

Desiring Under the Proper Guise

Michael Milona and Mark Schroeder

1.1 Introduction

According to the thesis of the guise of the normative, all desires are associated with an appearance or judgment to the effect that the object of desire, or at least something systematically related to the object of the desire, satisfies an important normative property or figures in an important normative relation. This thesis has held promise, according to many, to explain why motivation by desire *rationalizes* or *makes intelligible*, and doesn't merely explain, action.¹ And it holds promise to yield a non-mysterious source for normative epistemology² that is not debunked by empirical observations that affect has strong effects on our moral judgments.³

The guise of the normative can take different forms; in its most familiar guise, it is the thesis that the desire that p represents it as being *good* for it to be the case that p . This thesis is a version of the guise of the *good*. But another version of the guise of the normative holds that when you desire that p , facts about which actions would help to bring it about that p are represented to you as *reasons* to perform those actions. This thesis is a version of the guise of *reasons*. Importantly, *both* of these versions of the guise of the normative hold promise to make good on the key motivations for that thesis. In this paper we will be arguing for the *comparative* thesis that the guise of reasons thesis is more promising than the guise of the good. The guise of the normative should stand or fall with the guise of reasons.⁴

Many authors have discussed the guise of the good or the guise of reasons without explicitly taking on the comparative question of which view is to be preferred. For example, Anscombe (1963/2000), Davidson (1969/2000), Quinn (1993), Oddie (2005), and Tenenbaum (2007), just to mention a few, all defend forms of the guise of the good with little if any explicit commentary on why the guise of the good is to be preferred to the guise of reasons. In contrast, Scanlon (1998) and Bedke (2009) defend forms of the guise of reasons, with little if any explicit commentary on why the guise of reasons is to be preferred to the guise of the good.

¹ See particularly Anscombe (1963/2000), Quinn (1993), Tenenbaum (2007).

² Oddie (2005), Schroeder (2007).

³ Compare especially Greene (2008), Haidt (2001).

⁴ Many philosophers hold the similar view that *emotions* involve normative representations. See, for instance, Roberts (2003), Döring (2007), Tappolet (2012). The general strategy that we use to identify the most plausible form of the guise of the normative can also be used to get leverage in debates about the normative content of emotions.

But the fact that this question has not, in general, been tackled explicitly does not mean that it is not important. In particular, the guise of the normative cannot be true unless some *particular* thesis about what normative features are represented is true. And for many purposes, it will matter greatly which form of the guise of the normative is held true. For example, as Alex Gregory (2013) has argued, there are a number of purported objections to the guise of the normative which are really best understood as objections to the guise of the good. And for all that Gregory has shown, there could also be other objections to the guise of reasons that do not apply to the guise of the good. So in general, the guise of the normative cannot be adequately defended without showing that some particular version of the thesis is defensible.

Moreover, for many of the philosophical *applications* of the guise of the normative, it makes a great difference both for how those applications will go and for how plausible they will be, which form the guise of the normative takes. For example, we happen to be particularly interested in the *epistemological* applications of the guise of the normative. For if desires involve perceptions of the normative, then moral or more broadly, normative epistemology could have a perceptual basis, analogous to the justificatory story for ordinary empirical beliefs. But it makes a very important difference whether we think that our basic epistemic access to normative facts runs through reasons, or through the good.⁵

The structure of our argument is simple. In order for the normative to figure in any way at all in the guise under which we perceive things when we have desires, the thesis that desires are associated with *contents* about the corresponding normative property or relation must be defensible. But there

⁵ In correspondence, we have come across just one argument that it is ultimately a non-issue which of the guise of good or the guise of reasons is accepted. This argument begins with the assumption that reasons are analyzable in terms of the good or the good is analyzable in terms of reasons, and infers from this that if desires involve perceptions of reasons, then they will also, *ipso facto*, by virtue of the analyzability of the good in terms of reasons, involve perceptions of the good, and similarly if desires involve perceptions of the good. One important point that emerges from this argument, which we do not want to deny, is that it is *possible* that both some form of the guise of reasons and some form of the guise of the good are true – assuming that some analysis of reasons in terms of the good or the good in terms of reasons is suitably transparent, and that the *right form* of the guise of the normative is true.

But what the argument does *not* succeed at showing, is that if *any* form of the guise of reasons is true, some form of the guise of the good is true, or conversely. On many proposed analyses of the good in terms of reasons, for example, the analysis takes the form of a *fitting attitudes* account (Rabinowicz and Rønnow-Rasmussen (2004), for example), on which the good is analyzed in terms of reasons *to desire*, or for other attitudes. But on many forms of the thesis of the guise of reasons, desires are associated with representations of reasons for *action*, rather than for attitudes, in all or at least some cases. But when we put these two views together, it turns out that so far from entailing that the guise of the good turns out to be true after all, they actually entail that the guise of the good *cannot* be true – for the reasons that are represented by desires are the wrong kind to be involved in the analysis of the good.

are constraints on how it could be that mental states are associated with contents. And in particular, these constraints make it harder, other things being equal, for states to be associated with contents that are systematically mistaken. But on each of the most plausible ways of associating desires with representations of what is good, we will argue, desires *are* systematically mistaken, and other things are equal. This presents the thesis of the guise of the good with a serious challenge.

This challenge arises because *good* and *better* are properly predicated of propositions or *outcomes*, and this poses a *constraint* on how the contents represented by the desires of different agents could be true together. This same fact about *good* and *better*, we note, is responsible for historical puzzlement about the possibility of agent-centered constraints in moral philosophy. Reflection on the structure of agent-centered constraints, we argue, shows that what we ought to do is better explained in terms of *reasons* than in terms of what is *good*, and that reasons support *properties*, rather than outcomes. A similar lesson transfers over to the case of the guise of the normative, where it supports the conclusion that desires represent reasons in support of properties – *means* – to achieving the object of the desire. The thesis of the guise of reasons therefore answers our challenge to the guise of the good, and moreover it does so in the way that we should expect, given the proper diagnosis of the source of the difficulty. This, we argue, is what makes it a preferable view.

But before we can lay out this argument, we must make clearer what is at stake. So sections 1.2, 2.1, and 2.2 clear the ground for our main argument by clarifying how to think about the thesis of the guise of the normative in the most general case, and about the objects of both desire and goodness. In sections 3.1 and 3.2 we argue for an important high-level constraint on what it takes for any mental state to be associated with any particular content, and use that constraint to raise our central challenge for the thesis of the guise of the good. In sections 4.1 and 4.2 we turn to show that philosophical puzzlement about the possibility of agent-centered constraints has been driven by a closely analogous structural problem, and explain what we believe to be its solution. And finally, in section 5 we leverage this discussion of constraints to return to the cases that raise troubles for the guise of the good, where we will show that the thesis of the guise of reasons handles the cases straightforwardly – as should be expected. Our moral is that it should be no more surprise that these cases pose trouble for the guise of the good, than that agent-centered constraints pose trouble for traditional act-consequentialism. Each arises, we will argue, because desire and obligation are concerned with *properties* an agent might have – paradigmatically including *actions* she might perform – but goodness is a property of *propositions* or outcomes.

1.2 The Guise of the Normative

Traditionally, many distinct theses have fallen under the heading of ‘the guise of the good’ or, generalizing, ‘the guise of the normative’. But we have in mind a very particular target. According to the thesis of the guise of the normative as we conceive of it,

Generic Guise There is some particular normative property or relation N and some other relation R , such that every time an agent has a desire with object o , she perceives things that are R -related to o as being N .

This formulation is intended to be neutral about a number of things. It is neutral, obviously, about which normative property N is represented by desire. This allows the thesis to be common ground between those who endorse the guise of the good, and those who endorse instead the guise of reasons. It is also neutral about the relationship of priority between perceiving things as satisfying this normative property, and having the desire. So it can be common ground between those who hold that having a desire *consists* or *partly* consists in perceiving these normative things, and those who hold instead that there are independent facts about the nature of desire, of which it is a *consequence* that desires are associated with representations of these normative things.⁶

Our statement of Generic Guise is also intended to be neutral among conceptions of the object of desire. Some philosophers who have discussed the guise of the good have seemed to presuppose that the objects of desire can be individuals, such as when you desire a cup of coffee. On such a view, desire must involve a perception of some normative state that is systematically related to the object of desire – in this case, to the cup of coffee. On another, more general, view, the objects of desire are propositional. So on this view, in the most general case, every desire is a desire that p . Given this sort of view, desire must involve a perception of some normative state that is systematically related to the propositional object of desire.

