

The Teleological Conception of Practical Reasons

DOUGLAS W. PORTMORE

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. The aim could be nothing more than to bring it about that one performs the act. For instance, one could intend to run merely for the sake of bringing it about that one runs. The fact remains, though, that for every intentional action there is some end at which the agent aims.

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 what we have most reason to do is determined by which way we have most reason to want the world to go. To put things more precisely, an agent's

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.

reasons for action are a function of her reasons for preferring certain 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 want to be actual. 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. It's important in its own right, for asking what we have most reason to do is one of the most fundamental practical questions that we can ask. A theory about practical reasons won't necessarily tell us what we legally, morally, or prudentially ought to do, but it will tell us what's arguably even more important: namely, what we just plain ought to do (that is, what we ought to do, all things considered).² And, as noted above, the teleological conception of practical reasons may also have important implications for moral theorizing. Whether it does or not depends on whether *moral rationalism* is true. Moral rationalism is the view that an agent can be morally required to do only what she has most reason to do.³ If both moral rationalism and the teleological conception of practical reasons are true, then the correct moral theory will be some species of *rational-desire teleology*: the view that an act's deontic status is determined by the agent's reasons for and against preferring its outcome to those of the available alternatives, such that, if an

² Admittedly, if some reasons have no (rational) requiring strength (see GERT 2004 and DANCY 2004) or if some reasons with requiring strength are silenced, undermined, or bracketed off by other factors or considerations (see, e.g., SCANLON 1998), then it won't always be the case that agents are rationally required to do what they have most reason to do. But I'm using the phrase 'ought to do *x*' in a sense that doesn't imply a requirement to do *x*. In this sense, I may, for instance, tell my dinner companion that she ought to try the duck without thereby implying that she is in any sense required to try the duck. This sense of 'ought' can be called the *advisory sense* of 'ought'. Another example of the advisory sense of 'ought' comes from FERRY 2009: "[Y]ou really ought to visit your mother on her birthday, but you have got to at least send a card." See also MCNAMARA 1996. For more on the notion of what we just plain ought to do, see MCLEOD 2001.

³ The thesis that I call *moral rationalism* sometimes goes by other names. David Brink calls moral rationalism "the supremacy thesis" (1997, p. 255), Stephen Darwall calls it "supremacy" (2006b, p. 286), Samuel Scheffler calls it "the claim of overridingness" (1992, pp. 52-54), John Skorupski calls it "the principle of moral categoricity" (1999, p. 170), Sarah Stroud calls it the "overridingness thesis" (1998, p. 171), and R. Jay Wallace calls it the "optimality thesis" (2006, p. 130). For a defense of the thesis and for more on its implications for moral theorizing, see PORTMORE 2011a and PORTMORE 2011b.

agent is morally required to perform an act, then, of all the outcomes that she could bring about, she has most reason to desire that its outcome obtains.⁴

Below, I'll argue that we should accept the teleological conception of (practical) reasons—or 'TCR' for short. TCR is, on its face, quite plausible. Even its critics admit as much.⁵ Since TCR is *prima facie* plausible, my main task will be to show that it is not subject to the sorts of objections that critics have leveled against it. Thus, I'll dedicate most of this paper to rebutting objections and clearing up misconceptions. Nevertheless, I will, in the last third of the paper, try to offer some positive arguments for it.

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.⁶ 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 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

⁴ This contrasts both with *impersonal-value teleology* (the view that an act's deontic status is determined by the impersonal value of its outcome, such that, if S is morally required to perform x , then S 's performing x would maximize the good) and with *personal-value teleology* or egoism (the view that an act's deontic status is determined by the personal value of its outcome, such that, if S is morally required to perform x , then S 's performing x would maximize S 's good). Some hold that the term 'consequentialism' is best used as a label for only impersonal-value teleology. Others (e.g., PORTMORE 2011) argue that the term 'consequentialism' should be used broadly to encompass all three forms of teleology.

⁵ Scanlon, for instance, admits that this view "sounds plausible" (1998, p. 84).

⁶ 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.

be their corresponding outcomes, where an act's outcome is construed broadly as the possible world that would be actual were the act to be performed.⁷ 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 so in virtue of the fact that S has more reason to desire that o_i obtains than to desire that o_j obtains.⁸

More concisely, then, TCR is the view that S has more reason to perform a_i than to perform a_j if and only if, and because, S has more reason to desire that o_i obtains than to desire that o_j obtains.⁹ For my purposes, though, it will be useful to keep the three claims

⁷ Again, I'm assuming counterfactual determinism—see note 1 above.

⁸ Michael Smith's (2005, p. 17) formulation of the view is similar to mine. He says, "(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."

Of course, I've been assuming that the laws of nature are deterministic and, thus, that counterfactual determinism is true. If, instead, the laws of nature are indeterministic, then TCR will need to be reformulated so as to take the reasons to perform an action to be a function of the sum of the expected S-relative desirability values of each of the possible outcomes associated with that action—that is, a function of $\sum_i Pr(o_i/a_j) \times D_s(o_i)$, where a_j is the given action, $Pr(o_i/a_j)$ is the probability of o_i 's obtaining given S's performance of a_j , and $D_s(o_i)$ is the S-relative desirability value of o_i , which is just a measure of how much reason S has to desire that o_i obtains. Thus, if the laws of nature are indeterministic, we must replace TCR with the following:

TCR* (1) S has more reason to perform a_j than to perform a_k if $\sum_i Pr(o_i/a_j) \times D_s(o_i)$ is greater than $\sum_i Pr(o_i/a_k) \times D_s(o_i)$; (2) S has more reason to perform a_i than to perform a_j only if $\sum_i Pr(o_i/a_j) \times D_s(o_i)$ is greater than $\sum_i Pr(o_i/a_k) \times D_s(o_i)$; and (3) if S has more reason to perform a_j than to perform a_k , then this is so in virtue of the fact that $\sum_i Pr(o_i/a_j) \times D_s(o_i)$ is greater than $\sum_i Pr(o_i/a_k) \times D_s(o_i)$.

