PENULTIMATE DRAFT, PLEASE CITE PUBLISHED VERSION

Bernard Williams on the Guise of the Good

Francesco Orsi

University of Tartu, Estonia

orsi@ut.ee

Abstract

The guise of the good is the thesis that an agent can only want, or intentionally do or pursue something, if and because this seems good to the agent in some respect or other. Bernard Williams criticizes the guise of the good in *Ethics and the Limits of Philosophy*. In this paper I reconstruct and assess his hitherto unnoticed critical remarks. Williams's opposition is based on the idea that it takes an "extra step" to go from desiring or pursuing something to thinking of it as good. To show this, he employs an argument from elimination (with two strands: an argument from personal ends and an argument from immediate ends), and a separate argument which I call the one-desire argument. I review and reject two types of answers to Williams's objections: answers that explore different varieties of what it can mean to "think of as good", and answers that focus on different varieties of goodness under the guise of which one can desire something. I then consider whether a different, but nearby thesis, namely the guise of normative reasons, can accommodate Williams's objections, concluding that this alternative also fails. I end by looking at how this debate might move forward.

1. Introduction

Bernard Williams briefly discusses the guise of the good in *Ethics and the Limits of Philosophy* (2006 [1985]). As far as I know, Williams's critical remarks have not been taken up in the now rich literature on the guise of the good.[1] Addressing his remarks is useful both to shed further light on Williams's philosophy of action—something I will not attempt to do here—and to challenge advocates of the guise of the good to come up with convincing answers to his concerns—something I will do in the course of this article. By the guise of the good, I mean the thesis that an agent can only want, or intentionally do or pursue something, if and because this seems good to the agent in some respect or other.[2]

In section 2 I reconstruct Williams's objections to the guise of the good. Williams's opposition is based on the idea that it takes an "extra step" to go from desiring or pursuing something to thinking of it as good. To show this, Williams employs an argument from elimination (with two strands, which I call the argument from personal ends and the argument from immediate ends) and a separate argument, which I call the one-desire argument. In section 3 I highlight the novel features of Williams's objections when compared to common criticisms in the contemporary debate. In sections 4-6 I review and assess two types of answers to Williams's objections: answers that focus on what it is to "think as good" (sections 4-5), and answers that focus on the varieties of goodness (section 6). I argue that none of these replies are successful. In section 7 I consider whether a different, but nearby thesis, namely the guise of normative reasons, can accommodate Williams's objections. However, while being initially promising, upon closer inspection the guise of normative

[1] Nor have they been engaged with in the growing Williams scholarship, as can be seen from surveying collections such as: Altham and Harrison 1995, Thomas 2007, Heuer and Lang 2012, Chappell and van Ackeren 2019.

[2] I will not comment on the classical philosophical sources of this idea (Orsi 2023). Useful introductions (with plenty of references to the contemporary debate) are Tenenbaum 2013, Tenenbaum 2020, Orsi 2015.

reasons falls prey to similar problems as those faced by the guise of the good (section 8). I conclude by taking a quick look at how the Williams vs. guise of the good dispute might move forward (section 9).

2. Williams's arguments against the guise of the good

Williams's discussion of the guise of the good occurs in the context of his criticism of the Kantian project to ground morality on the nature of rational agency. After suggesting that, as rational agents, we must want a certain degree of freedom, Williams goes on to ask whether this implies or commits to thinking of our own freedom as something good. One argument for the latter would be as follows:

(1) "[W]hen an agent wants various particular outcomes, he must think that those various outcomes are good."
(2) Freedom is "involved in securing these outcomes".
(3) Therefore, an agent "would be bound to think that his freedom was a good thing" (Williams 2006: 58).

Premise (1) is understood by Williams as a version or an implication of the guise of the good, and I'll come to it shortly. Premise (2) is not targeted for criticism. One's own freedom is involved in securing the outcomes we want, presumably, simply because those outcomes come about as a result of our choosing among a reasonable amount of options. (Of course there are also outcomes, i.e. states of affairs, that we desire, which do not depend on our choice, and hence on freedom, but Williams here is only discussing what we want or are committed to wanting as agents.)