Our statement of Generic Guise includes one important complication – it allows, but does not insist, that a desire involves a perception of its object as satisfying some normative property. It allows this, because one possible candidate for R is the relation of identity. So one form of the guise

⁶ For example, Schroeder (2007, chapter 8) claims to differ from Scanlon’s (1998) account of attention-directed desire in virtue of explaining the perceptual content of desires in terms of an independent account of the nature of desire, rather than taking desire to be constituted by contents about reasons. We take no stand, here, on whether such an account might work; we only note that our question is of interest to both sides of this dispute, and so we have tried to frame it in a neutral way.

of the good is the thesis that the proper objects of desire are propositions, that R is the relation of identity, and that N is the property of being impersonally good – the property that we express by saying things like ‘it would be good if there were fewer conflicts’, and which is appealed to by traditional, agent-neutral act consequentialism. On this view, whenever you have a desire, you desire that p , and you perceive it as being good that p . But holding that the proper object of desire is an individual does not, on our formulation, prevent you from holding that desire is under the guise of the impersonal good. For you could hold that whenever you desire, you desire some individual x , and you perceive it as being good that you possess x . On this view, the objects of desire are individuals, N is the property of being impersonally good, and R is the relation of being the proposition that you possess o .

One dimension along which our formulation of Generic Guise is definitely not neutral, is that we specify that desire is associated with the *perception* that something has the relevant normative property.⁷ It is possible to hold similar views according to which desire is instead associated with something more akin to a *judgment* that something has a normative property.⁸ We believe that the examples that we will be discussing in section 3.2 can also be used, with care, to construct a related challenge to **judgmentalist versions of** the guise of the good, but that the issues involved are sufficiently different that the two views are worth being kept separate. In any case, we also believe both that the perceptual version of the view is more promising, and note that the perceptual version of the view is *prima facie* better positioned to play a novel role in normative epistemology, which we took to constitute part of the philosophical interest of the guise of the normative. So for each of these reasons, we will take it as given in what follows. Generic Guise is not committed, however, to the claim that you must in any *literal* sense perceive that something has the relevant normative property; the state it is committed to calling a ‘perception’ earns the title through its analogous phenomenological and epistemological role to ordinary perceptions, rather than by being a special case of vision, hearing, or the like.

A second dimension along which our formulation of Generic Guise is not neutral, is that it assumes that there is *uniform* relationship between desire and the representation of some normative property or relation. This is a substantive assumption, and one that it is possible to relax. On a strict generalization of Generic Guise, there is a function \mathcal{F}_1 from situations s to normative properties or

⁷ The perceptual model of the guise of the good arguably traces back to Aristotle. See Moss (2013) for careful discussion of Aristotle’s view.

⁸ Gregory (2013), for example, which otherwise offers an account particularly similar to ours, defends a view on which desires involve *judgments* about reasons, we will be focusing on views according to which desires are systematically related to *perceptions* of reasons.

relations N and a function \mathcal{F}_2 from situations s to relations R such that whenever someone has a desire with object o in situation s , she perceives things that are $\mathcal{F}_2(s)$ -related to o as being $\mathcal{F}_1(s)$. Later we will see that one of the best hopes for defending the guise of the good in response to our argument will be to distinguish between different *good* properties and generalize in this way, so this will be important, later.

So our question, in this paper, is not whether Generic Guise is true, but rather, if it is, what the appropriate values are for N and R , and hence, in addition, what the object is, of desire in general. In the next section, we will narrow this question by defending a general answer about the nature of the objects of desire.

2.1 Propertitional Desire

Ascriptions of desire can be divided grammatically into three classes: determiner-phrase (DP) desire, non-finite desire, and complementizer phrase (CP) desire. These three classes correspond to the following examples:

DP desire	Sue desires a coke.
non-finite desire	Sue desires to drink a coke.
CP desire	Sue desires that she drinks a coke.

Each of these grammatical distinctions corresponds to a distinct possible object of desire. In DP desire ascriptions, the verb ‘desire’ or its cognates, ‘desires’, ‘desired’, or ‘desiring’, is followed by a determiner phrase such as ‘a coke’. These ascriptions correspond to the idea that when you have a desire, you desire a thing. In non-finite desire ascriptions, ‘desire’ is followed by a non-finite clause, such as ‘to drink a coke’. These ascriptions correspond to the idea that when you have a desire, it is a desire to do some thing or to be some way. And in CP desire ascriptions, the verb ‘desire’ is followed by a complementizer phrase such as ‘that Sue drinks a coke’ or ‘for Sue to drink a coke’. These ascriptions correspond to the idea that when you have a desire, it is a desire for something to be the case.

Despite the fact that these three forms of desire ascription appear to make ‘desire’ out to take three different kinds of object, they cannot plausibly be maintained to be independent.⁹ In general, whenever a DP ascription such as ‘Sue desires a coke’ is true, there is a corresponding true non-finite

⁹ Contrast Brewer (2006), (2009) and Thagard (2006).

ascription, which says that Sue desires to bear some important *relation* to a coke. For example, ‘Sue desires a coke’ might be true because Sue desires to drink a coke, or because Sue desires to own a coke, so that she can give it to someone else. But it can’t be true unless there is some way in which Sue desires to be related to a coke. So every DP desire ascription has a corresponding non-finite desire ascription, with which it is equivalent. But the converse is not true. ‘Sue desires to fly’ does not correspond straightforwardly to any DP desire ascription.

This asymmetry, and the fact that non-finite desire ascriptions make explicit structure that is only understood from context in DP desire ascriptions, show that DP desire ascriptions should be understood *in terms of* non-finite desire ascriptions. There are not two kinds of desire – desires for things and desires to be some way. There is just desire to be some way, and among the ways that there are desires to be, are being related to things in certain contextually obvious ways.¹⁰ So it would be a mistake of insufficient generality to try to formulate a guise of the normative thesis in terms of objectual desire.¹¹

There is also a close correspondence between non-finite desire ascriptions and propositional desire ascriptions. If Sue desires to drink a coke, then her desire is not satisfied unless *she* drinks a coke. So she desires that something be the case – namely, that she drinks a coke. And under normal circumstances, if Sue desires that she drinks a coke, then she desires to drink a coke. So the orthodox view about the objects of desire – the view that desires are propositional attitudes – concludes that non-finite desire ascriptions can in turn be understood in terms of propositional desire ascriptions. On this view, since objectual ascriptions reduce to non-finite ascriptions, which in turn reduce to propositional ascriptions, it follows that all desires are essentially, at bottom, propositional. We gloss this claim as the view that the objects of desire, strictly speaking, are propositions.¹²

The view that the objects of desire are propositions is often misunderstood. It cannot be reported by saying that what we really desire, are propositions.¹³ To gloss the view in this way is to use an objectual desire ascription: ‘what I desire is a proposition’. But according to the view that desire is

¹⁰ This is *not* to say that there are not other attitudes that are fundamentally objectual. See Montague (2007), Grzankowski (2015).

¹¹ Brewer often formulates the guise of the normative in terms of objectual desire. For example, he says, “When we desire a person...the apparent goodness brought to light by the desire is the actual goodness of someone already wholly present” (2009: 62-3). He counsels against trying to understand many of our person-directed desires as being anything other than an attitude toward an individual. But for reasons we have already pointed out, we think that this is a mistake (see also Sinhababu (2015) for criticism).

¹² This view is encapsulated by the orthodox linguistic view that non-finite clauses contain a covert *referential* pronoun, PRO, which refers to the subject of the attitude ascription. For a textbook presentation of the motivations for postulating PRO, see Radford (2004).

¹³ Compare Rüdiger Bittner (2015, 8), Brewer (2009).

a propositional attitude, in order to understand objectual desire ascriptions, we have to understand them in terms of non-finite ascriptions. So it is true that I desire a proposition only if I desire to stand in some contextually salient relation to that proposition – such as, for example, owning it or holding it in my hands. But no one is claiming that all desire is desire to own or hold propositions, or the like. So it is extremely misleading to gloss the view that desire is a propositional attitude as the view that what we really desire, are propositions. So when we gloss this view by saying that the objects of desire are propositions, this is not what we mean.

