

Norms of Intentionality: Norms that Don't Guide

Abstract

More than ever, it is in vogue to argue that no norms either play a role or directly follow from the theory of mental content. In this paper, I present an intuitive theory of intentionality (including a theory of mental content) on which norms are constitutive of the intentional properties of attitude and content in order to show that this trend is misguided. Although this theory of intentionality—the teleological theory of intentional representation—does involve a commitment to representational norms, these norms are not problematic in the way critics have suggested they would be. In particular, these norms do not guide thinking by motivating intentional agents to (intentionally) accord with them; as a result, no obvious vicious regress threatens the theory. In the final section of this paper, I argue that accepting this teleological theory of intentionality need not commit one to thinking that intentionality is the product of natural selection.

Introduction

Is mental content inherently normative? Do mental states have the particular content they do because particular norms apply? What was, in the wake of Kripke's interpretation of Wittgenstein, once the heterodox view—that content is *not* inherently normative—may no longer be so.¹ In the last decade, a series of papers have contested the thesis that intentional properties (and corresponding linguistic properties) are, in some sense, inherently normative.² In this paper, I will show that the critiques advanced against this thesis do not apply to a very plausible theory that embraces it.

The thesis that content is inherently normative (CIN) is the thesis that norms either play a role in or follow immediately from the theory of mental content. In their recent article “Against Content Normativity,” Kathrin Glüer and Åsa Wikforss ultimately base their argument against CIN on the premise that any norms inherent to (bearing) content must be *regulative* in nature, i.e. norms that “guide our performances.”³ This premise is commonly accepted among advocates and opponents of CIN alike,⁴ but as Glüer and Wikforss (and others) show, the thought that norms inherent to content must be norms we intentionally use to guide our cognition leads to irresolvable difficulties for the view that content is inherently normative. After all, *intentionally* following these norms requires

¹ Kripke (1982). For examples of different defenders of CIN, see Dummett (1959), Boghossian (1989), (2005), Millikan (1990), Wright (1992), Tanney (1999), Brandom (2000), and Wedgwood (2007).

² See, for instance, Horwich (1995), (1998), (2005), (2006), Papineau (1999), Dretske (2000a), Hattiangadi (2006), and Glüer and Wikforss (2009).

³ Glüer and Wikforss (2009), p. 31.

⁴ Cf. Boghossian (2005) and Hattiangadi (2006). The critiques of CIN may be due at least partly to advocates only partially understanding the proper motivations of their own view.

that we antecedently have intentional states with content, and this requirement leads very quickly to a vicious regress.

Even without rehearsing the details of their arguments, it is plausible enough that the starting point of Glüer and Wikforss is not promising for CIN. Fortunately for advocates of CIN, they have no reason to accept this starting point. Indeed, I would suggest the primary motivation for CIN ought to press someone to *reject* the premise that the norms inherent to content guide performance in the relevant sense.⁵ Consequently, cogent arguments based on this premise are largely irrelevant to deciding whether content is inherently normative or not.

In Sections 1-6 of this paper, I will develop a view inconsistent with the Glüer and Wikforss premise that motivates CIN, and use that view to address some of the major critiques raised against CIN. In Section 7, I will show that the view I have developed is not obviously subject to any sort of vicious regress.

1 Why think that content is inherently normative?

CIN is properly motivated by three other ideas. These three ideas generate a very intuitive picture of mental representation that I will attempt to elucidate in this section. As my goal is merely to show that opponents of CIN have not addressed this picture, I will not attempt to argue affirmatively for it. Nonetheless, although I will not try to argue the point extensively in this paper, I do believe that this picture best fits many of our folk conceptions of intentionality, and, consequently, ought to be considered the default view even if it ultimately is not accepted.⁶

The first of these ideas is that the representation of artifacts that exhibit “aboutness,” e.g. maps, blueprints, etc., stems from having a representational *purpose* or *teleological function* (henceforth “telos”). The particular teloi of maps, blueprints, etc. fix that these artifacts are representational. They distinguish maps, blueprints, etc. from other non-representational artifacts, e.g. chairs and spoons. Thus, something is a map of Eugene, Oregon rather than a chair because it has the telos of bearing a certain structural resemblance to features of Eugene, Oregon rather than

⁵ I assume that “norms guiding performance” in the sense relevant to the Glüer-Wikforss premise requires that an agent intentionally follow the norms, i.e. that an agent follows the rule by using some sort of practical reasoning. In another sense of “norms guiding performance,” the performance of a well-designed coke machine is guided by the norms of its design; it does what it’s supposed to do (i.e. it distributes coke when the appropriate amount of money is inserted) because it was designed to do so. As far as I can see, the arguments Glüer and Wikforss deploy do not establish that norms inherent to content cannot “guide performance” in this sense, so either they have overlooked this sort of performance guidance (which occurs when something fulfills its telos because it has that telos), or the Glüer-Wikforss premise asserts something stronger. The most charitable interpretation, I think, is to assume the latter.

⁶ I argue for these points in works in progress. Because I think this picture of intentionality best fits our folk conception, I think it would be more charitable to interpret Kripke (1982) with this picture rather than the “regulative” version of the CIN thesis.

the telos of supporting a seated person. Teloi also distinguish maps, blueprints, etc. from each other according to differences in representational aspect. Thus, something is a map of Eugene, Oregon rather than a plan for development of Eugene, Oregon because it has the telos of matching up with Eugene rather than the telos of being matched up with by Eugene. Likewise, something is a map of Eugene rather than a map of Portland because it has the telos of matching up with Eugene rather than the telos of matching up with Portland. In this way, when it comes to maps, blueprints, and other representational artifacts, having some particular representational content—being about something or other—is a matter of having a particular telos.

The second idea is that intentional states are of a kin with representational artifacts in that they too have their representational properties in virtue of having particular teloi. To think that beliefs, desires, and other intentional states are of a kin with maps, blueprints, and other representational artifacts is to think that the representational features of intentional states arise in the roughly the same way as the representational features of artifacts.⁷ Beliefs and other doxastic states are evidently similarly to maps, paintings, and various detectors that represent by (metaphorically speaking) *reflecting* the way things are. Desires and other conative attitudes are likewise similar to objects like blueprints and designs that represent by (metaphorically speaking) *projecting* a way for things to be. According to this second idea, these are more than just superficial similarities. Just as the difference between reflecting and projecting for artifacts is due to differences in telos, so it is with intentional states. Likewise, just as maps, blueprints, etc. are about things because they have particular teloi, so too beliefs, desires, etc. are about things because they likewise have analogous teloi. In both cases, differences in representational content also stem from differences in telos.

The final idea is the observation that having a particular telos is a matter of being such that particular norms apply—the norms that are constitutive of having that telos.⁸ (I defend this claim in more detail in Sections 3 and 4). Just as maps, blueprints, etc. have particular teloi because particular norms apply, so beliefs, desires, etc. also have particular teloi because particular norms apply.

In the case of artifacts, teleological norms frequently *do* come to apply in virtue of the intentions of designers or users. Intending an object to be an artifact of a particular sort involves a tacit commitment to approve

⁷ See Stalnaker (1984), p. 6. Fred Dretske frequently compares artifact detectors of various sorts with mental states. See, for instance, Dretske (1986) and (1988).

⁸ Gibbard (2003) discusses the idea that norms might be constitutive of content.

or disapprove of the object according to how it comports with particular norms in a way that establishes these norms as applying to the object. Significantly, teleological norms *need* not come to apply to artifacts in this way; this generalization is subject to counterexamples. The members of a society might intend a particular substance to purify them before the gods so as to avoid their wrath, which they find evidence for in epidemics of disease. It might well turn out that (at least) a telos of this substance is to kill bacteria because the use of the substance is ultimately counterfactually sensitive to the fact that the society is thereby benefiting from lower rates of disease due to the anti-bacterial properties of the substance.⁹

In the case of intentional states, teleological norms frequently *could not* come to apply in virtue of the intentions of designers or users. (I will discuss how they might come to apply in Section 7.) Of course, this generalization is also prone to counterexample due to the possibility of artificial intelligence. Counterexamples in both directions strongly suggest that the teleology of intentional states need not be of a fundamentally different kind than the teleology artifacts exhibit.