⁹ It may be objected that complete possible worlds—or, in other words, total outcomes, such as o_i —are far too complex to be the objects of one's conscious desires. And, as the objection might run, if o_i is not the sort of thing that one can desire, then it's not the sort of thing that one can have reason to desire. But I don't see why one must have every aspect of o_i conscious before one's mind in order to desire it or to have a reason to desire it. It seems clear that there are reasons for individuals to desire certain total outcomes and to prefer them to others. For instance, the fact that Smith's child as opposed to some stranger's child would be saved if she performs a_1 as opposed to a_2 is clearly a reason for her to prefer o_1 to o_2 . And although Smith may be incapable of appreciating in total all the various reasons that she has for preferring o_1 to o_2 , this doesn't mean that she

separate; I'll refer to them as TCR-1, TCR-2, and TCR-3, respectively. And although not stated above, I take TCR to include the claim that S has a reason to perform a_i if and only if, and because, S has a reason to desire that o_i obtains.

Having stated the view as precisely as I can, I will now proceed to clarify it, explaining in the process how and why it differs from T. M. Scanlon's statement 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). Although my statement of TCR differs from Scanlon's in how it's worded, I don't believe that it differs in substance from the view that he intended to describe—or so I shall argue. I take my statement of TCR to differ from Scanlon's only in its degree of clarity and precision. There are, in fact, five separate points that need clarifying.

§§1.1 *Reasons for desiring as opposed to desirability*: Unlike Scanlon, I state TCR in terms of the agent's reasons for desiring various possible outcomes as opposed to the desirability of those 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 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 the right kind to desire that o_i obtains.¹⁰ What are the right kinds of reasons? They are all and only those reasons that are relevant to determining whether, and to what extent, o_i is desirable. Let's call these *fittingness reasons*.

can't have most reason to prefer o_1 to o_2 . After all, which total outcome Smith has most reason to prefer is simply a function of whatever the various specific reasons, on balance, supports her preferring. I don't see, then, why Smith must be capable of having every aspect of o_1 before her mind in order for her to have most reason to desire it or to prefer it.

¹⁰ 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 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 of the states of affairs in the relevant contextually supplied comparison class. For more on this, see SCHROEDER 2010 and SCHROEDER 2008.

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 controversial issues, and I have nothing new to add to the growing debate. Even so, I can plausibly claim that there are some clear cases of what 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, then these are clearly not fittingness reasons, for such pragmatic "reasons" for desiring that o_i obtains clearly have no bearing on whether, or to what extent, o_i 's obtaining is desirable.¹¹ 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 reasons only to want, and to act so as to cause oneself, to desire that o_i obtains. If, however, I'm wrong about this, then admittedly I'll need to revise TCR so as to exclude such pragmatic "reasons," for such reasons are no more relevant to whether one should act so as to bring it about that o_i obtains than they are to whether o_i 's obtaining is desirable.¹²

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 the outcome in which my child lives to the outcome in which her 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 desirable than the outcome in which her child goes on living. Other things being equal, the outcome in which my child lives is

¹¹ 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 RABINOWICZ AND RØNNOW-RASMUSSEN 2004, PARFIT 2011, and SCHROEDER 2010.

¹² The revised version of TCR would, then, read as follows: S has more reason to perform a_i than to perform a_j if and only if, and because, the set of all the *non-pragmatic reasons* that S has for desiring that o_i obtains is weightier than the set of all the *non-pragmatic reasons* that S has for desiring that o_j obtains. Pragmatic reasons for S to desire that o_i obtains are reasons that are provided by facts about the consequences of S's desiring that o_i obtains, and non-pragmatic reasons are just reasons that are not pragmatic—for a more careful account of the relevant distinction, see STRATTON-LAKE 2005. The distinction that I'm drawing between pragmatic and non-pragmatic reasons is closely related to Derek Parfit's (2001) distinction between state-given and object-given reasons as well as to Christian Piller's (2006) distinction between attitude-related and content-related reasons; it's not clear to me, though, that either is extensionally equivalent to mine.

As will be evident shortly, it's important to note that many agent-relative reasons for S to desire that o_i obtains, such as the fact that S's child will live if and only if o_i obtains, are non-pragmatic reasons and thus will not be excluded by the restriction to non-pragmatic reasons in the above revised version of TCR.

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.¹³ And, thus, agent-relative reasons for preferring one outcome to another are not fittingness reasons for preferring the one to the other.¹⁴

It is this last exclusion that makes trouble for stating TCR in terms of the value or the desirability of outcomes. If TCR is to be stated 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 as opposed to some stranger's child. Yet there is no reason why the teleologist should exclude the possibility of agent-relative reasons for action, as even the 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, p. 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 have more reason to bring about some state of

¹³ This is known as the *partiality challenge* to the fitting-attitude or buck-passing account of value. See OLSON 2009, SUKKANEN 2009, and ZIMMERMAN 2010 for some potential solutions to this particular problem.

¹⁴ 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 at some point after performing that 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 wish/prefer (at this point) that she had had the abortion and then later given 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 she had formed any special relationship with what was then just an early-stage fetus, she should have preferred the outcome in which she had had the abortion and then later given 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 if and only if, 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 preferable state of affairs in which that loved one doesn't exist, see HARMAN 2009.

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.

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). We can properly value/desire some outcomes more than others without claiming that they are better than, or preferable to, those others. This is because which is most preferable (i.e., best) is a function of only our agent-neutral reasons for preferring them to the others, whereas which possible world we have most reason to prefer is a function of both our agent-relative and our agent-neutral reasons for preferring them to the others.

As Jussi Suikkanen (2009, p. 6) argues, it is better (impersonally speaking) that o_i obtains than that o_j obtains if and only if it would be fitting for an impartial spectator to prefer o_i to o_j . An impartial spectator must be both impartial and a mere spectator. To ensure that she is impartial, we must assume that she has no personal relations with anyone involved. And to ensure that she is a mere spectator, we must assume that she is not involved either in bringing it about that o_i obtains or in bringing about that o_j obtains. Thus, the impartial spectator can have nothing but agent-neutral reasons for preferring one outcome to the other.

But, unlike the impartial spectator, situated agents can have agent-relative reasons for preferring one outcome to another given both their agential relations to those outcomes and their personal relations with those who would be better or worse off were those outcomes to obtain. So whereas it can, for instance, be fitting for me to prefer the outcome in which my child lives to the one in which some stranger’s child lives given my personal relations with my child, it would not be fitting—at least, not if other things are equal—for an impartial spectator to have the same preference. And thus it can be

appropriate for me to prefer the outcome in which my child lives even if this outcome is not better than (or preferable to) the one in which the stranger's child lives.