However, even if desire is, in some sense, a propositional attitude, CP-desire ascriptions are *not* the most general form for ascriptions of desire. This is because though in general, if Sue desires to drink a coke, Sue desires that *she* drinks a coke, the entailment does not go the other way around. Suppose, for example, that Sue is in a busy shopping center, looks across the crowd, sees someone glossy-eyed and grumpy-looking, and forms the desire that she – the person who Sue sees looking glossy-eyed and grumpy-looking – drinks a coke (so that she perks up). If it turns out that Sue is seeing herself in a mirror¹⁴, then though it is true that Sue desires that she drinks a coke, it is not true that Sue desires to drink a coke. In general, non-finite desire ascriptions truth-conditionally entail that their subjects conceive of themselves in the normal, first-personal, way in the content of the desire, but propositional desire ascriptions do not. What this shows is that non-finite desire ascriptions are more fine-grained than propositional desire ascriptions. So insofar as we are trying to formulate a *general* thesis about what is represented by desires, we should focus on the case of non-finite desire ascriptions.

But non-finite clauses, we believe, do not express propositions. Some verbs, such as ‘neglect’, take non-finite clauses as their arguments, as in ‘Jake neglected to come to the talk’. But ‘neglects’ does not express a relation between agents and propositions. Jake cannot neglect an arbitrary proposition; he can only neglect to do something. Since the proper object of ‘neglects’ is a property, we can infer that ‘to come to the talk’ expresses a property.¹⁵ So we may infer that ‘Jake desires to come to the talk’ expresses a relation between Jake and this same property – the property of coming to the talk. Desire, in the most general case, is a relation toward properties – a *propertitional* attitude, we might say, rather than a *propositional* attitude.¹⁶

The resulting view of desire resembles David Lewis’s (1979) claim that desire is an attitude toward properties, rather than propositions. We don’t, however, accept all of Lewis’s commitments, and particularly not his idea that attitudes like desire involve *locating* oneself in physical and logical

¹⁴ Compare the classic cases discussed by Mach (1897) and Perry (1979).

¹⁵ See Schroeder (2011), especially footnote 31.

¹⁶ Compare Chierchia (1989).

space. What is important for us in this paper is not the proper metaphysics of desire nor the proper semantics of ‘desire’ ascriptions, but just fixing on the most *general case* of desire. Because non-finite desire is always present when objectual or propositional desire ascriptions are true, but makes more discriminations than objectual or propositional desire ascriptions, it is what we must fix on in order to evaluate general claims such as Generic Guise.

Consequently, from here forward, we will assume that in the fully general case, the objects of desire are *properties*, and we will evaluate each implementation of the Generic Guise thesis in that light. So our question is, when an agent desires *to F*, what is the normative property or relation *N* and the relation *R* such that she perceives things that are *R*-related to *F* as being *N*? In section 3.1 we will argue for a very high-level constraint on the answer to this question, and in section 3.2, we will employ that constraint in order to argue against the thesis that *N* is the property of being *good*. But first, we must get clear on one more thing, in order to be able to make the thesis of the guise of the good precise.

2.2 The Structure of Good

What makes the guise of the good distinctive is that it claims that the relevant normative property *N* is the property of being *good*, or some closely related property. So in order to formulate a version of the guise of the good precisely, we need to know what the relation *R* is. In order to characterize the relation *R*, however, we need to know what sort of thing the property *good* applies to – *R* must be a relation that relates *that* kind of thing to properties (the objects of desire).

As with ‘desire’, attributions of ‘good’ can be grammatically distinguished as falling into several different classes. In fact, with ‘good’, there are many more possible surface forms of attributions, but they can be naturally classified as falling under three categories:

CP good	It is good that you came./It would be good for you to come.
good for	It is good for you to eat your vegetables.
attributive good	Those are good vegetables.

In attributive ‘good’ attributions, ‘good’ functions as a predicate modifier, transforming some predicate such as ‘vegetables’ into a complex predicate, ‘good vegetables’. Such uses of ‘good’ can be applied to individuals, but strictly speaking ‘good’ does not pick out a property of individuals, but only transforms kind terms like ‘vegetables’ into corresponding properties of individuals, *good vegetable*. It is implausible that the guise of the good theory should be construed in terms of attributive good, since it’s not clear

in general *which kind* it is such that desires should be associated with the perception that things are good things of that kind. So we set it aside.

Some other sentences attributing ‘good’ *do* seem to predicate ‘good’ directly of individuals. For example, consider the sentence, ‘vegetables are good’, which appears to predicate ‘good’ of vegetables. But this is a mistake; this sentence is systematically related to sentences like ‘vegetables are good for you to eat’, ‘it is good for you to eat vegetables’, and ‘vegetables are good to eat’. Robert Shanklin has shown that as used in all of these sentences, ‘good’ belongs to the syntactic category of what are known as *experiencer* adjectives.¹⁷ Another example is ‘fun’ for which we have ‘vegetables are fun’, ‘vegetables are fun to eat’, ‘vegetables are fun for you to eat’, and ‘it is fun for you to eat vegetables’.

The basic form of experiencer-adjective sentences is ‘it is good for you to eat vegetables’, in which ‘good’ expresses a relation between you and the proposition that you eat vegetables. The word ‘it’, in this sentence, does not refer to anything; it is simply included because English grammatically demands each sentence to have a subject. This same grammatical demand can be satisfied through *movement*: the word ‘vegetables’ can be moved up to the beginning of the sentence, to get the transparently equivalent, ‘vegetables are good for you to eat’.¹⁸ The sentences ‘vegetables are good to eat’, ‘vegetables are good for you’ and ‘vegetables are good’ are just elisions of this sentence. So though they all *appear* to predicate ‘good’ of vegetables, this is a grammatical illusion. Really, they all predicate ‘good’ of a relation between someone and some proposition. They belong to the category of *good for* sentences – what we might call *personal* good.

The final category of ‘good’ attributions are CP attributions. In these sentences, ‘good’ is grammatically predicated of some complementizer phrase, such as ‘that you came’ or ‘for you to come’ or ‘if you come’, to get sentences like ‘it was good that you came’, ‘it would be good for you to come’ or ‘it will be good if you come’. It is important to distinguish the role of ‘for you’ in these sentences from its role in *good for* attributions, and this distinction can be brought out by contrasting sentences like ‘it would be better for Jeffrey Dahmer to be punished for him not to be punished’ with sentences like ‘it would be better for Jeffrey Dahmer to not be punished than to be punished’, both of which intuitively express truths. In the former sentence, ‘for’ is a complementizer, and in the latter, it is a preposition. The latter sentence says that Dahmer and not being punished stand in the *good for* relation; the former says that the outcome of Dahmer’s being punished is *good*.

¹⁷ For arguments and discussion, see Shanklin (2011).

¹⁸ Radford (2004) is a good introduction to the relevant issues in syntax.

CP ‘good’ ascriptions ascribe ‘good’ to propositions – the referents of complementizer clauses.¹⁹ It is this reading of ‘good’ in which traditional consequentialists are interested, when they claim that each agent ought always to perform that action such that it will be best for them to perform it. On this view, CP ‘good’ ascriptions ascribe a monadic property, *good*, to propositions – the property of what we might call being *impersonally good*. It is important to note that on some views – most famously, that of Foot (1985) and also Thomson (2003) – CP ‘good’ ascriptions do not ascribe a monadic property, because they are really just a special case of *good for* ascriptions, with the person for whom the outcome is good elided or quantified out. But if ‘good’ denotes a simple property at all, CP attributions are the best case for making that out.

We have ruled out the idea that a plausible version of the guise of the good thesis will be formulated using attributive goodness. So that leaves the possibility of formulating it with *personal* good, or with *impersonal* good. The former is a relation between a person or subject and a proposition, and the latter is a property of a proposition (both, perhaps, relative to a standard, since ‘good’ is gradable). If an agent desires *F*, then if the impersonal view is true, there must be some related proposition *P* such that she has a perception as of it being good for it to be the case that *P*. In order to get *P* from *F*, we need to supply some individual, and the only obvious individual is the agent – the subject of the attitude. So the easiest way to make sense of this view is as the view that when an agent desires *F*, she has an associated perception as of its being good for it to be the case that she is *F*.

Guise Impersonal If *X* desires *F*, then she has an associated perception as of its being impersonally good that she is *F*.

