

*Aims as Reasons**

In recent discussions, one finds the claim that, despite first appearances, attitudes do not give us reasons for action.¹ More precisely, the “Thesis”—to give it a name—is that the fact that a person currently has or lacks an attitude does not give her normative reasons for action, where

* I am very grateful to private comments from Tim Scanlon, Sam Scheffler, Jay Wallace, and Wai-Hung Wong, as well as to public comments from participants in the Seminar on Ethics and Normative Theory (SENT) at Stanford; at the University of Valencia, where Eduardo Ortiz delivered prepared comments; and at the University of Stirling, where Kent Hurtig delivered prepared comments.

¹ For instances of this claim, focusing variously, and often selectively, on belief, intention, or desire, see Michael Bratman, *Intention, Plans, and Practical Reason* (Cambridge: Harvard, 1987) and “Intention, Practical Rationality, and Self-Governance,” unpublished; John Broome, “Normative Requirements,” *Ratio*, 12 (1999): 398–419; “Are Intentions Reasons? And How Should We Cope with Incommensurable Values?” in Christopher Morris and Arthur Ripstein, ed., *Practical Rationality and Preference: Essays for David Gauthier* (Cambridge: Cambridge University Press, 2001), pp. 98–120; and “Practical Reasoning,” in José Bermúdez and Alan Millar (eds.) *Reason and Nature: Essays in the Theory of Rationality* (Oxford: Oxford University Press, 2002) pp. 85–111; John Brunero, “Are Intentions Reasons?” *Pacific Philosophical Quarterly* 88 (2007): 424–444; Garrett Cullity, “Decisions, Reasons, and Rationality,” *Ethics* 119 (2008): ; 57–95; Jonathan Dancy, *Practical Reality* (Oxford: Oxford University Press, 2000); Richard Moran, *Authority and Estrangement* (Princeton: Princeton University Press, 2001), Ch. 3; Derek Parfit, “Reasons and Motivation,” *Proceedings of the Aristotelian Society*, suppl. vol. 71 (1997): 99–130; and *On What Matters* (unpublished); Joseph Raz, *The Morality of Freedom* (Oxford: Oxford University Press, 1986), Ch. 6; *Engaging Reason* (Oxford: Oxford University Press, 1999), Ch. 3; “The Myth of Instrumental Rationality,” *Journal of Ethics and Social Philosophy* 1:1 (2005); and “Instrumental Rationality: A Reprise,” *Journal of Ethics and Social Philosophy, Symposium 1* (2005): 1–19; T.M. Scanlon, *What We Owe to Each Other* (Cambridge: Harvard University Press, 1997), esp. Ch. 1, Sect. 9; “Reasons: A Puzzling Duality?” in R. Jay Wallace, Philip Pettit, Samuel Scheffler, and Michael Smith, ed., *Reasons and Value: Themes from the Moral Philosophy of Joseph Raz* (Oxford: Clarendon Press, 2004), pp. 231–246; and *Moral Dimensions: Permissibility, Meaning, and Blame* (Cambridge: Harvard University Press, 2008); Dennis Stampe, “The Authority of Desire,” *The Philosophical Review* 96 (1987): 335–381; Barry Stroud, *Metaphysics, Modality, and Value* (unpublished), Ch. 5; Judith Jarvis Thomson, *Goodness and Advice*, (Princeton: Princeton University Press, 2001), pp. 24–29; and R. Jay Wallace, “Normativity, Commitment, and Instrumental Reason,” *Philosophers’ Imprint* 1:3 (2001). I should note that Bratman allows that intentions can provide “framework” reasons. This may be compatible with the Thesis, however, if framework reasons are identified with requirements of structural rationality of the kind discussed in section 2. For discussion of a similar thesis about belief, see Hannah Ginsborg, “Reasons for Belief,” *Philosophy and Phenomenological Research* 72 (2006): 286–318.

by “attitudes” we mean what Scanlon would call “judgment-sensitive” attitudes.² These are attitudes that are responsive to their bearers’ judgments about the reasons for them: including beliefs, desires, and intentions, but excluding sensations, such as those of pain and pleasure.

Since the motivation and contours of the Thesis are somewhat obscure, I begin by exploring these. Next I consider a challenge to the Thesis, which Scanlon raises: that in many cases, particularly cases of underdetermination by reason, *having an aim* seems to affect one’s reasons for action. In the bulk of the paper, I suggest how the challenge might be met, by building on suggestions that Scanlon offers.

1. Clarifying the thesis

Let us, provisionally, take the *Thesis* to say:

No basic normative principles take the:

attitude-based form: *S* has at *t* reason to do *X*, where *X* is a function of just *S*’s attitudes at *t* (or disposition at *t* to have attitudes) and non-normative facts.

A *normative principle* is a principle of the form

One has reason to *phi* in *C1*.

Such a principle is *basic* iff it is not explained some other normative principle

One has reason to *psi* in *C2*

and the fact that

One’s *phi*-ing in *C1* makes it more likely that one *psi*’s in *C2*

and:

General Transmission (GT): If there is *pro tanto* reason for one to *E*, and if it is likely, conditional on one’s *M*-ing, that one *E*’s and that one’s *M*-ing (nonsuperfluously) helps to bring this about, then there is, because of this, *pro tanto* reason for one to *M*. There is

² See *What We Owe to Each Other*, pp. 18–22.

more reason the more likely it is—the more “effective” a means one’s M-ing is to one’s E-ing.³

This understanding of the Thesis would explain why its friends are willing to allow that, in countless cases, one’s having an attitude at a given time can affect one’s reasons at that time for action. For example:

- (1) The fact that I *believe* that everyone is out to get me usually means that my anti-psychotic medication is wearing off, and so makes it the case that I have reason to take another dose.
- (2) The fact that I *desire* a slice of pie means that I will enjoy a slice of pie, and so makes it the case that I have reason to have one.
- (3) The fact that I *desire* to be a doctor means that I am more likely to succeed in attempts to become a doctor, and so makes it the case that I have more reason to try.
- (4) The fact that I do not *want* to spend time with my aunt means that if I were tell her that I want to, I would be deceiving her, and so makes it the case that I have reason not to tell her that I want to.⁴
- (5) My *well-being* consists in the satisfaction of my desires. So the fact that I desire to *X* means that *X*-ing would promote my well-being, and so makes it the case that I have reason to *X*.⁵

³ I discuss a more precise, but more complicated form of *GT* (as well as its relation to the Facilitative Principle of Raz, “The Myth of Instrumental Rationality,” p. 6, to which it is indebted) in my “Instrumental Transmission,” unpublished. The reason for “by helping to bring it about” is to rule out possible cases in which one’s M-ing makes it more likely that one E’s merely by providing evidence that one E’s. The reason for “nonsuperfluously” is to rule out possible cases in which one would have E-ed even if one had not M-ed. In this paper, I will often drop “(nonsuperfluously)” and “that one E’s and…” for less cumbersome formulations. They should nonetheless be read as shorthand for the formulation here (which is, in turn, shorthand for the formulation in “Instrumental Transmission”).

Judgments about what there is reason to do and judgments about what is likely are relative to a body of information. The body of information need not be the agent’s. After all, when we have more information than the agent, our advice about what he has reason to do is typically in light of our information, not his. In general, I believe, the relevant body of information is selected by the context of the person *assessing* the reasons- or likely-claim for truth or falsity. See my and John MacFarlane’s “Ought: Between Objective and Subjective,” unpublished.

⁴ See, for example, Dancy, *Practical Reality*, pp. 40–41, 55–56; Ginsborg, “Reasons for Belief”; Parfit, *On What Matters*, Section XXXX; Raz, *Engaging Reason*, pp. 59–60; Scanlon, *What We Owe to Each Other*, Ch. 1, Sect. 9, and *Moral Dimensions*, Ch. 2; Thomson, *Goodness and Advice*.

⁵ Parfit, *On What Matters*, Section XXXX.

(6) There is a *virtue of “resoluteness”* which consists in doing what one intends. So the fact that I intend to *X* means that by *X*-ing, I would display this virtue, and so makes it the case that I have reason to *X*.⁶

(7) *Friendship* is partly constituted by its participants’ intentions, including their present intentions. The fact that I have certain intentions toward Kevin makes him my friend, and so makes it the case that I have reason to help him.⁷

In each of these cases, the effect of my attitude on my reason is explained by the combination of

(i) the *GT*, (ii) a normative principle (not necessarily basic) that is not of the attitude-based form,

such as:

- (1) I have reason to avoid psychotic episodes.
- (2) I have reason to enjoy my food.
- (3) I have reason to become a doctor.
- (4) I have reason not to deceive my aunt.
- (5) I have reason to promote my well-being.
- (6) I have reason to display the virtue of resoluteness.
- (7) I have reason to help my friends.

and (iii) the fact that that attitude makes it the case that I am more likely to (e.g.) avoid psychosis by (e.g.) taking medicine. Such explanations need not appeal to basic normative principles of the attitude-based form.

With some sense of the Thesis’s content, let us ask about its point. Why would anyone deny that any basic normative principles take the attitude-based form? A common objection to *attitude-based* theories,⁸ which claim that *all* basic normative principles are of this form, is that

⁶ Broome, “Are Intentions Reasons?” pp. XXXX.

⁷ Scanlon, *What We Owe to Each Other*, Ch. 1, Sect. 9, and *Moral Dimensions*, Ch. 4.

⁸ I am not understanding attitude-based theories as *metaphysical* accounts of relation of being reason for, which claim either (i) that the relation that we naively call “being reason for” is reducible to some nonnormative relation involving attitudes, or (ii) that what we naively call “beliefs that something is reason for something else” are reducible to attitudes, other than belief, toward nonnormative contents. Instead, I am taking attitude-based theories to be substantive, although broad, theories of *what* reasons we have. This is not to deny, of course, that some of the appeal of attitude-based theories, understood as substantive normative theories, may derive from their being confused with such metaphysical theories. See Stroud, *Metaphysics, Modality, and Value*, ch. 5; and Parfit, *On What Matters*, Sec. XXXX. Nor is it to deny that some deliberately present attitude-based theories as metaphysical theories. See Mark Schroeder,

they are *extensionally incorrect*. Consider the *Desire-Based Theory*, which claims that there is just one basic normative principle:

S has reason at *t* to do just what would serve *S*'s intrinsic desires at *t*.

The Desire-Based Theory implies, on the one hand, that given the right set of desires, one would have reason to do something that, intuitively, one would lack reason to do.⁹ If one had a desire served by committing an atrocity, and no desires served by refraining from it, then one would have conclusive reason to commit that atrocity. And the Desire-Based Theory implies, on the other hand, that given the right set of attitudes, one would lack reason to do something that, intuitively, one would have reason to do. If one were to have no desire served by avoiding suffering, then one would lack sufficient reason to avoid suffering.

However, the fact that some attitude-based theories are extensionally incorrect does not seem a compelling ground for denying that there are *any* basic normative principles with the attitude-based form. For one thing, other attitude-based theories are less clearly extensionally incorrect. There are Kantian views, for example, according to which all reasons for action derive

Slaves of the Passions (Oxford: Oxford University Press, 2007), esp. Ch. 4, 11. Nor am I addressing views, such as that of Michael Smith, *The Moral Problem* (Oxford: Oxford University Press, 1994), that understand the relation of being a reason for not in terms of *nonnormative* relations involving attitudes, but instead in terms of *rational* relations among attitudes.