Now, if we can properly value/desire some outcomes more than others without claiming that they are better than, or preferable to, those others, then we should ask, "Why does Scanlon state 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 in the context of TCR, he means for this to include agent-relative value or what he refers to as "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 since Scanlon equates 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 my stating TCR in terms of the agent's reasons for desiring. Indeed, given what he says, we should think that my statement of TCR in terms of the agent's reasons for desiring various outcomes is equivalent to his statement of TCR in terms of the value/desirability of outcomes, for, in his statement of TCR, he just means for his talk of the value/desirability of an outcome to stand for the agent's reasons for wanting that outcome to obtain.

So readers should note that TCR is to be understood in terms of reasons to desire and not necessarily in terms of (impersonal) value or desirability. With 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). So they are not using words such as 'valuable' and 'desirable' in their ordinary, impersonal senses. The reader should, then, assume that when I revert back to talk of desirability so as to engage TCR's critics on their own terms,

I'm using the word 'desirable' as they do, as a kind of short-hand for 'that which the agent has sufficiently weighty reasons to desire'.

§§1.2 *Total outcomes as opposed to intended effects*: Another way that my statement of TCR differs from Scanlon's is that I state 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. It seems that Scanlon meant to state TCR in terms of total outcomes, since, as he points out, the "result" of the action must take into account "the action itself" (1998, p. 84). The problem is that his statement of TCR, as quoted at the beginning of this section, is somewhat unclear on this point in that it refers to "*that* result [emphasis added]," which in this case refers back to the result at which the action was aimed. The problem, then, is that the reasons for performing an action may lie, in part, with the results at which the agent didn't aim. I assume that Scanlon would agree, and, thus, I have stated TCR in terms of total outcomes as opposed to intended effects.

Note, then, that, as I've stated TCR, it is not restricted to only the causal consequences of actions.¹⁵ Indeed, it would be odd for the teleologist (which is what I 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 an act. For instance, one way the world changes when a subject, S, performs an act, 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 perform an act of that type.¹⁶ Since all these ways in which the world might change could potentially make a difference as to whether or not S has a reason to desire that o_1 obtains 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 consequences.

¹⁵ 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 2010.

¹⁶ 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, in performing a_1 , she thereby fulfills her past promise to perform an act of that type. For more on this point, see STURGEON 1996, pp. 511-514, and ANDERSON 1996, p. 541. Of course, the teleologist must deny that the reason that S has to desire the outcome in which her promise is fulfilled is that she has a reason to fulfill her promise. For the teleologist, it's the other way around: she has a reason to fulfill her promise in virtue of the fact that she has a reason to desire that her promise is fulfilled.

§§1.3 *Reasons (to intend) to act*: Following Scanlon, 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 some bodily movement, and the end-product is, then, an intentional action. Nevertheless, the most immediate product 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.¹⁷

Note, though, 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 my 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 (call this 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 (call this 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 .¹⁸ If I'm wrong about this, if the truth of such an attitude-directed judgment does constitute a reason 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 now to drink some toxin tomorrow is certainly no reason to drink that toxin when tomorrow comes around.¹⁹

§§1.4 *The narrow as opposed to the broad construal of desire*: Although some philosophers (e.g., HEUER 2004, p. 48) take a desire to be nothing more than a disposition to act, where one desires that ϕ obtains if and only if one is disposed to act so as to bring it

¹⁷ As I see it, all talk of reasons for action is really just a somewhat sloppy but more idiomatic way of talking about reasons for intending to act. That is, I don't think that there is any distinction between what we ought, or have reason, to do and what we ought, or have reason, to intend to do. The latter is just a precisification of the former. If I'm wrong about this, then TCR should be taken to be a view about reasons for action.

¹⁸ See SCANLON 2007, esp. pp. 90-91

¹⁹ This example comes from KAVKA 1983.

about that o_i obtains (that is, to perform a_i), I will, following Scanlon, 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.²⁰

On this more narrow interpretation, desiring that o_i obtains is sufficient for being motivated (to some extent) to perform a_i , but being motivated to perform a_i is not sufficient for desiring that o_i obtains.²¹ 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 various 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).²²

Unlike Scanlon, though, I don’t think that “desiring something involves having a tendency to see something good or desirable about it” (1998, 32). This suggests that preferring o_i to o_j involves having a tendency to see o_i as better than (or preferable to) o_j . But I don’t think that can be right, as I can, even other things being equal, prefer the outcome in which I am saved to the outcome in which five others are saved without having any tendency to see the former as better than the latter. Of course, I might rightly think that the former is better *for me*, but that’s not the same as thinking that it would be better if I was the one whom was saved. I think, then, it’s more accurate to say that

²⁰ 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.

²¹ 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.

²² 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.

desiring something typically involves having a tendency to see something about it as providing one with a reason to desire it—likewise, for preferring.

§§1.5 *TCR as opposed to just the bi-conditional that it entails*: 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 truth of the bi-conditional stated by TCR-1+2. To illustrate, let ‘PER’ stand for ‘S has more reason to *perform* a_i than to perform a_j ’ and let ‘DES’ 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 ‘PER if and only if DES’—are: (i) ‘PER, because DES’, (ii) ‘DES, because PER’, or (iii) both ‘PER, because BET’ and ‘DES, because BET’—where, for instance, ‘BET’ might stand for ‘ o_i is *better* than o_j ’.²³ In defending TCR, I must not only defend TCR-1+2, but also argue that it is explanation *i* as opposed to either explanation *ii* or explanation *iii* 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.

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, pp. 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.²⁴ 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

²³ 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 Mark Schroeder and Shyam Nair for pointing out the need to consider such common “cause” explanations. And I thank Schroeder for suggesting that someone might take ‘ o_i is *better* than o_j ’ to be the common “cause.”

²⁴ This is true whether TCR is to be formulated in terms of how much reason the agent has to desire the available outcomes or in terms of how valuable/desirable (in the ordinary, agent-neutral sense) the available outcomes are.

view—the view that the fundamental bearers of intrinsic value are concrete entities (e.g., persons, animals, and things)—is called *value concretism*.²⁵

Contrary to what Scanlon and others (e.g., ANDERSON 1993) have claimed, there is no reason why the teleologist cannot accept value concretism.²⁶ 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;²⁷ (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 beings of consideration, virtuous ones of admiration, convenient things of use” —ANDERSON 1993, p. 11). TCR is compatible with all of (a)-(d), as I will now explain.

