

## Holism about Fact and Value

Kenny Walden

“Two Dogmas of Empiricism” has two projects, one negative and one positive. The negative project is an attack on Carnap’s distinction between internal and external questions, which Quine undertakes with an argument against the analytic/synthetic distinction. The positive project is to sketch an alternative picture of the dynamics of inquiry, a picture of how science looks without Carnap’s dichotomy. In place of Carnap’s bifurcated science of evidence and conventional framework, Quine gives us a holistic science suffused with both fact and convention. Knowledge, as he says elsewhere, is “a fabric of sentences [which] develops and changes, through more or less arbitrary and deliberate revisions and additions of our own, more or less directly occasioned by the continuing stimulation of our sense organs. It is a pale grey lore, black with fact and white with convention. But I have found no substantial reasons for concluding that there are any quite black threads in it, or any white ones.”<sup>1</sup>

When joined with a mild form of methodological naturalism, this “Quinean holism” leads to a distinctive approach to metaphysics, epistemology, and semantics. It does not necessarily entail particular first-order doctrines in these fields but prescribes a uniformity of regard. The thought is that once we appreciate this holism, we see that there are no good grounds for distinguishing between different parts of the “web of belief”—of saying that sector *A* of the web is substantive, while sector *B* is merely conventional, that sector *C* aims to represent the world while sector *D* is merely expressive, or that the objects of quantification for sector *E* genuinely exist, while those of *F* are mere posits.

The goal of this paper is to argue for a version of holism concerning fact and value. It is to argue that judgments of fact and judgments of value are part of the same “web” in Quine’s sense.

---

<sup>1</sup> “Two dogmas” is reprinted in *From a Logical Point of View* (Cambridge: Harvard University Press, 1953). The fabric quote is from “Carnap and logical truth” reprinted in *Ways of Paradox* (Cambridge: Harvard University Press, 1976), p. 132.

What is this sense, exactly? Quine's first criterion of holism—the notion that sits behind the textile rhetoric—seems to turn on the confirmation relationships of judgments. Say that a judgment *a* is confirmationally relevant to a judgment *b* if the acceptance or rejection of *a* could rationally affect our acceptance or rejection of *b*—not merely cause the acceptance or rejection, but be a reason for it. (This definition is consistent with different conceptions of what it is to accept a judgment—believing it, being in a certain planning state, etc.) One important feature of confirmation relationships is that they are transitive. If *a* is confirmationally relevant to *b*, and *b* is confirmationally relevant to *c*, then *a* is confirmationally relevant to *c*. If *a* is not confirmationally relevant to *b*, then say that *b* is confirmationally isolated from *a*. If *a* is confirmationally relevant to *b*, and *b* confirmationally relevant to *a*, then say that *a* and *b* are in a *mutual confirmation relationship*. Suppose that for every pair of judgments *a* and *b* in a set *S*, *a* and *b* are in a mutual confirmation relationship. Then *S* is a *holistic web* of judgments. Quine's claim is that scientific theories are holistic webs, not, as he reads Carnap as saying, separated into autonomous sets of conventional and empirical elements.

Confirmation will also be my criterion, and my goal is to show that judgments of fact and value are members of the same holistic web.<sup>2</sup> If successful, such an argument would undermine one conception of the fact/value dichotomy. According to this conception, there is a body of “pure” practical or value judgments that are confirmationally isolated from all theoretical or fact judgments. (Assume for the sake of argument that we know how to pick out these judgments, and bracket questions about just which judgments they are.) There is also a corresponding body of purely theoretical or fact judgments that are confirmationally isolated from all practical judgments. If we assume that these two bodies of judgments are in confirmational isolation from the outside, but are themselves holistic, then we have two

---

<sup>2</sup> Importantly, my criterion is not the semantic holism of fact and value, the view that we cannot distinguish fact judgments from value judgments in respect of meaning (and so, presumably, it makes no sense to talk about “fact” and “value” judgments at all). The two theses are obviously related, and one can argue from confirmation holism to semantic holism with the addition of auxiliary premises—I address such a corollary in the final section of the paper—but the views are distinct. For more on the relationship between confirmation and semantic holism, see Jerry Fodor and Ernie Lepore, *Holism: A Shopper's Guide* (Oxford: Blackwell, 1992), especially the first chapter on Quine.

holistic webs, a value web and a fact web. While each of these theories may have downstream confirmational effects on some hybrid judgments—the conjunction of “I ought to avoid danger” and “radioactive cockroaches are dangerous” produces the downstream hybrid judgment, “I ought to avoid radioactive cockroaches”—the webs themselves are confirmationally isolated from each other because the confirmation effect that each has on “I ought to avoid radioactive cockroaches” is unidirectional. Thus no facts inform our “pure” value judgments, and no values inform our “pure” fact judgments. This is how we can avoid having a superficial intermingling of fact and value judgments vitiate the dichotomy too easily. I think this is one reasonable gloss on fact/value dichotomy—not the only one, certainly—and it is this dichotomy that my argument targets. It is also, I will argue later, an implicit presupposition for some central views in metaethics, and so the collapse of the dichotomy so understood makes the articulation of these views problematic.

My argument is also modeled on one of Quine's. Not the argument of “Two dogmas”, but on Quine's argument for holism about mathematics and natural science. This argument is a bit scattered and obscure, so I am going to take an interpretive liberty and simply assert that Quine's argument for this holistic thesis is an instance of the following schema. It is the instance for  $A$  = physical science,  $B$  = mathematics, and  $P$  = general empirical phenomena.<sup>3</sup>

1.  $A$ -judgments are indispensable to explaining phenomena  $P$ . (First indispensability premise.)
2.  $B$ -judgments are also indispensable to explaining phenomena  $P$ . (Second indispensability premise.)
3. These explanations are holistic: which  $A$ -judgments will contribute to the best overall explanation of  $P$  depends on which  $B$ -judgments we include in our explanation and vice versa. (Holism of explanation.)
4. A judgment's being included in the best overall explanation of a phenomenon is a reason to accept that judgment—it confirms it. (Inference to the best explanation.)

---

<sup>3</sup> A better elaboration and defense of this argument can be found in Mark Colyvan's *The Indispensability of Mathematics* (New York: Oxford University Press, 2001).

5. Therefore, *A* and *B* judgments stand in mutual confirmation relationships.  
(From the holism of explanation plus inference to the best explanation.)
6. Therefore, *A* and *B* constitute a single holistic web.

In this paper I present another instance of this argument schema. In mine *A* is natural science, understood relatively broadly to include the special sciences and much of “common sense”. *A* will therefore consist of descriptive, theoretical claims about the world. *B* will be what we could variously call *normative* or *practical* or *value* claims—claims that do not just describe the world but evaluate things and recommend actions. And *P* is the phenomenon of human action. The conclusion I reach is that *A* and *B* stand in mutual confirmation relationships and so constitute a single, holistic web of value and belief—a web of “valief”, if you will.<sup>4</sup>

## 1. The context of the argument

I have some antecedents in this campaign, and so before setting out on my argument, I owe a brief explanation of what I find wanting in their accounts and arguments, and thus why I think it is worth giving mine.

