative value is just a species of value *simpliciter*, there are good reasons to suspect that this is a mistake. As Mark Schroeder has argued, the notion of goodness from a particular individual’s point of view—that is, the good-relative-to relation—would seem to be just a theoretical posit that has nothing whatsoever to do with the word ‘good’. And, thus, it would seem that what Scanlon calls ‘agent-relative goodness’ might just as well have been called ‘agent-relative orangeness’.¹² For these reasons, it is, I think, better to formulate TCR, as I have, in terms of reasons to desire rather than in terms of value/desirability.

¹¹ Given that Scanlon seems to equate what is “good [or desirable] from a particular individual’s point of view” with “the way she has reason to want things to go” (1998, p. 81), he should have no objection to this friendly amendment.

¹² See SCHROEDER 2007 and SCHROEDER 2008a.

Of course, doing so means that someone who holds that reasons for action are solely a function of the value of outcomes will not necessarily count as a teleologist (that is, as someone who endorses TCR). Whether they do or not will depend on whether or not they think that agents can ever have more reason to prefer a worse outcome to a better outcome than vice versa. If they think that they can (if, for instance, they think that an agent can have decisive reason to prefer the outcome in which her child is saved to the one in which some stranger's child is saved even if the latter is slightly better than the former), then they must reject TCR. But it is plausible to suppose that many who hold that reasons for action are purely a function of the value of outcomes will instead hold that the best outcome available to a given agent is necessarily the one she ought to prefer to the available alternatives, either because fittingness reasons are the only reasons or because fittingness reasons always override non-fittingness reasons. Regardless, it's clear that TCR, as I've formulated it, is distinct from the view that reasons for action are purely a function of the value of outcomes, and it is, I think, equally clear that it's a view worth considering.

So readers should take note that TCR is to be understood in terms of reasons to desire and not necessarily in terms of value/desirability. That said; I will occasionally revert back to talking about the value/desirability of states of affairs, since this is the language that TCR's critics so often employ. Keep in mind, though, that these critics mean for value/desirability to somehow include what's "good from a particular individual's point of view," which they equate with "the way she has reason to want things to go" (SCANLON 1998, p. 81).¹³ So the reader should assume that when I revert back to talk of

¹³ The other leading critic of TCR is Elizabeth Anderson. Although she uses the term 'consequentialism' as opposed to 'teleology', she defines 'consequentialism' so broadly (see 1993, 30-31) that it is, in spirit, equivalent to TCR. She says, for instance, "consequentialism specifies our rational aims, and then tells us to adopt whatever intentions will best bring about those aims" (ANDERSON 1996, p. 539), which is exactly what TCR tells us to do. Thus, as Anderson uses the term, 'consequentialism' refers not to a moral theory but to a conception of practical reasons that is roughly equivalent to TCR. In certain passages, Anderson, like Scanlon, defines 'consequentialism' in terms of value as opposed to reasons to desire. But, as with Scanlon, this is only because she talks as if intrinsic goods include both what's good for an individual and what's good relative to an individual. This way of talking, however, can quickly lead to nonsense. To illustrate, consider the following two quotes from ANDERSON 1993: (1) "I call any theory of rational or moral action consequentialist if it meets the following three conditions: First, it gives people the sole ultimate aim of maximizing intrinsic

desirability, I'm using the word 'desirable' as short-hand for 'that which the agent has sufficiently weighty reasons to desire'.

§§1.2 *Intended effects versus total outcomes*: Another way that my formulation differs from Scanlon's is that I have chosen to formulate TCR broadly in terms of reasons to desire total outcomes, and not narrowly in terms of reasons to desire only the results that the agent aims to produce. Note, then, that, as I've formulated the view, it is not restricted to only the causal consequences of actions.¹⁴ Indeed, it would be odd for the teleologist (which is what I'll call the proponent of TCR) to exclude in advance from consideration any of the ways that the world might change as a result of an agent's performing some act. For instance, one way the world changes when some subject, S, performs some act, say, a_1 , is that the world becomes one in which S has performed a_1 . And, as a result of S's having performed a_1 , it may also thereby be true that S has fulfilled her past promise to do so.¹⁵ Since all these ways in which the world might change could potentially make a difference as to whether or not S has any reason to desire o_1 and, if so, how much reason, we should formulate TCR, as I have, so that it does not exclude from consideration such possibly relevant non-causal effects.

§§1.3 *Reasons to act versus reasons for intending to act*: When I talk about reasons to perform an action, I am, strictly speaking, referring to reasons to intend to perform that action. Most immediately, practical reasoning gives rise not to bodily movements, but to intentions. Of course, when all goes well, these intentions result in, say, some bodily movement and then we have an intentional action. Nevertheless, the immediate product

value..." (p. 30), and (2) "My use of the term 'consequentialism' is more expansive than most people's. In my usage, egoism as a theory of rationality, which gives each agent the sole ultimate aim of maximizing his own personal welfare, is a consequentialist theory" (1993, p. 31). But the first and second quotes are inconsistent. If egoism gives Ted Bundy the sole ultimate aim of maximizing his own personal welfare and his fulfilling this aim would not maximize intrinsic value, then egoism is both consequentialist and nonconsequentialist according to these two quotes. And, of course, both conjuncts of the antecedent are true. So these two quotes are indeed inconsistent. It would be better, then, to talk about reasons to desire than to try to have value/desirability include *goodness for* and *goodness relative to*.

¹⁴ Even critics of TCR admit that teleologists can be concerned with more than just the causal consequences of acts. See, for instance, SCANLON 1998 (pp. 80 & 84) and WALLACE 2005 (pp. 26-27).

¹⁵ Thus TCR will not be exclusively forward-looking. The teleologist can hold that S has a reason to desire that o_1 obtains in virtue of the fact that it is true in that world that S has performed a_1 and that her having done so constitutes the fulfillment of her past promise to do so. For more on this point, see STURGEON 1996, pp. 511-514, and ANDERSON 1996, p. 541.

of practical reasoning is an intention to perform some act, not the act itself (SCANLON 1998, pp. 18-22). Having clarified this, I will, however, sometimes (when it seems not to matter) slip into the more customary way of speaking in terms of reasons for action when I really mean to be speaking of reasons to intend to perform some action. And note that I don't consider facts such as the fact that I will receive some reward or punishment if I intend to φ to constitute genuine reasons for or against intending to φ . The judgment that the intentional content of one's intention to φ , viz., φ , has certain consequences that one has reason to desire (i.e., a *content-directed judgment*) is the sort of judgment that can give rise to an intention to φ and that, if true, constitutes a reason to φ . But the judgment that one's having the attitude of intending to φ would have consequences that one has reason to desire (i.e., an *attitude-directed judgment*) is not the sort of judgment that can give rise to an intention to φ , nor is it, to my mind, the sort of judgment that, if true, constitutes a reason to intend φ .¹⁶ If I'm wrong about this, if the truth of such attitude-directed judgments do constitute reasons to intend to φ , then I will need to rethink my view that a reason to intend to φ is just a reason to φ , for the fact that I will be rewarded if I intend to drink some toxin tomorrow is certainly no reason to drink that toxin tomorrow.¹⁷

§§1.4 *The broad versus the narrow construal of desire*: Although some philosophers (e.g., HEUER 2004, p. 48) take a desire to be nothing more than just a disposition to act, where one desires that o_i obtains if and only if one is disposed to act so as to bring it about that o_i obtains (that is, to perform a_i), I will use 'desire' in the more narrow, ordinary sense, such that one desires that o_i obtains only if one finds the prospect of o_i 's obtaining in some way attractive or appealing.¹⁸ On this more narrow interpretation, desiring that o_i obtains is sufficient for being motivated (to some extent) to perform a_i , but being

¹⁶ See SCANLON 2007, esp. pp. 90-91

¹⁷ This is example comes from KAVKA 1983.