As rational *beings*, 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 (desirable in the ordinary, agent-neutral sense) by desiring their actualizations—at least, that’s how we appropriately

²⁵ I borrow the terms *concretism* and *abstractism* from TÄNNSJÖ 1999.

²⁶ Elizabeth Anderson is, as I see it, another leading critic of TCR. Although she uses the term ‘consequentialism’ as opposed to ‘teleology’, she defines ‘consequentialism’ so broadly (see 1993, pp. 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 (ANDERSON 1993, pp. 30-31).

²⁷ 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, p. 27).

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 bring about only outcomes. We cannot bring about concrete entities, for 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 a certain concrete entity 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.²⁸

§§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 provided by my friend and our friendship 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

²⁸ 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.

feelings.²⁹ 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.³⁰ 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

²⁹ It's important to note here both that TCR is concerned only with intentional acts and that the mental "acts" of caring about my friend, empathizing with her pain, and taking joy in her successes are not *intentional* acts. I don't come to care for my friend by intending to care for her anymore than I come to believe that the Earth is spherical by intending to believe that it is spherical. Rather, I come to care for someone who is worthy of care in the same way that I come to believe what my evidence supports—that is, by a non-voluntary process of responding to reasons for having these attitudes. See, for instance, SCANLON 1998, pp.20-22. See also my discussion of Anderson on blaming and believing in §§2.3 below.

³⁰ TCR would be incompatible with the existence of agent-relative reasons only if TCR were to be formulated in terms of how valuable/desirable (in the ordinary, agent-neutral sense) the available outcomes are. But, as I've shown, neither Scanlon nor Anderson thinks that this is the way to formulate TCR.

preventing you from betraying two of yours.³¹ And TCR certainly allows for the 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.

What's more, 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).³² 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.³³

³¹ Similarly, the teleologist can even hold that I have reason to prefer your betraying my friends to my betraying my friends. And such agent-relative reasons will, given TCR, generate agent-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.

³² As mentioned in note 14, 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_i if and only if, and because, S has more reason, at t_i , to desire that o_i obtains than to desire that o_j obtains.

³³ Take the following substantive view about what I have most reason to do: I have most reason to refrain from destroying someone else's friendship even in those circumstances in which my doing so would prevent five others from each destroying someone else's friendship. To say that this substantive view is compatible with TCR is not to say that both this substantive view and TCR will be true irrespective of what other claims are true. This is not so. Unless I have most reason to prefer the outcome in which five others each destroy someone else's friendship to the outcome in which I myself destroy someone else's friendship, TCR will not imply that that I have most reason to refrain from destroying someone else's friendship even in those circumstances in which my doing so would prevent five others from each doing the same. TCR is like utilitarianism in this respect. Utilitarianism doesn't provide a substantive account of what we morally ought to do absent some substantive account of what utility consists in (e.g., pleasure, achievement, desire satisfaction, etc.). Likewise, TCR doesn't provide a substantive account of what we have most reason to do absent some substantive account of what we have most reason to desire. The problem, then, is that critics of

§§2.3 TCR is compatible with the view that attitudes such as belief and blame are rationally justified on non-pragmatic 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 pragmatic 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” (ANDERSON 1993, p. 39). This is mistaken for at least three 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 .³⁴ Thus, it is the consequences of *the act*, not of *the intention*, that is relevant on TCR. And so there is no way to generalize 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 the teleologist would be committed to the more general claim

TCR, such as Scanlon, assume that TCR is, in principle, incompatible with certain substantive views about what we have most reason to do, for they say nothing about what we have most reason to desire.

³⁴ 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).

that this is true of *all* attitudes. Why couldn't the teleologist hold that what is true of intending to act is not true of other attitudes?³⁵

Third, it seems that Anderson is wrongly assuming that mental attitudes such as blaming and believing are actions in the relevant sense, as this seems to be the only possible explanation for why she thinks that the teleologist is committed to the view that blaming and believing are rationally justified on the grounds that performing these "actions" would have desirable consequences. But, unlike intentional actions, we don't "perform" the "act" of blaming (or believing) at will—that is, by intending to blame (or by intending to believe). Take belief, 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 beliefs as the result of our intending to form them. Since blaming and believing are not intentional actions, they are not the sorts of things to which TCR applies.

So the teleologist can accept, contrary to what Anderson claims, that whether a person has sufficient reason to believe that *p* depends on only what her evidence is for the truth of *p* and not on the desirability of the consequences of her believing that *p*. And the teleologist can accept that whether blaming someone is rationally justified depends only on whether there is sufficient reason to blame that someone. 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 (LEVY 2005, p. 2). One can blame a person for her wrongdoing without intending to act so as to criticize, punish, or otherwise censure her.

§§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

³⁵ Or if one prefers to talk about reasons for action as opposed to reasons for intending to act, why can't the teleologist hold that although reasons for action are a function of our reasons for desiring its consequences, neither reasons for believing nor reasons for blaming are a function of our reasons for desiring the consequences that stem from our having these attitudes?

domain of reasons is taken to be explanatorily prior to the domain of value (cf. WALLACE 2010)—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).³⁶ Nevertheless, TCR is compatible with the following more general formulation of the 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's 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 (in the ordinary, agent-neutral sense) 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.³⁷

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 that "many reasons bearing on an action concern not the desirability of outcomes but rather the eligibility or

³⁶ 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).

³⁷ The agent-neutral teleologist holds that the best outcome (in the ordinary, agent-neutral sense) 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.

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 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 §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 of *What We Owe to Each Other*, he provides examples in which 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.³⁸ 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,

³⁸ 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.

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 account certain considerations when deciding whether to perform a_1 .³⁹ It seems that when we consider whether or not 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 will now show.

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

³⁹ I am, perhaps, speaking a bit loosely when I treat taking into account (or ignoring) some consideration as itself an intentional act, for, admittedly, whether one takes into account (or ignores) some consideration is not always directly under one's volitional control. In those cases in which it is not directly under one's volitional control, though, one can still intend to do that which 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 and immediately diverting my attention to something else whenever a consideration of that sort creeps into my consciousness. So the expression 'taking into account (or ignoring) certain considerations' should, perhaps, be taken as a convenient shorthand for 'doing that which will cause oneself to take into account (or ignore) certain considerations'.

strategic shot might make my opponent feel crushed or disappointed just is not relevant to whether or not I should make the shot. I don't think that we should accept this claim.