Reductionist naturalism entails confirmation holism about fact and value. If all value properties are identical to natural properties, which are presumably at home in our web of factual judgments, then by the substitution of identicals, value judgments are also part of that web. But the argument for this view works rather differently from mine insofar as it depends on the specification of reductions of value properties to particular natural properties through (e.g.) Railton’s “reforming

---

<sup>4</sup> This Miriam Solomon’s coinage. She is critical of the idea. See her “The web of valief: an assessment of feminist radical empiricism” in S. Crasnow and A. Superson (eds), *Out from the Shadows: Analytical Feminist Contributions to Traditional Philosophy* (New York: Oxford University Press, 2012).

naturalistic definitions of non-moral goodness and moral rightness” or Boyd’s synthetic identities.<sup>5</sup> The argument for holism, however, does not depend on such reductions. The argument is not that value properties are identical to recognized natural properties, but that our best criterion of naturalness—inclusion in our most explanatory web of fact judgments about the world—already countenances value properties. As a result, the present argument eludes some of the classic problems facing this reductionist naturalism.

On the other hand, Hilary Putnam offers an argument more in the spirit of the present paper, but I find it inconclusive. He begins with a genealogy of the fact/value dichotomy. Our intuitive understanding of the theoretical/practical or fact/value dichotomy, he argues, is inherited from Hume’s idea that facts must be intimately connected to sense experiences. More proximately it derives from the modern elaboration of the Humean idea by logical empiricists like Carnap. But Carnap’s program proved unworkable. Quine’s argument against the analytic/synthetic distinction showed that we could not maintain a rigid separation between the content of a theory and its background framework, and the underdetermination of theory by evidence shows that something intuitively “non-factual” will inevitably find its way into our theories.<sup>6</sup> Of course, these reflections do not show that the dichotomy is untenable, just that one motivation for it fails. We need something more for an argument against the dichotomy itself.

Here Putnam and others turn to examples of the “entanglement” of fact and value. Helen Longino, for example, argues that there is no value-free notion of “evidence” that is capable of performing the methodological tasks we put to it, so we have to rely on values in constructing theories. Nicholas Sturgeon offers more specific examples of explanations with moral terms in them: opposition to slavery arose earliest in the Anglo- and Francophone worlds because slavery in these areas

---

<sup>5</sup> Railton, “Moral realism” reprinted in his *Facts, Values, and Norms* (New York: Cambridge University Press, 2003), p. 32. Railton characterizes his own view as reductionist in “Toward fin de siècle ethics”, with Stephen Darwall and Allan Gibbard, *Philosophical Review* 101(1), 1992, p. 174. Boyd, “How to be a moral realist”, in Geoff Sayre-McCord (ed), *Essays in Moral Realism* (Ithaca: Cornell University Press, 1989).

<sup>6</sup> This is the argument of the first chapter of *The Collapse of the Fact/Value Dichotomy* (Cambridge: Harvard University Press, 2002).

was *worse* than in others. Elizabeth Anderson cites recent research into divorce. This research is more successful than its rivals because it presupposes different value judgments about the family, and the success of these new theories of divorce amounts to a kind of confirmation of those values. Putnam himself offers an argument with a similar structure, only his example is the value-laden economics of Amartya Sen. Catherine Elgin says that the dichotomy founders on the ubiquity of thick concepts in our projects. And Morton White argues that there is no good reason to exclude emotional reactions from the range of things contributing to the confirmation of a theory, and once we allow this, we can justify the inclusion of value judgments in the holistic web on the ground that theories including them enjoy better confirmation by this liberal standard than those without.<sup>7</sup>

I am sympathetic to all of these arguments, and I believe they lead us to more or less the same holism of fact and value that I favor. Nonetheless, I don't think they are fully convincing on their own. They offer what we might call a "low-road" to holism. They marshal examples to show that there is entanglement between theory and practice, fact and value, but they do not give deep or principled reasons as to why there *must* be this kind of entanglement. This makes them more easily grasped, and potentially more persuasive, but it also leaves them vulnerable to parrying. If we reconstruct these examples, critics may say, we find that the practical content we thought was important is dispensable gilding. For example, they may say that Sturgeon's "worse" is the most convenient way to pick out a state of affairs that is explanatorily relevant, but that this does not mean that normativity of the word itself explains anything—the moral assessment could be a nomological dangler. Similarly, they could dispute the gloss that Anderson and Putnam put on their

---

<sup>7</sup> Longino, *Science as Social Knowledge* (Princeton: Princeton University Press, 1990); Sturgeon, "Moral explanations", in G. Sayre-McCord (ed.) *Essays on Moral Realism* (Ithaca: Cornell University Press, 1988). Anderson, "Uses of value judgments in science: A general argument, with lessons from a case study of feminist research on divorce" *Hypatia* 19(1), 2004, pp. 1-24. Elgin, "The relativity of fact and the objectivity of value" reprinted in her *Between the Absolute and the Arbitrary* (Ithaca: Cornell University Press, 1997). Elgin is of course preceded in this attention to thick concepts by *inter alia* Philippa Foot and Bernard Williams. I mention her in particular because she connects such concepts to holism. White, *What Is and What Ought to Be Done* (New York: Oxford University Press, 1981) and *A Philosophy of Culture* (Princeton: Princeton University Press, 2002), ch. 10. In a spirit similar to White, see Owen Flanagan argument that Quine's rejection of value claims is unwarranted by his own lights in "Quinean ethics", *Ethics* 93(1), 1982, pp. 56-74.

examples, saying that the success of, e.g., the new research on divorce is due not to its countenancing different values but to some other innovative feature of that work. Or they could maintain that these value judgments are helpful auxiliaries for working scientists but not strictly speaking indispensable, and so should be paraphrased away in the final analysis. Critics may also note, *pace* Elgin, that thick concepts can change their valence or lose it altogether (as “chaste” and “pious” have) and use this fact to argue that such concepts can always be factored into a fact component and a value component. And, finally, they can simply reject a proposed liberalization of the confirmation base of theories of sort White proposes. Maybe Quine had reason to accept such a liberalization, I can imagine them saying, but that doesn’t mean we do.<sup>8</sup>

I do not visit these debates here. What I want to offer is a supplement to this work, what amounts to a “high road” argument for holism about fact and value: an argument showing that a pervasive kind of entanglement follows from the relationship between explanation and action. That is, I try to give a more abstract but also more principled reason for thinking that facts and values *must* be in entangled in the manner of confirmation holism, instead of offering up particular cases where they appear to be.