Understanding intentional states on the model of representational objects and artifacts motivates the idea that the “correctness” relevant to beliefs is *genuinely* normative, and not merely a way of categorizing beliefs as true and false.¹⁰ What makes some piece of paper a map of Eugene is the fact that the piece of paper is *correct* (relative to some arbitrary scenario) if and only if (that scenario is one in which) Eugene has features that bear a particular structural resemblance to the actual features that are printed on the paper. In other words, the piece of paper is a map because there ought (in some sense) to be structural resemblances between its features and the features of some particular place—and it is a map of Eugene because that particular place is Eugene. It is not enough for the features of Eugene merely to bear a structural resemblance to features printed on the paper—it is not enough for the piece of paper merely to meet the standards of a good map of Eugene. Meeting the standards can happen by sheer accident (even if the correspondence is caused); just because some piece of paper might be successfully used as a map of

⁹ Indeed, there is even a temptation to think that this telos is *the* telos of the substance, especially if we suppose the substance is wholly useless as a means of purification before the gods due to the fact that these particular gods don’t even exist.

¹⁰ Here I engage directly with Glüer and Wikforss (2009), p. 36. Suggesting that correctness is normative is compatible with the idea that correctness directly entails some sort of permission rather than some sort of obligation. Cf. Whiting (2010). When I say that there ought to be a structural resemblance between the object that represents and what is represented, I intend to suggest that there ought to be a structural resemblance given that there is representation at all.

Eugene doesn't mean it is one. The standards must *apply* to the piece of paper in order for it to be a map; the norms are a prerequisite for the piece of paper to have the right sort of telos.¹¹

Analogously, what makes a particular cognitive state a belief with a particular content that *p* might just be (at least partly) the fact that the cognitive state is *correct* (relative to some arbitrary scenario) if and only if (that scenario is one in which) *p*. (Thus, the fact that a cognitive state is correct (relative to some arbitrary scenario) if and only if (that scenario is one in which) snow is white might significantly contribute to the cognitive state being a belief that snow is white.) Correctness here would be genuinely normative just as in the case of the map, and not just a way of categorizing which beliefs have “matched up with the actual world” in the sense of having content that is true. Just as structural resemblance between a piece of paper and Eugene is not enough for correctness, so the truth of the content of a cognitive state is not sufficient for the correctness under consideration. After all, the cognitive state might be, for instance, a disbelief, which is certainly not correct when its content is true. What is relevant to whether a cognitive state is a belief is not that the state “matches up with the actual world” in the sense of having a content that is true, but rather that the state (if it exists) ought to “match up with the actual world” in this way. Indeed, on this picture, a belief with content that is true is, in some particular sense, how it ought to be even if the belief is irrational. This result is analogous to the fact that a map of Eugene can be correct—and hence, can be as it ought, in some sense, to be—even if the map was drawn on the basis of poor evidence or using substandard cartographic methods.

On the picture under consideration, it is not just that a cognitive state is a belief (partly) because particular norms apply. Rather, the cognitive state is a belief *with content that p* because these norms apply.¹² (Thus, that some cognitive state is correct (relative to some arbitrary scenario) if and only if (that scenario is one in which) snow is white significantly contributes not only to its being a belief, but also being a belief that snow is white rather than a belief that grass is green.) On this picture, a cognitive state is a belief with content that *p* (at least partly) because the state is correct (relative to a scenario) if and only if (that scenario is one in which) *p*. If the state were correct (relative to a scenario) if and only if (that scenario is one in which) *q*, where the content that *p* and the content that *q* have different representational/truth-conditional content, then the state would not have the content that

¹¹ The distinction I draw between meeting standards and those standards applying is, of course, wholly analogous to the distinction often made between conforming to a rule and following it.

¹² Cf. Glüer and Wikforss (2009), pp. 37-9.

p . The state has the content that p precisely because its correctness depends on whether p . This result is completely analogously to the map case; the map has the (representational) content it does—it is *about* Eugene—precisely because the correctness of the map depends on how Eugene is.

It's important to highlight that the inherent normativity of content on this picture is not due to anything peculiar about belief.¹³ A cognitive state is a desire with content that p (at least partly) because the state is *correctly* realized (in an arbitrary scenario) if and only if (that scenario is one in which) p .¹⁴ Again, the norm of correctness not only fix that the cognitive state is conative rather than doxastic, they also fix the content of this conative state. On this picture, norms are constitutive of a telos that fixes (at least partly) both attitude and content. Again, this result is analogous to the situation with representational objects and artifacts. Norms constitutive of the teloi of maps and blueprints, and the teloi of maps and blueprints fixes both that they are maps rather than blueprints or blueprints rather than maps *in addition to* their content.

Let's call the picture I've just sketched the "*teleological theory of intentional representation*" (TTIR). The three ideas behind TTIR commit one to the thesis that norms are constitutive of having representational content. Even if representational content does not exhaust mental content, this conclusion would be enough to establish CIN. (On this picture, norms are (at least partly) constitutive of content.) In my mind, TTIR is the principal motivation for accepting the thesis that content is inherently normative.

How does TTIR bear on the Glüer-Wikforss premise that any norms inherent to content must be norms that guide performance (though intentional adherence to them)? The norms that apply to maps, blueprints, and other representational artifacts are not norms that guide performance in this way. After all, maps, blueprints, etc. are not agents, and hence, do not "perform" in the requisite sense.¹⁵ If one accepts—in accordance with TTIR—that the norms that apply to intentional states are just like the norms that apply to representational objects and artifacts, then

¹³ This undercuts the supposed difficulty for CIN spelled out in Glüer and Wikforss (2009), pp. 39-41.

¹⁴ One of the faults of Boghossian (2005) is that he fails to recognize that there is as much reason to think that there is inherent normativity in the case of desire as there is in the case of belief.

¹⁵ Of course, there might be a perfectly respectable sense in which we can judge the performance of maps if all this amounts to is judging how well they do as maps. There is also a perfectly respectable sense in which this sort of performance can be guided by representational norms. The fact that this map does well as a map might well be at least partly explained by the fact that the map is supposed to look that way. There will be an analogous sense in which intentional states perform, and which this performance can be guided by representational norms. See footnote 5.

we shouldn't expect that these norms will guide performance in this way either—quite the contrary, in fact. TTIR is inconsistent with the Glüer-Wikforss premise.

TTIR is likewise inconsistent with other similar principles marshaled to refute the thesis that content is inherently normative. For instance, given TTIR we shouldn't expect the "oughts" of norms inherent to content to imply "cans." Maps can't do anything intentionally. Neither can beliefs. Advocates of TTIR should reject the "ought"- "can" principle; it is inappropriate to the relevant sort of "oughts."¹⁶ The "ought"- "can" principle likewise applies to norms that govern agents rather than to norms that constitute *teloi*.

2: Telos and representation

In the last section, I argued that CIN is motivated by three ideas. The first idea was the observation that part of what it is to be a representational artifact, e.g. a map, a blueprint, etc., is to have a particular sort *telos*, and the third was that the existence of norms is constitutive of having such a *telos*. An opponent of CIN might be tempted to challenge these claims. As these claims are important both to motivate CIN and to motivate the rejection of the Glüer-Wikforss premise, I will take the next two sections to defend them.

Why think that being a representational artifact involves having a particular *telos*? The natural alternative view is presumably that being a representational artifact of a particular sort is a matter of nomological features rather than teleological ones. For instance, being a map of Eugene might be a matter of being caused in the right way so that there is counterfactual dependence between the features printed and the features of Eugene. Likewise, being a blueprint of a house might be a matter of having a tendency to cause a house to be so that there is a similar sort of counterfactual dependence.

A straightforward nomological view of representational objects and artifacts that requires counterfactual dependence is not promising.¹⁷ As has been noted, counterfactual dependence is not necessary for representational properties because maps, blueprints, and their kin can represent poorly.¹⁸ A map of Eugene can be very inaccurate, and a blueprint can be a blueprint even if nobody has any serious motivation to carry it out. These points are a strong strike against not just the straightforward nomological view, but nomological views generally. Very plausibly, there must be some sort of causal connection between Eugene and a piece of paper in order for the piece

¹⁶ The "ought"- "can" principle is the focus of a criticism of CIN found in Hattiangadi (2006).

¹⁷ See Prinz (2002), Chapter 9 for one discussion of problems with causal theories of mental representation.

¹⁸ Cf. Millikan (1989).

of paper to be a map of Eugene.¹⁹ However, the sort of causal connection required is not robust enough to distinguish between those pieces of paper that are maps of Eugene and those that aren't.