And this is intuitively what impersonal guise of the good theories hold. If Lauren desires to drink water, then she perceives it as being good that *she* drinks water.

Similarly, if an agent desires *F*, then if the *personal* view is true, there must be some related proposition or outcome, *P*, such that she has a perception as of it being good for *someone* for it to be the case that *P*. In order to make sense of this view, we need to fill in for *whom* she perceives it as being good for it to be the case that *P*, and in order to get *P* from *F*, we need to supply some individual who might have *F*. Again, the most obvious answer to each of these questions is the agent herself – the subject of the attitude. So the easiest way to make sense of this view is as the view that when an agent

¹⁹ See especially Finlay (2014) for discussion and defense of this point.

desires F , she has an associated perception as of its being good for *her* for it to be the case that *she* is F – that it is good for her to be F .

Guise Personal If X desires F , then she has an associated perception as of its being personally good for X that X is F .

And this is intuitively what personal guise of the good theories hold. If Lauren desires to drink water, then she perceives it as being good for her to drink water.

It is *possible*, of course, that a guise of the good theory could claim that personal good is represented, without it being the personal good of an individual, or without it being the personal good of the desirer. For example, she might instead perceive only that it is good for *someone* to drink water, or that it is good for *someone* that Lauren drinks water, or even that it is good for her that *someone* drinks water. Or she might even perceive that it is good for (for example) Larry that Lauren drinks water. All of these would count as versions of Generic Guise. But the quantified versions of the principle are all intuitively much too weak. And the version that appeals to another agent's personal good makes unclear what determines whose personal good is perceived, and why such perceptions do the right work in making action intelligible.²⁰

Similarly, the guise of the good theorist could claim that impersonal good is represented, without it representing that it is good for *the subject* to have the property. For example, Lauren could perceive that it is good for it to be the case that *someone* drinks water. But again, this is intuitively much too weak. None of these weaker views help to make sense, for example, of why desires rationalize the actions that they do. In consequence, we will assume in what follows, at least to get our argument started, that the best formulations of each Guise theory are Guise Impersonal and Guise Personal. These formulations will bring our challenge into particularly sharp focus. Alternative formulations will face obstacles with respect to being implausibly weak or implausibly strong, or with respect to making action intelligible, particularly in the same kinds of cases that constitute our challenge, in section 3.2.

²⁰ We take it that something like this explains why none of these alternative views seem to have been defended in the literature, to the best of our knowledge.

3.1 The Accuracy Constraint

According to Generic Guise, desires are systematically associated with perceptual states with normative contents. It is our aim in this paper to suggest that we can get a handle on the question of *which* normative contents are represented, by thinking more closely about the question of how any normative content could come to be represented, at all.

To better understand this question, consider it from the perspective of Guise Impersonal. According to this theory, if, say, John desires to win the race, John's desire is associated with a perception as of its being good that John wins the race. This theory entails, *inter alia*, that this perceptual state has a certain *content* about what is good. But we know, in general, from thinking about ordinary perceptual states, grounded in vision, hearing, smell, taste, touch, and proprioception, that it is highly non-trivial that any state counts as having the content that it does. Philosophers of mind have written many articles and books exploring different answers to the question of how it is that ordinary perceptual states come to have contents, and specifically, how they come to have the particular contents that they do. So the Guise Impersonal theorist owes a similar explanation of how it is that states associated with desires come to represent outcomes as good.²¹ Indeed, there is a greater debt of explanation, because it is much more controversial whether we have perceptions of what is good, than it is whether our ordinary perceptual states have ordinary, non-normative, contents about the world.

In the philosophy of mind, theories differ sharply over exactly what is required in order for a perceptual state to come to have a particular content. According to some theories, a covariation condition suffices; other theories are causal or have a causal component, and yet other theories require teleological conditions to be satisfied.²² Often combinations of various forms of each of these conditions are packaged together. Against the background of such disagreement, we cannot take the approach of relying on any particular theory of content determination. But it is still possible to formulate a *constraint* on content determination that is broadly acceptable to theorists of very different persuasions. Such a constraint will pose a necessary condition on a perceptual state coming to have a particular content, and can therefore be used to tell against possible normative contents that cannot plausibly be construed as satisfying that condition.

The intuitive idea behind the constraint that we will defend is that the presence (or absence) of a systematic covariation relation between an experience and what it is hypothesized to be about at least provides evidence for (or against) the hypothesis, whether or not covariation is ultimately what


²¹ Schroeder (2008) and Schafer (2013) note that this question has received very little attention.

²² Compare, for example, Dretske (1995), Neander (2013), and Millikan (1989).

explains content.²³ To give an illustration, consider the hypothesis that squarish perceptual experiences (i.e., perceptual experiences with the phenomenology as of squares) are about triangularity. We take it that this is an implausible view. On the assumption of this hypothesis, the squarish perceptual experiences of human agents with normal, well-functioning psychologies would be *systematically* in error, since as a matter of fact such agents tend to have such experiences in the presence of square objects, not triangular ones. That there would be such systematic inaccuracy tells against the initial hypothesis, and it tells against it *because* they would be inaccurate in this systematic way.

Of course, sometimes there are systematic sources of error in our perceptual experiences. Perceptual illusions, including famous cases such as the Müller-Lyer illusion, occur under systematic conditions. So it is too strong to say that *any* systematic source of error tells against the hypothesis that the perceptual state has a content that makes it inaccurate. But we can get some insight into *when* systematic error is possible by considering such ordinary cases of perceptual error or illusion.

It turns out that we actually know quite a lot about how perceptual illusion works in ordinary cases. Take the Müller-Lyer illusion, for example:



Most humans, who have been raised in an ordinary built environment, perceive one of two lines as longer than another, despite the fact that they are exactly the same length. This turns out to be easily explained by the fact that in a normal built environment, the line segment that visually appears to be longer actually *is* longer. Even though the images projected on the retina are identical in size, the line flanked by outward facing arrowheads corresponds in normal settings to something farther away; and if two objects project the same size retinal image, but one is perceived as farther away, we perceive the

²³ We take this constraint (which will be refined below) to be amenable to the explanations of perceptual content offered by Millikan (1989), Dretske (1995), Tye (2000), Prinz (2004), Burge (2010), and Neander (2013), among others. Not all of these theorists hold that covariation is part of the fundamental explanation of content. Neander, for instance, has a causal-teleological theory of content, but nevertheless she accepts that covariation is evidential. For example, Neander (2013: 29) says, “If our ancestors experienced color contrast illusions, or the light wasn’t always white light, with the result that [reddish visual experiences] were sometimes produced in the absence of red, certain pathways in the human visual system could nonetheless have been adaptive and could have been selected because *often enough* the [reddish visual experiences] they produced carried information about red.”

one that is apparently at a greater distance as larger.²⁴ Other illusions cut across the senses in interesting ways. To take a familiar example, ventriloquists make a living creating auditory illusions which are triggered by visual cues. In particular, a ventriloquist moves a puppet in a way that that corresponds to how we would expect someone to move who is speaking in the way that the ventriloquist is speaking, but the ventriloquist herself does not move her mouth at all. This helps to create the auditory illusion that the sounds are coming from a different direction than they in fact are.²⁵ In both cases, then, the dispositions which lead to the illusory perceptions are such as to produce accurate perceptions in typical environments.

Similar explanations go for other visual and auditory illusions. In fact, it should not be a surprise that this is so, because one of the standard methods of research into perceptual cognition presupposes that it is so. According to this methodology, one of the principle sources of insight into how perceptual cognition works is to pay special attention to perceptual illusions. We can learn a lot about how a perceptual process actually works by assuming that the process in an illusion is an extension of a process that is accurate in its core cases.²⁶

Here, then, is the constraint that we are looking for, formulated in a general way:

Accuracy Constraint

A hypothesis about the content of a perception is plausible only if perceptions of that sort arising out of normal, well-functioning psychologies are accurate, or else can plausibly be understood on the model of an illusion.

In the next section, we will use this Accuracy Constraint to argue that the two most natural formulations of the guise of the good – Guise Impersonal and Guise Personal – are not very plausible. This is because the desires of agents with normal, well-functioning psychologies are systematically associated with states which fail in typical contexts to covary with what is impersonally good and with what is good for the desirer.