The thought animating this argument is that such confirmation holism is an inevitable consequence of how we are situated as explainers. Quine’s argument about mathematics is premised on the idea that what which judgments we should countenance depends on what we call on to explain the phenomena we encounter. But explanation is a situated undertaking. What counts as a good explanation for us depends on features of us—on our cognitive architecture, the beliefs, hypotheses, and background principles we have going into the project of explanation, the concepts we employ, the social contexts in which explanations are offered, even our practical aims. True enough, it is a goal of inquiry to make our explanations as unparochial as possible, to give explanations that would be valid for agents who

---

<sup>8</sup> Versions of these critiques are pursued *inter alia* by Solomon, “The web of valief” and Simon Blackburn, “Through thick and thin” reprinted in *Practical Tortoise Raising and Other Essays* (New York: Oxford University Press, 2011).

were not situated as we are. But we should not confuse this regulative ideal with the suggestion that the only good explanations are ones equally effective for all possible epistemic agents. Mathematics is indispensable to *our* scientific explanations, not necessarily to those that we could share with Laplace's Demon. Nor should we attempt, as the logical empiricists did, to distinguish *a priori* those features of explanation that are fully objective from those that are mere artifacts of our situation.

One kind of situatedness of explanation underwrites the holism of fact and value I argue for: *explainers of actions are also performers of actions*. Given that these two endeavors—performing actions and explaining them—concern exactly the same objects, our default assumption should be that these two projects will become bound up with each other, that thinking about how to act will affect our thinking about why agents acted as they did, and vice-versa. This is the inchoate idea motivating my argument.<sup>9</sup> Once again, just to have it before us, that argument goes like so:

1. Fact judgments are indispensable to explanations of action. (First indispensability premise.)
2. Value judgments are also indispensable to explanations of action. (Second indispensability premise.)
3. These explanations are holistic: which fact judgments will contribute to the best overall explanation of an action depends on which value judgments we include in our explanation and vice versa. (Holism of explanation.)
4. A judgment's being included in the best overall explanation of a phenomenon is a reason to accept that judgment—it confirms it. (Inference to the best explanation.)

---

<sup>9</sup> One final point about context. As the reader will see, the principal move in this argument depends on a Davidsonian dictum about interpretation. I believe that the sort of argument I give here was available to Davidson, but so far as I can tell he nowhere explicitly embraces the conclusion. The closest he comes is suggesting that values share the “objectivity” of beliefs by dint of an isomorphic role in interpretation. But this is different from holism: one could hold that facts and values are on all fours according to this particular notion of objectivity, but still confirmationally separate and therefore distinguished along other dimensions. See “The objectivity of value” in *Problems of Rationality* (New York: Oxford University Press, 2004).



5. Therefore, fact and value judgments stand in mutual confirmation relationships. (From the holism of explanation plus inference to the best explanation.)
6. Therefore, fact and value judgments constitute a single holistic web.

The following four sections introduce and defend each of the non-derivative premises in this argument.

## **2. Premise 1: the indispensability of fact judgments**

The first premise of my argument is that any satisfactory and complete explanation of an action will inevitably include some judgments of fact—or, slightly differently, that there will always be some judgment of fact that is relevant to a question of the form, “why did A  $\phi$ ?” Let’s work with an example. We have an agent; call him The Butler. One evening we see him rummaging through the papers on the desk of his master, Lord Musgrave. Why does this The Butler do this? And more philosophically, what sorts of considerations are relevant to explaining this action?

I won’t belabor the point. The Butler’s rummaging is a concrete event in a physical world ordered by causes and effects, so *of course* fact judgments about the character of that world and its causal order will enter into our best and fullest explanation of that rummaging. No conjunction of value judgments about what the Butler ought to do, what it would be sensible for him to do, or what it would be good for him to do will determine his *doing* a thing absent judgments linking these to concrete events in the world—absent fact judgments. So we will inevitably find such judgments in our explanation.

## **3. Premise 2: the indispensability of value judgments**

The second premise is that value judgments are indispensable to action explanation in the same way that any satisfactory and complete explanation of an

action will inevitably include some judgments of value. This is, by a wide margin, the more difficult and controversial of the two indispensability theses.

Here I think it is useful to think along the lines of Davidson's theory of interpretation. Davidson maintains that the interpretation of actors and actions is constrained by a methodological precept that has come to be called a principle of charity, expressed most dashingly thus:

In our need to make him [the object of explanation] make sense, we will try for a theory that finds him consistent, a believer of truths, and a lover of the good (all by our own lights, it goes without saying).<sup>10</sup>

If Davidson is right in asserting this as a kind of constitutive principle of interpretation, then premise 2 follows. For if we must interpret The Butler as a "lover of the good", then the fact that something is indeed good (by our own judgment, it goes without saying) will be part of the reason that we see The Butler's rummaging in one light rather than another—as one kind of thing rather than another. And so the relevant value judgment will be part of the comprehensive explanation we offer when putting The Butler's rummaging in this light. It is worth emphasizing how strong this claim is. A common thought against value-laden explanation is that while it may be that someone's *holding* a value judgment enters an explanation of her action, the value judgment *itself* does not.<sup>11</sup> But what follows from the charity principle is the second thing, the thing we need: *that* something is good or choice-worthy (by our lights) is part of our interpretation—not merely that the agent *finds it so*.<sup>12</sup>

---

<sup>10</sup> "Mental events" reprinted in *Essays on Actions and Events* (New York: Oxford University Press, 2001), p. 222. This is just one formulation of the principle of charity. See Ernie Lepore and Kirk Ludwig, *Davidson: Meaning, Truth, Language, and Reality* (New York: Oxford University Press, 2007), pp. 198ff for discussion of Davidson's different statements and defenses of charity.

<sup>11</sup> Gilbert Harman emphasizes this contrast in his argument against the explanatoriness of moral facts in *The Nature of Morality* (New York: Oxford University Press, 1977), ch. 1.

<sup>12</sup> One might argue that this does not show that value judgments are *part of* the explanation itself, but only that they are part of the background that makes a certain explanation appropriate. (One could make the same kind of point in reply to the claims of the previous section.) This may be a good objection were the aims of my argument different, but as it stands I don't think this distinction affects it. I am out to show that value judgments are confirmationally relevant to fact judgments

My argument for premise 2 will therefore be an argument that something like Davidson's charity principle is correct, that it is a stricture of the interpretation of agents that we try to render them, as far as possible, as acting for genuine reasons, responding to things of real value, and following valid rules. If this is right then our judgments that certain things are indeed reasons, valuable, or valid will be part of our explanation.

I can't claim to fully understand Davidson's argument for the principle of charity, so I will try to cobble one together on my own. It goes like this.

- 2.1 We must adopt some methodological precepts in our explanatory activities to overcome the underdetermination of explanation by our immediate evidence.
- 2.2 Insofar as we are trying to explain actions, these precepts should aim to render an event *intelligible as an action*.
- 2.3 What we find intelligible as an action will depend on our own theory of action.
- 2.4 This theory will include judgments made in the mode of practical reasoning—value judgments.
- 2.5 Therefore, our methodological precepts in explaining action will make reference to our own value judgments.
- 2.6 Therefore, such value claims are indispensable to the explanation of action.