⁹ Broome, "Normative Requirements," and, following Broome, Brunero, "Are Intentions Reasons?" stress that, on some attitude-based theories, attitudes would be reasons for *themselves*. For example, if any belief is a reason to believe its logical consequences, then, since *p* is a logical consequence of *p*, any belief is a reason for itself. But Broome and Brunero never explain what the sting of this observation is supposed to be. Is it simply a way of dramatizing the problem of extensional incorrectness, the point being that the content of the belief, *p*, could be *anything*? Or is it supposed to have some independent force? Is the worry perhaps that if an attitude were a reason for itself, then it would be a reason that one could not go against, making deliberating about it, or offering it as advice, superfluous? But how troubling would this be? After all, some have explicitly argued that attitudes are self-justifying. For example, Gilbert Harman, *Change in View* (Cambridge: MIT Press, 1986) argues, not implausibly, that any belief has some provisional justification simply in virtue of being held.

from exercises of the will that respect the will's structural, self-imposed standards. These standards require that some things be willed (at least if anything is willed) and that other things not be willed.¹⁰ For another, even if all attitude-based theories are extensionally incorrect, it is less clear that a “hybrid” must be. A hybrid would say that *some, but not all*, basic normative principles are of the attitude-based form. Take, for example, the:

Intention-Based Hybrid: There is a basic normative principle that *S* has reason at *t* to do what *S* intends at *t* to do. But there are other basic normative principles, which are not of the attitude-based form.

This hybrid does not imply that if one did not intend to avoid suffering, then one would lack sufficient reason to avoid suffering. It does imply that if one intended an atrocity, then one would have *pro tanto* reason to commit it. However, it need not imply that one would have *conclusive* reason to commit it, since the intention-based reason might be overridden by other reasons, such as those provided by human suffering. To my mind, the mere fact that one intends an atrocity does not make it the case that one has even overridden *pro tanto* reason to commit it. So, by my lights, this hybrid is still extensionally incorrect. But I also find it hard to muster independent evidence for this claim,¹¹ or to feel confident that, on its own, it has much importance, either for practice or theory. The Intention-Based Hybrid would also seem to imply that, when two options are otherwise tied—when there is merely sufficient reason for option *X*, merely sufficient reason for option *Y*, conclusive reason for (*X* or *Y*), and not sufficient reason for (*X* and *Y*)—forming an intention for one of these options can tip the scales, making it the case

¹⁰ See, for example, Christine Korsgaard, *The Sources of Normativity* (Cambridge: Cambridge University Press, 1996).

¹¹ Moreover, Mark Schroeder has suggested a pragmatic explanation of our reluctance to accept that the intention makes it the case that one has *pro tanto* reason. We expect claims about reasons to be made only when the reason is strong, and this reason is weak. See his “Instrumental Mythology,” *Journal of Ethics and Social Philosophy, Symposium 1* (2005): 1–12; “The Scope of Instrumental Reason,” *Philosophical Perspectives* 18 (2004): 337–64; and *Slaves of the Passions*, Ch. 5.

that one has conclusive reason to pursue it. But this implication actually seems intuitive, as we will see.

Scanlon, one of few friends of the Thesis to explicitly consider the possibility of a hybrid,¹² rejects it not on extensional grounds, but instead because it would involve a “troubling” or “puzzling” “duality in the *sources* of reasons.”¹³ It is elusive, however, what is supposed to be puzzling. Scanlon already accepts a multiplicity of sources of reasons. The value of human life is one source, the grandeur of the Grand Canyon another. Why not yet another source, consisting of desire or rational willing? If we already accept a multiplicity of basic normative principles, with a multiplicity of forms, why would one more basic normative principle, of the attitude-based form, introduce a puzzling duality?

What puzzles Scanlon seems not so much a duality of *sources* as a duality in modes of *sourcing*: a difference between the *way* in which pleasure, friendship, etc. are “sources” of reasons, according to friends of the Thesis, and the *way* in which desire or the will is a “source” of reasons, according to attitude-based theorists. At several points, Scanlon suggests that he seeks to avoid the idea that attitudes (specifically intentions) “create” or “generate” reasons, or “confer” on certain considerations the status of a reason (231–32, 239–40). But these terms leave the contrast rather elusive. Scanlon himself agrees that intentions “make it the case” that one has reasons, at least in the sense that in coming to have an intention, one comes to have an attitude that one would not otherwise have had (240). It is not clear what the difference is supposed to be between *creating* a reason and “merely” *making it the case that there is* a reason

¹² Parfit, *On What Matters*, Sec. XXXX, considers a hybrid desire-based theory, but of a different form. According to this theory, desires give one additional reason for action, but *only when* one has reason to have those desires. Although Parfit finds such a theory implausible, he is not particularly troubled by it, finding it “fundamentally value-based.” This hybrid theory is structurally similar to the “further attitude-constituted value” that I describe in Section 9 below.

¹³ Scanlon, “A Puzzling Duality?” 246, see also 233, 237, 239.

when there was not before. One might try saying that when an attitude merely “makes it the case” that there is a reason, it does so only in virtue of fulfilling the condition of a normative principle that is “just” there, to be “discovered,” independent of one’s attitudes (231). But attitude-based theories will say the same thing. When a desire “creates” a reason, the desire-based theorist will say, it does so only by fulfilling the principle:

If *S* has a desire at *t*, then *S* has reason at *t* to do what will serve that desire.

This principle is already there to be discovered, independent of one’s attitudes.¹⁴

Perhaps the difference at issue is this. For friends of the Thesis, sources provide reasons in the way in which things that we think of as *values* or *things of value* provide reasons. They are sources in the sense that they are the things that we most fundamentally have reason to *promote*, or *to act for the sake of*. Accordingly, the sources of reasons that friends of the Thesis recognize—pleasure, friendship, etc.—are naturally described as values or things of value.¹⁵ I

¹⁴ One might suggest: “The attitude-based form is strictly:

If *S* has an attitude at *t*, then *S* has reason at *t* to do certain things.

My attitude ‘creates’ the reason in the sense that I did not have the reason prior to coming to have the attitude. Contrast this with a case in which, having had no friends, I make one, by coming to have the relevant attitude. Here the relevant principle is of the form:

S has reason to help *S*’s friends, whoever they may be.

In this case, my attitude does not ‘create’ the reason. Even prior to having the attitude, I *already* had reason to help every member of the set of my friends. It’s just that the set was empty.” (As Raz at one point writes, it is not problematic that adopting goals should change our reasons by “activat[ing] conditional reasons that we have anyway” (“The Myth of Instrumental Rationality,” p. 23).) One problem with this suggestion is that it seems at least as natural to interpret the relevant principle as:

If *S* has a friend at *t*, then *S* has reason at *t* to help him,

and to say that, prior to making a friend, I did not already have reason to help my friends.

Instead, it was true of me that, if I made a friend, I would then have reason to help him.

¹⁵ Raz, “Myth of Instrumental Rationality,” pp. 3–4, and elsewhere, is explicit about the grounding of reasons in value.

This is not necessarily to say that the concept of value is somehow explanatorily prior to the concept of providing reasons in this way. It is just to say that there is a way of providing reasons that is characteristically associated with what we are prepared to call values or things of value. So it is compatible with (although not committed to) Scanlon’s “buck-passing” proposal

won't try to give a positive description of the different, "attitude-based" way in which the attitude-based theorist thinks of desires or the will as "sources" of reasons. For our purposes, what matters is the negative point that the attitude-based theorist does *not* think of desires or the will as "sources" in the sense of:

Value-Provision: A type *T* (or, alternatively, instances of a type *T*) is a *source of*, or *provides*, reasons for action only if there is a basic normative principle that is of the form either:

(i) if one stands in a relationship of a certain kind to an instance of *T*, then one has reason to *bring it about*

or:

(ii) if one stands in a relationship of a certain kind to an instance of *T*, then one has reason to *respect* or *engage with*¹⁶ or *honor* or act in a certain way *for the sake of* it.

No desire-based theorist, for example, will say that one is supposed to *bring about* one's present desires (contrast the *objects* of these desires), or that one's present desires are things that one somehow *respects* by satisfying. Accordingly, no desire-based theorist will say that one's desires are *themselves* things of *value*. The desire-based theorist may say that *other* things are of value: in particular, that the *objects* of one's present desires are valuable ("for one"?) because one desires them. The point is that what the desire-based theorist takes as the *source* of reasons is not itself of value—or, to be less committal, is not a source of reasons in the way that values characteristically are.

This may seem, at first glance, to set Kantians apart from attitude-based theorists, since what Kantians see as the source of reasons—the rational will—is something that they also see as *being of value* and *calling for respect*. In the relevant regard, however, Kantianism is like the

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" (*What We Owe to Each Other*, p. 96).

¹⁶ These activities of "respecting" and "engaging with" things of value are elucidated by Joseph Raz, *Value, Respect, and Attachment* (Cambridge: Cambridge University Press, 2001). For our purposes, the intuitive connotations of the terms are enough.

Desire-Based Theory: the rational will is not a source of reasons in the way in which values are sources of reasons. Kantians do not advance a *basic* normative principle that says that one has reason to do what respects the rational will, because the rational will is something of value that calls for respect.¹⁷ Instead, their basic principle is something like:

If the rational will takes something as its object, then, because of this, it is valuable and one has reason to treat it in certain ways.

They then claim that the rational will cannot but take itself as its own object. From this they derive, as a kind of theorem, that the rational will is valuable and something that we have reason to respect.¹⁸

If this is right, then it suggests that Thesis is really:

*Thesis: All sources of reasons for action are sources in the way in which values characteristically are. That is, all reasons for action derive from such sources, or are value-based.*¹⁹

Accordingly, attitude-based theories should be understood as claiming:

Attitude-Based Theory: No sources of reasons are sources in the way in which values characteristically are. All are sources in the alternative, attitude-based way.

and hybrids should be understood as claiming:

¹⁷ Contrast the way in which Scanlon, *What We Owe to Each Other*, Ch. 3, says that we have reason to respect the value of human life, or Raz, *Value, Respect, and Attachment*, Ch. 4, says that we have reason to respect people.

¹⁸ Christine Korsgaard, “Kant’s Formula of Humanity,” in *Creating the Kingdom of Ends* (Cambridge: Cambridge University Press, 1996), pp. 106–32 at times suggests a shorter argument:

(i) When the rational will takes something as its object, then we have reason to treat it in a certain way. So it is a source of reasons.

(ii) Since the rational will is a source of reasons, it is itself something that we have reason to respect.

This argument just equivocates on the two senses of “source of reasons” that we have been distinguishing. In order to get the conclusion that the rational will is itself something to be respected, the Kantian needs the claim that the rational will must take itself as one of its objects.

¹⁹ This might explain why it is so natural to label the opposition to the Desire-Based Theory the “value-based” view. See Dancy, *Practical Reality*, pp. 29–31, who asks why the opposition to the desire-based view must involve any reference to value.

Hybrid: Some sources of reasons are sources in the way in which values characteristically are, whereas other sources of reasons are sources in the alternative, attitude-based way.

This makes it clearer why a hybrid would involve a puzzling duality. The Thesis represents practical deliberation as seeking, ultimately, to trace reasons for action back to the kind of source that values or things of value represent. Perhaps we can make sense of the suggestion of attitude-based theories that practical deliberation takes an altogether different form: that it is a kind of existential “self-legislation,” or calculation of means to ends set by groundless desire. What is harder to make sense of is the idea that practical deliberation should take both forms: that we should at once seek to ground our reasons in value and accept as reasons what lacks any such grounding. What puzzles me, and may also puzzle Scanlon, is the difficulty of conceiving of a single, unified deliberative viewpoint that could integrate these two very different stances toward our reasons.

This understanding of the Thesis is compatible with the cases that we considered earlier, in which it is granted that one’s attitudes do affect one’s reasons. In (5)–(7), my present attitude helps to *constitute* the *value* that is the source of the relevant reasons: my well-being, my resoluteness, my friendship. We might put this by saying that my attitude is part of an *attitude-constituted value*. In (1)–(4), the relevant value is constituted by something else: my mental health, my pleasure, my aunt’s interest in knowing the truth. My attitude simply affects whether a given action facilitates my doing what the relevant value provides me with reason to do. It plays a *purely facilitative* role. The crucial point is that in all of these cases, my attitudes affect my reasons only by affecting my value-based reasons.