Indeed, it is difficult to see how facts about how a piece of paper is and how it was caused to be could fix that the piece of paper is a map of Eugene, unless those facts also fix that the piece of paper has the telos of representing Eugene. Consider the following story:²⁰

Joe has a remarkable memory; he is also bored. Joe decides that he is going to make a detailed map of Eugene from memory. The difficulty of the project motivates Frank to bet Joe that, despite his amazing memory, his map will be wrong in at least some significant details. At the same time, José, a city planner, is drawing up a plan for Eugene. He uses current maps and photographs for his plan. His plan retains much of the city as is.

We might suppose (for dramatic effect) that by sheer coincidence, the map of Joe and the plan of José turn out to be molecule-for-molecule duplicates. Furthermore, José's plan is wholly correct in its depiction of the majority of Eugene that José wants to keep unchanged, and Joe's map is wholly correct in its depiction of that part of Eugene as well. Unfortunately, Frank proves to be right; Joe's map is not correct in all of its depiction of Eugene. Coincidentally, Joe's map misrepresents the very parts of Eugene that José is proposing to change. Obviously, José's plan is not incorrect in depicting Eugene as Joe's map does.

Both Joe's map and José's plan are causally sensitive to the features of Eugene. Indeed, to the extent that Joe's map would work as a map of Eugene, so would José's plan. Likewise, to the extent that José plan would work as a plan for changing Eugene, so would Joe's map. The two products are qualitatively identical, and thus could be swapped without making a difference. Nonetheless, the piece of paper Joe produces is a map while the piece of paper José produces is a plan.

What fixes that Joe's product is a map while José's is a plan is the respective teloi of the two products. The intrinsic properties of these products are largely irrelevant to whether the product is a "reflective" or "projective" representation, and the only causal properties are particularly relevant here are the etiological properties—in particular, having been created with certain intentions—that fix the teloi of the products.

¹⁹ Putnam (1975).

²⁰ Ironically, the same sorts of thought experiments that motivate the causal theory of reference/representation also undermine it. Cf. Putman (1981), Chapter 1.

The same sort of point might be made about content of representational artifacts. Both using Mary as a model, Tim and Jim might produce an identical paintings. Nonetheless, even though Tim's painting is of Mary, Jim's painting might be of Mary's identical twin sister Cary if Jim was intentionally using Mary as a stand-in for Cary. (We can even suppose Jim tries to adjust for subtle differences between Mary and Cary, and Tim happens to make mistakes in depicting Mary that exactly coincide with those adjustments.) Differences in content between the two paintings can exist even if both painting are duplicates of one another, and moreover, are causally sensitive to the same thing. The content of each painting is fixed by what the painting is supposed to be—the telos. As in the previous example, the only causal properties that are relevant in these cases are the etiological properties concerning the intentions of the creators.

Obviously, the counterexamples I've concocted may only apply to more crude nomological theory of artifactual representation. However, it's difficult to see how sophisticated theories could be developed so as to avoid all counterexamples if only because there are no strong constraints on how the causal relation between artifactual representation and what is represented can be realized. I can use information on the internet to draw a map of Eugene; alternatively, I could go to Eugene and make direct observations. A wide variety of causal connections between Eugene and me are sufficient for me to draw (at least a bad) map of Eugene (or alternatively find some object to be a map). The only natural way to single out genuine maps of Eugene from things that just happen to be causally connected to Eugene is by looking for teleological properties that in the most typical cases will stem directly from the intentions of users or designers.

I see no way, then, to avoid the conclusion that, at least when it comes to the sort of representation exemplified by representational objects and artifacts, representation is a matter of telos.²¹ Obviously, this conclusion does not imply that the “representation” of intentional states is a matter of telos. (Better: this conclusion does not imply that intentional states are genuinely representational even if they are, in some sense, about things.) However, it does imply that a commitment to understand intentional states on the paradigm of representational objects and artifacts does require understanding intentional states as inherently having teloi. If having a telos is inherently normative, then such a commitment ultimately motivates accepting CIN.

²¹ Richard Heck pointed out to me that this conclusion might explain why it may be a mistake to begin the search for mental representation by looking at information, cf. Dretske (1981), which is not teleological in nature. It's not obvious we should count information or Grice's “natural meaning”, cf. Grice (1989), as a form of genuine representation.

3: Norms and telos

But is having a telos inherently normative? Does someone attracted to the general TTIR picture whereon intentional states have their representational content due to having some particular telos really have to accept CIN?²² I think so.

What else could be constitutive of teloi besides norms? Teloi often emerge from the intentions of designers and users of objects, but as the “purifying substance” counterexample from Section 1 demonstrates, intentions are not necessary for teloi. (A substance can have the telos of an anti-bacterial even if nobody intends it to.) This same example also suggests intentions may not be sufficient either. The members of that imagined society may intend for the substance to purify them before the gods. However, if there are no gods to pacify, it’s not altogether obvious that the telos of the substance is to purify them before the gods. When the members of the society discover that there is no reason to believe in the gods of their folklore, it seems like they might reasonably judge that the substance isn’t really for purification before the gods after all. Perhaps people can be wrong about what they think the telos of an artifact to be, or if it really has a telos at all. Artifacts can serve different roles in our lives than we think they do. In any case, there is good reason to conclude that intentions are not even essentially tied to teloi—they are merely common associates.

On the other hand, norms are at least essentially (if not constitutively) tied to teloi. Suppose I am an anthropologist and I want to discover which objects present in a foreign society have the telos of being a map of the surrounding area. It is necessary and sufficient for my purposes to discover which objects are such that their features are correct or incorrect depending on how the surrounding area is. This project is not the same as discovering which objects bear structural resemblance to the surrounding area. Any object of sufficient complexity will bear structural resemblance to the surrounding area; moreover, any such object could become a map if somebody learned to interpret its features in the right way.²³ I’m not trying to discover whether there are any objects that could make for good maps; I’m trying to discover which objects are the maps (whether they are good maps or not). Consequently, what I ultimately need to discover is the objects that aren’t as they ought to be if there isn’t a structural resemblance. I need to discover which objects that the teleological norms of being a map apply to.

²² I am actually inclined to think that the inherent normativity of teleology is platitudinous, so that anybody who denies it is either confused or has some other concept of teleology in mind. Nonetheless, the burden rests on me to argue that this sort of teleology is generally exhibited by artifacts.

²³ Cf. Lewis (1984).

Another example may help to illustrate these points. Consider a picture taken from space of North America on a relatively cloudless night across the continent. Presumably, the illuminated points in the picture will non-accidentally coincide with urbanization. Consequently, there will be a non-accidental structural resemblance between illumination in the image and urban centers in North America. Noting this non-accidental structural resemblance, however, does not license our judging that this image has the telos of a map of the urban areas in the North America region. Of course, we could use the image as this sort of map; occasionally, we do use pictures of North America at night in this way for dramatic effect. However, we usually don't use images in this way because we artificially manufacture maps of this sort that are far more accurate. Although a picture of North America at night might become a map of North American urbanization if it were consistently used in the right way, pictures of this sort aren't generally considered to be maps of North American urbanization because they generally aren't used in that way. Generally, we just think that images of this sort are representations of the visual appearance of North America from space on some particular occasion. Suppose that when the picture under consideration is taken, Seattle is having a complete blackout due to power failures. Consequently, the part of the image that would be illuminated if Seattle were lit up is dark. We wouldn't typically judge that the image is incorrect on that account precisely because we don't think the image is a map of North American urbanization. If we were using the image as a map of urbanization, however, we would typically judge that it was incorrect on that account—at least for those purposes. This change in what we would judge if the telos of the image were different reflects our tacit commitment to the idea that, at least for artifacts, representational norms of correctness necessarily track teleology.

In general, something has a telos precisely because it ought to be a certain way—i.e. because being that way merits a certain sort of approval. Things don't have teloi because they do anything in particular. For instance, something can have the telos of a mousetrap even if it completely fails to catch mice. Mousetraps can be broken or poorly designed. What makes something a mousetrap is that it ought, in some sense, to catch mice. This is not to say that the world is how it ought to be if this thing catches mice; it is not to say that the thing merits all things considered approval if it does catch mice. If a mousetrap catches mice, it merits mousetrap-approval. Mousetrap-approval is wholly compatible with thinking that mousetraps are an abomination.

In other words, just because something is as it ought to be teleologically—so that it merits a sort of “job”-approval—does not imply that it is how it ought to be in any other sense. Indeed, something fulfilling its telos need not even be an objective good. An evil maniac might invent a child torture machine. What makes this object a child

torture machine is that it have the telos of torturing any nearby children. It has this telos precisely because when it tortures children in accordance with its design plan, it is working correctly—exactly as it ought (teleologically) to be working. Obviously, that the object ought to work in this way teleologically, does not imply that working in this way is even an objective good. (Meriting “job”-approval can be a bad thing if one has a bad job.)