We might have two slightly different questions about principles like Davidson's. First, why should we accept such methodological precepts at all; why should we see the explanation of action as constitutively aiming at any particular goal whatsoever? And, second, assuming that we do adopt such precepts, why should we accept ones that would have us explain agents as lovers of the good, rather than something else?

---

(and vice-versa), and "background conditions" on an explanation understood in this way are confirmationally relevant: they affect which judgments we adopt and reject, and so will be part of the larger holistic web. So whether or not certain judgments are part of the "explanation proper" or the "background conditions" doesn't matter.

The answer to the first question appeals to the need overcome the underdetermination of our explanations by the observations that are supposed to support them. For any body of evidence about someone's behavior there are innumerable competing explanations about that person and why they do what they do. If we are going to have any chance of understanding what an actor is up to, we need some principles for sorting through these rival explanations. For all we are told about The Butler in my little vignette, for example, we could postulate explanations centered on such claims as these:

- (1) The Butler is searching for papers containing a clue about where a treasure is hidden.
- (2) The Butler wants to go up in the aliens' spaceship and believes that rummaging is a way to signal them.
- (3) The Butler wants to know how many tokens of the letter "u" are in Musgrave's papers and is counting them in an unconventional way.

Strictly speaking, absent principles guiding our explanation, we could offer a story like this to explain The Butler's rummaging:

- (4) The Butler felt a bit ripe, and he thought a shower would fix him up nicely, so he decided to rummage through some papers.

These are all strictly compatible with the evidence we have about The Butler, as are many, many other stories. Some of them make The Butler look like rather an odd fellow, but that's exactly the point: we need some principles about the explanation of action that will motivate the demotion of these stories relative to more plausible ones.<sup>13</sup>

This answers the first question, but what of the second. Perhaps we need some *a priori* principles to overcome this underdetermination, but why Davidson's charity principle in particular? Why suppose that when we interpret The Butler we ought to prefer explanations that portray him as a "lover of the good", rather than a lover of the bad, the weird, or the Chicago Cubs?

---

<sup>13</sup> David Lewis also emphasizes the role of precepts like these in overcoming underdetermination. He goes so far as to say that we should accept enough precepts that we can reach a univocal explanation. Lewis accepts a version of charity that entails my premise 2. See his "Radical interpretation", reprinted in *Philosophical Papers Volume I* (New York: Oxford University Press, 1983).

When we are trying to explain something, we are trying to render it intelligible as the kind of thing it is, and this goal will be one source of principles of explanation like those we are interested. For example, when I want to explain why carbonic acid results from carbon dioxide being dissolved in water, I try to render this event intelligible as a chemical reaction. There are certain things it takes to be a chemical reaction, and these standards will guide my attempts at explanation. They will be methodological precepts of my explanatory endeavor. We may discover that some event simply doesn't fit the template and try to explain it in some other way, but so long as we are trying to explain it as a chemical reaction, we will be guided in that explanation by a conception of what it means to be a chemical reaction.

So it should also be with the explanation of action. To explain an action is to render it *intelligible as an action*. But what does it take to render something intelligible as an action? At this point we could assert some essentialist claim about action—that it aims at a particular thing or is performed under a certain guise. We could then argue that this feature should be reflected in our methodological precepts of interpretation. This isn't the way I want to go, though: any such claim would be controversial, and I don't think we need it anyway. We don't need it because each and every person already has a conception of what makes an action intelligible that is acquired from their understanding action, as it were, *from the inside*. Creatures like us have an understanding of what action is, and thus of what it means for something to be intelligible as an action because we are actors. (This is, in the terms I used above, part of how explanation is situated for creatures like us.) This understanding does not consist in knowledge of what is essential to or constitutive of action, but in a vast network of individual judgments made in the midst of our own practical reasoning. Given the aim of finding the treasure and the possibility of being hindered in doing this if one is found out, one ought to search for the treasure in secret. There is really no good reason to be interested in counting the number of tokens of "u" in a stack of papers. It would make no sense at all to go rummaging through papers because you have decided you need a shower. These are workaday value judgments, and they form part of our "theory" of action in the simple sense

that they are judgments *about* action. They are judgments about which actions are appropriate when, which actions make sense under which circumstances, which ends are sensible and senseless.

Such judgments naturally inform what we find intelligible in the way of action. To see why, return to the analogy with chemistry. When we are explaining the formation of carbonic acid, we are trying to render this event intelligible as a chemical reaction. But whence our conception of a chemical reaction? We don't come on it *a priori* as a clear and distinct axiom, or from the testimony of an oracle. It is an empirical standard that evolves alongside the rest of our chemical theory. As much as our conception of a chemical reaction guides our theorizing, *that conception* also depends the judgments we reach in such theorizing. That is, what we find intelligible as a chemical reaction is informed by the workaday judgments of chemistry. Indeed, this is precisely what we see in the history of the subject: the transformation of the concept of a chemical reaction in Lavoisier's critique of the phlogiston theory did not come out of nowhere. It was motivated by actual difficulties encountered in the laboratory, and in this way, the ensuing conception of chemical reaction was part of the same "web", to use the Quinean image, as these first-order chemical judgments.

The same should be true *mutatis mutandis* for what we find intelligible in the way of action. This conception will depend on the multitude of particular judgments we make about action. Only here this multitude will include those judgments we make *as actors*—value judgments about which actions make sense when, etc. The most obvious example of this influence involves the unintelligibility of actions premised on schemes of value that are simply too far out of sync with our own. This is why we find explanations like (3) and (4) inadequate. The idea that someone would see needing a shower as a reason to rummage through papers on a desk is so strange to us, so clearly nonsense that we cannot feel that we have rendered The Butler's rummaging intelligible by giving such an explanation. By the same token, though perhaps to a lesser degree, by attributing such a bizarre fixation to The Butler as discovering the number of "u" inscriptions in a stack of papers, we feel like we have not really made his doing this intelligible as an action—we judge that this

fixation *makes no sense*—and that’s why, other things being equal, we prefer a different explanation.