Now, Scanlon never says what consideration is the reason for me to ignore this other consideration (i.e., that my opponent might feel crushed) when deciding whether or not to make this strategic shot, but clearly 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 maybe it's 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 now appealing to the desirability of taking my opponent's feelings into account, which is precisely what the teleologist holds to be relevant. The fact that we are not appealing to the desirability of my making the strategic shot and weighing the desirability of the outcome in which I make the shot against the desirability of the outcome in which I don't 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 example, not as a counterexample, but rather as the first step in an argument against TCR. That is, Scanlon could be arguing as follows:

- P1)** In many cases, we do not, when deliberating about what to do, treat the fact that an act would have certain desirable/undesirable consequences as a reason for/against performing it. For instance, when playing to win, we do not treat the fact that taking a certain shot would embarrass one's opponent (an admittedly undesirable consequence) as a reason against taking the shot.
- P2)** Our deliberative experience in such cases fits better with (*a*) the contention that these are cases in which the agent treats as non-reasons facts that are

actually non-reasons than it does with (*b*) the contention that these are cases in which, on pragmatic grounds, the agent acts so as to cause herself to treat as non-reasons facts that are actually reasons (reasons with actual weight).⁴⁰

- C) Therefore, absent some other reason for accepting *b*, we should accept *a*, which supports the denial of TCR.⁴¹

Unfortunately, though, Scanlon never even considers *b* and so never explains why he thinks that the phenomenology of our actual deliberations supports *a* as opposed to *b*. Perhaps, he just thinks that this obvious. If so, then all I can do is cite my dissent. 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. To borrow another one of Scanlon examples: when I serve on a search committee and ignore how I would personally benefit from our hiring the candidate who specializes in moral philosophy, this feels like a case in which I am choosing to ignore a genuine, self-interested reason that I have to recommend the candidate who specializes in moral philosophy to the department. I choose to ignore this self-interested reason, because I recognize that I have better reason to want to live up to the standards associated with the role of being a good search committee member, and those standards require me to make my recommendation on the basis of what would be best for our program as a whole, and not on the basis of what would be best for me.

So when I consider the phenomenology in these cases, they seem to me to support *b*, and Scanlon gives me no reason to think that I'm mistaken in this, as he fails to even consider *b*, let alone explain why he thinks that the phenomenology supports *a* instead. And, even if I'm wrong about the phenomenology, Scanlon's argument fails to establish that TCR is false, for, as we'll see presently, there are reasons to accept TCR and thus *b* even if the phenomenology supports *a* instead. After all, things are not always as they

⁴⁰ 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.

⁴¹ Something along these lines was suggested to me by Peter de Marneffe.

seem. Thus, even if I'm wrong and the phenomenology fits better with a , we should accept a only if there are no better reasons for accepting b instead.

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 the reasons 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 the following is possible: an agent ought to prefer the outcome in which she has performed a_i (viz., o_i) even though a_i is not the alternative that she ought to perform. Given such a possibility, the agent ought to hope that she performs the alternative (viz., a_i) that she ought not to perform. And if, in spite of hoping that she'll perform a_i , she performs a_j , she ought to wish that she had performed a_i instead. I find this counterintuitive. Surely, agents shouldn't hope that they'll act as they shouldn't. Nor should they wish that they had acted as they shouldn't have acted. To avoid such counterintuitive implications, we should accept TCR-1.

To illustrate how the denial of TCR-1 can lead to such counterintuitive implications, imagine that I'm picnicking on a hill beside a lake and that I must decide whether or not to run as fast as I can down the hill to the lake's only boat dock. Suppose that Richard and his two daughters are complete strangers to me, that my daughter and Richard's two daughters are all drowning, that there's only enough time to rescue either my daughter or Richard's two daughters (as they're at separate ends of the lake), that there's only one boat tied to the boat dock, and that there's no other way to rescue the children except by using this boat. Further suppose that if I run as fast as I can down to the boat dock, I'll get there first and rescue my daughter. If, however, I don't run as fast as I can, Richard will get there first and rescue his two daughters. Assume that everything else is equal.

Let's stipulate that o_1 —the possible world in which I run as fast as I can and rescue my daughter—will obtain if and only if I perform a_1 —the act of running as fast as I can down to the boat dock. And let's stipulate that o_2 —the possible world in which I refrain from running as fast as I can and thereby allow Richard to rescue his two daughters—will obtain if and only if I perform a_2 —the “act” of refraining from running as fast as I can down to the boat dock.⁴² Lastly, let's assume, for the sake of argument, that o_2 is, impersonally speaking, better than o_1 , but that, in this instance, I ought not to do what would bring about the impersonally best outcome (viz. a_2). That is, let's assume that I should perform a_1 , thereby rescuing my daughter.

Interestingly, if TCR-1 is false, the following is possible: I ought, all things considered (that is, have most reason, all things considered), to prefer o_2 to o_1 even though I ought, as stipulated, to perform a_1 as opposed to a_2 . Given such a possibility, I ought to hope that I perform the act that I ought not to perform. After all, it seems that if I ought, all things considered, to prefer the outcome in which I have performed a_2 (viz., o_2), an outcome that will obtain only if I perform a_2 , then I ought to hope that I will perform a_2 . Yet a_2 is the alternative that I ought not to perform. Thus, on this possibility, I ought to hope that I will act as I ought not to act. Again, I find this counterintuitive. It seems much more plausible to accept TCR-1 and thereby deny such possibilities. We can deny such possibilities by accepting TCR-1 and claiming that, given that I ought to perform a_1 as opposed to a_2 , it can't be that I ought to prefer o_2 to o_1 .

Of course, I've stipulated that o_2 is, impersonally speaking, better than o_1 . And, given this, an impartial spectator ought to prefer o_2 to o_1 . But just because an impartial spectator ought to prefer o_2 to o_1 , we shouldn't infer that I ought to prefer o_2 to o_1 . After all, I differ from an impartial spectator in that I'm neither impartial nor a mere spectator. Unlike an impartial spectator, I bear certain agential relations to these two outcomes (i.e., I'm the agent whose actions determine which outcome will obtain), and I bear a special

⁴² For the sake of simplifying the discussion, I'm assuming that there are only two act-tokens available to me: one that falls under the description of my running as fast as I can down to the boat dock and another that falls under the description of my refraining from running as fast as I can down to the boat dock. If it's necessary, the reader can assume that these are the only two alternative act-tokens available to me, because God will immediately cause me to cease to exist should I attempt to perform any other act-token.

relation to my daughter, who will be better off if o_1 obtains. Given these relations and, in particular, the special relationship that I have with my daughter, I ought to prefer o_1 to o_2 despite the fact that o_2 is, impersonally speaking, better than o_1 .