²⁴ See Gregory (1997). This explanation has not gone uncontested (e.g., Rock (1995), Orlandi (2014)), but the important thing is not the details of the explanation but the idea that the dispositions which lead to the illusions are such as to be accurate in more typical cases, a point which is agreed on by all parties to the dispute about exactly how the Müller-Lyer illusion is to be explained.

²⁵ In general, when visual experiences influence the direction from which we perceive as sound as coming from, this is known as ‘the ventriloquism effect’ (Howard and Templeton (1966)). For a powerful version of this effect, see P.T. Young (1928).

²⁶ Textbooks on perception make liberal use of this strategy. See, for instance, Gregory (1995), Foley and Matlin (2010), Snowden et al (2012), and Goldstein (2013).

3.2 Problem Cases

We are now in a position to precisely formulate the problem for Guise Impersonal and Guise Personal – a problem to which we will later argue that the guise of reasons offers an elegant and natural solution. The easiest way to begin to see the trouble is by thinking about several of the ‘flavors’ that normal, well-functioning desires come in: egoistic, impersonal, and other-regarding but not purely impersonal. Consider Larry. Larry wants to help the homeless in his neighborhood. He doesn’t want to do this merely because doing so is a way of, say, impressing his peers or creating pleasure for himself. Of course, he’s not completely impersonal in his desires. He really does want to impress his peers, get pleasure for himself, and so on; and he doesn’t have those desires merely because he sees satisfying some impersonal desire. Other of Larry’s desires are not egoistic but also not completely impersonal; a prime example is his desire for his child to win a prestigious fellowship. Most of us are like Larry. We have a mixture of genuinely egoistic, purely impersonal, and other-regarding but not purely impersonal desires.²⁷ The difficulty for the guise of the good is that Guise Personal runs into trouble with purely impersonal desires, and Guise Impersonal struggles with egoistic desires. Moreover, other-regarding but not purely impersonal desires present similar obstacles for *both* views.

It’s easy to see how purely impersonal desires make trouble for Guise Personal. According to the theory, when an agent desires *F*, she has an associated perception as of it being good *for her* that she is *F*. So if Kunal sees that a friend of a friend is up for a promotion, and desires that she gets it, then, according to Guise Personal, he has an associated perception as of its being good for him that she gets a promotion. But unless there is some coincidence, in general it is *not* good for Kunal that friends of his friends get promotions. So the vast majority of such cases turn out to be inaccurate, on this view. Such cases can easily be multiplied.

This is not yet inconsistent with the Accuracy Constraint. To settle the question of whether Guise Personal runs afoul of the constraint, we need to answer: (i) Are desires like Kunal’s normal and well-functioning? and (ii) If so, can explain away inaccuracy in those cases on the model of an illusion? To begin, such impersonal, or more specifically, *altruistic* desires seem to arise out of normal, well-functioning psychologies.²⁸ There is nothing defective about the desiderative system of a person wants

²⁷ Some have challenged the claim that there are any desires which aren’t egoistic. The view that all desires are ultimately egoistic, or instrumental to the satisfaction of egoistic desires, is called *psychological egoism*. But we side with those who take psychological egoism to be misguided. See Sober and Wilson (1998) for forceful objections to psychological egoism.

²⁸ Altruistic desires are a paradigm case of impersonal desires, but not all impersonal desires are altruistic. Some desires have nothing to do with the welfare of another person or animal, and we may also have impersonal desires that another be harmed.

others to get promotions; we expect normal human agents to have such altruistic desires. But now consider the possibility of an illusion. If under normal evolutionary conditions, our ancestors would have seen someone up for a promotion only under conditions under which it would have been accurate to conclude that it was better for the viewer for that person to get a promotion, then that would be the right kind of thing to explain why it is that we experience an illusion when we see someone up for a promotion. But the assumptions required by this explanation seem to be straightforwardly false – and not just because promotions were rare under evolutionary conditions. Under normal conditions – and even under normal evolutionary conditions – not everything that is good for the people (or even friends of friends) one encounters is good for oneself. And similar problems seem to arise for similar styles of explanation.

It's equally straightforward why egoistic desires make trouble for Guise Impersonal. According to the theory, when an agent desires *F*, she has an associated perception as of its being good for it to be the case that she is *F*. The easiest way to begin to see the trouble is to consider egoistic desires as they arise in competitive contexts:

Race An Olympic runner, Shelly, is about to compete for the 400 meter gold. She desires to win the race and is averse to any of the other runners winning. The other runners, of course, also want to win the race and their opponents not to win.²⁹

As is often the case in such competitions, it makes no difference (or no significant difference) to the total goodness of different possible outcomes who wins the race. So if proponents of Guise Impersonal want it to turn out that Shelly's desires are associated with accurate perceptions, they will be forced to say that the other runners' desires are associated with inaccurate perceptions. This is for the simple reason that, according to the theory, Shelly represents it as good that she wins the race, and the other runners represent it as not good that Shelly wins. But it cannot be both impersonally good and not good that Shelly wins the race. Cases like Race are quite typical and could be multiplied.

Now let's think about such competitive desires in light of the Accuracy Constraint. Similar to before, we need to answer: (i) Are egoistic desires in competitive contexts normal and well-functioning? and (ii) If so, can we explain away inaccuracy in those cases on the model of an illusion? To begin, egoistic desires in competitive contexts seem arise out of normal, well-functioning psychologies. There

²⁹ This example is adapted from Oddie (forthcoming).

is nothing defective about the desiderative system of a runner who wants to win a race and is averse to other runners winning, for example; we expect normal human agents to have egoistic desires when in such contexts. Now, to clarify, one might think it would be better all things considered if we just never had egoistic desires; but we should not confuse what would be best all things considered with the notions of normal and well-functioning relevant to the Accuracy Constraint. For example, it might be better all things considered if our visual dispositions were slightly different so that the world looked a bit prettier than it in fact does, but even if this were so, there is nothing in the relevant sense defective about our visual faculties when we (say) represent dandelions as plain yellow rather than rainbow colored.

If the agents' desires in Race, and in similar cases, are well-functioning, then proponents of Guise Impersonal will need to insist that there is an illusion. That is, they will need to say that the runners' desires are well-functioning but nonetheless misrepresent the good. The trouble for such a move is that in competitive contexts, we do not seem to be responding to considerations which tend to covary closely with the impersonal good. A competitor's desires will normally be highly and systematically sensitive to thoughts of *her* glory, *her* status, *her* future endorsements and so on; and these egoistic considerations just aren't relevant to the impersonal good.

At this point, a proponent of the guise of the good may make the following maneuver. Instead of adopting *either* Guise Impersonal or Guise Personal, she may insist that somehow *both* are correct. The easiest way to make sense of this idea, is to suppose that there are actually *two* types of desire, egoistic desires which are associated with perceptions as of personal goodness and impersonal desires which are associated with perceptions as of impersonal goodness.

The problem with this view is that it fares no better with other-regarding but not purely impersonal desires. In fact, the vast majority of our impersonal desires are not *purely* impersonal at all but are rather what we call *competitively altruistic*. Such desires are centered on the desirer in a way that makes it implausible that they track the impersonal good. Consider the following, more harrowing kind of competition:

Hungry Jasmine and Jacqueline are members of the same tribal community, but they're not close family members. During one especially harsh winter, food is limited and survival is in question. On occasions when there is not enough food for everyone, Jasmine desires that her family get the food and that Jacqueline's not get it. Jacqueline's desires are similar.

In *Hungry*, as in *Race*, it makes no difference to the impersonal good who “wins” the competition, and so the guise of the good theorist, even on the disunified version of the theory, cannot have it that both Jasmine’s and Jacqueline’s desires are accurate. But at the same time, both of their desiderative systems seem to be functioning well. And, moreover, it is hard to see why we should expect illusions. Jacqueline’s desire is sensitive to the fact that a certain family is *hers*, and similarly for Jasmine. But these considerations are irrelevant to the impersonal good, and yet their desires are highly and systematically sensitive to such considerations. So once we have the Accuracy Constraint on board, it becomes very implausible not only that desires in general are systematically associated with perceptions of the impersonal good but even that *altruistic* desires are. The trouble is, we might say, with the *desirer-centeredness* of desire; even our altruistic desires are often centered on the desirer in a way that is difficult even for the disunified theory of desire to accommodate.