It is also worth noting that, in many of these cases, my present attitudes affect my reasons in an *attitude-independent* way: that is, in a way in which factors that are not my present attitudes can affect my reasons. For example, a *blood test* can tell me just as soon as my paranoid belief

that I need my meds, I can deceive my aunt about an *investment* just as soon as about my desires, someone can cease to be my friend because of *his* feelings just as soon as because of mine, and my *future* desires determine what constitutes my well-being just as soon as do my present desires. However, in theory—that is, in abstraction from further claims about what values there are and what reasons they provide—the Thesis does not require that my present attitudes affect my reasons (if at all) only in an attitude-independent way. Indeed, in theory, the Thesis is even compatible with some or all basic normative principles being of the attitude-based form.²⁰ For example, if one’s present desires were things of value to be respected—like human life, or the natural environment, or cultural achievements—and if respecting them consisted in doing what served them, then there might be a value-based basic normative principle that said that one had reason to do what served one’s present desires. However, we can set aside this abstract possibility, since it is uncontroversial that there are no such values. That is something that even attitude-based theorists accept.

2. Explaining the appearance that our attitudes affect our reasons

To explain why one might wish to defend the Thesis, of course, is not yet to defend it. In what follows, I want to focus on a particular challenge to the Thesis: that our attitudes *seem* to affect our *reasons*, even though they *seem not* to affect our *value-based* reasons. To defend the Thesis, we need to explain away either the first appearance—that our attitudes affect our reasons—or the second—that our attitudes do not affect our value-based reasons.

2.1 The explanatory and the normative

²⁰ Note that even if we understand the Thesis in the earlier way, as the claim that no basic normative principles are of the attitude-based form, it still may be compatible with a theory extensionally equivalent to a hybrid or attitude-based theory. In theory, a normative principle of the attitude-based form might be derived from a normative principle not of that form.

In some cases, the appearance that our attitudes affect our reasons arises from confusing explanatory and normative reasons. If we ask *why* someone did or believed something, it may be appropriate to answer that it was because he had some attitude: because he desired, intended, or believed something else. In this case, we ask for, and receive, an *explanatory* reason for his belief or action. But this is a different question from whether he had *normative* reason for, whether there was anything to be said in favor of, believing or doing what he did.²¹

2.2 *Patterns of Reasons*

In other cases, the appearance that our attitudes affect reasons arises from a different sort of confusion. When we have an attitude, we often have reason for it. And when we have reason for an attitude, we typically have reason for other attitudes or responses. Reasons, we might say, come in patterns. One such pattern is reflected in the *GT*. The *GT* implies that if I have an aim, and I have reason to have it, then I have reason to take means to it. So it may appear that my having the aim gives me reason to take the means. But what in fact gives me reason to take the means is my *having reason* to have that aim. I would have the same reason to take the means even if I did not have that aim.²²

2.3 *The evaluative and the deliberative*

In other cases, the appearance that our attitudes affect our reasons results from overlooking a distinction that Scanlon has done much to clarify: the distinction between an attitude's being

²¹ See Dancy, *Practical Reality*; Ginsborg, "Reasons for Belief"; Parfit, *On What Matters*, Sect. XXXX; Stroud, *Metaphysics, Modality, and Value*, Ch. 5; Thomson, *Goodness and Advice*.

²² Raz, *Morality of Freedom*, Ch. 6, (and, following Raz, Dancy, *Practical Reality*, Ch. 2; and Parfit, *On What Matters*, Sect. XXXX) makes a similar point about desires. Usually, when one desires something, one has reason to desire it. And usually when one has reason to desire something, one also has reason to act to achieve it. So it may seem that desiring something makes it the case that one has reason to achieve it, when in fact it is having reason to desire it that does so. For further discussion of patterns of reasons, see Niko Kolodny, "How Does Coherence Matter?" *Proceedings of the Aristotelian Society* 107 (2007): 229–63.

deliberatively significant and its being *critically* or *evaluatively* significant.²³ Reasons, at least in the sense with which we are concerned, are *deliberatively* significant. They have their home in the first-person standpoint of a deliberator trying to decide what to believe or do and, by extension, in the second-person standpoint of an advisor trying to help the deliberator reach that decision. Evaluations, by contrast, characteristically arise instead from the third-person standpoint of a spectator, who aims not to guide himself, or an advisee, but instead to appraise, after the fact, what someone else came to believe or do. One can be misled to think that one's having an attitude is a reason for one to do something—a consideration that carries weight within the first-person standpoint of deliberation—when in fact it only makes it the case that one's doing it would be subject to a certain kind of evaluation, from a third-person standpoint of assessment.²⁴

Although this evaluation can take other forms, the form most relevant for our purposes is what Scanlon calls “structural irrationality.”²⁵ One's attitudes can make it irrational for one to,

²³ See Scanlon, *Moral Dimensions*, Ch. 1.

²⁴ Speaking of “the first-person standpoint of deliberation” and “the third-person standpoint of evaluation” helps, I think, to indicate the relevant distinction. But it may be a ladder that one kicks away after ascending. The more one thinks about the distinction, the less confident one is of any neat correspondence with these standpoints. After all, we can make claims about the reasons that third parties have, purely hypothetically or after the fact. And we can consider before the fact, indeed in deliberation, how we ourselves would properly be evaluated for a particular performance. Compare Scanlon, *Moral Dimensions*, Ch 1. I am indebted to Sam Scheffler for pressing this point.

²⁵ See “Structural Irrationality,” in Geoffrey Brennan, Robert Goodin, Frank Jackson, and Michael Smith (eds.), *Common Minds: Themes from the Philosophy of Philip Pettit* (Oxford: Oxford University Press, 2007), pp. 84–103. Other modes of evaluation are especially relevant to the narrower thesis, which Scanlon defends in *Moral Dimensions*, that an agent's present intentions, or reasons for which he acts, do not affect the *moral permissibility* of his so acting. Opponents of this narrower thesis, who insist that intentions do affect permissibility, need not embrace attitude-based, or hybrid, theories of the kind that the Thesis rejects. Opponents of the narrower thesis may well accept that acting with a good (or bad) intention realizes a kind of value. Scanlon's reply, I take it, is that while that acting with a good intention can be said to realize a value, this value does not have *deliberative* significance. It does not provide one with

or to refuse to, form or revise other attitudes. Of particular interest to us is the putative rational requirement:

Means-End (ME): It is irrational of me (to intend at t to E , to believe at t that intending at t to M is a necessary means to my E -ing, but not to intend at t to M).²⁶

As Scanlon observes, we have no reason, in general, to avoid irrationality. Nevertheless, when we are irrational, we can be criticized for malfunctioning, or manifesting a kind of vice.²⁷

3. *The tie-breaking effect of intention*

In other cases, however, it is harder to deny that our attitudes affect our reasons. I will consider one such case, or cluster of cases, to which Scanlon draws our attention.²⁸ These are cases of underdetermination, in which, to start, I have merely sufficient reason to X , merely sufficient reason to Y , conclusive reason (to X or to Y), and lack sufficient reason (to X and to Y). Then I form an intention to X , and as a result, I come to have conclusive reason to take means to X and to lack sufficient reason to take means to Y . An example: I have merely sufficient reason to attend a conference in New York, merely sufficient reason to attend a conference in Boston, conclusive reason to attend one of these, and lack sufficient reason to attend both. As things stand, I haven't made up my mind. The New York organizers call me to ask whether I would

reason to realize it, or with any other reasons to act as one does. It has, instead, *critical* significance. It provides others with reasons to *respond* in certain ways to what one does. Less personally, it may give others reason, say, to admire one for displaying a virtue. More personally, it may give them reason to reaffirm, or alter, the attitudes and interactions that constitute their "relations" with one. This kind of value (or disvalue) constitutes what Scanlon calls "meaning." Even though malice, recklessness, and culpable ignorance are all impermissible ways of harming another, they give different meanings to what is otherwise the same action. In particular, they make the action blameworthy in different ways.

²⁶ For further discussion of *ME*, see Bratman, *Intentions, Plans, and Practical Reason*; Broome, "Practical Reasoning"; and Wallace, "Normativity, Commitment and Instrumental Reason."

²⁷ See also Niko Kolodny, "Why Be Rational?" *Mind* 114 (2005): 509–63; and Raz, "Myth of Instrumental Rationality." In "Why Be Rational?" I try to explain why considerations about one's irrationality can *seem* to play a role in deliberation and advice.

²⁸ See also Brunero, "Are Intentions Reasons?" and Cullity, "Decisions, Reasons and Rationality," p.p. 64–66.

like to pre-register, which, as they explain, is cheaper than registering on site. It would be reasonable for me to respond: “Sorry, I don’t know whether I’m going, so I don’t have reason to register yet.” Then I make up my mind to go to New York, and the organizers call back. Now, it seems, I do have sufficient reason to pre-register for New York, whereas I still lack sufficient reason to pre-register for Boston. Given that I have decided to attend the New York conference, the fact that pre-registering for New York would serve, in an efficient way, the aim of attending the New York conference is a stronger reason than it was before. What is the alternative? Would it be correct for me to think that I *continue to have just as much* reason to take steps to attend the conference that I have decided *against* attending? Should I just as soon pre-register for Boston, or book a flight there?

The resources of the last section seem not to explain this phenomenon:

Intention Effect, first pass (IE1): If I have merely sufficient reason to *E1*, merely sufficient reason to *E2*, conclusive reason (to *E1* or to *E2*), and lack sufficient reason (to *E1* and *E2*), then if I intend to *E1*, I have, because of this intention, stronger reason than I otherwise would to take means to *E1*-ing.

The *GT* does not explain *IE1* unless intending the end somehow gives one stronger reason for the end itself. But that would itself require explanation. Nor it is plausible that we accept *IE1* only because we confuse it with some principle of rationality. To begin with, this principle could not be *ME*. *ME* applies only to what I believe are *necessary* means, whereas in this case the means are not necessary. Pre-registration is not required; it is just cheaper than the alternative. Next, *ME* requires me only to *intend* the means. However, what seems out of place in the case described is not simply to refuse to intend the means, but also to refuse to *believe* that I have greater *reason* to *take* the means. Suppose that I have decided to attend the New York conference. Whether or not I intend to pre-register for New York, it would be odd of me to

maintain that, so long as I am decided on New York and against Boston, it is a matter of indifference whether I pre-register for New York or for Boston instead.

If we were confusing *IEI* with a principle of rationality, it seems, that principle would have to be:

Believed Intention Effect (BIE): If believe that (I have merely sufficient reason to *E1*, merely sufficient reason to *E2*, conclusive reason (to *E1* or to *E2*), and lack sufficient reason (to *E1* and *E2*)), intend to *E1*, and believe that (*M* is a means to *E1*-ing), then it is irrational of me to deny that, because of this intention, I have stronger reason than I would otherwise have to *M*.

The problem is not that there is no such principle of rationality. There may well be. The problem is instead that *BIE* seems to *presuppose IEI*, read literally, as a distinct fact about reasons. Suppose that *IEI* were false. Now consider a reflective agent who *knew* that *IEI* was false: that after deciding to *E1*, one does *not* come to have stronger reason. According to *BIE*, rationality would require her to ignore what she knows: require her to believe, falsely, that in deciding to *E1*, she comes to have reason that she lacked before. How can rationality, as a rule, require us to take a false view of our reasons?²⁹

There is a more basic problem with the suggestion that *ME*, or *BIE*, or any other principle of rationality might explain the phenomena that lead us to accept *IEI*. These principles require something of the agent when (provided the other psychological conditions are met) and only when he *believes* that *M* (or intending to *M*) is a means to *E*. Since an agent's rationality consists merely in the coherence of his attitudes, what matters for his rationality is merely what he believes, not what is the case. By contrast, *IEI* does *not* apply when the agent *falsely* believes that *M*-ing is a means to *E*-ing. If in fact there is no pre-registration for the conference, and I

²⁹ Scanlon, "Reasons: A Puzzling Duality?" p. 236 suggests such a principle of rationality. But he does not suggest that *IEI* is plausible only because it is confused with this principle. He agrees that this principle of rationality must be backed up by a truth about reasons, for much the same reasons as given in the text.

have simply misunderstood the situation, then spectators, realizing this, will *not* think that my decision makes it the case that I have reason to pay. They might advise me of this. And *IEI* *does* apply even when the agent *fails to realize* that *M*-ing is a means to *E*-ing. If I do not realize that I must register for the conference, then observers will *still* think that my decision makes it the case that I have reason to pre-register. Again, they might advise me of this. In sum, *IEI* appears to be a claim about reasons, not about structural irrationality.