¹⁸ To find o_i 's obtaining in *some way* attractive or appealing, one needn't have all that o_i 's obtaining entails conscious before the mind. Ana might know that the only way to ensure that her daughter excels in school is to hire Bill (a tutor) and, consequently, find the outcome resulting from her hiring Bill in this respect appealing.

motivated to perform a_i is not sufficient for desiring that o_i obtains.¹⁹ Thus, in Warren Quinn's famous example of a man who has a compulsive urge that disposes him to turn on every radio he sees despite his failing to see anything appealing about either these acts themselves or their effects (1993, p. 236), we do not have a genuine case of desire—at least, not in the sense that I will be using the term 'desire'. As I see it, then, having a desire involves a complicated set of dispositions to think, feel, and react in specified ways (SCANLON 1998, p. 21). A person who desires that o_i obtains will find the prospect of its obtaining appealing, will be to some extent motivated to perform a_i , and will, perhaps, have her attention directed insistently toward considerations that present themselves as counting in favor of o_i 's obtaining (SCANLON 1998, p. 39).²⁰

§§1.5 *TCR-1+2 versus TCR*: To be a teleologist, it is not enough to accept the bi-conditional that is entailed by the conjunction of TCR-1 and TCR-2 (call this bi-conditional "TCR-1+2"); the teleologist must accept TCR-3 as well. Of course, TCR-3 is but one of three possible explanations for the correlation stated by TCR-1+2. To illustrate, let 'P' stand for 'S has more reason to *perform* a_i than to *perform* a_j ' and let 'D' stand for 'S has more reason to *desire* that o_i obtains than to *desire* that o_j obtains'. The three possible explanations for TCR-1+2—that is, for the correlation between P and D—are: (1) 'P, because D', (2) 'D, because P', or (3) both 'P, because B' and 'D, because B'—where, for instance, 'B' might stand for ' o_i is *better* than o_j '.²¹ In defending TCR, I must not only defend TCR-1+2, but also argue that it is explanation 1 as opposed to either explanation 2 or explanation 3 that explains TCR-1+2; that is, I must defend TCR-3 in addition to both TCR-1 and TCR-2. But before I proceed to defend TCR, I will first try to clear up some actual and potential misconceptions about the view.

¹⁹ Because one can be motivated (to some extent) to perform a_i without being sufficiently motivated to perform a_i , desiring that o_i obtains does not necessarily result in an intention to perform a_i . After all, one can have conflicting motives.

²⁰ I am not sure whether the last of these three is essential to desiring, as the qualifier 'perhaps' is meant to indicate. For reasons to doubt that it is essential to desiring in the ordinary sense, see CHANG 2004, esp. pp. 65-66.

²¹ These three explanations are analogues of the three possible causal explanations for a correlation between events a and b : (1) a causes b , (2) b causes a , or (3) a and b have a common cause. I thank [redacted] for pointing out the need to consider such common "cause" explanations. And I thank [redacted] for suggesting that someone might take ' o_i is *better* than o_j ' to be the common "cause."

2. Clearing up some misconceptions about TCR

There are a number of misconceptions about TCR that have led philosophers to reject TCR for mistaken reasons. Below, I try to clear up some of these misconceptions.

§§2.1 *TCR is compatible with value concretism*: Although Scanlon (1998, 79-81) lumps the two together, TCR is distinct from, and independent of, *value abstractism*: the view that the sole or primary bearers of intrinsic value are certain abstracta—facts, outcomes, states of affairs, or possible worlds. On value abstractism, there is only one kind of value, the kind that is to be promoted, and so the only proper response to value is to desire and promote it, ensuring that there is as much of it as possible. The contrary view—the view that the fundamental bearers of intrinsic value are concrete entities (e.g., persons, animals, and things)—is called *value concretism*.²²

Contrary to what Scanlon and others (e.g., ANDERSON 1993; WALLACE 1995) have claimed, there is no reason why the teleologist cannot accept value concretism.²³ Indeed, the teleologist can accept all of the following: (a) that concrete entities—persons, animals, and things—are the primary bearers of intrinsic value; (b) that states of affairs generally have only extrinsic value in that they generally have no value apart from our valuing concrete entities;²⁴ (c) that “our basic evaluative attitudes—love, respect, consideration, honor, and so forth—are non-propositional...attitudes we take up immediately toward persons, animals, and things, not toward facts” (ANDERSON 1993, p. 20); and (d) that value and our valuations are deeply pluralistic, that there are many ways that we experience things as valuable (e.g., as interesting, admirable, beautiful, etc.) and that there are many different kinds of value as well as different modes of valuing that are appropriate to each (e.g., “beautiful things are worthy of appreciation, rational beings of respect, sentient

²² I borrow the terms *concretism* and *abstractism* from TÄNNISJÖ 1999.

²³ Again, Anderson uses the term ‘consequentialism’ as opposed to ‘teleology’, but see note 14 above.

²⁴ The reason for the qualifier ‘generally’ in item *b* is that Anderson does allow for the possibility that a state of affairs can have intrinsic value if it is one that is intrinsically interesting. Anderson says, “Interest does seem to be an evaluative attitude that can take a state of affairs as its immediate and independent object. This is an exception to the general rule that states of affairs have no intrinsic value” (1993, 27).

beings of consideration, virtuous ones of admiration, convenient things of use” — ANDERSON 1993, p. 11). TCR is compatible with all (a)-(d), as I will now explain.

As rational *creatures*, we appropriately respond to different sorts of things with different sorts of attitudes. We appropriately respond to beautiful objects by appreciating them, we appropriately respond to rational persons by respecting them, and we appropriately respond to desirable states of affairs by desiring their actualizations—at least, that’s how we appropriately respond to them when we don’t have weightier agent-relative or time-relative reasons to desire that they not be actualized. As rational *agents*, though, it is only the last of these three that is pertinent, for, as agents, we can effect only outcomes. We cannot effect valuable entities. A concrete entity is not the sort of thing that we can bring about or actualize through our actions. Of course, we can act so as to bring it about that, say, a rational person exists or that our actions express our respect for rational persons, but these are states of affairs, not concrete entities. As agents, then, we have the ability to actualize only certain possible worlds or states of affairs. Indeed, purposive action must aim at the realization of some state of affairs. So the teleologist can admit that we have reasons to have all sorts of different attitudes, including reasons to have certain non-propositional attitudes (such as, respect) toward various concrete entities (such as, rational persons). But the teleologist will insist that when it comes to the particular attitude of intending to act in some way, the reasons for having this attitude must always be grounded in the reasons that the agent has to desire that certain possible worlds or states of affairs be actualized. It is a mistake, however, to think that the teleologist is, in addition, committed to the denial of any of claims (a)-(d) above.²⁵

§§2.2 *TCR is compatible with appropriately valuing goods such as friendship*: Another common misconception regarding TCR is that it is incompatible with the thought that with respect to goods such as science and friendship, taking them to be valuable is not simply, or even primarily, a matter of promoting certain states of affairs (cf. SCANLON 1998, p. 88). Take friendship. TCR does not imply, for instance, that the only reasons

²⁵ This point is not particularly new, although it bears repeating given the stubborn persistence of this misconception. Others who have made essentially the same point include ARNESON 2002 and STURGEON 1996.

provided by my friend and our relationship are reasons to promote certain states of affairs. The teleologist can accept that I have reasons to care about my friend, to empathize with her pain, to take joy in her successes, etc., and that these are not reasons to promote certain states of affairs, but rather reasons to have certain non-propositional attitudes and feelings. TCR is, then, compatible with the thought that what lovers, friends, and family members value, fundamentally, is each other as opposed to certain states of affairs.