Consider a comparison with games. Winning *ipso facto* merits a certain sort of approval. What fixes that we are playing checkers rather than *misère* checkers is the fact that certain final states of game merit approval of a sort and others don't. Nonetheless, winning a game obviously doesn't merit approval all things considered. I can judge that Sharon has done well by winning the checkers game so decisively (and even played very well from a strategic standpoint), and yet wholly disapprove of her trouncing her little sister, Shirley, so thoroughly that it will discourage Shirley from really learning to play.²⁴

Meriting mousetrap-approval or mousetrap-disapproval is not the same as meeting or failing to meet the standards for a good mousetrap. The plastic spoon I used to eat breakfast and a broken mousetrap both fail to meet the standards for a good mousetrap, but only the latter (usually) merits mousetrap-disapproval on that account. Although neither the spoon nor the broken mousetrap functions well as a mousetrap, only the latter thereby functions *incorrectly*. It's not the “job” of the spoon to catch mice so it can't be incorrect for not being currently disposed to do so. On the other hand, it is the “job” of a mousetrap to catch mice, so it is incorrect for its failure in this regard. This point shows that the relevant incorrectness (and correctness) is not merely a matter of failing (or succeeding) to measure up to standards for mousetraps—after all, both the plastic spoon and the broken mousetrap fail to measure up. In the case of broken mousetraps, however, mousetrap norms actually apply, and, quite plausibly, it is because they apply that the items in question are mousetraps. To put the point more colloquially, something gets a “job” when norms are established to judge its functioning.²⁵

Indeed, it is very plausible to think in this way that teleological norms are not only essentially tied to teloi, but also constitutive of having them. Teleological norms provide the best explanation of what acquiring and having a telos consists in. How else does one intend for some object to be a mousetrap except by effectively committing

²⁴ It's worth saying that I'm not committed to thinking that game norms and teleological norms more generally can't derivatively give us reasons for action. They might in this way sometimes guide performance. I merely intend to suggest that they don't do so inherently. Thanks to Richard Heck for this point.

²⁵ As the functioning of this something can be explained by the fact these norms have been established, these norms cannot be purely evaluative norms. See footnotes 5 and 14.

oneself (and possibly others) to render mousetrap-approval when the object meets the standards for a good mousetrap and mousetrap-disapproval when it doesn't? What else would intending an object to be a mousetrap consist in? Intentionally making a mousetrap in part involves establishing that mousetrap norms apply to the object one creates. More generally, bestowing a telos on some object or artifact is wholly a matter of settling how the object or artifact is to be evaluated with regards to its functioning. When one establishes that an object is to be evaluated with approval or disapproval depending on whether it effectively catches mice, one doesn't have to do anything else to make that object a mousetrap. It's difficult to see how there could be anything more to a "job" than the norms that apply to something that has that "job."

4 Teleological representation, norms, and correspondence

We are now in a position to find fault with the following remarks from Dretske:

Beliefs and judgments must either be true or false, yes, but there is nothing normative about truth and falsity. What makes a judgment false (true) is the fact that it fails (or succeeds) in corresponding to the facts, and failing (or succeeding) in corresponding to the facts is, as far as I can see, a straightforward factual matter. Nothing normative about it.

An arrow (on a sign, say) can point to Chicago or away from Chicago. There is a difference here, yes, but the difference is not normative. Aside from our purposes in putting the sign there or in using the sign as a guide, there is nothing right or wrong, nothing that is supposed-to-be or supposed-not-to-be, about an arrow pointing to Chicago.²⁶

Here Dretske appears to commit himself to the idea that representation—both in the case of intentional states and in the case of representational artifacts—is *not* inherently normative. Putting aside the question of intentional representation for the moment, let us consider the case of representational artifacts. Does Dretske's sign represent that Chicago is in a particular direction given merely that it "points" to Chicago?

We should be very careful here about what we mean by "points." Consider the vectors beginning at the tail of the arrow and passing through the arrow's head. It is, no doubt, an entirely non-normative matter whether some such vector (roughly) ends in Chicago. If this fact settles the matter of whether the arrow "points" to Chicago, then

²⁶ Dretske (2000a) in Dretske (2000b), p. 247.

without question the matter is entirely non-normative. However, I see no reason to conclude on the basis of any such “pointing” that the arrow represents that Chicago is in some particular direction. It is compatible with the arrow “pointing to Chicago” in this sense that the arrow represents that Milwaukee is in that direction whether it be correctly—because, say, the arrow coincidentally also “points to Milwaukee”—or incorrectly—because, say, the maker of the arrow confused Chicago for Milwaukee.

Alternatively, “pointing to Chicago” might mean that the arrow represents that Chicago is in a certain direction. If so, it is no longer obvious that the arrow’s pointing to Chicago is entirely a non-normative matter (even though the question of whether the arrow represents that Chicago is in some particular direction is obviously settled). Indeed, I think careful consideration of Dretske’s arrow shows that the relevant representation would not, in fact, be entirely a non-normative matter. The arrow represents that Chicago is in a particular direction (if it does) because it has a telos—a telos bestowed upon it, no doubt, by designers and/or users of the arrow. Having this telos is settled by the arrow’s having certain conditions for correctness—where to say that the arrow is correct or incorrect is not merely to say that the arrow meets or fails to meet the standards we typically use for an arrow representing that Chicago is in some particular direction. After all, a plastic spoon and a broken mousetrap typically fail to meet these standards, but they are not typically incorrect on that score. Similarly, an arrow representing the direction of Milwaukee might meet the standards we typically use for an arrow representing that Chicago is in some particular direction by lining up appropriately with the direction of Chicago, but it is not thereby correct.

To say that the arrow is correct (incorrect) is to imply that certain norms of representing apply. To say that the arrow is correct (incorrect) is to imply that, in some sense, the arrow ought to be a certain way and that it does (doesn’t) comply. This is not to imply that the world is a better (worse) place for the arrow complying (not complying). An arrow can merit our approval *qua* representation even if we wish that it misrepresented and led people astray.

Having a particular telos—and, in particular, having a representational telos—is an inherently normative matter. Something is a representational artifact ultimately because of the teleological norms that apply to it. To understand intentional states on the model of representational artifacts (in accordance with TTIR), does, therefore, commit one to the view that intentional properties—content and attitude—are inherently normative. (At this point, I take it that I have established that TTIR does imply CIN.)

Dretske's confusion about the inherently normative aspect of artifactual representation no doubt stems from the misguided thought that, when it comes to artifacts, we can make sense of correspondence to the facts independent of any representational norms. In many cases, we are fooled into thinking we can make sense of correspondence independently because we confuse it with resemblance or structural resemblance. A painting may well resemble the visual likeness of a valley in southeast China taken in from a certain perspective, but the painting need not correspond to the valley in the relevant sense.²⁷ Confusion can arise because we often use resemblance as a measure of correspondence. Nonetheless, the two can evidently come apart. Consider a decent painting representing the Yangtze River Valley and the negative of a photograph taken of the same valley. Because the negative of the photograph inverts the color spectrum, the painting may well better resemble the visual likeness of the Yangtze River Valley from a certain perspective even though the negative of the photograph better corresponds with the visual likeness of the Yangtze River Valley from that very same perspective. A digital photograph of the Yangtze River Valley can produce data on a flash drive that corresponds quite well with the Yangtze River Valley even though it doesn't resemble a visual likeness of the Yangtze River Valley in the slightest. Resemblance is incidental to correspondence. Correspondence is settled by what would be correct—it is a matter of things lining up as they (in some sense) *ought* to. It can be that they ought to line up so as to resemble, but it need not be.

No doubt there is likewise structural resemblance in many cases of correspondence, but structural resemblance is ubiquitous, and generally not indicative of correspondence. Similarity in structure without correspondence is the rule, not the exception.²⁸ Structural resemblance only indicates correspondence when it's a structural resemblance that's supposed to exist.