If we cannot explain away the first appearance—that my intention affects my reason—then we can defend the Thesis only by explaining away the second appearance—that my intention does not affect my value-based reason. But it is not clear how my intention could affect my value-based reason. It would be a strange kind of action at a distance for my *mere decision* to go to New York to affect the value of attending that conference, say, by making it more important or illuminating.

What Scanlon suggests, more plausibly, is that the decision to attend the New York conference affects my reasons by affecting my *relation* to the value of attending the New York conference. This is entirely in keeping with the value-based view. The reasons that something of value provides us with depend, in general, on our *relation* to that thing of value.³⁰ As Scanlon puts it, adopting a goal changes one's reasons by putting one in a new relation to the goal, just as, say, developing specialized skills would put one in a new relation to it.

³⁰ As the wording of Value-Provision suggests, the relevant relation to the thing of value need not be the purely facilitative relation of being more likely to do what it gives one reason to do. Rather, what it gives one reason to do in the first place may depend on one's relations to it. For example, my friendship gives me reasons to do certain things that it does not give strangers. This is part of what is meant by saying that friendship gives rise to "agent-relative" reasons. (In this case, however, the thing of value cannot be neatly distinguished from my relation to it, since what constitutes my relation to my friendship also partly constitutes that friendship itself.)

It is clear how acquiring skills tailored to a particular goal might change one's reason. Those skills make it less costly for one to pursue that goal going forward, or make it more likely that one will accomplish it. But what new relation to the goal does merely *intending* it bring about, and how does this new relation affect one's reason? Scanlon writes:

What a person has reason to do in such a situation is to pursue one of these eligible goals. To do this, she must employ some procedure for selecting one of the goals and then pursue that goal, rather than any of the others. Once she has selected a goal, she is no longer in this situation. This might be put by saying that adopting a goal puts one in a different relation to that goal, and this makes a difference to what a given consideration gives one reason to do.³¹

But, again, what is it about this change in one's "situation" that changes one's reason? Unless more is said, one worries that this just abstractly describes the kind of explanation that the Thesis needs, rather than provides such an explanation.

Indeed, one might worry that *IEI* is straightforwardly explained by the intention-based hybrid, which says, recall, that in addition to whatever value-based reason to *E* one may have, one's intending to *E* provides non-value-based reason to *E*. In most cases, the explanation would say, this non-value-based reason is negligible. Either it is superfluous, because one has conclusive value-based reason to pursue *E*, or it is outweighed, because one otherwise has conclusive value-based reason not to pursue *E*. But when one's value-based reasons otherwise underdetermine the choice, one's non-value-based reason breaks the tie. So my intention to attend the New York conference makes it the case that I have conclusive reason to attend the

³¹ Scanlon, "Reasons: A Puzzling Duality?" p. 239. See also *Moral Dimensions*, pp. 93–94. Similarly, Raz, "The Myth of Instrumental Rationality," p. 23 suggests that intentions to pursue certain goals can affect our reasons to take certain subsidiary actions, but not by changing the value of the goals themselves. "Crucially, the way goals acquire their normative relevance is by being conditions on the applicability or stringency of reasons. Therefore, they can have that effect only if the goals are worth pursuing in the first place."

New York conference, and in turn, by the *GT*, makes it the case that I have conclusive reason to take the means to it. Doesn't that—the hybrid theorist will say—fit the facts just about right?

4. *The cost explanation*

It is often noted that if one *takes* some means to *E*, or makes some “investment” in achieving *E*, this can make it the case that one has greater reason, on balance, to *E*. Having already taken some means toward the end, one can realize it, going forward, with less cost. So, on balance—that is, taking into account the costs and benefits—there is greater reason to *E*.³² Since this mechanism is predicated on actually taking means to *E*, it is usually thought to be irrelevant to *IEI*. But this may be an oversight, as Michael Bratman once pointed out to me in conversation. Strictly speaking, intending to *E* is taking *some* means to *E*-ing. Having intended to attend the New York conference, I can realize the end of attending the New York conference with one fewer step than it would take to realize the end of attending the Boston conference. Boston requires, whereas New York does not, that I take the step of changing my mind. This step requires deliberation, which typically carries some cost. It taxes my attention and mental energy, and it involves the irritation of an unsettled future. Recall, moreover, that there is no offsetting benefit. Since, as I already know, neither choice is better than the other, thinking about it more will not help me arrive at a better choice. Of course, I and others are rarely so articulate about the fact that changing our minds is a pointless burden. But we naturally think and say things that suggest that we are in fact responding to such factors: “You could spend all day going back and forth on this. But you've already made a decision, and you have other things to do.”

³² See, for example, Broome, “Are Intentions Reasons?” pp. XX; Cullity, “Decisions, Reasons, and Rationality,” p. 64; Raz, “The Myth of Instrumental Rationality,” p. 22.

Scanlon suggests something very close to this explanation. He rightly observes that the cost of deliberation provides one with “second-order purely pragmatic” reasons *not to reconsider*” the decision (unless some evidence appears to cast doubt on the choice that one made). Reconsideration is a waste; to reopen the decision is to bear needless cost.³³ However, Scanlon seems not to take the further step of saying that the cost of deliberation provides one with reason *to pursue* what one now intends. But whether or not he says it, it seems to follow from what he says. If pursuing another option requires one to reconsider, and if to reconsider is to bear needless cost, then to pursue the other option requires one to bear needless cost.³⁴

This “cost explanation” supports something stronger than *IE1*, in two respects.

Intention Effect, second pass (IE2): When there is *reason* for me to *E*, if I intend to *E*, then I have stronger reason, *on balance*, than I would otherwise have *to E* and thus, by *GT*, to take means to *E* (which means presumably includes *sustaining* the intention to *E*).

First, intending *E* makes it the case that one has stronger reason, on balance, not only to take *means to E*, but also to *E itself*. The cost of *E-ing* as a whole is lower than it was, and so, to that extent, it has become the better bargain.

³³ The point is sometimes put by saying that, if someone is disposed to reopen decisions, then he will achieve less of what he has reason to achieve over the long run. This may be true, but there is no need to appeal to dispositions, or to what will happen over the long run. If the agent reopens this very decision, he thereby incurs unnecessary cost in this very instance.

³⁴ See “Reasons: A Puzzling Duality?” pp. 240–46. The following line of thought may make it seem as though the costs of deliberation provide reasons only *not to reconsider*, not also *to pursue* what one intends. “Suppose that one reopens the decision and, having done so, revises one’s intention. Granted, one did something wrong in *reopening* the decision. But, *having* reopened it, one did nothing wrong in then revising one’s intention. This shows that while one had conclusive reason not to reopen the decision, one did not have conclusive reason to pursue the option one had decided on.” But what is actually shown is that, *once one reopens the decision*, one no longer has conclusive reason to pursue the option one had decided on. But this is precisely what the cost explanation predicts. By reopening the decision, one erases the cost advantage. Now, going forward, pursuing either option will require some deliberation and choice. This does not show that, *before* reopening the decision, one did not have conclusive reason to pursue the option one had decided on.

Second, the effect is not limited to tie-breaking situations. Suppose the New York conference was the only option, so that I had conclusive reason to attend it. It would still be the case that, having decided to attend it, I no longer need to take that step. To that extent, then, the prospective costs of attending are lower than they were, and so, again, it is an even better bargain than it was. Even if I lacked sufficient reason to attend New York, the lower prospective cost of attending would give me greater, although perhaps still not sufficient, reason, on balance, to attend New York. Of course, it is clear why this effect may be overlooked in non tie-breaking situations. It is most often, if not always, in tie-breaking situations that the modest change in cost makes it the case that one has conclusive reason to pursue the option that one has decided on.

On this explanation, the effect of intending to *E* is purely facilitative. It makes it the case that an end can be achieved at lower cost. The general principle is:

Cost: If some condition lowers the cost, going forward, of achieving some end, *E*, then, because of this, there is less reason *against E-ing*. As a result there is more reason, on balance, to *E* and hence, by *GT*, to take means to *E*.³⁵

The cost explanation is compatible with the Thesis, therefore, so long as the reasons to pursue the end and to avoid the cost are value based. The effect of intention here is also attitude independent. Factors other than my having some attitude can lower the cost of achieving some end. The thing that I intend to buy, for example, can be put on sale.

One worry about the cost explanation is that revising one's intention is not always costly. There are situations, such as being stuck in an elevator or lying awake in bed, in which one has

³⁵ See "Instrumental Transmission" for discussion of why and when (since there are exceptions) the fact that an action requires some cost provides one with reason against that action.

“nothing better to do” than to revisit the decision.³⁶ The cost explanation would imply that, once one enters into such a situation, one no longer has greater reason to pursue the option one intends. But this seems like the right implication. If I have nothing better to do, then presumably I do not have the opportunity to take any means to my ends other than sustaining or revising my intentions themselves. So the test is whether, if I find myself in such a situation, reopen the decision to go to New York, and change my mind,³⁷ I have done something that I lack sufficient reason to do, or failed to do something that I have conclusive reason to do. It seems, as the cost explanation would predict, that I have not.

A different worry is that even when revising one’s intention is costly, this cost seems, in most cases, too modest to explain the more pronounced change that one’s intention seems to work on one’s reasons. So we turn to another effect of intention, which, together with the effect that we have been discussing, may help to explain this more pronounced change.

5. The effectiveness explanation

In general, one’s reason to intend to *X* depends on the epistemic probability that other conditions, that affect one’s reason to *X*, or whether one will succeed in *X*-ing, will or will not obtain. I have more reason to intend to picnic tomorrow, for example, if it is less likely to rain. These conditions are sometimes the actions of a person. I have more reason to intend to attend this conference if you are more likely to attend. And sometimes the person is the agent himself. I have more reason to intend to find accommodation where the conference is being held, if I am more likely to attend the conference. And I am more likely to attend if I intend to. When I

³⁶ Scanlon, “Reasons: A Puzzling Duality?” p. 241; Parfit, *On What Matters*, Sect. XXXX.

³⁷ However, as suggested in note 34, even in a case in which one does have something better to do, *once* one has reopened the decision, changing one’s mind may well not be against reason.

intend something, this changes what the future is likely to hold. And, in general, changes in what the future is likely to hold, whether or not those changes are brought about by my intentions, change what I have reason to intend and do.

This is the phenomenon that Scanlon calls the “predictive significance” of intention.³⁸ Why, Scanlon asks, is it wrong for me to buy rat poison with the intention of using it to kill my wife, but not wrong to buy it with the intention of using it to kill only rats? His initial answer is that when I intend to kill my wife, but not when I intend to kill only rats, I am likely to perform the other acts (such as putting it in her soup) that together with my buying rat poison will be jointly sufficient for killing her. Thus, when I intend to kill my wife, but not when I intend to kill only rats, buying rat poison puts her at risk. In buying rat poison for myself, I am, so to speak, an accomplice to murder, just as I would be if I gave my neighbor rat poison, knowing that he was likely to use it to kill his wife.