TCR is also compatible with the thought that a person who values friendship will see that what she has reason to do, first and foremost, is to be a good friend to her current friends and that these reasons are weightier than whatever reasons she has to cultivate new friendships or to foster good friendship relations among others (cf. SCANLON 1998, pp. 88-89). The teleologist can even hold that my friendships generate agent-centered restrictions on my actions (cf. ANDERSON 1993, pp. 73-74), such that I have more reason to refrain from betraying one of my own friends than to prevent more numerous others from betraying theirs. This is all possible given that TCR allows for agent-relative reasons. If there were only agent-neutral reasons, (e.g., agent-neutral reasons to promote friendships and to prevent friends from betraying one another), then I would often have sufficient reason to neglect one of my current friendships if I could thereby cultivate two or more new ones, and I would often have sufficient reason to betray one of my own friends if I could thereby prevent more numerous others from betraying theirs. But, as even critics of TCR admit (e.g., SCANLON 1998, p. 81), TCR is compatible with the existence of agent-relative reasons and thus with the idea that whereas *you* will have more reason to prefer that *your* friends are not betrayed, *I* will have more reason to prefer that *my* friends are not betrayed. And such agent-relative reasons to prefer the possible world in which your friends as opposed to my friends are betrayed will, given TCR, generate agent-relative reasons for me to refrain from betraying one of my own friends even for the sake of preventing you from betraying two of yours.²⁶ And TCR certainly allows for the

²⁶ Similarly, the teleologist can even hold that I have more reason to prefer that my friends are betrayed by *you* than to prefer that they are betrayed by *me*. And such agent-relative reasons will, given TCR, generate agent-

possibility that such reasons will be decisive and thereby generate an agent-centered restriction against betraying one's own friends even for the sake of preventing more numerous others from betraying theirs.

Moreover, TCR is compatible with the claim that I should not abandon my current friends even for the sake of cultivating more numerous new friendships, for the teleologist can hold that I currently have good time-relative reasons for preferring the preservation of my current friendships to the creation of otherwise similar new friendships given the shared history that I have with my current friends and the lack of any shared history (at present) with those possible future friends (HURKA 2006, p. 238).²⁷ And not only can the teleologist accept that one should not destroy one's own current friendships for the sake of creating more numerous future friendships for oneself or for others, but also that, out of respect for friendships generally, one should not destroy someone else's friendship for the sake of preventing numerous others from doing the same. Again, because the teleologist can hold that there are agent-relative reasons for preferring one possible world to another, the teleologist can hold that I should prefer the state of affairs in which, say, five others each destroy someone else's friendship to the state of affairs in which I myself destroy someone else's friendship.

§§2.3 *TCR is compatible with the view that attitudes such as belief and blame are rationally justified on non-instrumental grounds:* Elizabeth Anderson has claimed that the teleologist is committed to the implausible view that all attitudes (blaming, believing, intending, etc.) are rationally justified on instrumental grounds—that is, on the grounds that the agent's having the given attitude would have desirable consequences. For instance, she claims that the teleologist (or what she calls the “consequentialist”) must hold that beliefs “are justified to the degree that they bring about better states of affairs into existence”

relative reasons for me to refrain from betraying one of my friends even for the sake of preventing you from betraying two of my friends.

²⁷ As mentioned in note 10, TCR should, then, be formulated as follows: S has more reason to perform a_i at t_i than to perform a_j at t_j just when, and because, S has more reason, at t_i , to desire that o_i obtains than to desire that o_j obtains.

(ANDERSON 1993, p. 39).²⁸ This is mistaken for at least two reasons. First, note that the teleologist does not even hold that whether one is rationally justified in having the intention to perform a_i is a function of the desirability of the consequences of one's having this attitude. TCR does not imply, for instance, that the fact that an evil demon has threatened to produce undesirable consequences unless you intend to perform a_i gives you a reason to intend to perform a_i . TCR implies that it is the fact that your *performing* a_i will have desirable consequences (or, as I would prefer to say, consequences that you have sufficiently weighty reasons to desire), not the fact that your *intending to perform* a_i will have desirable consequences, that provides you with a reason to intend to perform a_i .²⁹ Thus, it is the consequences of *the act*, not of *the intention* to perform the act, that is relevant on TCR. And so there is no way to generalize, as Anderson apparently does, from the claim that intending to do a_i is justified on the grounds that *performing* a_i is instrumental in bringing about certain desirable consequences to the claim that attitudes, such as believing that p , are justified on the grounds that *having this attitude* is instrumental in bringing about certain desirable consequences. It is just not analogous.

Second, even if it were true that the teleologist thought that the rationality of intending to perform a_i depends on the desirability of the consequences of having this attitude, it is not clear why Anderson thinks that the teleologist is committed to the more general claim that this is true of *all* attitudes. Why can't the teleologist hold that what is true of intending to act is not true of other attitudes? In either case, then, the teleologist can accept, contrary to what Anderson claims, that whether a person has sufficient reason

²⁸ The mistake, here, is to assume both that mental attitudes such as believing and blaming are actions in the relevant sense and that the teleological view should be formulated in terms of reasons for *action* as opposed to reasons for *intending to perform actions*. Clearly, though, mental attitudes such as believing and blaming are not (or, at least, are not typically) actions in the relevant sense. For instance, we often form beliefs involuntarily in response to our perceptual experiences. These belief formations are not intentional actions, for we do not form these beliefs as a result of our intending to form them. By making explicit that, when I talk about reasons to perform an action, I am, strictly speaking, referring to reasons to intend to perform that action, I hope to avoid such confusions.

²⁹ To see that these can come apart, consider that one can, say, knock over a bucket without forming the intention to do so and that one can form the intention to do so without succeeding. On TCR, the reasons that one has to intend to knock over a bucket depend on what the world would be like if one knocked over the bucket (whether or not one does so intentionally), not on what the world would be like if one formed the intention to do so (whether or not one succeeds).

to believe that p depends on what her evidence is for the truth of p and not on the desirability of the consequences of her believing that p .

Anderson is also mistaken in thinking that the teleologist must hold that: “Blame should be meted out to wrongdoers only if it would deter future wrongdoing and to anyone else unlucky enough to be in a position where being made an example of (however unjustly) would deter further wrongdoing” (ANDERSON 1993, p. 40). There is no reason why the teleologist cannot hold that whether blaming someone is rationally justified depends only on whether that someone is blameworthy. To blame someone is to have a certain attitude toward that someone, but it is not the same attitude about which TCR has something to say, viz., intending to act. Of course, the teleologist is committed to the view that whether one is rationally justified in intending to act so as to criticize, punish, or otherwise censure someone does depend on the consequences of such an act. But this is distinct from the attitude of blaming someone for her actions, i.e., that of feeling guilt, resentment, or indignation in response to her actions. One can blame a person for her wrongdoing without intending to act so as to criticize, punish, or otherwise censure her. And the teleologist is not committed to even the view that one should intend to act so as to punish someone if and only if doing so would deter future wrongdoing, for the teleologist could hold that there are generally decisive reasons to desire all and only those worlds in which one acts so as to punish only those who are blameworthy.

§§2.4 *TCR is compatible with passing the normative buck from value to reasons*: Lastly, it should be noted that TCR is compatible with a buck-passing account of value, where the domain of reasons is taken to be explanatorily prior to the domain of value (cf. WALLACE 2005)—thus the normative buck is passed from value to reasons. There is potential for confusion here, for Scanlon sometimes formulates the buck-passing account of value specifically in terms of reasons for *action*, and TCR does treat evaluative reasons (e.g., reasons to desire) as explanatorily prior to practical reasons (i.e., reasons for action).³⁰ Nevertheless, TCR is compatible with the following more general formulation of the

³⁰ Scanlon says, for instance, that “to call something valuable is to say that it has other properties that provide reasons for behaving in certain ways with regard to it” (1998, p. 96).

buck-passing account of value: x 's being good or valuable is just the purely formal, higher-order property of having other properties that provide sufficiently weighty reasons of the right kind to respond favorably toward x . This leaves open what the relevant favorable response will be, and, arguably, it will vary depending on the kind of thing that is being evaluated. In the case of a rational person, the relevant response might be to have respect for that person. In the case of a valuable state of affairs, the relevant response might be to desire that it obtains. If this is right, then TCR is perfectly compatible with the buck-passing account of value. An agent-neutral teleologist could, for instance, hold that reasons for actions are a function of the value (or the desirability) of the states of affairs that those actions produce, but still ultimately pass the normative buck back to reasons by claiming that what it is for a state of affairs to be valuable/desirable is for it to have the purely formal, higher-order property of having other properties that provide sufficiently weighty reasons (of the right kind) to desire it.³¹

3. Scanlon's putative counterexamples to TCR

Besides these misconceptions, some philosophers are led to reject TCR given putative counterexamples. Scanlon, for instance, uses examples to argue, as he puts it, that "many reasons bearing on an action concern not the desirability of outcomes but rather the eligibility or ineligibility of various other reasons..., and judging that a certain consideration does not count as a reason for action is not equivalent to assigning negative intrinsic value to the occurrence of actions based on this reason" (1998, p. 84). Unfortunately, it is not clear what he is trying to establish by this. One possibility is that he is pointing out that there are sometimes reasons for *believing* (he says "judging") that a certain fact does not count as a reason to perform a given act and that these *epistemic* reasons (i.e., reasons for belief) do not concern the desirability of outcomes. But, then, this is no counterexample to TCR. To provide a counterexample to TCR, Scanlon must provide an example in which we are compelled to think that there are *practical* reasons (reasons for

³¹ The agent-neutral teleologist holds that the best outcome available to a given agent is necessarily the one she ought to prefer to all other available alternatives, either because fittingness reasons for preferring one outcome to another are the only reasons for preferring one outcome to another or because fittingness reasons always override non-fittingness reasons.

action) that do not concern the desirability of outcomes, not an example in which we are compelled to think that there are *epistemic* reasons (reasons for belief) that do not concern the desirability of outcomes. As we saw in §3.2.3, the teleologist can accept that reasons for belief do not concern the desirability of outcomes. This, then, is clearly not the most charitable way to interpret Scanlon.