So, although I have good agent-neutral reasons for preferring o_2 to o_1 , it seems that I have even better agent-relative reasons for preferring that my daughter is saved. To deny this is to allow that I can have decisive reason to act so as to save my own daughter even though I have decisive reason to prefer the outcome in which I have instead acted so as to allow Richard to save his two daughters. Such a claim is, I've suggested, counterintuitive. And if we want to avoid such counterintuitive implications, then we should hold that an agent's reasons for performing an act must always track her reasons for preferring the outcome in which she has performed that act to the outcome in which she has refrained 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*: Similar considerations count in favor 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. To deny TCR-2 is to hold that the following is possible: an agent ought to perform a_i even though the outcome in which she has performed a_i (viz., o_i) is not the outcome that she ought to prefer. Given such a possibility, the agent ought to perform a_i even though she ought not to hope that she will perform a_i . And if, as it turns out, she fails to perform the act that she ought to perform (viz., a_i), she ought not to wish that she had. I find this counterintuitive as well. Surely, agents should hope that they will act as they ought to act and should wish that they had acted as they should have acted. To avoid such counterintuitive implications, we should accept TCR-2.

To illustrate how the denial of TCR-2 can lead to such counterintuitive implications, imagine that I must decide whether or not to lie to Henri about Ida's whereabouts. Assume that Henri is looking to kill Ida, that Ida is completely innocent, that Henri will fail to kill Ida only if I lie to Henri, and that I know all of this. Let's stipulate that o_3 —the

possible world in which I lie to Henri and thereby prevent him from killing Ida—will obtain if and only if I perform a_3 —the act of lying to Henri. And let's stipulate that o_4 —the possible world in which I refrain from lying to Henri and thereby allow him to kill Ida—will obtain if and only if I perform a_4 —the “act” of refraining from lying to Henri.⁴³ Assume that everything else is equal. Lastly, assume, for the sake of argument, that I ought to prefer o_3 to o_4 , for the agent-neutral reasons that I have for preferring the outcome that's least bad (in this case, the one in which I have lied to Henri, thereby preventing him from killing Ida) decisively oppose the agent-relative reasons that I have for preferring the outcome in which I have refrained from treating anyone as a mere means (in this case, the one in which I have refrained from lying to Henri as a means to preventing him from killing Ida).

Interestingly, the denier of TCR-2 can say that although I ought to prefer the outcome in which I have lied to Henri, I ought, as Kant claimed, to refrain from lying to Henri. But if we really believe that I ought to prefer o_3 to o_4 , why would we insist that I ought to perform a_4 as opposed to a_3 ? If we insist that such is the case, then we must hold that I ought to perform the alternative that I ought to hope that I do not perform. It seems much more plausible to accept TCR-2 and accept that if I ought to prefer the outcome in which I have lied to Henri, then I ought, contrary to Kant, to lie to Henri.

So to preserve the intuitions both that an agent ought to hope that she will perform an act if and only if it's the act that she ought to perform and that an agent ought to wish that she had acted differently if and only if she didn't act as she should have, we should accept both TCR-1 and TCR-2. That is, we should accept TCR-1+2: S has more reason to desire that o_i obtains than to desire that o_j obtains if and only if S has more reason to perform a_i than to perform a_j .

§§4.3 *In defense of TCR-3*: As I noted in §§1.5, the teleologist is committed not only to TCR-1+2, but also to the right-hand side of that bi-conditional having explanatory priority. Thus, the teleologist must defend TCR-3: if S has more reason to perform a_i than

⁴³ Here too, I'm assuming, merely for sake of simplifying the discussion, that there are only two act-tokens available to me: one that falls under the description of my lying to Henri and another that falls under the description of my refraining from lying to Henri.

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 'PER' stand for 'S has more reason to *perform* a_i than to perform a_j ' and let 'DES' stand for 'S has more reason to *desire* that o_i obtains than to desire that o_j obtains', then TCR-3 is the view that, in every instance, 'PER, because DES'. And the three possible alternatives to this view are: (1) in every instance, 'DES, because PER'; (2) in some but not all instances, 'PER, because DES', and, in every other instance, 'DES, because PER'; and (3) in every instance, both 'PER, because BET' and 'DES, because BET' — where, for instance, 'BET' might stand for ' o_i is *better* than o_j '.

(1) *The First Alternative*: So the first alternative to TCR-3 is to hold that what explains the truth of the bi-conditional 'PER if and only if DES' is that, in every instance, 'DES, because PER', not 'PER, because DES', as TCR-3 supposes. This is to hold that when S has more reason to desire that o_i obtains than to desire that o_j obtains, this is *always* because S has more reason to perform a_i than to perform a_j . But this is clearly false, as one simple example is sufficed to show. Suppose that a_5 is the act of putting my money into a savings account that yields 5% annually and that a_6 is the act of putting my money into a savings account that 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_5 than to perform a_6 is the fact that I have more reason to desire that o_5 obtains than to desire that o_6 obtains. The same would seem to hold for any other purely prudential choice, for surely no one would argue that the reason I have for preferring the outcome in which I am prudentially better off is that this is the outcome in which I did what was prudent (i.e., what I had most prudential reason to do). We should, then, reject the first alternative.

Perhaps, this is too quick. After all, Elizabeth Anderson has argued that 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. For instance, on her view, I have reason to want, and to aim to bring about, the state of affairs in which I spend time alone with wife only because it's rational for me to act in ways that express my love for my wife. On Anderson's view, then, whether I have reason to want, and to adopt as my end, the state of affairs in which I spend time alone with my wife depends on whether it is rational for me to act so as to adequately express my love for her. And, perhaps, Anderson would similarly argue that whether I have reason to want, and to adopt as my end, the state of affairs in which I'm better off depends on whether it is rational for me to act so as to make myself better off.

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, not vice versa. Despite appearances, though, Anderson's view 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 . For instance, it may be that my wanting and pursuing the state of affairs in which 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 in which 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, for if this is right, then what ultimately accounts for the rationality of performing a given action is the rationality of wanting some end to obtain, not vice versa. The issue, then, is whether or not Anderson's injunction to act only in ways that express one's rational attitudes toward persons, animals, and things is extensionally equivalent to the injunction to act so as to promote some end—i.e., some state of affairs.