These examples illustrate the general thought of the argument. Particular value judgments made in the midst of our own practical deliberations are relevant to the explanation of action because they inform our conception of what is intelligible in the way of action (in the same way, I suggested, that individual judgments about chemistry will inform the chemist’s ideas about what is intelligible as a chemical reaction), and rendering an event intelligible as an action is a methodological precept that guides our explanation of action.<sup>14</sup>

Now I just said that judgments about action made amidst our own deliberations will naturally inform our conception of what is intelligible as an action, and will for this reason figure in our explanation of the actions of others. But this does not mean that it is strictly *impossible* to quarantine one’s own value judgments from the business of explaining others’ actions. And I don’t think it is impossible. But there is a substantial cost, I believe, and to complete the argument I want to give a sense of what that cost would be. One specific cost will be the inability to eliminate zany-yet-coherent alternative explanations of an agent’s action. For most explainers value judgments are an essential point of friction for overcoming the underdetermination of our theory of behavior by the evidence available. Thus in The Butler’s case we set aside potential explanations like (3) and (4), or at least demote them relative to others, because they do a poor job of making The Butler intelligible. We do this because they portray him as acting in blatant violation of fundamental norms of rationality or valuing the valueless. (Similarly, we set aside explanation (2) since it involves an uncharitable construal of The Butler’s beliefs.) If we did not rely on such value judgments, we could not do this, and the underdetermination we face in explaining action would be all the more acute. More generally, the costs of quarantine will approximate those of segregating the judgments of two different branches

---

<sup>14</sup> Besides Davidson and Lewis, similar sentiments about the normative aspects of rendering an action intelligible are articulated by Anscombe, *Intention* (Cambridge: Harvard University Press, 1957/2000), §§5-11, John McDowell, “Functionalism and anomalous monism” reprinted in his *Mind, Value, and Reality* (Cambridge: Harvard University Press, 1998), especially on p. 389, Thomas Nagel, *The View From Nowhere* (New York: Oxford University Press, 1986), pp. 142-3, and Jonathan Dancy, *Practical Reality* (New York: Oxford University Press, 2000), pp. 94-97.

of science, for example refusing to employ what we know about chemistry in human physiology or mathematics in physics. Physiology would not grind to a halt without chemistry. Our ancestors had a passable physiology without anything resembling modern chemistry. Yet it is undeniable that the additional leverage in our explanations provided by acknowledging that our bodies are made up of the same stuff that chemists study has paid dividends. For example, the manifest phenomena of inheritance are compatible numerous different theories of how traits are transmitted from parent to child. But only some of these are realizable by the substances that make up our bodies. Fewer still are compatible with what we know about nucleic acids in particular. What chemistry offers physiologists, then, is an increase in the points of theoretical friction a theory faces across the board. I have argued that relying on value judgments in the explanation of action offers the same benefit, and refusing to do so would entail analogous costs.

Before moving to the remaining premises of the argument, I want to say a few in words in response to some natural objections to this argument.

*First objection.* Are value judgments the only thing that can do this work for us? Mightn't we set these explanations aside on other grounds, for instance that they are statistically unlikely—that people tend not to care about number of “u”s? There are two problems with this. First, we can cook up examples where the statistics will not help, cases where the explanation that renders the agent most intelligible is one that attributes to him a statistically unusual motive. Second, it's not clear that assimilating an action to a statistical regularity is enough to make it intelligible as an action. The statistical regularity itself that seems to call out for some further explanation. Noting that 93% of cinereous powders create blue flames when ignited does not make that ignition intelligible as a chemical reaction; it just makes it statistically likely. By the same token noting that 93% of people behave thus in these circumstances does not necessarily make that behavior *intelligible as an action*.

*Second objection.* Does this argument entail (falsely) that it is impossible to explain other agents as acting contrary to our values? It doesn't because the methodological goal of charity must be balanced against other constraints on the explanation of action: capturing the evidence, charity principles concerning true belief,



and general explanatory desiderata like simplicity, fecundity, and the avoidance of ad hoc stipulations. So the argument does not entail that we cannot depict people as doing wicked things for crummy reasons; sometimes the weight of evidence will mandate such explanations. What the argument does entail is that we should attribute errors in valuing as shallow and localized as possible consistent with these other aims. And this seems appropriate. We will better understand the racist, for example, if we see him as sharing many of our values—tradition, maybe—but being blind to some, misunderstanding the contours of others, and mistaking the weights of still others, than we do if portray him as simply Satanic. Thus even in cases where we are explaining someone as acting on errant values, the charity principle plays a role in shaping (more) intelligible explanations.

*Third objection.* I have been speaking quite generally of value judgments and their contributions to rendering an action intelligible. But one might worry that in doing this I am eliding over important distinctions within the web of value judgments. For example, one might object that it is sufficient for rendering an event intelligible as an action to show why an agent undertook that action *given* her beliefs, desires, intentions, and so on. So it is both unnecessary and distorting to project our own values on a subject when all we need secure is some manner of procedural rationality. The first thing to notice about this objection is that it doesn't deny the indispensability of value judgments to explanation altogether. Rather, it suggests that only a certain subtype of value judgment is indispensable, namely judgments of procedural rationality about what would make sense given certain attitudes. This makes the consequences of the objection for my argument unclear.

But there is a deeper problem too. The suggestion that charity only extends to rendering an agent procedurally rational depends on a conception of interpretation that we should reject. It assumes that we first ascribe intentional attitudes and so have those ascriptions in hand when trying to make the agent's behavior intelligible as an action. But this is not how interpretation works. The explanation of action and the ascription of intentional attitudes are part of the very same interpretative project. The predicament of the action explainer is never one of making sense of why someone did something *given* a comprehensive picture of his psychology

because that very picture is worked out in a tapestry of overlapping explanations of the actor's behavior. It is true that for the explanation of individual actions we might treat certain intentional attitudes as fixed points, but this merely amounts to the ascription being more entrenched relative to other claims. Moreover, there are several types of case that give the lie to the notion that charity demands we always preserve procedural rationality, even if it means attributing eccentric values to an agent. These are the familiar cases where it is *appropriate* to interpret a person as suffering from weakness of will, compulsion, addiction, or some other failure of rationality.

#### **4. Premise 3: holism about explanation**

I have so far argued that both fact judgments and value judgments are indispensable to the project of explaining action. This does not yet give us fact/value holism, for we have not shown that these explanations are themselves holistic. We need to show that there is a confirmational give and take between fact and value in these explanations, that explaining action sometimes involves adjusting what we think on one front in light of what we say on the other.

There is a natural way to maintain a segregation of fact and value despite the two indispensability premises. Value judgments contribute to the explanation of one aspect of our actions: their intentional aspect. They are done for reasons, as part of a plan, and so on. Fact judgments contribute to a different aspect, the physical. They involve the moving of various bodies in space and time. But this only shows that we must explain these two things in parallel, not necessarily in concert. We can use value judgments to produce an explanation of the intentional aspects of an action, and fact judgments to produce an explanation of the physical aspects, but there is no reason that what goes on in one of these explanations would affect what goes on in the other. The total explanation of an action could be no more than a conjunction of two separate and independent sub-explanations.

The picture of action that this possibility relies on is incompatible with a very basic postulate, what I will call the unity of action. An action is something that

is at once intentional and physical, not something intentional *adjoined* to something physical. That is, an action is not just an attempt coinciding with a movement, but an insoluble unity of these two things. This is why it does not make sense to ask what is left over when we subtract the intention from the movement.<sup>15</sup> An explanation of an action must capture this unity; it must make sense of *this* deliberation corresponding to *that* movement.

Ordinary action explanation accomplishes this by going back and forth between intentional explanation and physical explanation with the goal of reaching equilibrium between the two. Here are two stylized examples.