At least when it comes to artifacts, genuine correspondence is what happens when world and representation are *correctly* aligned where correctness here is normatively loaded. We can't make sense of whether Dretske's arrow corresponds to the world in the relevant sense unless we know what correspondence would consist in, and we can't know what correspondence would consist in until we know what it is for the arrow to be as it ought to be. There's no principled reason why Dretske's arrow can't represent that Chicago is in a particular direction by

²⁷ I would be willing to concede that there is a sense of "correspondence" in which it just means resemblance or even structural resemblance. In this latter sense of "correspondence," we might say that the details of a painting correspond with the features of the valley even if the painting does not represent the valley. Henceforth, I will only use correspondence in the relevant sense, that is, correspondence as it relates to representation.

²⁸ Cf. Lewis (1984).

“pointing” (in the formerly described non-representational sense) in the opposite direction or some other direction entirely. Despite illusion to the contrary, arrows don’t represent direction intrinsically.

Let’s move our focus from artifacts to intentional states. Suppose that, as I have suggested, representational artifacts have representational teloi in virtue of constitutive representational norms. Opponents of TTIR can nonetheless reject the claim that intentional states “represent” as artifacts do. They can insist that intentional states have truth-conditional content, and hence are, in some sense, about things, but we need not assume on that basis that they “represent” in the same way artifacts do. However, if so-called intentional “representation” is merely a matter of having truth conditions, then correspondence in the intuitive sense between intentional states and world won’t necessarily make much sense. We can characterize a shallow notion of “semantic correspondence” whereby a cognitive state “semantically corresponds” just in case it has a content that is true. Nonetheless, “semantic correspondence” isn’t correspondence in any intuitive sense. On our characterization, a doubt that p “semantically corresponds” to the world in the event that it is true that p , but, intuitively, a doubt that p does not correspond to the world in the event that it is true that p . If beliefs genuinely correspond in the intuitive sense when they are true, it is very plausibly only because unlike doubts, beliefs by their very nature ought to be held only when they “semantically correspond.”

There is a longstanding tradition in the philosophy of mind that attempts to give the intentional representation of a cognitive state at least partly by attempting to *independently* characterize what it would be for a cognitive state to correspond to the world in non-normative terms.²⁹ (A dependent, shallow characterization of semantic correspondence is always available by identifying semantic correspondence with truth and tacking a deflationary theory of truth onto the theory of mental content.)³⁰ I believe that there are good reasons to think this tradition is misguided. Certainly, no attempt to accomplish this project has met with obvious success.³¹ One of the principal obstacles seems to be that we need to have some independent idea of when cognition is proceeding *correctly* and when it isn’t in order to know whether the resulting cognitive states exhibits genuine correspondence or not. In the end, I strongly suspect that this sort of obstacle is insurmountable. Ultimately, representation probably does not follow upon metaphysically prior relations of correspondence; rather, correspondence follows

²⁹ A clear example of this tradition is Fodor (1990), but any non-normative “referentialist” will do.

³⁰ Cf. Field (1994).

³¹ Fodor (1990) does admirably well considering the challenge, but I don’t ultimately find his asymmetrical dependence view very convincing. See the essays of Loewer and Rey (1991) for some criticism.

upon representation with a normative aspect. Although significant, this insight is widely underappreciated. In failing to appreciate it, philosophers such as Dretske are tricked into thinking that we can make sense of robust representational properties without postulating normativity.³² Making sense of robust, non-deflationary representation probably hinges on the truth of TTIR precisely because, in general, representational norms are conceptually and metaphysically antecedent to correspondence relations (so that the former fixes the latter rather than the other way round).³³ TTIR makes best sense of our intuitions that intentional states and reality sometimes do and sometimes don't correspond.

5 A confusion TTIR does not make

Opponents of CIN often accuse advocates of CIN of confusing the norms that apply to acts with norms that apply to states and objects that are the products of those acts.³⁴ Norms that apply to production should not be confused with norms that apply to the thing produced. Perhaps we ought to put an arrow here with this or that orientation, but whatever our moral or practical obligations might be in this regard, these obligations apply to us, not to the arrow. The complicated norms that guide our use of language are not part of the theory of linguistic content.³⁵ The epistemic norms that illuminate the proper production of judgments are (allegedly) not part of the theory of mental content.³⁶ Opponents of CIN often suggest that once we attempt to distinguish all the norms that apply to our acts from norms that apply to the product of the acts, we will see that CIN is no longer appealing precisely because we will see that there are no norms of the latter sort—or at the very least, none that are part of the theory of content.

The advocate of TTIR disagrees. She will agree, of course, that we should distinguish between norms that apply to acts (broadly conceived to include mental acts like judging) and norms that apply to the products of those acts. She will insist, however, that her agreement on these points in no way concedes anything to her opponent, but, in point of fact, is critical to her own view. She does not confuse practical or moral norms with the teleological norms requisite for representation, nor does she confuse these latter norms with norms of production (although the two are not entirely unrelated).

³² Cf. Austin (1950).

³³ I discuss this question further in work in progress.

³⁴ This is particularly clear in Dretske (2000a), but the same might be said of Horwich (1998), Chapter 8.

³⁵ Cf. Horwich (1998), Chapter 8.

³⁶ Cf. Horwich (2005), Chapters 5-6.

Whatever norms apply to the production of representational artifacts, e.g. maps, blueprints, etc., these norms are certainly distinct from the norms that apply to the representational artifacts themselves. Consider again Joe who has bet Frank that he can draw up a correct map of Eugene, Oregon from memory. What regulative norms apply to Joe's act of production? These norms depend almost certainly on Joe's moral obligation and pragmatic interests. It might be best for Joe, all things considered, to intentionally draw the map with significant inaccuracies. Frank may be a sore loser, and winning the bet might not be worth Joe's trouble. Alternatively, Joe might want to give Frank some money without hurting Frank's pride. Assuming Joe has no overriding moral obligation to do the best he can, either one of these reasons might make it the case that all things considered Joe ought to draw the map with significant inaccuracies.

Of course, this normative fact has nothing to do with whether or not the resulting map will be correct—as *it* ought to be—or not. As I established in Sections 2 and 3, the piece of paper that results from Joe's work is a map because *it* ought to be a certain way. If the piece of paper is not that way, then it is incorrect whether or not Joe was, all things considered, right to produce the piece of paper in the way that he did.

Obviously, in some sense, Joe incorrectly acts—he has drawn incorrectly—if he produces a map that is incorrect. This is not to say that Joe acted wrongly all things considered—or even that he had a *pro tanto* practical normative reason to act otherwise. Joe acts wrongly in the role of mapmaker, where the norms of mapmaking are derivative from the norms that are constitutive of maps. As a mapmaker, Joe is effectively “playing a game” in which the “object of the game” is to produce a correct map of Eugene. However, norms that apply to the players of games generally need to be distinguished from norms that apply to objects that figure in the games. The norms that apply to people in their roles as chess players should be distinguished from the norms that are constitutive of being a knight. (Some particular hunk of wood is a knight precisely because of the way it is permissible to move it.) Likewise, although the two sets of norms are related (and perhaps intimately so), any norms that apply to Joe as part of this “mapmaker game” should be distinguished from the norms that apply to the product of the “game”—i.e. the map. There are norms that apply to maps in their role as maps even if there are also norms that apply to mapmakers in their roles as mapmakers.

In any case, both the norms that apply to representations and the norms that apply to people in their role as representation-makers are strictly separate from other norms that apply to people. “Playing the game of mapmaking” need not give Joe even a *pro tanto* practical normative reason to play it at all well. Perhaps Frank

knows that Joe can draw an accurate map of Eugene, and makes a bet with him that he won't just to see if Joe will do what Frank wants. If Joe recognizes Frank's intentions, he may realize that it's in his interest to draw the map inaccurately. In this sort of case, Joe may not have any motivation whatsoever to draw a correct map precisely because he has no *pro tanto* practical normative reason to do so. The product of Joe's efforts is a map of Eugene nonetheless if he and Frank agree (perhaps tacitly) on how the product is to be evaluated, i.e. on what its correctness conditions are. This sort of (perhaps) tacit agreement only gives one a practical normative reason for action if one has an additional reason to bring about a product that merits a good evaluation according to the agreement.

Just as there are distinct norms that apply to representational artifacts themselves, so the advocate of TTIR will assert that there are norms that apply to the belief states and desire states themselves. These representational norms are not to be confused with intimately related norms of rationality that apply to the production of these states (unintentional as it may be and generally is) or other norms that may apply to the production of these states (unintentional as it may be and generally is). A good advocate of TTIR is not guilty of confusing norms that apply to acts of production with norms that apply to the products; in fact, her view crucially depends on making that distinction.