So far, we have not ruled out every possible version of the guise of the good. Instead of distinguishing between two kinds of desire, which represent different *good* properties, we might instead distinguish between *three*. Or we might relax the assumption that desire must be associated with a single normative property or relation, and instead assume that the way that desire represents is *context-sensitive* – there is some function which tells us, in any given situation *s*, which form of the guise of the normative is correct *for that context*. There are too many possible variations on this theme to refute each one in detail, but we believe that every such variation fails either by running into the same sorts of trouble in variants on cases of other-regarding but not purely impersonal desires, by being too weak, or by failing to play the same role in rationalizing action that we expect of a form of the guise of the normative.

For example, it is worth paying attention to *just how hard* it is to characterize a form of the guise of the good that yields plausible results about cases of competitive altruistic desires, let alone fits into a general account that makes sense of all cases of desire. The problem is that competitive altruistic desires don’t track the good of *any* individual or group – rather, at best they track a *weighted* balance of the goods of individuals. For example, Jasmine and Jacqueline may desire that each other’s families get food more than they desire that some other family outside of their tribe gets food. So their desires track the individual good of the *same* group of individuals, but weight those goods differently.

So if you want a version of the guise of the good to make Jasmine’s and Jacqueline’s competitive altruistic desires to come out as associated with *approximately true* representations of some facts about what is good for some people, these facts cannot just be facts about what is good for some *group* of people; they must be complex facts about the weighted balance of the good for some group of

people – and the weighting may need to turn out to be different for different people’s desires, and for the desires of the same person at different times! It is very hard to see how to give a non-*ad hoc* account of this kind, and it is even more puzzling how it could give us anything helpful for a perceptual moral epistemology. If our basic justification for normative judgments is for a complicated weighing of facts about what is good for individuals, and the weighing itself is different in different situations in ways that are not transparent to us, then it is extremely hard to see how this is going to help us to support the apparently basic normative judgments that we actually go on to make.

In this section, we have been focused on raising our challenge to the two most natural formulations of the guise of the good that we settled on in section 2.2, and to sketching how this challenge generalizes to more complicated variations on such views. In the remainder of this paper we turn to an indirect diagnosis of where our challenge comes from, and why it was to be expected. This diagnosis explains why the guise of reasons is the sort of view that we should *expect* to evade the challenge.

4.1 Constraints in Moral Philosophy

We have just been arguing that the guise of the good thesis has trouble making sense of competitive desires. This is because *good* and *better* are properly predicated of *propositions*, so getting one agent’s competitive desires to come out as *accurate* in how they represent the good forces other agents’ desires to come out as *inaccurate*. But the problem is not solved, either, by the form of the guise of the good theory which appeals to personal value, because psychological egoism is false – many, if not most, desires are for things that are *not* good for the desirer. We will now argue that this *very same structural problem* is familiar from consideration of the paradox of deontology, which is a philosophical puzzle about the possibility of agent-centered constraints in moral philosophy. And there, we believe, the puzzles have a familiar and straightforward answer: what makes actions right is the *reasons* that support them, rather than the goodness of the state of affairs that results if they are performed. A similar moral will show that what desires represent is what we have *reasons* to do.

An agent-centered constraint is an action that an agent ought not to perform, even if her performing it would result in a better outcome.³⁰ According to common-sense, deontological intuitive

³⁰ Scheffler (1982) defines agent-centered constraints slightly differently, as actions an agent ought not to perform, even if her doing so would prevent more agents from doing the very same thing. This definition is defective, because even traditional act-consequentialism holds that there are agent-centered constraints, in this sense. Act-consequentialism says that it is wrong to produce sub-optimal results, even if by doing so you can prevent more others from producing sub-optimal results (for example, by eliminating better alternatives). Scheffler’s definition

judgment, most intuitively forbidden actions are like this. It is not okay to kill an innocent person just to prevent two deaths, and that is not just because killings are worse than deaths – because under ordinary circumstances it would be wrong to kill an innocent person even to prevent two killings of innocent persons. It is impersonally better for there to be only one innocent person who is killed rather than two, but it is not, intuitively, okay for you to kill one innocent person who would not otherwise be killed, just in order to prevent the worse outcome in which two innocent people are killed.

Some philosophers have periodically struggled to understand how this could be so, and the puzzle of how it could be so is sometimes called the ‘paradox of deontology’. The reasoning behind the puzzle is simple, and goes like this: any explanation of what is wrong with killing the innocent person must appeal to something that is bad about your killing that innocent person. But by parity of reasoning, the same thing must be bad about the killings of innocent people that you prevent. And so unless you are special, we must conclude that there is *even more* of whatever tells against your killing telling in *favor* of it. But you are not special. And even if you were, it would not be okay for other people to kill, just in order to prevent you from killing. So it is hard, on this view, to understand how it could be wrong to kill, in this situation. Hence the idea that deontological intuitions are paradoxical.

There is much more to be said about this reasoning than we can say, here. But it is important to emphasize that the reasoning relies essentially on the assumption that whatever explains why it is wrong for you to steal is that there is something *bad* about you stealing. As we have argued in section 2.2, impersonal good is properly predicated of propositions or outcomes, and similar reasoning goes for badness. So this is the assumption that what counts against your killing is that some outcome that results if you kill – such as *that you kill* – is bad. And this outcome concerns *you*. But you are not special in any way. If it is bad *for you to kill*, then unless you are special in some way, it must also be bad *for others to kill*. And from this it seems inescapable that if you face a choice between killing and allowing two others to kill, you should kill.³¹

is also defective because cases involving constraints never involve the agent preventing the *very same thing* as she performs. For example, if you kill in a way that prevents two killings, it is true that ‘killing’ is an act-description under which you prevent more than you do. But your act is an act of killing that prevents killings, and the killings that you prevent are not killings that prevent killings – so you do not prevent two of the very same thing, after all. Compare Oddie and Milne (1990) and Nair (2014).

³¹ Unless, of course, it is even worse to kill in a way that prevents others from killing than it is to kill in a way that does not prevent others from killing. But in addition to being implausible, if this were true, then this would not be an example of an agent-centered constraint, after all, by our definition, because not killing would result in the best outcome.

Cases involving constraints are analogous to cases of competitive desires. There is no single account of what is impersonally good that makes each of a set of interlocking competitive desires accurate. Similarly, there is no single account of what is impersonally good – or at least, no account that satisfies some basic compelling constraints – that explains the possible agent-centered constraints facing each agent. And just as there is nothing special about you, such that we should explain the constraints facing you but not those facing other people, similarly there is nothing special about Shelly, such that we should explain why her perceptions about impersonal goodness are accurate, but not those of her competitors. So the reasons why we can't explain the wrongness of actions in terms of impersonal goodness, if we want to make sense of constraints, are closely analogous to the reasons why we can't make competitive desires turn out to be accurate as representations of impersonal good.

In the case of competitive desires, one apparent alternative was to suppose that they are instead associated with perceptions as of what is good for the agent. This, we argued, fails for similar reasons, because normal and well-functioning agents have altruistic desires. And exactly similar points go for the idea that we could explain agent-centered constraints by appeal to what is good for the agents in question. It is true that if we explain why it is wrong for you to kill an innocent person by assuming that it is bad *for you* in some way, then the only conclusion that we are led to by parity of reasoning is that it is bad *for others* for them to kill innocent people, rather than that it is bad *for you* to do so. But just as it is implausible that the objects of altruistic desires systematically happen to be good for the agents who desire them, it is implausible that it is always bad *for an agent* for it to be the case that she kills an innocent person, or the like. The fact that these are wrong things to do does not depend, in any obvious way, on whether it is plausible that doing them (or not) detracts from (or contributes to) the agents well-being.³²

4.2 Reasons Support Properties

What we've just argued, is that the paradox of deontology is shaped by the assumption that what explains why you should not do something is some fact about what is *good* or *bad*. As we argued, such

³² One strategy for resolving the paradox of deontology appeals to *agent-relative value*, which is supposed to be distinct from impersonal and personal goodness. According to this picture, we explain what an agent ought to do in terms of what is valuable-relative-to the agent. See Dreier (1993), Louis (2004). Additionally, and most importantly for our purposes, one might think that we can salvage the guise of the good by taking desires to involve perceptions as of agent-relative value. But this would be a mistake, for there is a serious question of what agent-relative value is even supposed to be, and the most promising ways of making sense of it analyze it in terms of *reasons* (e.g., what an agent has reason to desire). See Portmore (2011: 125). The resulting theory, then, would best be conceived as a version of the guise of reasons.

facts are always facts about goodness or badness as ascribed to *propositions* or *outcomes*, and this is why explanations of why *you* should not steal carry over to commitments that equally tell against your not preventing the killings of others. The simple answer to this reasoning is that the wrongness of actions is *not* explained by facts about what is good or bad. It is explained, in contrast, by *reasons* against performing those actions.