Suppose we observe The Butler kick Lord Musgrave in the shin. Our first thought will be that The Butler kicked Musgrave to hurt him. But we know antecedently that The Butler did not bear any ill will toward Musgrave, so The Butler's kicking Musgrave to hurt him wouldn't make sense. The rejection of our initial hypothesis on the grounds that it fails as an intentional explanation then leads us to an alternative hypothesis about the physical features of The Butler's kick. We might think that it was caused by his patellar reflex, and so was quite involuntary. But we check The Butler's knee with our best scientific methods and find that it was not. So we offer another hypothesis: perhaps The Butler kicked Musgrave out of play. We take out our value theory once more and hypothesize that this would make sense if he wanted to flirt with Musgrave. Life in the countryside can be boring after all. But we return to our theoretical theory and find that The Butler's kick was far too ferocious to be flirtation. (Another value-laden conclusion.) So we start over again and hypothesize that The Butler meant to disable Musgrave in order to escape somewhere. But this raises further practical questions. And so on.

Now suppose we observe The Butler sneaking into Musgrave's study. We are poised to explain this by saying that he wanted to rob Musgrave to pay off his gambling debts. This does make some sense given what we know about The Butler. But this is hard to square with some of the physical aspects of The Butler's action, like his going for old family documents instead of money. So we go back and revise

---

<sup>15</sup> As Robert Jaeger argues in "Action and subtraction", *Philosophical Review* 82(3), 1973, pp. 320-9.

our hypothesis: The Butler snoops through Musgrave's papers out of a prurient fascination with the lineage of well-bred men. But when we check this hypothesis against the facts gleaned from our theoretical methods we are disappointed, for we find that The Butler's skin conductance doesn't increase while he paws through these documents—usually a sign of arousal. So we ask once more when it would make sense for The Butler to snoop through Musgrave's documents in absence of any intrinsic interest in them. We then come up with the thought that he hopes to find clues to the location of the treasure. And so on.

In both of these examples our explanation oscillates between fact and value judgments, the intentional side of action and the physical, trying to reach a stable equilibrium between the two, a state where the intentional and physical pictures we paint are mutually supportive. We offer an intentional account, revise it in the face of some recalcitrant fact judgment, find this revised judgment wanting by the lights of our intentional account (which is informed by our value judgments) and so revise it again. We repeat this process, going back and forth, until our intentional and physical explanations are brought into equilibrium.

What these examples illustrate is that in ordinary action explanation we capture the unity of action by explaining holistically: by adjusting fact judgments in light of value judgments and vice versa until we reach a comprehensive equilibrium between both sorts of judgment that renders the event intelligible as a *unified action*. Theoretically there may be other ways of capturing this unity, but these are an odd lot, and it is enough to establish premise 3 that in fact our action explanations capture the unity of intentional and physical aspects of action through this holistic interaction of judgments of fact and value.

## **5. Premise 4: inference to the best explanation**

Just because certain judgments are used side by side in holistic explanations does not mean those judgments are in a mutual confirmation relationship. That is, holism of explanation does not necessarily entail holism of confirmation. But it is the latter thesis we are after. The gap between what we have and where we want to

go is a single premise, that inference to the best explanation (IBE) is a valid form of confirmation. I will simply assume that IBE is valid for fact judgments. This claim has its detractors, but these days I take it to be relatively uncontroversial. The harder case is the practical sphere. Why suppose that a value judgment's figuring in our best explanation is any reason at all to accept it in a practical context? That is, why should the explanatory utility of a claim like " $p$  is a reason to  $\phi$ " or " $p$  is good" or " $\phi$ -ing is reasonable" have any effect on my taking  $p$  to be a reason to  $\phi$ ,  $p$  to be good, or  $\phi$ -ing to be reasonable *while deciding what to do*?

To answer this question, we should get clear about the alternative. No one can seriously deny that giving a good explanation of The Butler that includes the fact that " $p$  is a reason to  $\phi$ " has *some* confirmational implications: my agreeing that  $p$  is a reason to  $\phi$  when explaining The Butler should make me more willing to cite that judgment when explaining what The Maid or The Chamberlain do. So the explanatoriness of the judgment *does* confirm the judgment " $p$  is a reason to  $\phi$ " at least in this sense—it at least lends it "theoretical" confirmation. What a person *can* deny, however, is that this judgment, the one used in explanation, has anything at all to do with his deliberations about whether to  $\phi$ . So someone who wants to resist this inference wants to keep two books. For him there are *theoretical* value judgments: judgments about what one ought to do, what action makes sense, etc. that he puts to the theoretical purposes I have outlined—the categorization of mental states and the explanation of behavior. And then there are *practical* value judgments: judgments about what one ought to do, what is good, what action makes sense, etc. that we make while deliberating and thereby guide our actions. In a more tendentious tone, such a person might dub the first category as *pseudo*-value judgments. They look value-laden because they use words like "reason", but their function and the logic that governs their acceptance is entirely classificatory and explanatory. If this view is correct, then the inference I am recommending would be

a *pun* on words like “reason” and “good”, for it would mean slipping from a theoretical sense useful categorizing mental states and explaining behavior to a practical sense that has a direct bearing on our deliberations (or back in the other direction).<sup>16</sup>

So this sort of double book-keeping is the alternative to accepting my final premise. I think that we should reject it for a few reasons. According to the position, the judgments about reasons and goods that I make in deliberating and acting, and the judgments about reasons and goods that I cite in explanations of others are in fact different judgments belonging to different spheres. One justifies or recommends, the other explains. It follows that when I say that The Butler acts for a reason or in recognition of some good, the “reason” and “good” I invoke in this explanation do no refer to the same kind of entity as those I encounter in my own deliberations—when I think that its being The Butler’s birthday is a good reason for me to give him a present or that it would be good for me to help The Butler with his financial problems. But now assume that we can offer adequate explanations of The Butler’s action, explanations that cite the real considerations for which he acts. If this is the case, then it follows that The Butler cannot really *act for* reasons or the sake of goods, at least not for the full-bodied “practical” reasons and goods that I act for by when I respond to them in my deliberations, since these cannot appear in our explanations of action. This is what Nagel calls practical solipsism: the view that the sole domain of genuine, normative reasons is my own deliberations.<sup>17</sup> Any view that entails this arrangement should be rejected.

In fact, the consequences will be even worse, since we can always take up the explanatory stance toward ourselves. If the double book-keeping strategy is appropriate, then I must view my own actions as explained by “pseudo”-value judgments, just as I would anyone else’s. And this creates an intolerable alienation. When deliberating I make genuinely practical value judgments about what it makes

---

<sup>16</sup> Here compare Timothy Schroeder’s argument that Davidson’s theory of mental content is, despite appearances, non-normative because even though it uses normative terms like “rational” in its categorization scheme, the fact that these terms have a certain normative *force* plays make no contribution to this scheme. “Donald Davidson’s theory of mind is non-normative”, *Philosophers’ Imprint* 3(1), 2003, pp. 1-14.