Conclusion (3), if correct, would play into the Kantian project, presumably because once we, as rational agents, take (or are committed to taking) our own freedom as a good, then it is a short step to recognizing others' freedom as a good, too, on pain of some kind of inconsistency. With the latter, morality (or at least an important portion of it) would get off the ground purely based on an argument from rational agency. So refuting this argument is crucial for Williams's dismissal of the Kantian project.[3]

Williams rejects premise (1). It is helpful to quote the passage in full:

> In any ordinary understanding of *good*, surely, an extra step is taken if you go from saying that you want something or have decided to pursue it to saying that it is good, or (more to the point) that it is good that you should have it. The idea of something's being good imports an idea, however minimal or hazy, of a perspective in which it can be acknowledged by more than one agent as good. An agent who merely has a certain purpose may of course think that his purpose is good, but he does not have to. The most he would commit himself to merely by having a purpose would presumably be that it would be good *for him* if he succeeded in it, but must even this much be involved? Even this modest claim implies a perspective that goes somewhere beyond the agent's immediate wants, to his longer-term interests or well-being. To value something, even relatively to your own interests, as you do in thinking that it would be better "for me," is always to go beyond merely wanting something. I might indeed come to put all the value in my life into the satisfaction of one desire, but if I did, it

[3] Admittedly, however, "I must regard my own freedom as a good" does not entail the claim that "my freedom constitutes a good, period" (ibid.: 59). It is only the latter claim which would commit us, as rational agents, to recognizing the value of others' freedom.

> would not simply be because I had only one desire. Merely to have one desire might well be to have no value in my life at all; to find all the value in one desire is to have just one desire that *matters* to me. (ibid.: 58-9)

Williams's argument against (1) can be formally summed up thus:

(i) According to the guise of the good, to want (or intentionally pursue) something implies thinking it as good.

(ii) Thinking something as good always involves an "extra step" beyond merely wanting or pursuing it.

(iii) If (ii) is true, then to want (or intentionally pursue) something cannot imply thinking it as good.

(iv) Therefore, to want (or intentionally pursue) something cannot imply thinking it as good.

(v) Therefore, the guise of the good is false.

The crux of the argument appears to be premise (ii): Thinking something as good always involves an "extra step" beyond merely wanting or pursuing it. In a sense, what Williams does in the quote above is simply illustrate, rather than argue for, premise (ii): it might be self-evident that to want or pursue something and to think it good in any way are different things. The burden is on the guise of the good to show otherwise, rather than on the opponents to provide arguments against the guise of the good. And if the passage does not really contain an argument for (ii), it is no wonder that the passage has not attracted attention in the debate over the guise of the good.[4]

[4] I thank a referee for inviting me to consider this reading of the text.

However, I believe that arguments supporting premise (ii) can be usefully *extracted* from the quote above, whether or not Williams's intention was to provide any of these arguments. As I show in the next sections, articulating and assessing these arguments is going to advance the debate on the merits of the guise of the good, because they pose novel and significant challenges to the guise of the good. Moreover, articulating the arguments is necessary in order to make sure that no plausible interpretation of "to think of as good", as can be given by the guise of the good, can accommodate Williams's objections. As I discuss in section 4, "to think of something as good" can mean different things (judging that it is good, having an appearance of it as good, having the attitude taking-as-good). For ease of exposition, I will continue to attribute the arguments to Williams, while remaining agnostic on whether he intended them as such.

There are two arguments that can be extracted in support of Williams's view that thinking something as good always involves an "extra step" beyond merely wanting or pursuing it. First, Williams uses what is, in effect, an argument from elimination, which itself branches out into two arguments, each covering a possible variety of value: impersonal value and prudential value. First, if "good" in "thinking as good" is the kind of value predicate which "imports an idea, however minimal or hazy, of a perspective in which it can be acknowledged by more than one agent as good", then Williams thinks it evident that not everything we want or pursue is thought of as good in this way. Williams doesn't in fact explain why, or what exactly he means by the interpersonal acknowledgment alluded to. I'll come back to this below. The claim, at any rate, is that at least sometimes we want or pursue things *without* any assumption that other agents should or may acknowledge the thing we want or pursue as

good, and this fact speaks against the guise of the good, for one ordinary sense of "good". I call this the *argument from personal ends*.

Second, in another ordinary sense of "good", "good" in "thinking as good" may mean "good for me" in the sense of "in my long-term interest". In this case, we need not be assuming acknowledgment of our ends as good by other agents. The evaluative perspective now shifts from impersonal value to prudence or the agent's long-term well-being. For example, we may know that our ends are morally bad, or harmful to most other people, but still pursue them as overall good for us. But even this much is too much to require of merely wanting something: there are "immediate wants" which need not be referred to any long-term interest of ours, and therefore need not be referred to the good (our good) in this second sense. I call this the *argument from immediate ends*.