Similarly, when I intend to attend the New York conference but not the Boston conference, my registering for the New York conference makes it more likely that I will attend the New York conference—is a *more effective* means to attending the New York conference—than registering for the Boston conference would make it likely that I will attend the Boston conference. That is, if I intend to attend the New York conference, but do not intend to attend the Boston conference, then:

the probability, conditional on registering for the New York conference, that registering for the New York conference helps to bring it about that I attend the New York conference

³⁸ See *Moral Dimensions*, Ch. 2. In “The Myth of Practical Consistency,” *European Journal of Philosophy* 16 (2008): 366–402, and “Instrumental Transmission,” I suggest that predictive significance explains many of the phenomena that taken to be evidence for rational requirements of consistency and means-end coherence in intention. If we take predictive significance seriously, in other words, we may have no need to posit such requirements.

is greater than:

the probability, conditional on registering for the Boston conference, that registering for the Boston conference helps to bring it about that I attend the Boston conference.

The reason is that if I don't intend to attend the Boston conference, it's all but certain that I will not take other means that, along with registration, would be jointly sufficient for attending the Boston conference. I won't book a flight to Boston, I won't board a Boston-bound plane, etc. So pre-registering for Boston will do next to nothing to increase the chances that I attend the Boston conference. It will be a waste. By contrast, if I do intend to attend the New York conference, it's likely that I will take the other means that, along with registration, will be jointly sufficient to attending the New York conference. Such is the "predictive significance" of that intention.³⁹

More generally put:

- (i) When I intend to *E*, I am more likely to take some insufficient means *M** to *E*-ing.
- (ii) When I am more likely to take some insufficient means *M** to *E*-ing, there is a higher probability, conditional on my taking some other insufficient means *M*, that my *M*-ing helps to bring it about that I *E*: that is, *M*-ing is a *more effective means* to *E*-ing.
- (iii) *Effectiveness* (implied by *GT*): If some condition makes my taking some step, *M* (which different from this condition) a more effective means to some goal, *E*, that there is reason for me to pursue, then, because of this, there is more reason for me to *M*.

In our case, the relevant condition is my intending to attend the New York conference, which makes pre-registration for New York a more effective means. My intention plays a purely facilitative role. So this explanation is consistent with the Thesis, so long as the reasons to pursue the goal are value based. The effect of my intention is also attitude independent. Other

³⁹ Brunero, "Are Intentions Reasons?" arrives independently at a similar explanation (although he suggests, arbitrarily in my view, that the effect is restricted only to *necessary* insufficient means). I have learned a great deal from his illuminating paper.

factors—such as my legs’ being rested, my gas tank’s being full, my bank account’s being flush, and so on—can also play this role, by making some means more effective.

The effectiveness explanation supports:

Intention Effect, third pass (IE3): When there is *reason* for me to *E*, if I intend to *E*, then I have stronger reason than I would otherwise have to take *insufficient* means to *E*-ing.

This differs from *IE2* in that intending to *E* does not affect my reason to *E*, or to take *sufficient* means to *E* (which entails *E*-ing). This is because, even if I intend to attend the New York conference,

the probability, conditional on attending the New York conference, that I will attend the New York conference

and:

the probability, conditional on attending the Boston conference, that I will attend the Boston conference

are equally one. Instead, it affects only my reason to take *insufficient* means to *E*: e.g., to pre-register for the conference.

IE3 also differs from *IE1* in that the effect is not limited to underdetermined cases.

Suppose that I have conclusive reason to attend the New York conference, but I haven’t decided to. Do I have sufficient reason to pre-register? Probably not. If I don’t decide to attend the New York conference, then I won’t take other means. So I won’t attend the New York conference even if I pre-register. So whatever reasons I might have for attending it do not translate into reasons for pre-registering. And there may well be reasons against pre-registering, such as that the money could be spent elsewhere. Thus, while the best option may be to pre-register and to take other means, the second-best option, assuming that I will not take other means, is not to pre-register at all. Or to put it in terms that will be further explained in Section 8, while I have most reason (to intend to attend New York and to pre-register), I have more reason (not to intend to

attend New York and not to pre-register) than (not to intend to attend New York and to pre-register).⁴⁰

Nevertheless, it is easy to overlook this effect in determined cases. The *IE3* effect is more pronounced in underdetermined cases. In most cases, one is more likely to do something that one hasn't yet decided to do if one has conclusive reason to do it than if one has merely sufficient reason to it, other things equal. Therefore, deciding to do something that one has conclusive reason to do has a smaller effect on the probability of one's taking other means than has deciding to do something that one has merely sufficient reason to do.

I turn now to some objections to the effectiveness explanation. First, it is not always true that intending the end makes insufficient means more effective. On the one hand, one's intending the end might not make it significantly more likely that one takes other means, because, say, one will refuse to take those other means. On the other hand, intending the end might not make it significantly more likely that one takes other means, because it was already so likely that one would take those means. In these cases, according to the effectiveness explanation, one's intending the end does not affect one's reason to take the means in the way it otherwise would.

⁴⁰ These are structurally similar to cases discussed by Frank Jackson, "On the Semantics and Logic of Obligation," *Mind* 94 (1985): 177–95; and Frank Jackson and Robert Pargetter, "Oughts, Options, and Actualism," *Philosophical Review* 95 (1986): 233–55. The claims in the text do not deny (taking a helpful example from Skip Schmall) that failing to intend (say) to save a drowning child is *against reason* if one has conclusive reason to intend it, and *irrational* if one believes that one has conclusive reason. It would also be *dishonest* to try to justify oneself by treating one's lack of intention as an impediment beyond one's control, or by saying, "Alas my diving in would be of no use. But I would help to save the child, if only there was some way that I could help." For there is a way one could help: namely, by intending to save the child. But this does not affect the point: that I have more reason (not to intend to save the child and not to dive in) than (not to intend to save the child and to dive in).

Viewed from the outside, this seems true and unproblematic. If observers believe that I won't ever board the plane, say, then it will seem to them a waste for me to pre-register. And if they are convinced that I will opt in the end to go to the New York conference, however sincere my professions of indecision, then it will seem to them that I ought to pre-register now, before the fees increase. It will probably seem to them that I ought to decide to attend now too; I'm wasting my time turning it over in my mind, simply forestalling the inevitable.

When we consider my point of view, however, things may be more puzzling. It would be strange for me to view my own intentions for the end as making no difference to whether I should take the means. On the contrary, we would expect me, upon making my decision, to see myself as having reasons that I did not have before. The explanation, I think, is that it would be difficult or impossible for me to believe what observers know, so long as my intentions are as described. If, in the first case, I knew what observers know—namely, that I will not take other means—then I could not sustain my intention for the end. One does not intend what one believes one will not do.⁴¹ If, in the second case, I knew what observers know—namely, that I will decide on the end—then I would likely feel pressed to decide on it myself. Once one knows how one will decide, it is bound to strike one as idle to continue wrestling with it as though it were a live question.

The second objection asks why premise (i) in the Effectiveness explanation is true when it is. *Why* am I more likely to take other insufficient means to attend New York when I intend to attend New York? Given that I am a rational agent, the explanation why I will do something is, in general, that I believe that I have reason to do it. So it is natural to answer: Because I believe that, having intended to attend New York, I have stronger reason to take those other means to

⁴¹ See Bratman, *Intention, Plans, and Practical Reason*, pp. 37–41.

attend New York. But why do I believe this? If I believe it because it is *true*, then we seem to be appealing to the very fact that we sought to explain: that intending an end gives me stronger reason to take insufficient means to it. How can we break out of this circle?⁴²

One reply would be to deny that I am more likely to take other insufficient means to attending New York because I believe that, having intended to attend New York, I have *stronger reason* to take those other means. Instead, it might be suggested, it is because I believe that, having intended to attend New York, it would be *irrational* of me to refuse to intend those insufficient means to attend New York, and when I believe that I am rationally required to do something, I am likely to comply.⁴³ One problem with this reply is that, if *ME* is the only relevant requirement of rationality, then I am rationally required to take insufficient means to what I intend only insofar as I believe that intending those insufficient means is itself a necessary means. However, I need not always believe this. The more serious problem with this reply is that, as I suggested in Section 2.3, we do not consider what rationality requires of us in deliberating what to do. Considerations of our own rationality are deliberatively inert; they have only evaluative significance.

So, if I am more likely to take other insufficient means to attending the New York conference, it must be because I believe that, having decided to attend the New York conference, I have stronger reason to take those other means. The reason why I believe this lies, I think, in the cost explanation. Having decided to attend the New York conference, I correctly believe that I have stronger reason to take insufficient means to attend the New York conference, because it is now the cheaper option. And because I expect that I will continue to believe this, I expect that I will take insufficient means to attend the New York conference. This appeal to the cost

⁴² I am particularly grateful to Jay Wallace for pressing me on this.

⁴³ Brunero, “Are Intentions Reasons?” suggests this.

explanation does not make the effectiveness explanation superfluous. The increase in the effectiveness of insufficient means gives one further reason, in addition to the cost of reconsideration, to take insufficient means. Recall that the main objection to the cost explanation was that it accounted for only a modest effect on one's reason to take means to one's end. The addition of the effectiveness explanation accounts for the more pronounced effect that decision seems to have.

6. Defending predictive significance: too "theoretical"?

Some may protest that the idea of predictive significance, which underlies the effectiveness explanation, attributes to agents an overly "theoretical" stance toward their own present intentions and future actions. To caricature the idea of predictive significance: Agents are supposed to observe that they intend certain things, to base predictions of their own future behavior on those observations, and then to "work around" that future behavior and its effects, as though they were obstacles heaved up by brute nature or some alien intelligence. That is fine for predictions of inanimate events, or the behavior of others, but it makes no sense in one's own case, when the intentions and actions are under one's control.

The first distortion of this caricature is that the idea of predictive significance requires that our knowledge of what we intend be somehow observational. The idea of predictive significance is compatible with the fact that our knowledge of what we ourselves intend is radically different from our knowledge of what others intend.⁴⁴ The second distortion is that

⁴⁴ How exactly to understand and explain the difference is, of course, a longstanding problem. In this connection, Moran, *Authority and Estrangement*, suggestively emphasizes the fact that we adopt a deliberative stance toward our own intentions, assessing the reasons for and against them with an eye to forming, retaining, or revising them in light of these assessments. One wonders, though, whether this explains knowledge of our intentions in underdetermination cases. In such cases, it seems, one might have intended the other thing, without having assessed the reasons any differently.

there needs to be, or usually will be, explicit reflection on what one intends. I suspect that our expectations about what we will do in the further future, and perhaps even our awareness of what we are currently intentionally doing, are, insofar as they are reasonable, based on inferences from the fact that we intend what we do.⁴⁵ But these inferences, like many other simple inferences, do not require us to reflect explicitly on their premises: the facts that we intend this or that.

The last distortion is that the idea of predictive significance asks the agent to treat his present intentions or future actions as not under his control, or not his responsibility, or not things for which he can have or lack reasons. Predictive significance asks the agent to treat his present intentions or future actions like natural occurrences or the actions of others in *one* respect: namely, as factors that affect the likelihood of, or determine, what will happen in the future, which in turn affects what he has or lacks reason to intend or do now. But nothing in the idea of predictive significance requires the agent to treat his present intentions or future actions like natural occurrences or the actions of others in *another* respect: namely, that they are not under the agent's control, or not his responsibility, or not things for which he has or lacks reasons.⁴⁶ Presumably, the agent needs to view his own agency in both lights. On the one hand,

⁴⁵ Sarah Paul, *How We Know What We're Doing* (Stanford University dissertation, 2008) gives several persuasive arguments for this conclusion and against the claims of G.E.M. Anscombe, *Intention* (Oxford: Basil Blackwell, 1957); Kieran Setiya, *Reasons without Rationalism* (Princeton: Princeton University Press, 2007); and J. David Velleman, *Practical Reflection* (Princeton: Princeton University Press, 1989).