<sup>17</sup> *The Possibility of Altruism* (New York: Oxford University Press, 1970), p. 114.

sense for me to do, but a moment later, if I stand back from myself and try to understand what I did, I will cite “pseudo”-practical versions of the same value judgments as the explanation. I will see myself not as acting for the reasons I recognized in deliberation, but for something else, perhaps described with similar language, but a very different kind of entity. Thus the reasons I must *see myself acting on* (in my explanation book) and the reasons that I *act upon* (in my deliberation book) will be of completely different species.<sup>18</sup>

My argument for the final premise of my argument, the validity of IBE, goes like this. Given the general validity of IBE and the foregoing premises of my argument, if we want to resist the idea that explanatoriness has confirmational implications for practical judgments we need to maintain that value judgments are in fact equivocal between two different sorts of judgment: a genuinely practical sense tied to practical reason, and a theoretical sense tied to explanation and classification. But this “double book-keeping” pushes us into a corner we should not want to crouch in: practical solipsism and a deep alienation from our own actions. This is enough to reject the double book-keeping strategy and accept that a value judgment’s explanatory usefulness has some effect on its “practical confirmation”: its aptness for being put forward in deliberation. I emphasize “some” because all I have established is a weak *pro tanto* relationship: for each of us the explanatory project and deliberation are not completely insulated from each other, and so the explanatory suitability of a value judgment gives us some *pro tanto* reason to accept that value judgment in the mode of practical reason.

## **6. Holism as a metaethical view**

This completes my defense of each of the premises in the master argument for the thesis that fact and value judgments stand in mutual confirmation relationships. Before closing, I want to attend to the metaethical consequences of this con-

---

<sup>18</sup> Dancy also emphasizes the implications for the explanation of action of the axiom that a person must see herself as acting on genuine, normative reasons. See *Practical Reality*, pp. 98ff.

clusion. There are two types of consequences: ones that we reach by joining confirmation holism about fact and value with further claims, and those that follow directly unaided. I'll begin with the first.

If we accept Quine's brand of methodological naturalism about the questions of metaphysics—that to be is to be the value of a bound variable in a successful scientific theory—then we get the conclusion that those values we commit ourselves to in our explanations are just as real as electrons and whales and all the other baubles of science. That is, we get a realism about value of the same variety as Quine's realism about mathematics. And importantly we get it without the problems that attend other arguments for ethical naturalism. For example, because this naturalism is not premised on the synthetic identity of normative and natural properties, we do not have to defend any such identity claim in the face of examples like those found on moral twin earth. And because we have not hitched our account to particular explanatory claims—the explanatory power of Passed Midshipman Selim Woodworth being no damned good—we are not vulnerable to having those examples parried away or redescribed in the fashion I suggested earlier. Instead we have a general argument that there *must* be examples of this type because of the general predicament of action explanation. (Of course, in the style of “high” and “low” road defenses, I think my case and examples like Woodworth's being no damn good are best pressed in concert.)

Second, if we accept a view that grounds the semantic content of judgments in the role they play in our theoretical projects (viz. verificationism, some varieties of functional role semantics) then we also get a kind of *semantic* or *meaning* holism of fact and value. That is, we get the conclusion that the semantic content of fact judgments depends on the content of value judgments in such a way that a change in the former will result in a change in the latter (and vice-versa). This result would seem to undermine the most natural way of drawing the fact/value dichotomy at the level of meaning. The meanings of our putative value terms would be just as entangled with those of putative fact terms as the judgments containing them are with respect to confirmation. For my argument this would have something like the effect of kicking away the ladder. I began by accepting a prima facie distinction between



fact judgments and value judgments and argued that these judgments occupy the same holistic web of confirmation. But if semantic holism about fact and value is true, then even this *prima facie* distinction is suspect. Naturally, we could still talk informally of a rough division between value and fact talk, but this division would not be grounded in anything principled enough to do philosophical work for us.

These are the two corollaries of fact/value holism that follow with the addition of further (controversial) theses. As for consequences that we can reach without such aid, these all take the form of complicating the arguments for and articulations of particular metaethical positions. I will mention three: constructivism, non-cognitivism, and a brand of realism.

Metaethical constructivism is often articulated in terms that rely on a version of the fact/value dichotomy. For example, Christine Korsgaard writes:

The moral realist thinks of practical philosophy as an essentially theoretical subject. Its business is to find, or anyway to argue that we can find, some sort of ethical knowledge that we can apply in action. According to constructivism, the only piece of knowledge that could be relevant here is knowledge that the problems represented by our normative terms are solvable, and the only way we can find out whether that is so is by trying to solve them. So for the constructivist practical philosophy is a practical subject.<sup>19</sup>

Similarly, Sharon Street understands metaethical constructivism as the view that “normative truth consists in what is entailed from within the practical point of view.”<sup>20</sup> For both Street and Korsgaard, then, the distinctness of constructivism as a metaethical view appears to rely on some antecedent distinction between theoretical and practical. So if there is nothing distinctive about the practical “point of view” *as against the theoretical*, then there is nothing special about normative judgments.

---

<sup>19</sup> “Realism and constructivism in Twentieth Century moral philosophy” in *The Constitution of Agency* (New York: Oxford University Press, 2008), p. 325.

<sup>20</sup> “What is constructivism in ethics and metaethics?”, *Philosophy Compass* 5(5), 2010, p. 367.

Korsgaard is explicit about this dichotomy in points of view. “When we look at our actions from the theoretical standpoint,” she says at one point, “our concern is with their explanation and prediction. When we view them from the practical standpoint our concern is with their justification and choice.”<sup>21</sup> But here is where holism proves problematic for this project. If my argument is correct, then the web of judgments associated with explanation and prediction is the same as the web of judgments associated with choice and justification. So we cannot distinguish Korsgaard’s two points of view in the obvious way—by associating them with two different bodies of judgments performing two different sorts of function. Insofar as the characterization of constructivism depends on this distinction, then, the doctrine is in danger of being ill-founded. Of course we can locate a *rough* distinction between two types of judgment and two points of view, but the metaethical work that constructivists need that distinction for demands a bright line—there is supposed to be a fundamental difference in the status of fact and value judgments—and that is what holism denies us. This does not mean that holism entails that constructivism is false or incoherent, just that a natural way of formulating the view is at odds with the thesis.