With the argument from personal ends and the argument from immediate ends, it looks as if Williams has covered all the ordinary senses of "good" which could serve the guise of the good, and concluded that none of them makes the thesis stand. In addition to this argument from elimination, Williams goes on to support premise (ii) with a further, rather condensed argument (beginning from "I might indeed come to put all the value in my life into the satisfaction of one desire…" in the quote above), which I'll call the *one-desire argument*. Consider the starting hypothesis of someone putting *all* the value of her life into the satisfaction of *one* of her desires. As per the guise of the good, this scenario logically entails that this desire is the only desire she can actually have. Why? Because if she had a second desire, then by the guise of the good she would "put some value into" or regard as valuable also the object (or the satisfaction) of this second desire, and the starting hypothesis would then be contradicted. Therefore, a person who puts all value into one desire *must* have only

one desire. But, Williams assumes, this consequence is absurd or, at best, conceptually unwarranted: we certainly can make sense of someone putting all value into one desire despite having other desires too. In turn, if it is possible for someone to put all value into one desire while having also other desires, it follows that, *pace* the guise of the good, not all objects of desire need to be regarded as valuable by the agent. Whenever an object of desire is regarded as valuable, this must involve an extra, conceptually optional "step", which goes beyond the fact that we desire it. In Williams's preferred wording, this extra step can be described as one (or more) of our desires, but not others, *mattering* to me. Whether "mattering to me" itself involves seeing what matters to me under the guise of the good is a further, and by this point in the argument, unimportant issue—although it is open to Williams to run again the previous two arguments (from personal ends and from immediate ends) in order to show that things may matter to people also in personal or immediate ways, hence rejecting the guise of the good also as far as "mattering" is concerned.

In the course of the one-desire argument Williams adds a final remark. The guise of the good entails that, if your desiderative life in fact does consist of one desire only, then this one desire *must* confer (at least apparent) value to your life—since you would necessarily regard the object of that one desire as valuable—and this, again, is for Williams an unwarranted conclusion: "Merely to have one desire might well be to have no [even apparent] value in my life at all" (ibid.: 59). This is a potentially separate argument from the one-desire argument previously described. I will briefly come back to it in the penultimate section.

3. The novelty of Williams's arguments

What is striking about Williams's arguments is that they constitute challenges to the guise of the good that have seldom, if ever, appeared in the literature. Debates on the guise of the good have centered on nearby, but rather differently motivated, avenues of criticism. For example, a long-standing objection is that the guise of the good overintellectualizes desire (and intentional action), by requiring possession of evaluative concepts for something as simple as, say, the desire to drink when thirsty (and for the related intentional actions). This objection is typically supported by appeal to agents such as small children or non-human animals: since they do not possess evaluative concepts, it follows they cannot desire anything—not in the same sense as we, human adults, desire something. But, it is argued, clearly these agents have desires in the same sense as we, human adults, have desires. So the guise of the good must be false.[5]

Williams's arguments from personal ends and from immediate ends, however, do not make any such assumptions either regarding the conceptual capacities of young children and animals vs. those of adults, or regarding the supposed continuity in "natural kind" between the desires of humans and the desires of other animal species. Williams's arguments refer quite clearly to the agency of human adults, especially given the dialectical context—a critique of the Kantian project to ground morality on rational agency. Moreover, the one-desire argument seems to refer to the rather intellectually mature phenomenon of "putting value into" something, as a psychological activity intuitively distinct from merely desiring something. If anything, on that argument it is the guise of the good which makes much of our evaluative life too shallow, by viewing evaluation as easy to come by as merely desiring something (since all desire implies evaluation, then evaluation is there whenever we desire something).

[5] See Velleman 1992 for the criticism, and Tenenbaum 2007: 6.2, for a response.

Another common objection to the guise of the good is that it appears to be both conceptually possible and psychologically real that some people desire and intend to do something while seeing absolutely nothing good about it (as in some cases of compulsion or addiction), or even while being attracted by their action being bad *qua* bad.[6] By contrast, Williams's arguments make no assumption of conceptual possibility or psychological reality about any such controversial, and perhaps ultimately marginal, cases of intentional agency. On the contrary, Williams's focus is on acting from what I have called personal and immediate ends, which appear to be perfectly ordinary cases of agency. As for the one-desire argument, it does not ask us to imagine any *realistic* scenario where a person puts all value into the satisfaction of only one desire (or, as in the added remark, where a person only has one desire). The argument works rather as a thought experiment to tease out the implications of the guise of the good.