⁴⁶ There is a kind of Sartrean "bad faith," illuminated by Moran, *Authority and Estrangement*, Ch. 3, that consists in a similar slide from treating one's present intention, like a natural event or the mind of another, as an object of belief—a "facticity"—to treating it, like a natural event or the mind of another, as something not under one's control—a *mere* facticity. Moran has no objection, at least not here, to treating our intentions as objects of belief, or evidence of what we will do (see p. 79). The problem lies in *not* treating them as things that we sustain or abandon by our own deliberation. I would add that while there is bad faith in treating one's *present* intention as something not under one's control, there is not the same bad faith in treating the *influence* of one's present intention on one's *future* behavior as something not under one's *present* control—although it is less clear to me whether Moran would agree. What one

he needs to see his present intentions and future actions as things under his control, to consider the reasons that he has for them, and to respond accordingly by forming, revising, or abandoning his present intentions. On the other hand, he also needs to be sensitive to what those formations, revisions, and abandonments mean for the future, and to what this future then means, by a kind of feedback, for his reasons.

What, after all, is the alternative? Suppose that we could not base any expectations of our future behavior on what we currently intend. Practical reasoning, which relies pervasively on such expectations, would be crippled. It is not clear how we could ever take means to our ends, or coordinate our actions.⁴⁷ Theoretical reasoning would also be crippled, since many of the future events about which we form expectations depend on what we will do. Fissures would appear elsewhere. Our generalizations about peoples' behavior would break down, without explanation, in our own case. And others' predictions of our own behavior would seem to us groundless.

7. Defending predictive significance: Scanlon on self-facilitation

In motivating the idea of predictive significance, I invoked an explanation of Scanlon's, which we might call the "appeal to self-facilitation." This is the explanation that, when I intend to

now intends is under one's present control, but what one will do in a week, say, is not under one's present control, except via one's present intentions (or self-manipulation). If one is convinced that one's present intention will fail to influence one's future behavior, then I think there is nothing amiss in judging that present intention futile and abandoning it. Of course, much may be amiss in the conviction itself, which may be the product of self-deception. And much may be amiss in the back-sliding that it anticipates, which will be akratic unless one changes one's mind about whatever reasons underlie one's present intention.

⁴⁷ It might be replied that we accomplish this by directly conforming to rational requirements on our intentions, such as *ME*, in a way that bypasses expectations about what we will do. For example, when we intend the end, we come to intend what we believe are necessary means, without any reflection on how we are likely to behave. Among other problems with this suggestion, our deliberations often display finer-grained sensitivities, to our chances of success and to non-necessary means, which the suggested mechanism cannot explain.

poison my wife, my buying rat poison is wrong, for the same reason that, when my neighbor intends to poison his wife, my giving him rat poison is wrong: because it facilitates a murder. The difficulty is that Scanlon later distances himself from the appeal to self-facilitation, for reasons that might seem to threaten the idea of predictive significance itself. Scanlon observes that, whereas in the neighbor case it seems natural to say that what I am doing is wrong because it facilitates a murder, it seems odd about to say this when I myself am the prospective murderer, at least if I expect that I will *control* whether or not I actually poison my wife. If instead I expect that I will not be able to control my murderous urges, then buying the rat poison for my future self is much like giving it to my neighbor. But if I do expect that I will control whether I poison my wife, then “the idea of facilitation involves an odd relation between [me] and [my] own future conduct.”⁴⁸ I agree that there is something odd about the appeal to self-facilitation. But I do not think that the explanation of what is odd undermines the idea of predictive significance, or the suggested uses of it.⁴⁹

What’s odd about the appeal to self-facilitation, at least in the first instance, seems to lie in how the appeal expects the agent to reason. The initial parts of his reasoning, however, seem unproblematic:

(1A) I intend to poison my wife.

(2A) If I intend to poison my wife, I have reason to believe that if I get rat poison, I will poison her (even though I know that this will be under my control).

⁴⁸ Scanlon, *Moral Dimensions*, Ch. 2. I take it that the point is not simply: “An intention to murder one’s wife would have to be akratic, or impulsive. So, if it remains under the agent’s control, there is little reason to expect that he will carry it out. Sooner or later he will come to his senses. Therefore, manipulation is unnecessary, unless it will not be under his control.” For one thing, this would not provide any contrast with the third-person case.

⁴⁹ It may not even undermine the truth of the appeal to self-facilitation itself: the claim that buying the poison is wrong, precisely because the agent thereby facilitates his own murderous plans. See note 54.

(3A) If I have reason to believe that if I get rat poison, I will poison my wife (even though I know it will be under my control), and if I have no other compelling reason to get rat poison, then because of this, it is wrong for me to get rat poison.
(4A) So, it is wrong for me to get rat poison.

One might worry that (1A) and (2A) are “too theoretical,” but, if the argument of the last section is correct, this worry has little merit. At worst, (1A) and (2A) are unnaturally explicit. (3A) follows from applying to the facts of the case the following plausible moral principle:

Control Principle: If I have no compelling reason to *X*, and I have reason to believe that *X*-ing will significantly increase the risk of harm to someone, then it is wrong for me to *X*.

Of course, (3A) would not follow if we rejected the Control Principle in favor of:

Non-control Principle: If I have no compelling reason to *X*, and I have reason to believe that *X*-ing will significantly increase the risk of harm to someone, *but not through events that I control*, then it is wrong for me to *X*.

But the Control Principle seems more compelling morally, especially within a contractualist framework. On the one hand, the objections to permission are weaker against the Control Principle.⁵⁰ From the victim’s point of view, it seems irrelevant whether the route from *X*-ing to the harm depends on events that the agent controls. What matters is simply that *X*-ing increases the risk of harm. On the other hand, there seems no valid objection to prohibition against either principle, since the agent has no compelling reason to *X*.⁵¹

So what makes the reasoning in the appeal to self-facilitation so odd? Amplifying this question is the fact that many structurally similar explanations, which Scanlon himself describes,

⁵⁰ See Scanlon, *What We Owe to Each Other*, Ch. 5 for further explanation of “objections to permission” and “objections to prohibition.”

⁵¹ Even if we assume the agent has compelling reason to *X*, the objections to prohibition against the Control Principle do not seem significantly stronger. If *X*-ing leads to harm only through events that the agent will control, then the agent has a way to *X* without violating the Control Principle. He only needs to intend now to control future events in such a way that *X*-ing will not lead to harm. Then *X*-ing no longer increases the risk of harm, and so does not violate the principle.

do not seem odd in the same way. Withdrawing from the appeal to self-facilitation, Scanlon suggests the fallback position that, although *buying* the poison is not wrong, *refusing to abandon my intention* to kill my wife is wrong.⁵² But why is it wrong? Scanlon suggests that it would *not* be wrong to stick pins in a voodoo doll with the intention of killing my wife, if I had no reason to believe that it would actually kill my wife.⁵³ This suggests that it would not be wrong to *have the intention* to kill my wife *unless* I had reason to believe that the intention would actually lead to my killing her. But if so, then my reasoning about keeping the intention is similar to my reasoning about buying the poison:

(2C) I have reason to believe that if I intend to poison my wife, I will poison my wife (even though I know that it will be under my control to revise this intention).

(3C) If I have reason to believe that if I intend to poison my wife, I will poison my wife (even though I know that it will be under my control to revise this intention), then, because of this, it is wrong for me to intend to poison my wife.

(4C) So, it is wrong to intend to poison my wife.

Moreover, when explaining how intentions can make certain actions *permissible*, rather than, as in the rat poison case, *impermissible*, Scanlon suggests that one might correctly reason:

(1D) I intend to use embryos for therapy.

(2D) If I intend to use embryos for therapy, I have reason to believe that, if I create an embryo, I will use it for therapy.

(3D) If I have reason to believe that, if I create an embryo, I will use it for therapy, it is permissible, whereas it would otherwise be impermissible, to create an embryo.

(4D) So, it is permissible to create an embryo.

(1E) I intend to fight the fire.

(2E) If I intend to fight the fire, then I have reason to believe that, if I cross your land, I will use this crossing to fight the fire.

(3E) If I have reason to believe that, if I cross your land, I will use this crossing to fight the fire, then it is permissible, whereas it would otherwise be impermissible, to cross your land.

(4E) So, it is permissible to cross your land

⁵² Scanlon actually writes that I “should abandon” this intention, but it’s hard to see what else “should” could mean in this context.

⁵³ See Scanlon, *Moral Dimensions*, Ch. 2.

These analyses are even closer to the appeal to self-facilitation. They involve an agent's judging something about the permissibility of a present action on the grounds that it would facilitate a future action of his that he has reason to believe that he will perform because of what he currently intends. And yet they don't seem odd in the same way. Why?

The answer, I think, lies in the steps the reasoning takes next:

(4A) So, it is wrong for me to get rat poison, since putting my wife at risk is wrong.

(5A) So, I won't get rat poison.

(6A) But, of course, I still intend to kill my wife, even though it puts my wife at risk.

This further stretch of reasoning is incoherent, because it involves simultaneously taking and not taking the risk to my wife as a reason. That is, I judge in (4A) that the risk to my wife means that I lack sufficient reason to *buy the poison*, and I respond accordingly in (5A) by not buying the poison. But in (6A) I fail to *drop the intention* to poison her, either because I fail to judge that the risk to my wife means that I lack sufficient reason to *intend to poison her*, or because I fail to comply with that judgment. Why do we assume (6A): that I do not drop the intention to kill my wife? Because if I *did* drop the intention, then I could not sustain the judgment in (4A) that my getting rat poison puts her at risk. It would be obvious to me that (1A) and (2A), on which (4A) depends, were no longer true. To put it paradoxically, my intention is in tension with the very reasoning that it supports.

The general structure of the incoherence is:

- (i) One intends to *X*.
- (ii) One believes that, since one intends to *X*, it is likely, conditional on one's taking some insufficient means, *M*, to *X*-ing, that one's *M*-ing will help to bring it about that one *X*'s.
- (iii) One concludes that, because it is likely to help to bring it about that one *X*'s, one lacks sufficient reason to *M*, and so refuses to *M*.
- (iv) One does not similarly conclude (or does not respond to the similar conclusion) that, because it is likely, conditional on one's intending to *X*, that intending to *X* will help to bring it about that one *X*'s, one lacks sufficient reason to intend to *X*—as is evidenced by one's continuing to intend to *X*.

Notice that the same incoherence does not arise in the other, contrasting cases that we have considered. The reasoning in (2C)–(4C), in which I abandon the intention to kill my wife, because I believe that that intention puts her at risk, leaves no opening for such incoherence, since there is only one judgment about my reasons, and I comply with it. Consider next the neighbor case. I can coherently (a) judge that I lack sufficient reason not to give my neighbor rat poison because, given that he intends to poison his wife, my giving it to him puts her at risk, (b) judge that *he* lacks sufficient reason to intend to poison her, because his having that intention puts her at risk, (c) comply with (a), and (d) judge that *he* fails to comply with (b). Or suppose that I expect that I will not be able to control my future actions. No matter what I intend now, if I get rat poison, I will poison my wife. In this case, I can coherently judge that I lack sufficient reason to intend to poison my wife, because this puts her at risk, and that I lack sufficient reason to buy the poison, because that puts her at risk. And responding to the first judgment *now* does not falsify the second judgment. Even if I don't intend to kill her now, buying rat poison still puts her at risk, because I will intend it later.

Cases with the structure of the conference decision also escape the incoherence. My intending to attend the New York conference fits the following coherent pattern:

- (i) One intends to *X* (e.g., to attend New York).
- (ii) One believes that, since one intends to *X*, it is likely, conditional on taking some insufficient means, *M*, to *X*-ing, that *M*-ing will help to bring it about that one *X*'s.
- (iii) One concludes that, because it is likely to help to bring it about that one *X*'s, one has sufficient reason to *M* and so *M*'s.
- (iv) One does not reject the conclusion that, because it is likely, conditional on intending to *X*, that intending *X* will help to bring it about that one *X*'s, one has sufficient reason to intend to *X*—as evidenced by one's continuing to intend *X*.