In fact, Anderson argues that her injunction, which she calls "E," to act only 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 in which one's actions adequately express one's rational attitudes. She argues that E is no more extensionally equivalent to P than the injunction, E' (i.e., E prime), to make only 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 allows me to express my love for her more fully than before (ANDERSON 1996, p. 545). Anderson shows, therefore, 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 in which 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 express 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* directs an agent to violate E if her doing so will minimize her own violations of E. But, as Nicholas Sturgeon (1996, p. 521) has pointed out, it is hard to imagine any but the most fanciful of cases in which an agent's violating E now will prevent her 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 intuitively plausible results.⁴⁴ But rather than debate the issue as to 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 in which one's

⁴⁴ For more on this issue, see PORTMORE 1998 and PORTMORE 2010 (Chapter 4).

current actions do not fail to adequately express one's rational attitudes toward persons.⁴⁵ P** and E yield the exact same prescriptions, and so we see that there is nothing that E can account for that P** cannot account for. And since the proponent of TCR-3 can adopt P** and thereby account for all the intuitive judgments that Anderson appeals to, there is no reason for us to think that Anderson's view, even if correct, is incompatible with TCR-3.

I've argued, then, that in many instances what an agent has reason to do depends on what she has reason to desire, not vice versa. For instance, it seems that I have reason to put my money in a higher-yielding savings account, because I have reason to prefer the state of affairs in which I earn more interest on my savings. And I've shown that, contrary to initial appearances, Anderson's view, even if correct, gives us no reason to deny this.

(2) *The Second Alternative:* This brings us to the second alternative to TCR-3, which is to claim that what explains the truth of the bi-conditional 'PER if and only if DES' is that, in some instances, 'PER, because DES', in other instances, 'DES, because PER', 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. Until such an explanation is forthcoming, we should, I believe, accept the more unified and systematic TCR-3 instead.

(3) *The Third Alternative:* The third and final alternative to TCR-3 is the view that what explains the truth of the bi-conditional 'PER if and only if DES' is that, in every instance, both 'PER, because BET' and 'DES, because BET'. We should, I think, reject this third alternative. Consider what BET might be. That is, consider what might possibly explain both PER and DES? The most obvious candidate is that 'BET' stands for ' o_i is better than o_j '. But, if 'BET' stands for ' o_i is better than o_j ', then we should reject the claim that 'DES, because BET', for it seems that sometimes agents 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 DES is BET.

⁴⁵ 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, PORTMORE 2007 and PORTMORE 2009 and LOUISE 2004.

To illustrate, consider again the case from §§4.1. In this case, I have more reason to desire o_1 (the possible world in which I run as fast as I can and rescue my daughter) than I have to desire o_2 (the possible world in which I refrain from running as fast as I can and thereby allow Richard to rescue his two daughters) even though o_2 is, impersonally speaking, better than o_1 . In this case, it seems that my agent-relative reasons for preferring that my own daughter is saved outweigh my agent-neutral reasons for preferring the better outcome. Since, in this case, I have reason to prefer the worse outcome (viz., o_1), it can't be that what explains why I have more reason to desire that o_1 obtains than to desire that o_2 obtains is that o_1 is better than o_2 , for that's simply false.

So, if the only likely candidate for 'BET' is ' o_i is *better* than o_j ', and if, on this interpretation of 'BET', we should reject 'DES, because BET', as I've argued, then we should reject this third alternative to TCR-3. Of course, these might seem like two big ifs. So the argument here is only presumptive. Given the prima facie plausibility of these antecedents, there is a presumptive case to be made against this third alternative to TCR-3. But even if the case is only a presumptive one, it's important to note that the burden is on the critic of TCR to supply some more plausible interpretation of 'BET'.

§§4.4 *In defense of TCR on the whole:* Admittedly, the arguments for TCR-1, TCR-2, and TCR-3 are not absolutely decisive. For one thing, the arguments for TCR-1 and TCR-2 relied on certain intuitions that perhaps not everyone will share. For another, I haven't canvassed every possible interpretation of 'BET', but have considered only the one that seemed most obvious to me. So let me explain why even those who don't find my arguments for each of TCR's individual three claims absolutely decisive, nevertheless, have good reason to accept TCR on the whole.

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 in at least 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 if and only if, and because, you have more reason

to desire the outcome in which you have experienced more pleasure than to desire the outcome in which you have experienced less pleasure. No one thinks that it's the other way around: that you have more reason to want the outcome in which you have experienced more pleasure because you have more reason to choose the more pleasure-inducing meal 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, this is a misconception: there are no substantive views about values and reasons that TCR cannot accommodate. TCR is an extremely ecumenical and accommodating view. Thus, another reason to accept TCR, apart from both its *prima facie* plausibility and the 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 preserve the intuitions both that an agent ought to hope that she will perform an act if and only if it's the act that she ought to perform and that an agent ought to wish that she had acted differently if and only if she didn't act as she should have. Moreover, I have argued that TCR, in contrast to its alternatives, is superior in its ability to systematize our considered convictions about practical reasons.

Perhaps, some will think that I have been too successful in defending TCR. That is, some might think that if TCR is really as ecumenical as I've claimed, then it rules nothing out and becomes almost trivial. This is not so. Although TCR is compatible with any substantive view about what we have most reason to do as well as with any substantive view about what we have most reason to desire, it does rule out certain positions. For one, it rules out the view that S has most reason to do a_i , but does not have most reason to desire that o_i obtains, for, according to TCR, S has most reason to do a_i if and only if S has most reason to desire that o_i obtains. For another, it rules out the view that S has most reason to desire that o_i obtains in virtue of the fact that S has most reason to do a_i , for, according to TCR, it's the other way around.