If anybody confuses norms that apply to production with norms that apply to what is produced, it is someone who accepts the Glüer-Wikforss premise that any norms inherent to content must be regulative norms that (through intentional adherence) guide performance (or similarly someone who thinks that norms inherent to content are subject to a "ought"- "can" principle). (Proponents and opponents of CIN alike are guilty of this confusion.)³⁷ If there are norms that are inherent to content, they are norms that apply to intentional states, not to intentional agents. Hence, there is no reason to think they have anything to do with guiding performance in the requisite sense.

6 Can there be norms that don't guide performance?

If TTIR is the principal motivation for accepting the thesis that content is inherently normative and TTIR is inconsistent with the Glüer-Wikforss premise that any norms inherent to content must be regulative norms (that guide in the requisite sense), then why would Glüer and Wikforss take this premise as a starting point?

"Prescriptions," Glüer and Wikforss tell us, "involve genuine 'oughts'."³⁸ This statement leaves the reader with the impression that Glüer and Wikforss do not think there are interesting genuine "oughts" other than those that are

³⁷ Cf. Boghossian (2005).

³⁸ Glüer and Wikforss (2009), p. 37.

“prescriptive” in the requisite sense of guiding. (Note: Where Glüer and Wikforss use “prescriptive,” I have been using “regulative.”) Obviously, if that were so, the Glüer-Wikforss premise would be well motivated. In previous sections, I have presupposed that it is not so.

Certainly, not all *genuine* norms are *regulative*.³⁹ A regulative norm is an applicable binding prescription that one can't felicitously beg-off and thereby treat as irrelevant to one's course of action, most especially if one admits that the norm is applicable. We would have trouble understanding someone making a sincere judgment whose content was normative in the *regulative* sense if they do not have any even *pro tanto* motivation to comply with the judged norm. Such a judgment seems to be impossible if the person making the judgment is rational. A regulative norm is, by its very nature, a norm that someone could attempt to comply with intentionally.

In this paper, I have essentially conceded that representational norms—norms of correctness for representational artifacts—are not regulative. To do so is not necessarily to concede that representational are not *genuine* norms. *Genuine* norms contrast with fake norms where normative language can be paraphrased away. Here are some examples of fake norms:

- (1) If one's aim is to harass the neighbors, one ought to play music loudly.
- (2) If one's overriding goal is to win the lottery, then one ought to buy a lottery ticket.

It's not clear that these examples say anything above and beyond what we might say in purely descriptive terms. (1*) and (2*) might serve as adequate replacements for the examples above:

- (1*) An effective means for harassing the neighbors is playing music loudly.
- (2*) Buying a lottery ticket is necessary for winning the lottery ticket.

These replacements neither use evidently normative language/concepts nor seem to entail any other statement that uses evidently normative language/concepts. To the extent that (1*) and (2*) are adequate paraphrases, we have reason to believe that (1) and (2) are actually descriptive facts rather than genuine norms.

We can readily distinguish fake norms from genuine norms by considering what we do when we judge that one or the other is true. To judge that a fake norm is true is merely to evaluate something with respect to some

³⁹ Regulative norms are governed by what is sometimes called the “internalist” constraint. For some discussion see Korsgaard (1986) or Wedgwood (2007), pp. 25-7.

considered standard or end without actually embracing the standard or end. When I judge that is (1) true, I merely evaluate playing music loudly with respect to the end of harassing the neighbors; I do not embrace the end of harassing the neighbors. When I judge that some particular fork isn't constituted how something with a chair-*telos* ought to be constituted, I am evaluating this fork with respect to chair standards, but I am not embracing the chair standard for it. Because in both cases I do not embracing any standard, the "ought" is inessential to the content of these judgments. I am not making judgments of genuine norms.

In contrast, judgments about representational correctness and, more generally, teleological "oughts" have genuinely normative contents. Judging that a teleological "ought" statement is true requires embracing the standard or end. When I judge that some particular piece of furniture ought not fall apart when sat on (because it is a chair), I *am* embracing a standard for that piece of furniture in its role of chair. I am not merely evaluating the item according to the chair standard; I am taking this standard to apply to the item (in a way I did not do with the fork). As a result, the teleological "ought" cannot be paraphrased away.

Nonetheless, we need not think that teleological "oughts" are regulative. We can draw a distinction between those genuine norms that are regulative and those that aren't. To make a genuine normative judgment is to embrace a standard for something (perhaps even oneself). To make a regulative normative judgment is to embrace a standard for oneself *qua* agent.⁴⁰ One need not do the latter to do the former. Genuine norms are not necessarily regulative.⁴¹

The fact that teleological norms are not regulative does not imply that these norms are not deontological. Just because they are not regulative, we need not assume that teleological norms are purely evaluative norms about which states of affairs would be better or worse. Teleological norms are not closely tied to *intentional* action, but it is their very nature to *govern* and *explain* outcomes.

For example, even though artifacts are not agents, they plainly do have "jobs" that can explain how they function. Artifacts are not *motivated* (in the sense of having desire-like attitudes) to comply with the teleological norms corresponding to their "jobs"; they do not comply with the teleological norms corresponding to these "jobs" intentionally. Nonetheless, there is a perfectly good sense in which teleological norms do govern the workings of

⁴⁰ Cf. Korsgaard (1996).

⁴¹ See Wedgwood (2007), pp. 119-120 for some discussion.

artifacts.⁴² The fact that some particular coke machine dispenses coke is always (somehow) explained partly by the fact that it ought (teleologically) to dispense coke. In typical cases, a token coke machine works the way it does because of the way it was built, but the machine was built with sensitivity to what coke machines ought to do, so ultimately, the way the machine works is partly explained by the fact that coke machines ought to work in a certain way.⁴³ This example generalizes: regular conformity with teleological norms is almost always explained (at least partly) by the teleological norms themselves (in at least some way or other).

Thus, suppose in accordance with TTIR that there are teleological norms of representational correctness that inherently apply to belief. These norms will, in many cases, explain why particular correct beliefs exist while potential incorrect beliefs were prevented from existing. The teleological norms of representational correctness will be (external) norms that govern (rather than guide) the (clearly non-intentional) production of beliefs. TTIR is a commitment to the idea that teleological norms that facilitate representation correctness must apply to and generally explain the workings and resulting products of any cognitive architecture that is capable of producing beliefs.

Note that the fact that the workings of states and objects *are* governed (in this way) by the teleological norms that apply to them makes it perfectly natural to make definitive ascriptions of responsibility. For example, when a coke machine fails to dispense coke, we can often isolate the part of the machine that is responsible for this failure. We might blame some particular cog for the failure even when the atypical working of this cog might have been compensated for by atypical workings elsewhere. The cog is responsible (rather than any other part) because it is the cog (rather than some other part) that is failing to fulfill its telos.

The fact that conformance to teleological norms is typically *not* accidental strongly suggests that teleological norms are, in at least some perfectly good sense, deontological rather than purely evaluative. The fact that teleological norms also allow for ascriptions of responsibility confirms this suggestion.

7 The Glüer-Wikforss vicious regress returns?

In Section 2 of their paper, Glüer and Wikforss argue that advocates of CIN are subject to a variety of vicious regresses. The specific formulation of their argument depends on the Glüer-Wikforss premise that norms inherent to content are regulative in such a way as to guide performance (in the requisite sense). In this paper, I have established that this premise is largely unacceptable for TTIR supporters of CIN. Nonetheless, one might worry that

⁴² See footnotes 5 and 14.

⁴³ Cf. Dretske (1988).

the threat of a similar vicious regress remains for such supporters of CIN even after they reject the Glüer-Wikforss starting point.

Ultimately, a Glüer and Wikforss style argument for the conclusion that CIN is subject to a vicious regress only depends on the premise that intentional states must be conceptually antecedent to any norms putatively inherent to intentional properties of attitude and content. If for norms to apply to beliefs, desires, etc., some subject must already have beliefs, desires, etc., then CIN is in trouble. Postulating that there are beliefs, desires, etc. would require antecedently postulating that there are beliefs, desires, etc. The pre-conditions for having intentional states could never be met.

An advocate of TTIR might worry that her view is subject to such a regress, not because she adopts the Glüer-Wikforss premise, but rather because she is committed to thinking of intentional states on the model of representational artifacts, and in paradigm cases, representational norms come to apply to representational artifacts in virtue of the intentional states of their designers. (For instance, Joe's map of Eugene comes to be a map of Eugene because of Joe's intention that it be so.) If this paradigm feature of representational artifacts generalizes to all of their kin, then TTIR is doomed to fail by way of a vicious regress.⁴⁴ Fortunately for the advocate of TTIR, there is no reason to think that this feature of representational artifacts generalizes to all of their kin.