So let us consider his actual examples. In Chapter 1, Section 10, he provides examples where an agent “judge[s] one consideration, C, to be a reason for taking another consideration, D, not to be relevant to... [his or her] decision to pursue a certain line of action” (1998, p. 51). What he gives us, then, are not reasons that concern the eligibility or ineligibility of various other reasons, but rather reasons that concern whether or not one ought to *take* various other reasons into account in one’s deliberations, and taking other reasons into account in one’s deliberation is itself an action.³² Moreover, the reasons for or against taking other reasons into account in one’s deliberations do seem to concern the comparative desirability of the outcomes associated with doing so versus not doing so. Indeed, in the sorts of examples that Scanlon provides, the reason not to take a certain consideration into account in one’s deliberations is that doing so would have an undesirable effect (or, at least, an effect that one has good reason to want to avoid). And so, Scanlon’s putative “counterexamples” are not counterexamples after all.

It seems, then, that Scanlon may be conflating reasons for and against performing an act, a_i , with reasons for and against taking into account various considerations when deciding whether to perform a_i . But the act of taking into account (or ignoring) various considerations when deciding whether to perform a_i is not the act of performing a_i , but instead the act of performing some other act, say, a_2 . So let a_2 be the act of taking into

³² The verb ‘take’ has many senses. In one sense, to take something to be irrelevant is to regard or treat it as irrelevant with the implication that this may be contrary to fact. This is how I am interpreting Scanlon here. Thus I am assuming that, by the phrase “taking another consideration, D, not to be relevant,” he means “treating another consideration, D, as being irrelevant (whether or not it is, or is believed to be, relevant).” If instead Scanlon is using his phrase to mean “believing D not to be relevant,” then he is not talking about reasons for actions (or, as he says, “reasons bearing on an action”), but is instead talking about reasons for belief. If he is talking about such epistemic reasons, then, as I explained above, this poses no problem for TCR, which instead concerns practical reasons.

account certain considerations when deciding whether to perform a_1 .³³ It seems that when we consider whether to perform a_2 , we should, as TCR implies, consider the agent's reasons for and against desiring that o_2 obtains. Thus TCR does not deny that the reasons for and against performing a_2 (i.e., taking into account certain considerations when deciding whether to perform a_1) may have nothing to do with the reasons for and against desiring that o_1 obtains. What it denies is only that the reasons for and against performing a_2 may have nothing to do with the reasons for and against desiring that o_2 obtains. And none of Scanlon's examples repudiate this, as I'll show presently.

Consider one of Scanlon's main examples, one where I have met someone for a game of tennis. Assume that I have determined that there are no strong reasons for or against my playing to win so that whether I have reason to play to win just depends on what I would enjoy doing at the moment. And assume that what I would enjoy most at this moment is playing to win and so this is what I have decided to do. Given all this, Scanlon claims, contrary to TCR, that the fact that my succeeding in making a certain strategic shot might make my opponent feel crushed or disappointed just is not relevant to whether or not I should make the shot. But this seems to be the wrong conclusion to draw. Scanlon never says what consideration here is the reason for ignoring (or doing what will cause me to ignore) this other consideration (that my opponent might feel crushed) into account when deciding whether or not to make this strategic shot, but it is clear that it is some pragmatic consideration such as the fact that it is very difficult, if not impossible, to play to win while taking these sorts of considerations into account, or it may be the fact that one cannot enjoy a good competitive game if one is constantly worrying about whether one's opponent's feelings might get hurt. But notice that we are

³³ I'm speaking a bit loosely when I treat taking into account (or ignoring) some consideration as itself a voluntary act, for, admittedly, whether one takes into account (or ignores) some consideration is very rarely, if ever, directly under one's volitional control. So, one doesn't so much intend to take into account (or ignore) some consideration as intend to do what will cause oneself to take into account (or ignore) that consideration. For instance, if I want to take into account certain considerations, I will intentionally focus my attention on considerations of that sort. And if, instead, I want to ignore those considerations, I will intentionally focus my attention elsewhere, purposely diverting my attention to something else whenever a consideration of that kind creeps into my consciousness. So the expression 'taking into account (or ignoring) certain considerations' should be taken as a convenient shorthand for 'doing what will cause oneself to take into account (or ignore) certain considerations'.

now appealing to the desirability of taking one's opponent's feelings into account, which is precisely what the teleologist takes to be relevant. The fact that we are not appealing to the desirability of making the strategic shot and weighing the desirability of the outcome where one makes the shot against the desirability of the outcome where one does not make the shot is neither here nor there. So it seems that this example shows only that the desirability or undesirability of embarrassing my opponent is not relevant to whether or not I should ignore such considerations. It does not, however, show that the desirability or undesirability of embarrassing my opponent is not relevant to whether or not I should make the shot.

Perhaps, even this is not the most charitable way of interpreting Scanlon. Perhaps, we should interpret Scanlon to be offering this sort of example, not as a counterexample, but rather as the first step in an argument against TCR. Scanlon could be arguing as follows: (1) As this example shows, we often do not treat the desirability/undesirability of certain features of an act's outcome as a reason for/against performing it. For instance, when playing to win, we do not treat the undesirability of embarrassing one's opponent as a reason against taking a certain shot. (2) In such cases, we are reasoning correctly in the sense that the considerations that we are not treating as reasons are, in fact, non-reasons.³⁴ (3) Therefore, TCR is false.³⁵ But, of course, premise 2 just begs the question. Therefore, Scanlon must do more than just offer this question-begging argument; he must offer some compelling reason for thinking that the examples that he appeals to (e.g., the tennis example) are best interpreted as (A) cases where the agent treats as non-reasons considerations that are, in fact, non-reasons as opposed to (B) cases where, on pragmatic grounds, the agent acts so as to cause herself to treat as non-reasons considerations that are, in fact, reasons (reasons with actual weight). Unfortunately, though, Scanlon offers no such argument. Perhaps, Scanlon just thinks that the phenomenology of our actual deliberations supports his interpretation: interpretation A. If so, then all I can do is cite my

³⁴ In calling these considerations 'non-reasons', I mean to allow that these considerations may count as genuine reasons in other contexts. The assumption, here, is only that due to their being silenced, undermined, or bracketed off, these considerations have ceased to be reasons in the context at hand.

³⁵ This interpretation was suggested to me by [redacted].

dissentation. When I reflect on my first-person deliberations in such cases, it feels to me like I'm deciding on pragmatic grounds to ignore certain relevant reasons. In any case, it would seem that Scanlon has failed to provide those sympathetic to TCR with any good reason to reject TCR, for those who are sympathetic to TCR will likely find interpretation B more plausible than interpretation A.

4. Arguments for TCR

Having argued that many of the reasons that philosophers have given for rejecting TCR are not in fact good reasons, I will now argue that there are good reasons to accept each of TCR's three claims. Furthermore, I will argue that there are good reasons for accepting TCR that go beyond whatever reasons there are for accepting each of its three claims.