The problem, then, with many of the criticisms that have been leveled against TCR is that they've proceeded on the assumption that TCR is incompatible with certain substantive views about what we have reason to do (such as, the view that we have more reason to refrain from betraying one of our own friends than to prevent more numerous others from betraying theirs) irrespective of what the correct substantive view about what we have reason to desire is. This is not so. If, for instance, we hold that whereas *others* have more reason to prefer that *their* friends are not betrayed, *each of us* has more reason to prefer that *our* friends are not betrayed, then TCR will imply that we have more reason to refrain from betraying one of our own friends than to prevent more numerous others from betraying theirs. The lesson, then, is that if one wants to reject TCR on the grounds that it rules out certain substantive views about what we have reason to do, then one must demonstrate that whatever substantive view about what we have reason to desire

that would render TCR compatible with such views about what we have reason to do is false. None of the critics have done so. And, until one of them does, I think that we should accept TCR, for it offers an intuitively compelling and systematic approach to understanding practical reasons.⁴⁶

References

- 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. (1997). "Kantian Rationalism: Inescapability, Authority, and Supremacy." In G. Cullity and B. Gaut (eds.), *Ethics and Practical Reason*, pp. 255-291. New York: Oxford University Press.
- 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.
- DANCY, J. (2004). "Enticing Reasons." In R. J. Wallace, P. Pettit, S. Scheffler, and M. Smith (eds.), *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, pp. 91-118. Oxford: Oxford University Press.

⁴⁶ For helpful discussions on *PEA Soup* <<http://peasoup.typepad.com/peasoup/>>, I thank Michael Almeida, Campbell Brown, Brad Cokelet, Arto Laitinen, Francesco Orsi, and David Sobel. For helpful comments on earlier drafts, I thank Thomas Baldwin, Stephen Darwall, Peter de Marneffe, Josh Glasgow, G. Shyam Nair, Sven Nyholm, Ben Sachs, Mark Schroeder, Jussi Suikkanen, and a number of anonymous referees. Lastly, I thank participants in my spring 2007 seminar entitled "Reasons and Values," especially Andrew Khoury, Errol Lord, and G. Shyam Nair, for valuable comments and discussions.

- DARWALL, S. (2006). "Morality and Practical Reason: A Kantian Approach." In D. Copp (ed.), *The Oxford Handbook of Ethical Theory*, pp. 282-320. Oxford: Oxford University Press.
- FERRY, M. (2009). "Does Morality Demand Our Very Best? On Moral Prescriptions and the Line of Duty." Retrieved August 9, 2009, from http://www.colorado.edu/philosophy/center/RoME_2009_full_papers.shtml.
- GERT, J. (2004). *Brute Rationality: Normativity and Human Action*. Cambridge: Cambridge University Press.
- HARMAN, E. (2009). "'I'll Be Glad I Did It' Reasoning and the Significance of Future Desires." *Philosophical Perspectives* 23: 177-199.
- HURKA, T. (2006). "Value and Friendship: A More Subtle View." *Utilitas* 18: 232-242.
- KAVKA, G. (1983). "The Toxin Puzzle." *Analysis* 43: 33-36.
- LEVY, N. (2005). "The Good, the Bad and the Blameworthy." *Journal of Ethics & Social Philosophy*, www.jesp.org, vol. 1, no. 2.
- LOUISE, J. (2004). "Relativity of Value and the Consequentialist Umbrella." *Philosophical Quarterly* 54: 518-536.
- MCLEOD, O. (2001). "Just Plain 'Ought'." *The Journal of Ethics* 5: 269-291.
- MCMAMARA, P. (1996b). "Must I Do What I Ought (or Will the Least I Can Do Do)?" In M. A. Brown and J. Carmo (eds.), *Deontic Logic, Agency and Normative Systems*, pp. 154-173. Berlin: Springer-Verlag.
- OLSON, J. (2009). "Fitting Attitude Analyses of Value and the Partiality Challenge." *Ethical Theory and Moral Practice* 12: 365-378.
- PARFIT, D. (2011). *On What Matters, Volume 1*. Oxford: Oxford University Press.
- . (2001). "Rationality and Reasons." In D. Egonsson, J. Josefsson, B. Petersson, and T. Rønnow-Rasmussen, eds., *Exploring Practical Philosophy*. Aldershot: Ashgate.
- PILLER, C. (2006). "Content-Related and Attitude-Related Reasons for Preferences." *Royal Institute of Philosophy Supplement* 81: 155-182.

- PORTMORE, D. W. (2011a). *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Forthcoming with Oxford University Press.
- PORTMORE, D. W. (2011b). "Consequentialism and Moral Rationalism." Forthcoming in *Oxford Studies in Normative Ethics*.
- PORTMORE, D. W. (2009). "Consequentializing." *Philosophy Compass* 4: 329-347.
- PORTMORE, D. W. (2007). "Consequentializing Moral Theories." *Pacific Philosophical Quarterly* 88: 39-73.
- PORTMORE, D. W. (1998). "Can Consequentialism Be Reconciled with Our Common-Sense Moral Intuitions?" *Philosophical Studies* 91: 1-19.
- 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.
- 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*, pp. 84-103. Oxford: Oxford University Press.
- . (1998). *What We Owe to Each Other*. Cambridge, Mass.: The Belknap Press.
- SCHEFFLER, S. (1992). *Human Morality*. New York: Oxford University Press.
- SCHROEDER, M. (2010). "Value and the Right Kind of Reason." Forthcoming in the *Oxford Studies in Metaethics*.
- . (2008). "Value Theory." *The Stanford Encyclopedia of Philosophy (Spring 2008 Edition)*, Edward N. Zalta (ed.), URL = <
<http://plato.stanford.edu/archives/spr2008/entries/value-theory/>>.
- SKORUPSKI, J. (1999). *Ethical Explorations*. Oxford: Oxford University Press.
- SMITH, M. (2005). "Meta-Ethics." In F. Jackson and M. Smith (eds.), *The Oxford Handbook of Contemporary Philosophy*, pp. 3-30. Oxford: Oxford University Press.
- 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. (2009). "Consequentialism, Constraints and the Good-Relative-to: A Reply to Mark Schroeder." *Journal of Ethics & Social Philosophy*, www.jesp.org, March 2009.
- TÄNNSJÖ, T. (1999). "A Concrete View of Intrinsic Value." *The Journal of Value Inquiry* 33: 531-536.
- VAN ROOJEN, M. (2010). "Moral Rationalism and Rational Amoralism." *Ethics* 120: 495-525.
- WALLACE, R. J. (2010). "Reasons, Values, and Agent-Relativity." Forthcoming in *Dialectica*.
- . (2006). *Normativity and the Will: Selected Papers on Moral Psychology and Practical Reason*. Oxford: Oxford University Press.
- ZIMMERMAN, M. J. (2010). "Partiality and Intrinsic Value." In B. Feltham and J. Cottingham (eds.), *Partiality and Impartiality: Morality, Special Relationships, and the Wider World*, pp. [not yet available]. Oxford: Oxford University Press.