The land-crossing and embryo-creation cases, (1D)–(4D) and (1E)–(4E), share this coherent structure. I can coherently judge that fighting the fire is sufficient reason to cross your land

given that I intend to fight the fire and that fighting the fire is sufficient reason to intend to fight the fire in the first place. And my responding to the second judgment—that I have sufficient reason to intend to fight the fire—verifies, rather than falsifies, the first judgment—that I have sufficient reason to cross your land. Thus, the oddity of the appeal to self-facilitation does not threaten the idea of predictive significance or the applications of it that we have considered.⁵⁴

8. *What matters to us*

So far, I have been asking how having an intention, for whatever object, might affect one's reason for action. However, there is something else that often comes under the label "having an aim." To "have an aim" in this sense is to "identify" with, to be "committed" to, to "care" about, to "value" some life-structuring pursuit. It is for some cause, career, or calling to be "important" or "matter" to one.⁵⁵ "Having an aim" in this sense differs from having an intention at least in taking a narrower range of objects. While one can intend to pick up dry cleaning or to order a sandwich, these are not activities that can be "important to" one or that one can "care about" in

⁵⁴ The objections to predictive significance in this and the previous section have been claims that an agent cannot *take* certain considerations as reasons. It is a further question whether, if these claims are true, this means that the considerations *are* not reasons. On the one hand, I have myself suggested that reasons are to be understood by reference to their role in deliberation. So there must be some connection between reasons and deliberation. On the other hand, it seems a mistake to construe this connection so tightly that *any* psychological impediment to taking them into account deprives them of standing as reasons. My inability to face a certain terrible fact, for example, should not, it seems, disqualify it as a reason. Similarly, the fact that it would be incoherent to think about the fact that I lack sufficient reason to continuing thinking about my reasons ("Stop dithering and act already!") should not mean that there is no such fact. Perhaps the sort of inability alleged in the previous section—the alleged inability to form predictions about or on the basis of what one believes one controls—would have been sufficiently basic and structural to constrain what reasons we have. But the kind of inability that we have acknowledged in this section seems more derivative and local. So I am not persuaded that we should conclude that, because responding to the judgment that buying the rat poison is wrong would involve the sort of incoherence that we have discussed, buying the rat poison is not in fact wrong.

⁵⁵ These expressions no doubt have different shades of meaning. For example, as Samuel Scheffler, "Valuing" XXXX points out, valuing involves, whereas caring need not, believing that the thing valued is valuable. My topic is the area where they overlap.

the relevant sense. “Having an aim” in this second sense also differs from having an intention in the psychological orientation it involves. Merely deciding to pursue a career is not the same as “caring about” or “valuing” it. For one thing, when one has to decide whether (say) to stick with a career, one’s decision often is preceded by an attempt to figure out what one really cares about or values. For another, the decision is something that one can get right or wrong, by sticking with, or giving up, what really “matters” to one. When the judge’s son opts for law school over a promising future as a jazz musician, we may think that he has given in to expectations, or shrunk from the risk of failure, instead of pursuing what is truly “important to” him. It is not clear how either of these things could be so if “caring” consisted in the decision itself.

What is it to “care” about something? The most plausible account, I think, is that it is to be *emotional vulnerable* to it: to be disposed to experience positive emotions when it fares well or is properly regarded by others, and to experience negative emotions when it fares poorly or is not properly regarded by others.⁵⁶ If I care about my work, I will feel elated when it advances and dejected when it is set back. I will feel affirmed when others admire my work, crestfallen when they call its flaws to my attention, and resentful when they dismiss it too hastily.

It is natural to think that having an aim, in this sense, gives me reason that I would not otherwise have. If I care about my life’s work, if it matters to me, then I have reason to pursue it that someone else, who does not have that relationship with it, lacks. And if I care about my life’s work, I have reason to pursue it that I do not have to pursue other aims, to which I am not attached in the same way. To be sure, the work itself must be independently worthwhile for my

⁵⁶ See Elizabeth Anderson, *Value in Ethics and Economics* (Cambridge: Harvard University Press, 1993), Ch. 1; Niko Kolodny, “Love as Valuing a Relationship,” *Philosophical Review* 112 (2003): 135–89; and Scheffler, “Valuing.” Scheffler argues persuasively against other proposals: that caring, or rather the relevant component of valuing, is desiring, desiring to desire, believing valuable, or having a particular feeling.

being invested in it in this way to provide me with reasons. The work must be something that everyone, whether or not he or she cares about it, has reason at very least to respect, if not also to support. I have no reason to pursue the project of counting blades of grass on the asylum grounds, for example, no matter how much, in my madness, I may care about it. Nevertheless, when a pursuit is independently valuable, my caring about it seems to give me *further* reason beyond those that its independent value gives everyone.

The challenge for the Thesis is to explain how caring about a pursuit can affect one's value-based reasons in this way. It might first be suggested that the fact that an aim matters to me can make it more likely that I will succeed in it. The more likely I am to succeed in an aim, other things equal, the stronger my reason to pursue it. As Scanlon writes:

[B]eing drawn to a pursuit is (at least under favorable conditions) a condition for having a good reason to undertake it as a career. For one thing, if one takes up a career which "leaves one cold," then one is unlikely to succeed in it.⁵⁷

One problem with this suggestion is that it is not always true that I am more likely to succeed in what matters to me more. First, it is a familiar, if tragic, phenomenon that one can be less likely to succeed at something precisely because it matters more to one. Because one cares about an aim so much, one may be overly anxious about it, or perversely inclined to self-sabotage. Second, people sometimes reliably pursue aims that they don't care much about. They need to earn a living, or are creatures of habit. Third, one may simply have greater talent for an aim that matters less to one. Even so—even though I am less likely to succeed in what matters to me more—this does not seem to extinguish the force of my reason to do what matters to me more. Of course, the fact that I am more likely to succeed in something else may *outweigh* this reason.

⁵⁷ Scanlon, *What We Owe to Each Other*, p. 49. See also Parfit, *On What Matters*, Sect. XXXX.

But that is compatible with the present point: namely, that the reason that arises from its mattering to me does not itself depend on whether I am more likely to succeed in it.

A second suggestion appeals to an implication of our account of caring: that when I care more about a pursuit, I will experience more intense positive emotions when it goes well. The prospect of these welcome experiences, the suggestion runs, is what gives me reason to pursue aims about which I care. One problem is that caring about a pursuit also makes me vulnerable to more intense negative emotions when it goes badly. It is enough to call to mind the cliché of the tormented artist, who “suffers for his work,” to see that pursuing what matters to us need not make our lives more pleasant. And yet the prospect of such suffering seems not to extinguish the force of the reason to do what one cares about, although, of course, it may outweigh it. Moreover, this suggestion seems too extrinsic. Even if pursuing art would be pleasant for a talented, well-adjusted artist, it would be oddly alienated for him to pursue his artwork rather than some no less valuable project, on the grounds that, given his psychological make-up, it promised more pleasant experiences.

A similar problem recurs for a third suggestion: that I have reason to pursue an aim that I care about because pursuing it enhances my well-being. It is plausible that one’s well-being consists, *inter alia*, in the successful pursuit of independently valuable aims that one cares about.⁵⁸ What is less plausible is that I have reason to enhance my own well-being, or at least

⁵⁸ Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984), p. 502, for example, suggests an account of well-being along these lines, in which it consists in the combination of “objectively” valuable activities and the right psychological orientation to them. “Pleasure with many other kinds of object has no value. And, if they are entirely devoid of pleasure, there is no value in knowledge, rational activity, love, or the awareness of beauty. What is of value, or is good for someone is to have both; to be engaged in these activities, and to be strongly wanting to be so engaged.” It is worth asking, though, why well-being should consist in caring about (or being otherwise psychologically attuned to) independently valuable activities. One explanation is, first, that well-being consists in engaging with things of value

that the reasons that I have to pursue what I care about are reasons to enhance my well-being. Again, this seems too extrinsic. As both Scanlon and Raz have argued, agents typically view their aims themselves, and not those aims' contribution to their well-being, as the source of their reasons to pursue them.⁵⁹

This is thrown into relief by an example of Raz's, in which, as he puts it, "a person's life loses nothing by a forced change of career, and yet where this person is acting reasonably in trying to stave off the change, being willing to pay a considerable price in the process."

Let us imagine a ballet dancer enjoying a reasonably successful career with a small provincial ballet company, the only one within hundreds of miles. Then dwindling audiences threaten the future of the company. If it is forced to disband our dancer will have to abandon ballet, to change career and look for something else to do. It does not surprise us that he regards the prospect as a great personal disaster. I think that it would be agreed that it is reasonable for him to try to prevent the collapse... Let us assume that the company has to close. Our dancer looks for other possibilities and starts a new career as a theatre director with the local theatre company where he remains until his retirement. He quickly comes to like his new work, enjoying a success comparable to his success in his first career.

Even if we idealize the case, so that the dancer knows that this is what will happen, that his life will be no worse when he abandons ballet, Raz observes:

it is still reasonable for him to regard the threat to his dancing career as a disaster for him. He wants to be a dancer not a theatre director. He knows that if he is forced into theatre he will come to like it and be reasonably successful. But that is not what he wants and he is willing to try hard to stay with dancing. These admissions of his entail, on the assumption that his feelings and actions are rational, that he is not concerned with his well-being. His concern is to have the life he wants to have, meaning the life he has become committed to.

The dancer's reasons to stay with dancing neither derive from the prospects for his well-being, nor from his greater likelihood of success at dancing, nor from any supposed greater value of

and, second, that caring about an independently valuable pursuit constitutes a further value. If this explanation is correct, then the present suggestion, that caring matters because of its effects on well-being, actually depends on the idea, discussed below, that caring constitutes a further value.

⁵⁹ Scanlon, *What We Owe to Each Other*, Ch. 3; and *Moral Dimensions*, pp. 94–95; Raz, *Engaging Reason* (Oxford: Oxford University Press, 1999), Ch. 13. The following quotations are from pp. 316–17.

dancing. Instead, his reasons seem to derive from the simple fact that dancing, not directing, is the life he “wants to have.”

This may seem a telling admission: a concession that, after all, we need desire-based reasons, or attitude-based reasons of some kind, to account for the reasons that the dancer takes himself to have.⁶⁰ But this is not, I take it, what Raz has in mind. Consider these two passages from his earlier book, *The Morality of Freedom*:

Previously, I have argued that wanting something is not a reason for doing it. We can see now that, while fundamentally right, in one respect that claim was exaggerated. Saying “I want to...” can be a way of indicating that one is committed to a project, that one has embraced a certain pursuit, cares about a relationship (389).

In embracing goals and commitments, in coming to care about one thing and another...[o]ne creates values, generates... reasons which transcend the reasons one had for undertaking one’s commitments and pursuits (387).⁶¹

One might put it this way: The fact that the dancer cares about dance means that his pursuing dance constitutes something of further value—a value that would be not constituted if he did not care about it: if he pursued it coldly, say, from a sense of obligation or inertia. This further value is what provides him with reason to stave off a change, even a change for something in which he would be no less happy or successful. The suggestion, put generally, is that a certain kind of attitude-constituted value is realized by one’s pursuing an independently worthwhile aim that one

⁶⁰ Ruth Chang, “Can Desires Provide Reasons for Action?” in Wallace, Pettit, Scheffler, and Smith, *Reasons and Value*, pp. 56–90, suggests something like this at p. 86.