§§4.1 *In defense of TCR-1*: According to TCR-1, if S has more reason to desire that o_i obtains than to desire that o_j obtains, then S has more reason to perform a_i than to perform a_j . To deny TCR-1 is to hold that it is possible for an agent to have more reason to perform a_i than to perform a_j even though she has more reason to prefer the world in which she performs a_j (viz., o_j) to the world in which she performs a_i (viz., o_i) than to prefer the world in which she performs a_i to the world in which she performs a_j . Thus, the denial of TCR-1 entails that the following will be possible: an agent ought (i.e., has most reason, all things considered) to perform a_i even though she ought to prefer the world in which she refrains from performing a_i to the world in which she performs a_i . This seems quite counterintuitive. If this were right, then, prior to action, the agent ought to hope that she doesn't do what she ought to do, and if she does somehow succeed in doing what she ought to do despite her hoping that she won't, she should wish that she hadn't. After all, it seems both that she should hope for the world that she has most reason, all things considered, to want to be actual and that, if things fail to go as hoped, she should wish that they had.³⁶ So, to deny TCR-1 is to accept that agents sometimes ought to hope that they refrain from doing what they ought to do and regret having done what they ought to

³⁶ For more on this point, see [redacted].

have done. Again, this seems quite counterintuitive. It would be better, then, to accept TCR-1 and deny such possibilities.

Given this, why would anyone reject TCR-1? One thing that might lead someone to reject to TCR-1 is the thought that although agents can have good (agent-neutral) reason, say, to prefer a world in which there has been only one murder to a world in which there has been five murders (other things being equal), agents ought not to commit murder even so as to prevent five others from each committing murder. The thought is that agents ought never to act so as to fail to *express* respect for rational persons and that, by murdering one person to prevent five others from doing the same, they would be doing just that. But why not also think that although agents have good *agent-neutral* reasons for preferring a world in which there has been only one murder to a world in which there has been five murders, they have even better *agent-relative* reasons for preferring a world in which five others commit murder to a world in which they themselves commit even one murder? To deny that an agent's agent-relative reasons for preferring that she not commit murder decisively oppose her agent-neutral reasons for preferring that there be fewer murders is to suggest that an agent can have decisive reasons not to commit murder even though she has decisive reasons to prefer that she does so. Such a suggestion can, I've claimed, seem quite counterintuitive. And if I'm right about this, then we she should hold that an agent's reasons for or against performing an act must always correlate with her reasons for or against preferring the world in which she performs that act to the world in which refrains from performing that act. And, thus, we should think that it will never be the case that an agent has more reason to desire that o_i obtains than to desire that o_j obtains but more reason to perform a_j than to perform a_i . That is, we should accept TCR-1.

§§4.2 *In defense of TCR-2*: According to TCR-2, if S has more reason to perform a_i than to perform a_j , then S has more reason to desire that o_i obtains than to desire that o_j obtains. We should accept TCR-2, I will argue, so as to accommodate the following two plausible claims:

- IN** *The Irreducible Normativity of Practical Reasons:* Practical reasons are irreducibly normative. Any purely reductive account of practical reasons will be unable to capture the normativity of practical reasons.
- ME** *The Motivational Efficacy of Practical Reasons:* An agent's normative reason for performing a_i must be capable of becoming her motivating reason for performing a_i . More specifically, if S has most reason to perform a_i , knows all the relevant facts, and is fully rational, then S's reasons for performing a_i will themselves be sufficient to move her to perform a_i . And if S has more reason to perform a_i than to perform a_j , knows all the relevant facts, and is fully rational, then these reasons will themselves be sufficient to move her to perform a_i rather than a_j when faced with the choice of performing one or the other.³⁷

I'll start by explaining why we should accept both IN and ME, and then I'll explain how TCR-2 enables us to accommodate the two. Consider ME first. ME is a fairly weak claim. It does not presuppose the Humean theory of motivation, according to which no belief could motivate us unless combined with some *independent desire*—a desire that is not itself the product of that belief (PARFIT 1997, p. 105). As far as ME goes, it may be, for instance, that the belief that some end is worth pursuing can itself produce a desire for that end. Nor does ME presuppose internalism about reasons, according to which S has a reason to perform a_i only if S would necessarily be motivated to perform a_i were she to deliberate on the relevant facts in a procedurally rational way (PARFIT 1997, p. 108). ME differs from internalism about reasons in allowing that it may not be enough for an agent to merely deliberate properly to ensure that she will be motivated to do what she has a normative reason to do; she may also need to desire what she has reason to desire. Thus, on ME, judgments about reasons for action are capable of motivating fully rational agents to act, but, to be fully rational, it may not be enough to simply follow some deliberative

³⁷ This is very close to what Christine Korsgaard call the "internalism requirement"—see KORSGAARD 1986, p. 297. And this something that rational intuitionists, Kantian internalists, and Humean internalists can all agree on. See DARWALL 2006, p. 298.

procedure in reflecting upon one's given beliefs and desires; one may also need to undertake the deliberative process with the desires and beliefs that one has reason to have—that is, one may also need to be substantively rational. Compared to internalism about reasons, then, ME takes the connection between (normative) reasons for action and motivation to be possibly more tenuous. How tenuous the connection will be depends on whether we take being fully rational to be merely a procedural matter (as internalists about reasons do) or also a substantive matter of having certain desires and beliefs (as externalists about reasons do). Yet whether we are internalists or externalists about reasons, it seems that we must, as ME does, preserve some necessary connection between reasons for action and motivation, for if some putative reason is not something that could motivate an agent even if she were fully rational, then it is hard to see in what sense it would be a reason for *her* at all.

We should also, I believe, accept IN. Recently, a number of leading philosophers (e.g., Parfit and Scanlon) have argued for IN. Below, I will try to briefly rehearse some of their arguments. One of the main reasons they give for accepting IN is the systematic failure of reductive accounts of 'a reason for S to perform a_i ' (SCANLON 1998, p. 57). Such accounts fall victim to the open question argument, for questions such as, "Act a_i is N ('N' standing for some non-normative property), but do I have a reason to perform a_i ?" always have an open feel to them. The reason seems to be that judging that the fact that a_i is N provides me with a reason to perform a_i entails drawing a normative conclusion that is not drawn if I judge merely that a_i is N. The normative claim that the fact that a_i is N is a reason to perform a_i seems to be a non-trivial addition to the non-normative claim that a_i is N. Non-normative facts may provide the grounds for reaching such normative conclusions, but judging that these facts obtain need not involve explicitly drawing these normative conclusions (STRATTON-LAKE AND HOOKER 2006, p. 151). So even if some non-normative facts are often or even necessarily practically relevant, the fact that such non-normative facts are practically relevant is still a further fact about them: a normative fact (DANCY 2006, p. 137). And this normative fact is not identical to the non-normative fact on which it is based.

Of course, the sorts of considerations cited above may seem to rule out only the possibility of analytic reduction. On the prospects for non-analytic reduction, we do well to consider the following from Derek Parfit:

Reductive views can be both non-analytical and true when, and because, the relevant concepts leave open certain possibilities, between which we must choose on non-conceptual grounds. But many other possibilities are conceptually excluded. Thus it was conceptually possible that heat should turn out to be molecular kinetic energy. But heat could not have turned out to be a shade of blue, or a medieval king. (1997, p. 122)

And just as heat could not turn out to be a certain shade of blue, it seems that a truth about what there is reason to do could not turn out to be a certain kind of purely naturalistic truth. Given the depth of difference between these two types of truths, a non-analytic reduction of one to the other is conceptually excluded. Here, an analogy might be helpful, but, as Parfit explains, it is hard to come by a close analogy for the irreducibility of normative facts to naturalistic facts. Here is Parfit:

Since normative facts are in their own distinctive category, there is no close analogy for their irreducibility to natural facts. One comparison would be with proposed reductions of necessary truths—such as truths of logic and mathematics—to certain kinds of contingent truths. Given the depth of the difference between these kinds of truth, we can be confident, I assume, that such reductions fail. There is a similar difference, I believe, between normative and natural truths. (PARFIT 1997, p. 122)

The thought is, then, that the difference between normative facts and naturalistic facts is as substantial as that between the truths of mathematics and contingent truths, and, thus, substantial enough that we can be confident that all reductions (analytic or non-

analytic) of the normative to the natural will fail.³⁸ Moreover, it seems that the difference between normative facts and naturalistic facts is no more substantial than the difference between normative facts and other non-normative facts, such as facts about the supernatural. So we can further conclude that all reductions (analytic or non-analytic) of the normative to the non-normative fail. When we attempt to identify normative facts with non-normative facts of any sort (natural or supernatural), we change the subject matter and thus abolish their normativity. The fact that a_i is N does not have the same subject matter as the fact that there are reasons to perform acts that are N. The latter fact is a normative fact about the former (DANCY 2006, p. 140). There is, then, a strong presumptive case to be made in favor of the view that any adequate account of practical reasons will need to take IN into account.