If ethical non-cognitivism is the view that one class of statements, value-laden ones, have a different content than those of another class, purely factual ones, then holism may pose a similar problem for that view. How are we to individuate the classes of statements that get this separate treatment in a way that can justify that treatment? In some contexts at least, Allan Gibbard gives an answer that invokes the fact/value dichotomy in the form targeted by my argument. “G. E. Moore argued that oughts don’t form a part of the natural world that empirical science can study. The picture I have sketched has the upshot that Moore was right. The scientific picture tells us why organisms like us would have questions whose answers can’t be made a part of science.” The scientific picture can’t answer these questions because “questions of what I ought to do and what it would be wrong to do or not

---

<sup>21</sup> “Personal identity and the unity of agency: a Kantian response to Parfit” in *Creating the Kingdom of Ends* (New York: Cambridge University Press, 1996), pp. 377-8.

to do aren't questions amenable to science. They are, I have been saying, questions of whether to help and of how to feel about not helping.”<sup>22</sup> So Gibbard thinks there are two classes of questions and judgments answering them. The most natural way to understand these classes is as two separate “theories” of connected judgments, one aimed at answering questions about how the world is and one aimed at answering questions about what to do. And the natural way of individuating these theories, I think, is in terms of the confirmation relationships of those judgments constituting them. Science is one body of confirmationally connected claims aimed at explaining the natural world. Ethics is the body of claims about what one ought to do. If the argument for holism is correct, however, then this separation is not possible. There are not two classes of judgments, but just one. Once again, a rough division is possible, but not one that can support the semantic dichotomy that non-cognitivism would like to offer. Just as was the case with constructivism, we don't have a conflict between holism and non-cognitivism itself, but rather between holism and a congenial way of understanding the demarcation of judgments that non-cognitivism relies on.<sup>23</sup>

Holism has more direct relevance to T. M. Scanlon's realism. Scanlon argues that naturalistic doubts about the existence of normative entities like reasons are misplaced because natural science and normativity make up separate, largely independent “domains”. Questions about existence are domain-specific: whether we have a reason to do something is settled not by the canons of natural science, but by the principles specific to the normative domain. This view, Scanlon himself notes, has “obvious similarities” with Carnap's view, the one Quine undertook attack in “Two dogmas”.<sup>24</sup> Put succinctly fact/value holism is to Scanlon's division of domains as Quine's general holism was to Carnap's multiplication of frameworks. If fact/value holism is true, then the individuation of domains—a normative domain and several non-normative domains—that Scanlon's metaphysical position

---

<sup>22</sup> *Reconciling Our Aims* (New York: Oxford University Press, 2011), p. 17.

<sup>23</sup> If, however, we accept the semantic holism of fact and value, the conflict with non-cognitivism (and constructivism for that matter) would appear to be more direct.

<sup>24</sup> *Being Realistic about Reasons* (New York: Oxford University Press, 2014), pp. 16-30. Scanlon's self-comparison to Carnap is in footnote 2.

is premised on cannot be drawn in what seems to be the natural way—that is, in a way based on the confirmation relationships of the relevant judgments. There are no domain-specific existence questions because the domains in question cannot be adequately individuated.

This conclusion is a little too quick, though, as Scanlon parts ways with Carnap on one important issue. He claims that his domains are not entirely autonomous. “Even pure statements in one domain can entail or presuppose claims in some other domain, and when this happens these claims need to be reconciled, and some of them modified or given up.” For example, we might have a domain of claims about witches and spirits, some of which entail propositions about the conditions under which people get sick or cows stop giving milk. “These claims conflict with claims of physics and other empirical sciences,” Scanlon says, “and this conflict provides decisive reason to reject the idea that there are witches and spirits.” In general, Scanlon says, “there can be meaningful ‘external’ questions about the adequacy of the reasoning in a domain, and about the truth of statements, including existential statements, that these modes of reasoning support.”

At first blush this concession would appear to endanger Scanlon’s whole project, since we can imagine someone like Mackie insisting that the claim the proposition that I have a reason to take my medicine has entailments and presuppositions that conflict with the claims of natural science in precisely the way that the claim that witches make people sick does. It can appear that by sacrificing the autonomy of domains, by admitting that there are legitimate “external” questions arbitrated by other domains, Scanlon may in effect be admitting that the answers to existence questions do not turn on the specific demands of their native domain, but on overall demands of all domains taken at once. Scanlon’s reply is that the domain of witches and ghosts is actually rather atypical in its entailing claims that really do conflict with physics. And the domain he is most interested in, the normative domain, is not entangled in this way. He explains:

Some claims about gods may involve claims of this kind, such as claims about the creation of the universe, or about what is happening when lightning occurs. An interpretation of moral claims, or claims

about reasons for action that takes these claims to be straightforwardly true would be “incompatible with a scientific view of the world” if these normative claims entailed, or their supposed significance presupposed, claims about the natural world that science gives us good reasons to reject. But pure normative statements do not involve or presuppose such claims.<sup>25</sup>

Here is where my argument for holism comes into conflict with Scanlon. I have argued that normative claims *do* entail—not in the narrow sense of logically entailing, but in the wider, more interesting sense of confirming—claims in the domain of natural science. Claims about what reasons an agent has, or what it would be rational to do are part of holistic explanations of the behavior of that agent, explanations that include perfectly humdrum empirical claims. Because these explanations are holistic, and inference the best explanation is generally valid, they furnish us with examples of mutual confirmation relationships between fact and value judgments. This is by no means a prelude to throwing out reasons statements because they fit poorly with our “scientific picture of the world”. On the contrary, it is an argument that these statements have a place in that picture for the same reason everything else does: because of the explanatory work they do for us.

Scanlon doesn’t address this kind of argument. He says at one point later in the chapter that he does not think that psychology is normative in a sense that imperils the fact/value dichotomy, but elects not to argue for such a position.<sup>26</sup> Without a rebuttal of this kind of argument, however, Scanlon seems to be in the same position as Carnap after “Two dogmas”. He can say that there are separate, semi-autonomous domains, and that existence claims are settled in a domain-specific manner, but it is far from obvious how to carry out the separation.

## 7. Conclusion

---

<sup>25</sup> pp. 21-2.

<sup>26</sup> p. 36.

I have argued for confirmation holism about fact and value judgments. I am not the first advocate for this thesis. What I take my argument to add is a “high road” complement to more familiar “low road” cases for such holism. The latter arguments proffer examples of the “entanglement” of fact and value judgments, usually in the form of ways that value judgments can be explanatorily useful. I find many of these examples convincing, but like all example-driven arguments they are vulnerable to parrying: maybe it’s not the value judgment itself that explanatorily useful, but something else; the relevant value judgments may be useful, but they are not really indispensable; a given judgment isn’t really a value judgment in the appropriate sense; etc.

What I have offered is a more abstract and principled argument as to why fact and value not only *are* entangled in the ways suggested by these examples, but why they *must be* holistically unified. Because of how we are situated as explainers—because we are explainers who also act—there is a presumption in favor the relevance of our own value judgments to the business of explaining action. In each step of the argument we see that a different way of overcoming this presumption—denying the indispensability of fact or value judgments, denying the holism of explanation, denying the validity of IBE—leads to untenable consequences. As a result we should accept this presumption and agree that fact and value judgments ultimately inhabit a single web of “valief”. This argument works best, perhaps, as a supplement to “low road” arguments, for it gives us reason to think that the examples marshaled by these arguments are not solitary, exceptional, or anomalous, but rather reflect a pervasive and ineliminable feature of our situation as actors and explainers that manifests itself in the holism and fact and value.