In general, a reason is always a reason *for someone* to *do some thing*. The reason relation holds between considerations or facts (which are what we identify as ‘the reason’), agents (who are the people *for whom* it is a reason), and *actions*, (as expressed by phrases like ‘to not kill’). But actions in this sense, we believe, are just a special case of properties. The action of riding a bike, for example, is the property that is shared by everyone who is riding a bike. And as properties, actions are the right sort of thing to be denoted by non-finite clauses like ‘to ride a bike’, and hence to be reported by reason ascriptions like ‘the fact that it is good exercise is a reason for you to ride a bike’.

Since reasons tell in favor of *actions*, rather than *outcomes*, and actions are *properties*, rather than *propositions*, we get a very simple explanation of why the reasoning behind the paradox of deontology fails. The explanation of why you ought not to kill appeals to a *reason* for you not to kill. This reason is a reason *for you* not to kill, but it does not support the proposition *that you do not kill*; it just supports the action of *not killing*. So by parity of reasoning, based on the observation that you are not special, leads us to conclude only that there are also reasons for other people in support of the very same action – *not killing*. Their reasons don’t concern your killing any more than your reasons concern theirs. In both cases, your reasons are just reasons not to kill.

So the idea that reasons count in favor of properties is not only well-motivated in its own right, but it offers a clean and elegant explanation of where the reasoning behind the paradox of deontology gets into trouble. Organizing the explanation of wrongness around propositions, rather than around properties, distorts the phenomena. This motivates the thought that the thesis of the guise of reasons might offer similarly better prospects for satisfying the Accuracy Constraint. We will now argue that this is right. Organizing the perceptual contents associated with desires around normative properties of propositions, rather than of properties, distorts the phenomena.

5.1 Desire and the Guise of Reasons

In order to see how the guise of reasons does better in light of the Accuracy Constraint, we first need to get clear about what the view is. According to *any* version of the guise of reasons, desires are systematically associated with perceptions as of reasons. On the version that we favor, these reasons

will be considerations which favor actions which are *means* to the satisfaction of the desire. For example, if you want to drink a coke and there are cokes in the fridge, then you will have a perception as of a reason to go to the fridge. Of course, if you don't *believe* that there are cokes in the fridge, then it *won't* seem to you that you have a reason to go to the fridge. So we need to be careful in formulating the theory. The theory says when you desire to drink a coke, what appears to you to be a reason will depend on your beliefs; if given your beliefs, some action would help bring it about that you drink a coke, then that will appear to you to be a reason. Put generally,

Guise Reasons If X desires *F* she perceives herself as having reasons to do actions which, given her beliefs, would help to bring it about that she is *F*.

Now let's work through how the guise of reasons helps with competitive desires, focusing on the case which made trouble for Guise Impersonal, Race. In this scenario, each of the competitors wants to win the race. And so according to Guise Reasons, each of the runners will see herself as having reasons do actions which, given her beliefs, would help to bring it about that she wins the race. For example, if the race is two weeks away, each runner will see herself as having reasons to eat well, train hard, and avoid distractions, among other things. And these perceptions are entirely accurate; the runners really do have reasons to take such means to help bring about the object of their desire. By contrast, Guise Impersonal ran into trouble with Race, since (*ex hypothesi*) it makes no significant difference with respect to impersonal goodness who ultimately wins. So even though the runners' desires are normal and well-functioning, it cannot be that they are each associated with accurate normative perceptions (as the Guise Reasons theorist can say). And, moreover, it is hard to see how the desires (in this case and all similar ones) can be understood on the model of an illusion.

Guise Reasons also does well when our competitive desires are altruistic rather than egoistic. In Hungry, for example, there is not enough food for all of the children and Jacqueline wants *her* children to get the food. She will thus see herself as having reasons to take steps which ensure that her children, rather than the other children, get the food. And to allow that her desires involve accurate normative perceptions, we are not forced to say that there is anything objectionably special about Jacqueline's children, for other parents in the tribe have analogous desiderative profiles to Jacqueline and accurately perceive themselves as having reasons to bring it about that *their* children get the food.

(The reasonable assumption at work here is that, when survival is at stake, we normally have reason to favor those with whom we have the closest of bonds.)

Guise Reasons also does well with altruistic desires which are non-competitive – what we have earlier called purely impersonal desires. To take a simple case, suppose that Alejandro sees someone in pain and desires for her pain to stop. According to the theory, Alejandro will see himself as having reasons to bring about the cessation of her pain (e.g., a reason to call an ambulance). And, of course, Alejandro does have such reasons. Other cases of normal, well-functioning impersonal desires will receive similar treatment. Guise Reasons, then, apparently does quite well in light of the Accuracy Constraint. It gives a simple, unified treatment of competitive desires, whether altruistic or egoistic, as well as purely impersonal desires.

That Guise Reasons has these advantages shouldn't be terribly surprising, given what we have already observed about how explaining obligations in terms of reasons offers an elegant solution to the paradox of deontology. To reiterate, what we're obligated to do (ought to do, should do, etc.) is always a property (in particular, an action), and reasons tell in favor of properties. By contrast, goodness and badness are ultimately ascribed to propositions, and so trying to explain obligations in terms of goodness and badness is much less natural. Matters are structurally similar with desire, since desires, like obligations, are *property-oriented*. In many cases, desires will have as their object properties which are also actions. For example, we might desire to win a race, ride a bike, or drink a smoothie. But even when a desire is oriented to a property which is not an action – say, a desire to be such that there is world peace – the desire still motivates us (all else being equal) to behave in ways which, given our beliefs, would satisfy the desire. Guise Reasons adds that her desire involves perceptions as of reasons to do these actions; and so the normative perception is neatly integrated with the motivational character of desire.

But there is more to the story. The action-oriented character of desire helps us to understand *why* a guise of the good theory, but not a guise of reasons theory, is likely to run into trouble with the Accuracy Constraint. A guise of the good theorist needs to find some kind of perception of goodness which parallels in the right way the diverse range of actions that our desires orient us toward. But as we have seen, the way in which our normal, well-functioning desires fluctuate between motivating egoistic and altruistic behavior makes this task very difficult. (Matters become even more complicated for a guise of the good theorist when we remember that we should want to preserve the initial motivations in action-theory and moral epistemology for having a guise of the normative theory.) Our solution to this difficulty is straightforward. Just as we can solve the paradox of deontology by taking reasons to explain obligations, so too can we solve our trouble with the Accuracy Constraint by taking

desires to involve perceptions as of reasons. The key is that reasons (in contrast with goodness) support properties, and both obligation and desire are property-oriented.

Finally, one of the concerns that has interested us throughout, is whether the thesis of the guise of the normative will have any sort of payoff in moral epistemology. One of the important points that we made in response to the increasingly complicated sorts of variants on the guise of the good that one might propose in response to the Accuracy Constraint was that such views make it increasingly puzzling how what they yield could be the basis for a plausible moral epistemology – particularly if the content represented by desires varies between situations in ways that are not transparent or are difficult to unpack. In contrast, if the contents that are represented by desires just concern the agent’s own reasons, then this fits smoothly into a plausible perceptual moral epistemology, as it is highly plausible that agent’s judgments about what it is okay or not okay to do in different situations are sensitive to her perception of her reasons in those situations.

5.2 Conclusion

The guise of the normative is an important theory which promises exciting results in action-theory and moral epistemology. But it is underappreciated how many different forms the theory can take, and it is an important task to determine which version is the most promising. In this paper, we have argued that the guise of reasons approach to the guise of the normative is more promising than the guise of the good approach.

We first argued against two of the most attractive versions of the guise of the good – Guise Impersonal and Guise Personal. In particular, we argued that both theories run afoul of the Accuracy Constraint, which is a useful tool for determining whether a state is really associated with certain contents. We then argued that the same structural problem occurs in moral philosophy when we attempt to explain all of our obligations in terms of the goodness/badness of outcomes, which is the ‘paradox of deontology’. And just as we can resolve the paradox of deontology by explaining obligations in terms of reasons, so too can we resolve our trouble with the Accuracy Constraint by taking desires to involve perceptions as of reasons.