Now, to accommodate both IN and ME, it seems that we will need to provide an account of reasons for action in terms of reasons for possessing some motivational state. Of course, not just any motivational state will do, as some motivational states are not sensitive to our judgments about reasons. Mere urges, for instance, are motivational, but they are not judgment-sensitive, that is, they are not sensitive to our judgments about reasons (SCANLON 1998, p. 20). Desires, by contrast, are judgment-sensitive. And, contrary to Hume, many philosophers now think that reason is the master of, not a slave to, desires (see, e.g., SCANLON 1998, chap. 1). If this is right, motivation can be the consequence of recognizing reasons to desire (BRINK 2008). All this leads me to conclude that the best account of reasons for action will be in terms of reasons for desiring. I see no other better way of accommodating both ME and IN. Other potential accounts of practical reasons seem to run afoul of ME, IN, or both.

To illustrate, consider the following schema:

S has a reason to intend to perform a_i if and only if a_i is F.

³⁸ You find this same sort of argument in HUEMER 2005, pp. 94-95.

If, on the one hand, we take F to be something normative but not intrinsically motivating, such as “an action that S has sufficient reason to believe is utility maximizing,” then we will ensure IN but only at the cost of forsaking ME, for it seems that the judgment that one has sufficient reason to believe that a_i is utility maximizing is not the sort of judgment that, by itself, is sufficient to motivate the judger to perform a_i even when both fully informed and fully rational. Even those who think that some beliefs are *intrinsically motivating* (i.e., capable of motivating even absent some independent desire) do not think that this sort of belief is intrinsically motivating.³⁹ The sorts of beliefs that are typically held to be intrinsically motivating are normative beliefs—for instance, the belief that a_i is what one has most reason to do, not the belief that a_i is utility maximizing. Moreover, it is only a subset of normative beliefs that could possibly motivate S to perform a_i absent some independent desire: normative beliefs about a_i and its consequences, not normative beliefs about, for instance, beliefs about a_i . Thus, normative beliefs, such as, “ a_i is morally required,” “ a_i is what there is most reason to do,” and “ a_i has an outcome that there is sufficient reason to desire,” are all, perhaps, intrinsically motivating. But it is clear that the normative belief “there are reasons to believe that a_i is utility maximizing” could motivate only when conjoined with some independent desire, e.g., the desire either to maximize utility or to do what there is sufficient reason to believe is utility maximizing. Thus if F is something normative but not intrinsically motivating, then ME will not hold—that is, it will not be the case that S, even if fully informed and fully rational, will be moved by F *alone* to do a_i (i.e., without some independent desire needing to be present).

If, on the other hand, we take F to be something non-normative but intrinsically motivating, such as “an action that will fulfill at least one of S’s present desires,” then we will have to forsake IN, for the fact that doing a_i will fulfill one’s present desires is non-normative. Moreover, the fact that performing a_i will fulfill one’s present desires is not even a fact of normative significance, since such facts do not, as Parfit has argued, provide

³⁹ It seems to me that those who take beliefs to be intrinsically motivating should concede that a desire will need to be present to move the agent but argue that the desire will itself be a product of the belief and, thus, be a dependent desire.

one with any reason to perform a_i .⁴⁰ To ensure that both ME and IN are met, then, we should accept TCR-2, where reasons for intending to perform a_i are explained in terms of reasons for desiring that o_i obtains.

Admittedly, I have, in this section, relied on a number of contentious metaethical claims, and I have had the space here only to briefly sketch some arguments in support of these claims. At best, then, the case presented here is only a presumptive case in favor of TCR-2. Nevertheless, it is important to note that these claims have been defended more extensively elsewhere (SCANLON 1998 and PARFIT 1997) and by at least one of the leading critics of TCR (i.e., Scanlon). So this presumptive case should have persuasive force against even some of TCR's leading critics.⁴¹

§§4.3 *In defense of TCR-3*: As noted above, the teleologist is committed not only to the bi-conditional relationship between reasons for acting and reasons for desiring entailed by the conjunction of TCR-1 and TCR-2, but also to the view that reasons for desiring are explanatorily prior to reasons for acting. And, thus, the teleologist must defend TCR-3: if S has more reason to perform a_i than to perform a_j , then this is in virtue of the fact that S has more reason to desire that o_i obtains than to desire that o_j obtains.

In defense of TCR-3, I will argue that it is more plausible than its alternatives. If we let 'P' stand for 'S has more reason to *perform* a_i than to *perform* a_j ' and let 'D' stand for 'S has more reason to *desire* that o_i obtains than to *desire* that o_j obtains', then the three possible alternatives to TCR-3 (i.e., to the view that, in every instance, 'P, because D') are: (1) in every instance, 'D, because P'; (2) in some instances, 'P, because D', in other instances, 'D, because P', but, in every instance, one or the other; and (3) in every instance, both 'P, because B' and 'D, because B' — where, for instance, 'B' might stand for ' o_i is *better* than o_j '.

⁴⁰ See PARFIT 2008 for an extended discussion of why such desire-based theories of practical reasons fail. For more on the distinction between a normative fact (a fact about reasons) and a fact of normative significance (a non-normative fact that provides reasons), see DANCY 2006.

⁴¹ Although I've relied on IN in arguing for TCR, it should be noted that TCR is compatible with the denial of IN. TCR is, for instance, perfectly compatible with ethical naturalism, for one could give a reductive naturalistic account of reasons for desiring.

So the first alternative to TCR-3 is to hold that what explains the correlation between P and D is that, in every instance, 'P, because D', not 'D, because P', as TCR-3 supposes. This is to hold that when I have more reason to desire that o_i obtains than to desire that o_j obtains, this is *always* because I have more reason to perform a_i than to perform a_j . But this is clearly false, as one simple example is sufficed to show. Take the case where a_i is the act of putting my money into savings account A, which yields 5% annually, and where a_j is the act of putting my money into savings account B, which yields 3% annually. Assume that both savings accounts are otherwise equal and that, other things being equal, I am better off putting my money in a higher-yielding savings account. In this case, it is clear that what explains the fact that I have more reason to perform a_i than to perform a_j is that I have more reason to desire that o_i obtains than to desire that o_j obtains. The same would seem to hold for any other prudential choice without moral ramifications, for no one would argue that the reason I have for preferring the world in which I am prudentially better off is that this is the world in which I did what was prudent (i.e., what I had most prudential reason to do). We should, then, reject the first alternative.

The second alternative is to claim that what explains the correlation between P and D is that, in some instances, 'P, because D', in other instances, 'D, because P', but, in every instance, one or the other. The problem with this view is that it is fragmented, and thus its proponents owe us some explanation as to why sometimes reasons for acting are explanatorily prior and other times reasons for desiring are explanatorily prior. Unless there some explanation for this, unless there is, for instance, some set of intuitive judgments that only this fragmented view could account for, we should accept instead the more unified and systematic TCR-3.

To see whether there are any such intuitive judgments that only this fragmented view could account for, we do good to consider Elizabeth Anderson's view (1993). On Anderson's view, an action is rational only if it adequately expresses one's rational evaluative attitudes (e.g., respect, consideration, and appreciation) toward persons, animals, and things. And, on her view, it is rational to adopt the aim (or the end) of

bringing about some state of affairs only if the act that would bring it about is itself rational. To illustrate, consider that it is, on Anderson's view, rational for me to act in ways that express my love for my wife and that it is in virtue of this that I have reason to want, and to aim to bring about, the state of affairs in which we spend time alone with each other. The ground, then, for my wanting and pursuing this state of affairs is that in doing so I express my rational evaluative attitudes towards her (i.e., I express my love for her). On Anderson's view, then, whether I have reason to adopt as my end the state of affairs in which we spend time alone together depends on whether it is rational for me to act so as to adequately express my love for her.