To begin with, although artifactual teloi paradigmatically derive from the intentional states of designers, biological teloi (assuming there are any) obviously do not.⁴⁵ For instance, if hearts have the telos to circulate nutrients throughout the body by pumping blood—as they very plausibly do—this telos did not arise through any agent's design. Although not entirely uncontroversial, a number of philosophers of biology have defended the legitimacy of genuinely normative biological function, and Ruth Millikan has very prominently argued that the existence of genuinely normative biological function can provide the resources needed for a version of TTIR.⁴⁶

Millikan and her allies do not stray far from the idea that telos is closely connected to design. The now established tradition they operate within—the etiological tradition of telos—traces back to Larry Wright's landmark discussion of (teleological) function.⁴⁷ The guiding idea of this tradition is roughly that objects have teloi when their

⁴⁴ See Block (2007), p. 24.

⁴⁵ I am not necessarily committed to the idea that biological function is inherently teleological.

⁴⁶ Cf. Millikan (1984), (1993), Neander (1991a), (1991b), Godfrey-Smith (1994), and the essays of Ariew, Cummins, and Perlman (2002).

⁴⁷ Wright (1973). See Godfrey-Smith (1993) for a useful summary of the literature.

existence is explained by their ability to do something. (Thus, Joe’s map of Eugene, for instance, has the teloi of accurately corresponding in a specified way to Eugene because it came into existence with the features it has only because it well enough corresponded to Eugene in this way—Joe wouldn’t have drawn it if it hadn’t well enough corresponded.) This idea is frequently refined so that the telos of an object, mechanism, state, trait, etc. is a matter of having been selected by a natural or artificial process because of what the object, mechanism, state, trait, etc. could do. Processes of selection are, by nature, processes of design. Artificial design occurs when agents select features for their artifacts according to what those features can do. Natural design occurs when organisms are selected by survival and reproduction for what they can do.

Is the etiological theory of telos sufficient to satisfy the CIN skeptic that a vicious regress might be avoided? Perhaps not. Although the contemporary etiological tradition of telos is both prominent and important, it is not without challenges and challengers. Wright’s analysis of function was immediately subject to various counterexamples (at least if we understand it as an analysis of *teleological* function).⁴⁸ In a vicious deflationary cycle, the stock market may decline below some level causing fear that inhibits consumption and investment that ultimately keeps the stock market below that level.⁴⁹ Surely in this case, the teleological function of the stock market is not to stay below that level despite the fact the stock market existing at that level is explained by it having that property. There are cases where a process of “natural selection” does not result in telos—clay crystals have “heritable” features that promote their persistence and replication.⁵⁰ Finally, there seem to be cases of objects having teloi when they have been in no way designed. An unmodified rock can come to have the telos of a paperweight even though the rock was not designed for anything at all.⁵¹ Etiology and design may be paradigmatic associates of teloi, but perhaps these associates are neither necessary nor sufficient.

I, myself, find it plausible that there may be genuine cases of having telos that are not etiological in nature. Consider the case of Sally:

⁴⁸ Boorse (1976). I’m inclined to think that there is a perfectly respectable notion of etiological function that is distinct from teleological function. Counterexamples to Wright’s analysis understood as an analysis of teleological function would not be counterexamples to it understood as an analysis of etiological function.

⁴⁹ Precisely just this sort of vicious cycle seems to have played out in Japan over the last twenty years.

⁵⁰ See Bedau (1991).

⁵¹ Allen and Bekoff (1995).

Sally hates her family. Her father is a drunk and her mother is a nag. Unfortunately, Sally is self-deceived, and is convinced that she loves her family. For many years, she goes home every holiday season to see them.

One year, Sally receives extra assignments at work that make it impossible for her to travel home for the holidays. Working through the holidays serves to promote her well-being; without the distraction of her co-workers (who have taken the holidays off) Sally becomes more productive, and she avoids the unhappiness associated with seeing her folks.

Working through the holidays is a self-sustaining behavior. By doing it once, Sally develops the habit of working through the holidays. Sally doesn't develop the habit intentionally—she tells herself that next year she will return to be with her family for the holidays. She doesn't even develop the habit through some unconscious intention (if there be any such things). Working through the holidays turns into a habit merely because it promotes certain aspects of her well-being. It is an addictive behavior precisely because of the way it ultimately makes her feel.

Very plausibly, Sally might discover years later (perhaps through psycho-therapy) that her working through the holidays—even that very first time—had the teleological function of helping her avoid her (horrible) family (at minimal cost) while getting ahead in her job. This telos was not had—at least not in the first case—by design, as the product of some selection process, or because the ability to help her avoid her family while getting ahead in her job explains the existence of her working through the holidays. Her working through the holidays the first time was (we might suppose) a chance occurrence.

Considerations raised in the last two paragraphs may motivate drawing a distinction between having an *etiological* function—having a function in roughly Wright's sense—and having a *teleological* function in the sense of having a telos or “job” due to the establishment of teleological norms. Even if the two sorts of functions overlap in many cases, having an etiological function may not be sufficient for having a telos. (Assuming biological function is a species of etiological function, having a particular biological function may not be sufficient for having a telos either even if in particular cases, e.g. the human heart, biological function and teleological function overlap.)

Besides etiological function, there is another sort of function commonly discussed in the philosophical literature—the Cummins function.⁵² Very roughly, having a Cummins function is a matter of playing a certain role in a causal explanation of some outcome. As a result, having a Cummins function is relative to the goals of explanation. For instance, relative to the goal of explaining how plants grow, rain over land masses may have the Cummins function of providing a source of water for plants. On the other hand, relative to the goal of explaining how evaporated water ends up in the oceans, rain over land masses may have the Cummins function of putting water into the system of creeks, rivers, and lakes (which, due to the Earth’s gravity, typically run into oceans eventually). Quite clearly, having a Cummins function has very little to do with having a teleological function, i.e. a telos.⁵³ It is not the telos of rain to provide a source of water for plants or to put water into the system of creeks, rivers, and lakes.

If teleological function is not identical to etiological function or Cummins function, where does that leave us?⁵⁴ If teleological functions do not arise merely in virtue of etiological or causal properties, how do they arise? Let’s return to the example of Sally. As I see it, that Sally’s behavior is *self-sustaining because it is beneficial* appears to be sufficient for the behavior of working through the holidays to have precisely this telos. This observation appears to suggest that teloi arise when properties of an object persist and hence are ultimately explained due to promoting some real or perceived good. In other words, a correct theory of teleological function may well emphasize the importance of value in roughly the way Mark Bedau has suggested.⁵⁵ There’s no reason to think that this theory will have as a prerequisite of teleological function that there be a designer or user with the relevant sort of intentions. For instance, even if it fails to have any biological function, Swampman’s duplicate of a heart might have the teleological function of pumping blood because pumping blood promotes some good (for him)—Swampman’s survival and well-being—which in turn promotes the continued pumping of blood.⁵⁶

It seems to me—although I won’t try to argue it here—that the teleological aspect of belief-desire psychology might be grounded in the fact that belief-desire psychology is self-sustaining because it is beneficial, i.e.

⁵² Cummins (1975).

⁵³ Cf. Kitcher (1993).

⁵⁴ The thought that these various sorts of functions might all be distinct is consonant with those philosophers who suggest we ought to be pluralists about function. See Godfrey-Smith (1993).

⁵⁵ Bedau (1992).

⁵⁶ Swampman is the famous example of Davidson (1987).

because it promotes the good of intentional agents.⁵⁷ Quite plausibly, belief-desire psychology promotes our survival and well-being. Indeed, belief-desire psychology may well be self-sustaining because by promoting our survival and well-being, it promotes the survival and upkeep of our cognitive architecture. Thus, this self-preservation of belief-desire psychology clearly serves a purpose that is beneficial to us. As I see it, these features alone are sufficient for belief-desire psychology to have a telos.⁵⁸ Whether or not its telos involves representation akin to that of maps and blueprints—whether TTIR is true—depends entirely on whether beliefs and desires benefit us by acting as representations of this sort. It is plausible to me that they do, but I will not try to argue the point here.