This result is not surprising, for reasons support properties and both obligations and desires are property-oriented. Furthermore, the property-oriented character of desires helps to explain *why* a view such as Guise Reasons does well in light of the Accuracy Constraint and why guise of the good alternatives struggle. What makes us confident in this conclusion, therefore, is not simply that it easily and elegantly avoids problems that force the guise of the good into complications and qualifications.

What makes us so confident, rather, is that it proceeds so immediately from what we have argued is a deep *diagnosis* of where those difficulties come from.³³

References

- Anscombe, Elizabeth (1963/2000). *Intention*, 2nd edition. Cambridge, MA: Harvard University Press.
- Bedke, Matt (2009). "The Iffiest Oughts: A Guise of Reasons Account of End-Given Conditionals." *Ethics* 119.4: 672-698.
- Bittner, Rüdiger (2015). "A Desirer's Reason." in *Action, Reasons, and Reason*. ed. Marco Iorio and Ralf Stoecker. Berlin: De Gruyter, pp. 5-26.
- Brewer, Talbot (2006). "Three Dogmas of Desire." in *Values and Virtues: Aristotelianism in Contemporary Ethics*. ed. Timothy Chappell. Oxford: Oxford University Press, pp. 253-285.
- _____ (2009). *The Retrieval of Ethics*. New York: Oxford University Press.
- Burge, Tyler (2010). *The Origins of Objectivity*. New York: Oxford University Press.
- Chierchia, Gennaro (1989). "Anaphora and Attitudes "De Se"." in *Semantics and Contextual Expressions*. eds. Bartsch, van Benthem, and van Emde Boas. Dordrecht: Reidel.
- Davidson, Donald (1969/2000). "How Is Weakness of the Will Possible?" in *Essays on Actions and Events*. New York: Oxford University Press, pp. 21-42.
- Döring, Sabine (2007). "Seeing What to Do: Affective Perception and Rational Motivation." *Dialectica* 61.3: 363 – 394.
- Dretske, Fred (1995). *Naturalizing the Mind*. New York: Cambridge University Press.
- Dreier, Jamie (1993). "The Structure of Normative Theories." *Monist* 76: 22-40.
- Finlay, Stephen (2014). *A Confusion of Tongues*. New York: Oxford University Press.
- Foley, Hugh and Margaret Matlin (2010). *Sensation and Perception*, 5th edition. New York: Routledge.
- Foot, Philippa (1985). "Utilitarianism and the Virtues." *Proceedings of the Aristotelian Society* 94.2: 196-209.
- Goldstein, Bruce (2013). *Sensation and Perception*, 9th edition. Belmont, CA: Wadsworth Cengage Learning.
- Greene, Joshua (2008). "The Secret Joke of Kant's Soul." in *Moral Psychology, vol. 3: The*

³³ Special thanks to (omitted for blind review).

- Neuroscience of Morality: Emotion, Disease, and Development*. ed. Walter Sinnott-Armstrong. Cambridge, MA: MIT Press, pp. 35-79.
- Gregory, Richard (1995). *Eye and Brain*, 5th edition. Princeton: Princeton University Press.
- Gregory, Alex (2013). "The Guise of Reasons." *American Philosophical Quarterly* 51.1: 63-72.
- Grzankowski, Alex (2015). "Not All Attitudes Are Propositional." *European Journal of Philosophy* 23.3: 374-391.
- Haidt, Jonathan (2001). "The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment." *Psychological Review* 108: 814-834.
- Howard, Ian and William Templeton (1966). *Human Spatial Orientation*. New York: John Wiley and Sons.
- Lewis, David (1979). "Attitudes De Dicto and De Se." *The Philosophical Review* 88.4: 513-543.
- Louis, Jennie (2004). "Relativity of Value and the Consequentialist Umbrella." *Philosophical Quarterly* 54: 518-36.
- Mach, Ernst (1897). *Analysis of Sensation*. Chicago: Open Court Press.
- Millikan, Ruth (1989). "Biosemantics." *Journal of Philosophy* 86: 281-287.
- Montague, Michelle (2007). "Against Propositionalism." *Noûs* 41.3: 503-518.
- Moss, Jessica (2013). *Aristotle on the Apparent Good: Perception, Phantasia, Thought, and Desire*. New York: Oxford University Press.
- Nair, Shyam (2014). "Reasoning with Conflicting Obligations." *Mind* 123.491: 753-790.
- Neander, Karen (2013). "Toward an Informational Teleosemantics." in *Millikan and Her Critics*. eds. Justin Kingsbury, Dan Ryder and Kenneth Williford. Malden, MA: Blackwell, 21-40.
- Oddie, Graham and Peter Milne (1991). "Act and Value: Expectation and the Representability of Moral Theories." *Theoria* 57.1-2: 42-76.
- Oddie, Graham (2005). *Value, Reality, and Desire*. New York: Oxford University Press.
- _____ (forthcoming). "Value Perception, Properties and the Primary Bearers of Value." In *Evaluative Perception: Aesthetic, Ethical, and Normative*, eds. Anna Bergqvist and Robert Cowan. Oxford: Oxford University Press.
- Orlandi, Nicoletta (2014). *The Innocent Eye: Why Vision is Not a Cognitive Process*. New York: Oxford University Press.
- Perry, John (1979). "The Problem of the Essential Indexical." *Noûs* 13.1: 3-21.
- Portmore, Douglas (2011). "The Teleological Conception of Practical Reasons." *Mind* 120.477: 118-153.

- Prinz, Jesse (2004). *Gut Reactions: A Perceptual Theory of Emotions*. New York: Oxford University Press.
- Quinn, Warren (1993). *Morality and Action*. Cambridge, MA: Harvard University Press.
- Rabinowicz, Wlodek, and Toni Rønnow-Rasmussen (2004). "The Strike of the Demon: Fitting Pro-Attitudes and Value." *Ethics* 114(3): 391-423.
- Radford, Andrew (2004). *Minimalist Syntax*. Cambridge: Cambridge University Press.
- Roberts, Robert (2003). *Emotions: An Essay in Aid of Moral Psychology*. New York: Cambridge University Press.
- Rock, Irvin (1995). *Perception*. New York: Scientific American Library.
- Scanlon, T.M. (1998). *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Scheffler, Samuel (1982). *The Rejection of Consequentialism*. Oxford: Oxford University Press.
- Schroeder, Mark (2007). *Slaves of the Passions*. New York: Oxford University Press.
- _____ (2008). "How Does the Good Appear to Us?" *Social Theory and Practice* 34.1: 119-130.
- _____ (2011). "Ought, Agents, and Actions." *Philosophical Review* 120.1: 1-41.
- Shanklin, Robert (2011). *On Good and 'Good'*. PhD Dissertation, University of Southern California. <http://digitallibrary.usc.edu/cdm/ref/collection/p15799coll127/id/632791>.
- Sinhababu, Neil (2015). "The Advantages of Propositionalism." *Pacific Philosophical Quarterly* 96.2: 165-180.
- Snowden, Peter, Peter Thomson, and Tom Troscianko (2012). *Basic Vision: An Introduction to Visual Perception*, 2nd Revised Edition. Oxford: Oxford University Press.
- Sober, Elliot and Robert Sloan Wilson (1998). *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA: Harvard University Press.
- Tappolet, Christine (2012). "Emotions, Perceptions, and Emotional Illusions." in *Perceptual Illusions: Philosophical and Psychological Essays*. ed. Clotilde Calabi. New York: Palgrave MacMillan, pp. 205-222.
- Tenenbaum, Sergio (2007). *Appearances of the Good: An Essay on the Nature of Practical Reason*. New York: Cambridge University Press.
- Thagard, Paul (2006). "Desires are Not Propositional Attitudes." *Dialogue* 45.1: 151-156.
- Thomson, Judith (2003). *Goodness and Advice*. Princeton: Princeton University Press.
- Tye, Michael (2000). *Consciousness, Color, and Content*. Cambridge: MIT Press.
- Young, P.T. (1928). "Localization with Acoustical Transposition of the Ears." *Journal of Experimental Psychology* 11.6: 399-429.