Anderson's view, thus, seems contrary to TCR-3 in that it holds that the rationality of desiring and pursuing some end, e_1 , can depend on the rationality of performing the act that brings it about. Despite appearances, though, this needn't be contrary to TCR-3. Whether it is or not depends on whether there is some *further* end, e_2 , in virtue of which it is rational to perform the act that brings about e_1 . Perhaps, my wanting and pursuing the state of affairs where my wife and I spend time alone together is rational only because it is rational for me to want and pursue the state of affairs where my actions adequately express my rational evaluative attitudes toward her. If this is right, then there is no conflict between Anderson's view and TCR-3. The issue, then, is whether or not Anderson's injunction to act only in ways that express one's rational attitudes towards persons, animals, and things is extensionally equivalent to some injunction to act so as to promote some end—some state of affairs.

In fact, Anderson argues that her injunction, which she calls "E," to only act in ways that *express* one's rational attitudes toward persons is not extensionally equivalent to the injunction, which she calls "P," to act so as to *promote* the state of affairs where one's actions adequately express one's rational attitudes. She argues that E is no more extensionally equivalent to P than the injunction, E' , to only make logically valid inferences is extensionally equivalent to the injunction, P' , to act so as to promote the state of affairs where one makes logically valid inferences. In both cases, the latter (P or P') tells one to do something that the former (E or E') does not: in the case of P, to perform more

acts that adequately express one's rational attitudes toward persons and, in the case of P', to make more valid inferences. Moreover, unlike E, "P tells me to violate E, if by doing so I can bring about more events containing E than if I never violated E" (ANDERSON 1996, p. 544). Thus, unlike E, P tells me to betray my wife by committing adultery if my doing so will create an opportunity for me to more fully express my love for her than I otherwise would have, as where this betrayal will enable a reconciliation that enables me to express my love for her more fully than before (ANDERSON 1996, p. 545). So Anderson does show that E and P yield different prescriptions. But these differences in what E and P prescribe do not demonstrate that E is not equivalent to *some* injunction to promote a state of affairs; they only demonstrate that E is not extensionally equivalent to P.

So consider P*: act so as to *promote* the state of affairs where as few of one's actions as possible fail to adequately express one's rational attitudes. Unlike P, P* does not tell one to act so as to promote greater instances of acting in ways that express one's rational attitudes toward persons. And, like E, it functions as a constraint, telling one not to violate E (i.e., not to act in a way that fails to adequately expresses one's rational attitudes toward persons) even in order to bring about more events containing E (i.e., more events where one adequately expresses one's rational attitudes toward persons). Of course, P* would tell us to violate E if doing so will minimize one's violations of E. But, as Sturgeon (1996, p. 521) has pointed out, it is hard to imagine any but the most fanciful of cases where violating E now will prevent me from committing more numerous, comparable violations of E in the future. Nor is it obvious that, in such fanciful cases, it is E rather than P* that gets the more plausible result.⁴² But rather than debate the issue of whether E or P* is more plausible, it is clear that E is extensionally equivalent to P**: act so as to *promote* the state of affairs where one's current actions do not fail to adequately express one's rational attitudes toward persons.⁴³ P** and E yield the exact same prescriptions, and so we see that there is nothing that E can account for that P** can't account for. And since the

⁴² For more on this issue, see [redacted].

⁴³ For any deontological injunction prohibiting agents from performing a certain type of action, there is, as I and others have argued, an equivalent consequentialist injunction requiring agents to bring about a certain state of affairs. See, for instance, [redacted] and LOUISE 2004.

proponent of TCR-3 can adopt P** and thereby account for all the intuitive judgments that Anderson appeals to. It remains unclear, then, why we should opt for the second alternative to TCR-3 when TCR-3 can explain everything that it can, only more systematically.

This brings us to the third and final alternative to TCR-3: that what explains the correlation between P and D is that, in every instance, both 'P, because B' and 'D, because B'. We should, I think, reject this third alternative. Consider what B might be. That is, consider what might possibly explain both P and D? The only likely possibility that comes to mind is the one where 'B' stands for '*o_i is better than o_j*'. But, if 'B' stands for '*o_i is better than o_j*', then we should, I think, reject the claim that 'D, because B'. As many philosophers have come to think, this just gets the order of explanation wrong, for to say that *o_i*'s obtaining would be good is just to say that *o_i* has the purely formal, higher-order property of having other properties that provide sufficiently weighty reasons of the right kind for desiring that *o_i* obtains.⁴⁴

But even if we don't accept this, the buck-passing account of value, we should still reject the claim that, in every instance, 'D, because B', for it seems that sometimes an agent will have more reason to desire that *o_i* obtains than to desire that *o_j* obtains even though *o_i* is *not* better than *o_j*. And so it can't be that, in every instance, what explains D is B. To illustrate, consider the following case in which it is plausible to think D but not B (that is, the plausible to think that the agent has more reason to desire that *o_i* obtains than to desire that *o_j* obtains even though *o_i* is *not* better than *o_j*). Suppose that both my child and some stranger's child are drowning and that I must choose either to save my child or to enable the stranger to save hers. Further suppose that the possible world in which I save my child is just as good as the possible world in the stranger saves her child—the stranger's child would live just as good a life as my child would, the stranger would lament the loss of her child just as much as I would lament the loss of my child, etc. It still seems that I could have more reason to prefer the possible world in which I save my child

⁴⁴ Again, this is called the fitting-attitude or buck-passing account of value—see, for instance, RABINOWICZ AND RØNNOW-RASMUSSEN 2004, SCANLON 1998, and STRATTON-LAKE AND HOOKER 2006.

to the possible world in which the stranger saves her child. This seems possible, because although which world is better is simply a function of agent-neutral reasons for preferring the one to the other, which I have most reason to prefer, *all things considered*, is a function of both my agent-neutral reasons and my agent-relative reasons for preferring the one to the other. And, in this case, it seems that I have compelling agent-relative reasons to prefer that I save my child, for it is my child whom I have loved and nurtured since she was a baby. In particular, the fact that it would be *my* child who be saved if I do the saving is a compelling agent-relative reason to prefer the possible world in which I save my child to the possible world in which the stranger saves her child.

So, if the only real candidate for 'B' is ' o_i is *better* than o_j ', and if, on this interpretation of 'B', we should reject 'D, because B', as I've argued, then we should reject this common-"cause" alternative to TCR-3. Of course, these might seem like two big ifs. So, again, the argument here is only presumptive. Given the prima facie plausibility of these antecedents, there is a presumptive case to be made against this alternative to TCR-3. And, importantly, we find, again, that the argument should have persuasive force against some of TCR's leading critics, for critics such as Scanlon accept my assumption that it is D that explains B, and not vice versa.

§§4.4 *In defense of TCR on the whole*: Admittedly, the arguments for TCR-1, TCR-2, and TCR-3 are each far from decisive. For one thing, the argument for TCR-1 relied on certain intuitions that may not be shared by everyone. For another, the arguments for TCR-2 and TCR-3 each relied, in part, on certain controversial metaethical assumptions about reasons and values, which I haven't had space here to adequately defend.⁴⁵ But, again, it is worth noting that these controversial assumptions are ones that at least some of TCR's leading critics endorse. Of course, my aim is not just to convince TCR's critics, but to convince others as well, and yet not everyone will both share these intuitions and accept these assumptions. So let me explain why I think that even those who don't do, nevertheless, have good reason to accept TCR.

⁴⁵ It's important to note that although my arguments for TCR have relied on these controversial metaethical assumptions, one needn't accept these assumptions either to find TCR attractive or to endorse it. For instance, one may think that reasons are not irreducibly normative and yet still want to hold that TCR is true.