I make these suggestions not because I am here interested in assaulting the etiological tradition of telos, but rather because I want to illustrate that there are alternatives to that tradition that may be more palatable to some depending on their theoretical commitments.⁵⁹ The important point is really that, in the current philosophical literature, there is an emerging consensus that having a telos is possible even absent a pre-existing intentional agent to bestow that telos.⁶⁰ This consensus transcends the debates over the etiological tradition. In fact, it may even transcend debates over whether putative biological teloi are legitimate.⁶¹ Even if what biologists refer to as “biological function” is not a species of genuine *telos* (and is instead some other sort of non-teleological function), it still may seem very plausible to accept that there are teloi that are not the product of pre-existing intentional states. I created the Sally example—which is by design not an example relevant to the philosophy of biology—to suggest that very point.

The emerging consensus that teloi can arise without a pre-existing intentional agent undercuts the worry that TTIR might be subject to the sort of vicious regress raised at the beginning of this section. Absent further argument to the contrary, I see no immediate reason for concern.

Conclusion

Proper motivation for CIN offers independent reason for rejecting the Glüer-Wikforss premise that any norms inherent to content are regulative—that these norms guide performance. Because opponents of CIN similarly tend

⁵⁷ Cf. Bedau (1992). Thus, I think psychology is likely to be autonomous of biology.

⁵⁸ This is not to say that less might not also be sufficient.

⁵⁹ See Braddon-Mitchell and Jackson (2007), pp. 210-215 for a brief discussion.

⁶⁰ Cf. Bedau (1990).

⁶¹ I have some sympathy with skeptical dissenters. See Bedau (1991).

to deploy putative normative constraints that do not generally apply to teleological norms, these opponents likewise generally fail to show that belief-desire psychology is not inherently teleological. Absent any such demonstration, we have no reason to reject CIN.

Acknowledgements

I would particularly like to thank Richard Heck, Christopher Hill, Joshua Schechter, Jamie Dreier, Jonathan Ichikawa, Michael Young, and Katherine Rubin for their comments and discussion on this or other closely related projects. This paper was presented at a research seminar jointly hosted by Arché Research Centre (The University of St. Andrews) and Centre for the Study of Mind in Nature (The University of Oslo). I offer my gratitude to the participants of that seminar.

References

- Allen, C. & Bekoff, M. (1995). Function, natural design, and animal behavior: Philosophical and ethological considerations. In N. S. Thompson (Ed.), *Perspectives in Ethology 11: Behavioral Design* (pp. 1-46). New York: Plenum.
- Ariew, A., Cummins, R., and Perlman, M. (Eds.). (2002). *Functions*. Oxford: Oxford University Press.
- Austin, J. L. (1950). Truth. *Aristotelian Society Supplementary Volume* 24, 111-129.
- Bedau, M. (1990). Against mentalism in teleology. *American Philosophical Quarterly* 27, 61-70.
- Bedau, M. (1991). Can biological teleology be naturalized? *Journal of Philosophy* 88, 647-55.
- Bedau, M. (1992). Where's the good in teleology? *Philosophical and Phenomenological Research* 52, 781-806
- Block, N. (2007). *Consciousness, function, and representation*. Cambridge, MA: MIT Press.
- Boghossian, P. (1989). The rule-following considerations. *Mind* 98, 507-49.
- Boghossian, P. (2005). Is meaning normative? In C. Nimtz & A. Beckermann (Ed.) *Main lectures and colloquia of GAP.5, Fifth International Congress of the Society for Analytical Philosophy, Bielefeld, 2003*. Mentis.
- Braddon-Mitchell, D. & Jackson, F. (2007). *Philosophy of mind and cognition*. Malden, MA: Blackwell Publishing.
- Brandom, R. (2000). *Articulating reasons*. Cambridge, MA: Harvard University Press.
- Boorse, C. (1976). Wright on functions. *Philosophical Review* 85, 70-86.
- Cummins, R. (1975). Functional analysis. *Journal of Philosophy* 72, 741-65.
- Davidson, D. (1987). Knowing one's own mind. *Proceedings and Addresses of the American Philosophical Association*, 441-58.
- Dretske, F. (1981). *Knowledge and the flow of information*. Oxford: Blackwell Publishing.
- Dretske, F. (1986). Misrepresentation. In R. Borgan (Ed.), *Belief: Form, content, and function*. Oxford: Oxford University Press.
- Dretske, F. (1988). The explanatory role of content. In R. Grimm & D. Merrill (Eds.), *Contents of thought*. Tucson, AZ: The University of Arizona Press.
- Dretske, F. (2000a). Norms, history, and the constitution of the mental. In Dretske (2000b), pp. 242-58.
- Dretske, F. (2000b). *Perception, knowledge, and belief*. Cambridge, England: Cambridge University Press.
- Dummett, M. (1959). Truth. In his *Truth and Other Enigmas*. Cambridge, MA: Harvard University Press, 1978.
- Field, H. (1994). Deflationary views of meaning and content. *Mind* 103, 249-285.
- Fodor, J. (1990). *A theory of content and other essays*. Cambridge, MA: MIT Press.
- Gibbard, A. (2003). Meaning and normativity. *Philosophical Issues* 13, 83-98.
- Glüer, K. & Wikforss, Å. (2009). Against content normativity. *Mind* 118, 31-70.

- Godfrey-Smith, P. (1993). Functions: Consensus without unity. *Pacific Philosophical Quarterly* 74, 196-208.
- Godfrey-Smith, P. (1994). A modern history theory of functions. *Noûs* 28, 344-62.
- Grice, H. P. (1989). *Studies in the ways of words*. Cambridge, MA: Harvard University Press.
- Hattiangadi, A. (2006). Is meaning normative? *Mind and Language* 21, 220-40.
- Horwich, P. (1995). Meaning, use, and truth. *Mind* 104.414, 355-68.
- Horwich, P. (1998). *Meaning*. Oxford: Clarendon Press.
- Horwich, P. (2005). *Reflections on meaning*. Oxford: Oxford University Press.
- Horwich, P. (2006). The value of truth. *Noûs* 40.2, 347-60.
- Kitcher, P. (1993). Function and design. In P. French & H. Wettstein (Eds.) *Midwest Studies in Philosophy XVIII* (pp. 379-97). Oxford: Blackwell Publishing.
- Korsgaard, C. (1986). Skepticism about practical reason. *Journal of Philosophy* 83, 5-26.
- Korsgaard, C. (1996). *The sources of normativity*. Cambridge, England: Cambridge University Press.
- Kripke, S. (1982). *Wittgenstein: On rules and private language*. Cambridge, MA: Harvard University Press.
- Lewis, D. (1984). Putnam's paradox. *Australasian Journal of Philosophy* 62, 221-36.
- Loewer, B. and G. Rey (Eds.). (1991). *Meaning in mind: Fodor and his critics*. Oxford: Blackwell Publishing.
- Millikan, R. (1984). *Language, thought, and other biological categories*. Cambridge, MA: MIT Press.
- Millikan, R. (1989). In defense of proper functions. Reprinted in Millikan (1993), pp. 13-30.
- Millikan, R. (1990). Truth rules, hoverflies, and the Kripke-Wittgenstein paradox. Reprinted in Millikan (1993), pp. 211-240.
- Millikan, R. (1993). *White queen psychology and other essays for Alice*. Cambridge, MA: MIT Press.
- Neander, K. (1991a). The teleological notion of function. *Australasian Journal of Philosophy* 69, 454-68.
- Neander, K. (1991b). Functions as selected effects: The conceptual analyst's defense. *Philosophy of Science* 58, 168-84.
- Papineau, D. (1999). Normativity and judgment. *Aristotelian Society Supplementary Volume* 73, 17-43.
- Prinz, J. (2002). *Furnishing the mind*. Cambridge, MA: MIT Press.
- Putnam, H. (1975). The meaning of "meaning." In his *Philosophical Papers, vol. 2: Mind, Language, and Reality*. Cambridge, England: Cambridge University Press.
- Putnam, H. (1981). *Reason, truth, and history*. Cambridge, England: Cambridge University Press.
- Stalnaker, R. (1984). *Inquiry*. Cambridge, MA: MIT Press.
- Tanney, J. (1999). Normativity and thought. *Aristotelian Society Supplementary Volume* 73, 45-61.
- Wedgwood, R. (2007). *The nature of normativity*. Oxford: Oxford University Press.
- Whiting, D. (2010). Should I believe the truth? *Dialectica* 64(2), 2123-224.
- Wright, C. (1992). *Truth and objectivity*. Cambridge, MA: Harvard University Press.
- Wright, L. (1973). Functions. *Philosophical Review* 82, 139-68.