⁶¹ See also *Engaging Reason*, pp. 63–64. I here suppress, or downplay, language in Raz’s surrounding discussion that might suggest two ideas that we have questioned. The first is that a mere choice—the formation of a bare intention—creates a value of the relevant kind. (Recall the judge’s son.) The second is that the reasons really flow from the value of living a successful life and that the effect of caring is to determine what counts for one as a successful life. (Recall the dancer.)

cares about. It is this attitude-constituted value that, compatibly with the Thesis, provides one with further reasons to pursue such an aim.⁶²

There are, however, two straightforward worries about this proposal. The first worry, to which I will return, is that it may seem not to explain Raz's dancer after all. Shouldn't the proposal lead us to say that by bringing about the change in career and coming to care about his new life as a theater director, the dancer will realize another instance of the same kind of attitude-constituted value, no worse than the first. If so, doesn't it follow that, as far as the value is concerned, he has no more reason to stave off the change than to accept it?

The second worry, more immediately pressing given the concerns of this paper, is that the proposal does nothing more than to affix the honorific title "value" to a more straightforward hybrid explanation: that some of our reasons to pursue an aim derive from the independent value of that aim, whereas other reasons derive, independently of value, from our present caring about it. Why think that caring ushers a new value onto the scene?

One reason is that the hybrid explanation, on closer inspection, is unconvincing. First, as we noted earlier, caring about a pursuit appears to provide one with reason only if the pursuit is independently valuable. The hybrid theorist might try to explain away this appearance. Even when the pursuit is not independently valuable, caring about it provides one with reason. It is just that this reason is usually outweighed by the opportunity cost: the value-based reason to devote one's resources to a pursuit that is independently valuable.⁶³ This is similar to the earlier

⁶² Compare Susan Wolf, "Meaning in Life and Why it Matters," Tanner Lectures on Human Values at Princeton University, November 2007, who suggests that such a value, fusing "objective" and "subjective" elements, is what makes for meaning in life.

⁶³ I am implicitly assuming here that the caring-based reasons that the hybrid theory recognizes would simply be "added" to the value-based reasons that it recognizes. Caring and value would not interact in any more significant way to determine what reasons an agent has. This rules out, for example, the possibility that caring provides reason, but only when the thing

suggestion of the intention-based hybrid that the non-value-based reasons that arise from our intentions are so weak that we notice them only in tie-breaking contexts. But this explanation seems less plausible in the present case, for the simple reason that caring about a pursuit, at least when the pursuit is independently valuable, seems to provide one with stronger reasons than would a mere intention to pursue it. One might think, for example, that one has reason to stick with a career that one cares deeply about rather than to abandon it for a career that would be, in the grand scheme of things, more valuable. But it is not clear that a bare decision to pursue the former career would have this effect.

Second, the fact that one cares about an independently valuable pursuit typically also gives *others* reasons: for example, not to undermine one's pursuit of it. The hybrid theorist might suggest that this is because undermining one's pursuit of what one cares about also undermines one's *well-being*. But this need not always be so, as Raz's example suggests. Or the hybrid theorist might suggest that respect for one's *autonomy* calls for us not to undermine one's pursuit of what one cares about. But what we respect, in such cases, is not, or not simply, one's self-regarding choices. We would not have the same reason to respect a choice to pursue something, such as counting blades of grass, that is not independently valuable. Indeed, our reason may not be to respect a *decision* at all, since the pursuit about which the agent cares, and whose claims we recognize, may be one that he himself has chosen, for bad reasons, to give up. On closer inspection, the relevant kind of "autonomy" seems little different from the attitude-

cared about is independently valuable. Given the fundamental difference in kind between caring-based and value-based reasons, it is hard to see how they could interact in any more significant way.

constituted value that we have proposed. Perhaps it just is that value, viewed, as it were, from the outside.⁶⁴

Finally, even when the pursuit is independently valuable, *currently* caring about a pursuit provides one with reason only if one has also had a *history* of pursuing it, with the right sort of concern, up until now. Someone who just found himself, out of the blue, caring about pursuing a life of dance—even if he had the same athleticism and feel for music—would not have the same reasons as Raz’s dancer, who has, as Raz puts it, “become committed to” that life over time. Indeed, it would seem out of place—“premature,” as we might say—for the former to invest dancing with the same importance in the first place. (Of course, people have reason to embark on pursuits to which they are attracted. But being attracted, in this sense, is something different from caring about the pursuit, in the way that Raz’s dancer does. Indeed, the normative significance of being attracted to a pursuit may be just that it indicates that one might, in due course, come to care about that pursuit, and so constitute the relevant value.) The attitude-constituted value in question, when fully spelled out, is that of pursuing an independently valuable aim, that one cares about, where one has a history with that aim of a kind that makes sense of that caring.⁶⁵

⁶⁴ See Raz, *Morality of Freedom*, p. 382.

⁶⁵ See Niko Kolodny, “Which Relationships Justify Partiality? General Considerations and Problem Cases,” in Brian Feltham and John Cottingham, eds, *Partiality and Impartiality: Morality, Special Relationships and the Wider World* (Oxford: Oxford University Press, forthcoming) for further discussion.

If the history provides reasons for one’s present caring, then one might wonder whether one’s present caring actually contributes anything to the value. Perhaps the history suffices for the value, and then this value provides one with reasons both for present caring and for pursuing the aim. This may be the case with interpersonal relationships, such as one’s relationship to one’s child. The history gives one reason to care about the child, but the history would provide one with reasons for action even if one did not care about the child. But the kind of aims—causes, careers, callings—that we have been discussing seem different in this respect. If the

Whatever the merits of the alternative hybrid explanation, it should not be a foreign idea that there is a value with this structure. We are familiar with other such values, such as friendship. One might “go through the motions” (or at least some of the motions) of friendship, to save the feelings of a pathologically sensitive acquaintance, without actually caring about that relationship. One would not have a friendship, or the reasons that a friendship provides. One has a friendship only if one currently cares about one’s friend and one’s relationship with him, and has the sort of history with him that would make sense of this. Perhaps, if one has a friendship, then one has additional reason for interacting with one’s friend, because his well-being and one’s own are now hostage to it, or because one is more effective in ministering to his needs, or because one takes greater pleasure in doing so. But surely the more fundamental and important reasons flow from the friendship itself. Friendship just is a distinctive kind of value, which is constituted by a present emotional vulnerability to someone and to relations with him and by the kind of history that makes sense of this vulnerability. The present suggestion is just that a distinctive kind of value is similarly constituted by pursuing something of independent value, when one is currently emotionally vulnerable to its pursuit, and when one has a history of pursuing it that makes sense of this vulnerability.

The analogy to friendship in turn helps answer the first of our two worries: that the proposal might not explain why Raz’s dancer has more reason to stave off the change in career than to accept it. After all, the alternative is to come to pursue and to care about pursuing being a theater director, which would instantiate the same kind of attitude-constituted value. But just as the value of friendship need not give me as much reason to cultivate a new friendship as it gives me to cherish an existing one, likewise this attitude-constituted value need not give me as much

dancer finds that he no longer cares about dancing, then that settles the matter. I am grateful to Sam Scheffler for raising this issue.

reason to cultivate a new pursuit as it gives me to sustain my existing one. The mistake, put broadly, is to think that things of value are sources of reasons only in the sense that, when we are able to bring about something of value, we have reason to do so. But as Value-Provision reflects, and as Scanlon has cautioned, this is overly narrow. Things of value can provide us with reasons when we stand in relations to them other than being able to bring them about—such as being currently engaged with them—and the reasons that they provide us with may be to do things other than to bring them about—such as to honor or respect them in suitable ways.

This brings to the foreground, however, my assumption that the attitude-constituted value can play a role in the agent's own deliberation, which might seem open to question. People who are fully engaged in pursuits that they care about are usually focused on the features that make it independently valuable. Thus, we would expect the dancer to be concerned with the aesthetic and expressive qualities of choreography and execution, and not to be given to retrospection on his history with dance, or introspection on his current feelings about it. It seems likely that this special value rises to deliberative salience only in from what Jay Wallace calls the standpoint of “eudaimonistic reflection,” where we step back from our relationships and pursuits, and survey our lives and their sources of meaning as a whole.⁶⁶ When we are contemplating some life change, for example, then it is entirely natural to reflect on our past with some project and to engage in “soul-searching”: an attempt to get clear on whether it is what we really care about. Much the same is true of friendship. When we spend time with friends, we focus on their charms and the pleasure of their company, not on our histories with them or our own present feelings.

⁶⁶ See R. Jay Wallace, “Caring, Reflexivity, and the Structure of Volition,” in Monika Betzler and Barbara Guckes, eds., *Autonomes Handeln* (Berlin: Akademie Verlag, 2000), pp. 213–34; and “The Rightness of Acts and the Goodness of Lives,” in Wallace, Pettit, Scheffler, and Smith, eds., *Reason and Value*, pp. 385–411.

But, in other contexts, such as when our loyalties are tested, considerations about what we have been through together and how our friend matters to us regain prominence.⁶⁷

9. Conclusion

I have been addressing only one version of the challenge that our attitudes seem to affect our reasons without affecting our value-based reasons. There are other versions of this challenge, of course, arising from the perceived effects, not of aims, but instead of beliefs and desires.⁶⁸ And

⁶⁷ Do the resources that we have canvassed—cost, effectiveness, the further attitude-constituted value—explain Raz’s observation, noted in 31 above, that “goals acquire their normative relevance by being conditions on the applicability or stringency of reasons”? I think so, although, of course, other explanations may also be possible.

Raz’s main example is the following: “For someone intent on running a marathon every day during August, running a marathon today, the 20th of August, is crucial to the realization of his ambition. For me it is just an opportunity to know what running a marathon feels like — a matter of much less moment.” So described, the example may be complicated by factors independent of our different goals. If I haven’t run a marathon on each of the first 19 days of August, then, whatever my goals are, my running a marathon today does not bring me any closer to achieving the end of running a marathon every day in August. If the August marathoner has run a marathon on each of the first 19 days of August and remembers what running them felt like, then, whatever his goals, his running a marathon today does not bring him any closer to achieving the end, long since achieved, of acquiring knowledge of what running a marathon feels like. So, to split hairs, suppose that the date is August 1, neither of us knows what running a marathon feels like, and yet each of us is capable, and equally capable, of running a marathon every day during August. Why does the August marathoner have stronger reason to run a marathon today than I have?

To begin with, one imagines that the August marathoner cares about this goal, and has been preparing it for some time. If so, then the attitude-constituted value that we have been discussing will be in play. But even if we set this aside, the effectiveness explanation seems to account for the difference. Given that I have no plans to run a marathon every day in August, there’s no chance that I will run a marathon every other day in August, even if I run one today. And if I won’t run a marathon every other day in August, then running one today does not make it more likely that I run one every day in August. By contrast, given that the August marathoner does have plans to run a marathon every day in August, there is a chance that he will run a marathon every other day in August, if he runs one today. And if he will, then running a marathon today does make it more likely that he will run one every day in August. Since my running one today is not a means to my running one every day in August, whereas his running one today is a means to it, I lack the reasons that flow from realizing that value, whereas he has them. (My running one today is still a means to something of value, namely knowing what running a marathon feels like. But then so too is his running one today.)

⁶⁸ For a recent revival of this challenge, see Chang, “Are Desires Reasons for Action?”

there are other challenges, which stem from the perceived “metaethical,” or “second-order” advantages of attitude-based theories. Only such theories, it is thought, can explain the metaphysics, epistemology, or motivating force of reasons.

If I have not dealt with such challenges, it is partly from a conviction that the resilience of attitude-constituted theories owes largely to their “first-order” or “substantive” plausibility. They would not enjoy their perennial appeal were it not obvious—an utterly familiar part of human experience—that our beliefs, wants, and decisions make a difference to what we have reason to do. Of course, it is also obvious—an unavoidable consequence, I think, of taking up the deliberative point of view—that we see our reasons as flowing from what is valuable, or good, or worthwhile. As I hope to have suggested, insistence on reconciling these commonplaces, as well as grounds for hope that reconciliation is possible, are among Scanlon’s many contributions to practical philosophy.