Compared to its alternatives, TCR does a far superior job of systematizing our various substantive convictions about what we have reason to do. Everyone seems to admit that, at least, in some instances an agent's reasons for performing a given action derive from her reasons for desiring its outcome. For instance, we think that, other things being equal, you have more reason to choose the more pleasure-inducing meal option to the less pleasure-inducing meal option just when, and because, you have more reason to desire the world in which you experience more pleasure than to desire the world in which you experience less pleasure. No one thinks that it's the other way around: that you have more reason to want the world in which you experience more pleasure because you have more reason to choose the more pleasure-inducing option. The only point in contention, then, is whether reasons for action always derive from reasons for desiring. Some have thought not, for they have thought that there are certain substantive views about value and practical reasons that TCR cannot accommodate—e.g., the view that it is certain concrete entities, and not states of affairs, that are the primary bearers of intrinsic value and that, when playing to win at a game of tennis, there are reasons not to take certain other considerations into account when deciding whether to make a certain strategic shot. But, as I've shown above, this is a misconception: there are no substantive views about values and reasons that TCR cannot accommodate. TCR is an extremely ecumenical view. Thus, one good reason to accept TCR, apart from whatever reasons there are to accept each of its three constitutive claims, is that TCR provides a systematic account of our various substantive views about reasons for action.

5. Conclusion

Whenever we face the choice of how to act, we also face the choice of which of various possible worlds to actualize. It is through our actions that we affect the way the world goes. The teleological conception of practical reasons holds that our reasons for acting are a function of our reasons for preferring some of these possible worlds to others. On this view, if I ought to prefer the possible world that will be actualized by my performing a_i at

t to the possible world that will be actualized by my performing a_j at t , then, given the choice of performing either a_i or a_j at t , I ought to perform a_i at t .

I have defended this view against both putative counterexamples and many actual and potential misconceptions. I have shown that this conception of practical reasons is quite ecumenical, and so we do not have to give up any of our considered convictions about values or reasons to accept it. I have also offered positive arguments in favor of each of TCR's three claims, arguing, among other things, that we must accept TCR so as to accommodate both the irreducible normativity and motivational efficacy of practical reasons. Moreover, I have argued that TCR is, in contrast to its alternatives, superior in its ability to systematize our considered convictions about practical reasons.⁴⁶

References [redacted]

- ANDERSON, E. (1993). *Value in Ethics and Economics*. Cambridge, Mass.: Harvard University Press.
- . (1996). "Reasons, Attitudes, and Values: Replies to Sturgeon and Piper." *Ethics* 106: 538-554.
- ARNESON, R. (2002). "The End of Welfare as We Know It? Scanlon versus Welfarist Consequentialism." *Social Theory and Practice* 28: 315-336.
- BRINK, D. O. (2008). "The Significance of Desire." *Oxford Studies in Metaethics* 3.
- BROOME, J. (1991). *Weighing Goods*. Oxford: Basil Blackwell.
- BYKVIST, K. (2003). "Normative Supervenience and Consequentialism." *Utilitas* 15: 27-49.
- CHANG, R. (2004). "Can Desires Provide Reasons for Action?" In R. J. Wallace, P. Pettit, S. Scheffler, and M. Smith, eds., *Reason and Value*. Oxford: Clarendon Press.
- CRISP, R. (2000). "Review of *Value...and What Follows*. By Joel Kupperman." *Philosophy*, 75: 458-462.

⁴⁶ [redacted] For helpful comments on earlier drafts, I thank [redacted].

- DARWALL, S. (2006). "Morality and Practical Reason: A Kantian Approach." In D. Copp, ed., *The Oxford Handbook of Ethical Theory*. Oxford University Press.
- . (2002). *Welfare and Rational Care*. Princeton: Princeton University Press.
- . (1983). *Impartial Reason*. Ithaca: Cornell University Press.
- DANCY, J. (2006). "Nonnaturalism." In D. Copp, ed., *The Oxford Handbook of Ethical Theory*. Oxford University Press.
- DREIER, J. (1993). "Structures of Normative Theories." *The Monist* 76: 22-40.
- HARMAN, E. (2005). "The Mistake in 'I'll Be Glad I Did It' Reasoning: The Significance of Future Desires." Working manuscript of 10/26/05.
- HUEMER, M. (2005). *Ethical Intuitionism*. New York: Palgrave Macmillan.
- HURKA, T. (2006). "Value and Friendship: A More Subtle View." *Utilitas* 18: 232-242.
- KAVKA, G. (1983). "The Toxin Puzzle." *Analysis* 43: 33-36.
- KORSGAARD, C. (1997). "The Normativity of Instrumental Reason." In G. Cullity and B. Gaut, eds., *Ethics and Practical Reason*. Oxford: Clarendon Press.
- LOUISE, J. (2004). "Relativity of Value and the Consequentialist Umbrella." *Philosophical Quarterly* 54: 518-536.
- MCAUGHTON, D. AND RAWLING, P. (1991). "Agent-Relativity and the Doing-Happening Distinction," *Philosophical Studies* 63: 167-185.
- OLSON, J. (2008). "Fitting Attitude Analyses of Value and the Partiality Challenge." Working manuscript of 5/16/08.
- PARFIT, D. (2008). *Climbing the Mountain*. Retrieved January 29, 2008, from http://individual.utoronto.ca/stafforini/parfit/parfit_-_climbing_the_mountain.pdf.
- . (2001). "Rationality and Reasons." In D. Egonsson, J. Josefsson, B. Petersson, and T. Rønnow-Rasmussen, eds., *Exploring Practical Philosophy*. Aldershot: Ashgate.
- . (1997). "Reasons and Motivation," *Proceeding of the Aristotelian Society Supplementary Volume* 71: 99-130.
- QUINN, W. (1993). "Putting Rationality in its Place." In his *Morality and Action*. Cambridge: Cambridge University Press.

- RABINOWICZ, W. AND RØNNOW-RASMUSSEN, T. (2004). "The Strike of the Demon: On Fitting Pro-Attitudes and Value." *Ethics* 114: 391-423.
- RIDGE, M. (2005). "Review of Pleasure and the Good Life." *Mind* 114: 414-417.
- SCANLON, T. M. (2007). "Structural Irrationality." In G. Brennan, R. Goodin, F. Jackson, and M. Smith, eds., *Common Minds: Themes from the Philosophy of Philip Pettit*. Oxford: Oxford University Press.
- . (1998). *What We Owe to Each Other*. Cambridge, Mass.: The Belknap Press.
- SCHROEDER, M. (2008a). "Value and the Right Kind of Reason." Working manuscript of 8/10/2008.
- . (2008b). "Value Theory," *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition), Edward N. Zalta (ed.), forthcoming URL = <http://plato.stanford.edu/archives/fall2008/entries/value-theory/>.
- . (2007). "Teleology, Agent-Relative Value, and 'Good'." *Ethics* 117: 265-295.
- SKORUPSKI, J. (1999). *Ethical Explorations*. Oxford: Oxford University Press.
- SMITH, M. (2006). "Two Kinds of Consequentialism." Forthcoming in *Philosophical Issues* 19. Working manuscript of 10/12/06.
- . (1994). *The Moral Problem*. Oxford: Blackwell.
- STURGEON, N. L. (1996). "Anderson on Reason and Value." *Ethics* 106: 509-524.
- STRATTON-LAKE, P. AND HOOKER, B. (2006). "Scanlon versus Moore on Goodness." In T. Horgan and M. Timmons, eds., *Metaethics after Moore*. Oxford: Clarendon Press.
- STROUD, S. (1998). "Moral Overridingness and Moral Theory." *Pacific Philosophical Quarterly* 79: 170-189.
- SUIKKANEN, J. (2008). "Consequentialism and the Good-Relative-to: A Reply to Mark Schroeder." Working manuscript of 9/10/04.
- TÄNNSJÖ, T. (1999). "A Concrete View of Intrinsic Value." *The Journal of Value Inquiry* 33: 531-536.
- VAN ROOJEN, M. (2007). "Moral Rationalism and Rational Amoralism." Working manuscript of 9/24/2007.

WALLACE, R. J. (2005). "Reasons, Values, and Agent-Relativity." Draft of June 8, 2005, retrieved August 22, 2006, from <http://philosophy.berkeley.edu/file/19/Buck-passing.pdf>.