If there is any current objection to the guise of the good which resembles Williams's arguments, it is probably the idea that the guise of the good offers a less economical account of agency via an unnecessary detour through evaluation. Kieran Setiya has argued that nothing normative or evaluative is needed to understand how one's own reasons for action can explain one's action (Setiya 2007, Setiya 2010). In particular, we do not need to see the agent as positively evaluating certain considerations (e.g. as good for her), in order for these considerations to be the reasons for which she acts. In a similar spirit, Williams appears to suggest that invoking any notion of good is unnecessary, in particular when our action is motivated by what I have called personal or immediate ends. If and when the agent acts on an

[6] See Stocker 1979, Velleman 1992 for the criticism, and Tenenbaum 2007: 6.3, Sussman 2009, Raz 2010, Tenenbaum 2018 for different responses to it.

idea of something as being good (or good for her in terms of self-interest), this will tell us something about her—i.e., about what matters to her—but no such thing is required by the bare notions of desire, intentional action, or acting for a reason.

However, Williams's objections are different. First, unlike what Setiya does (in particular, in Setiya 2007), Williams does not contrast the guise of the good with an alternative, supposedly more economical, account of agency. His critical remarks are supposed to stand on their own feet—they are not, on their face, arguments from the better explanation, and this aspect makes Williams's arguments relatively neutral and thus more widely adoptable. Second, Williams's stance towards the guise of the good may be described as critical but not as *antagonistic*: at least on my reconstruction of his arguments, Williams might concede that, if a sufficiently ordinary and non-trivial notion of goodness (and hence of "thinking as good") could be found, then the guise of the good might be able to accommodate personal and immediate ends. In this respect Williams does not entirely close off the chances for the guise of the good. The problem, as I show in the next sections, is that no such notion of goodness (and of thinking as good) is forthcoming. Third, and finally, the one-desire argument has nothing in particular to do with Setiya's or others' concerns in the philosophy of action.

4. Replies: The varieties of thinking as good

It might seem that Williams's arguments against the guise of the good have gone unnoticed for a good reason, namely, because there are easy replies to them. There are two kinds of replies that advocates of the guise of the good might be tempted to give: the first focuses on what it means to "think of x as good" (sections 4 and 5), the second focuses on the notion of

goodness (sections 6 and 7). In what follows, I will explain how these replies are supposed to answer Williams's arguments, and argue that neither is convincing.

The first kind of reply is that Williams wrongly assumes the guise of the good to be committed to the *judgment view*. On this view, in order to think of something as good, we need to be making an evaluative judgment with the content "x is good" or "x is good for me". If *every* desire needs to carry along such a judgment, one can see why the guise of the good looks implausible regarding both personal and immediate ends. Likewise, on the judgment view there would be no room for other desires than the ones the agent "puts value into", i.e. judges as worth satisfying. But we surely seem to have desires we do not judge to be worth satisfying.

However, this is at best an optional, and in fact rather unpopular, view of "thinking x as good" among advocates of the guise of the good.[7] In the recent literature at least two alternative views to the judgment view have been advanced. One is the *appearance view*, whereby to think of something as good is just for it to appear or seem good with an immediacy analogous to that perceptual appearances. Evaluative contents like "x is good" and similar are indeed entertained in any case of desire, but not as the object of judgments which commit to the truth of what is judged.[8]

On the appearance view, personal or immediate ends are not (or need not be) referred to the good (or to my good) by an act of judgment, which explains why it might seem that thoughts

[7] Joseph Raz does consistently talk of the agent's *belief* about value being required for intentional action (Raz 2010). But it is not clear whether value belief and value judgment should amount to the same thing.
[8] See Stampe 1987, Oddie 2005, Tenenbaum 2007.

about good may not be involved at all in having those ends: we are usually aware of our own acts of judgment. But if thoughts about good are instead understood as value appearances, then Williams's arguments from personal and immediate ends seem to lose their intuitive force, because they don't rule out that personal and immediate ends might simply appear good to the agent without being *judged* as such.

As for the one-desire argument: when Williams talks about an agent "putting all value into one desire", and concludes that, on the guise of the good, such an agent must possess only one desire, he must be assuming that any other desire of the agent would likewise be tantamount to an act of "putting value" into some other object, hence contradicting the starting hypothesis. But the appearance view offers an alternative: the agent who puts all value into one desire *judges* that only that one desire is worth satisfying, but she may possess many other desires that she doesn't put value into, because their objects merely appear good to her without being further judged to be worth satisfying. The guise of the good is thus not forced to ascribe only one desire to this type of agent.

The second alternative to the judgment view is the *attitude view*, whereby "thinking of x as good" in desiring x should not be understood as a representational or contentful state at all, whether that be a value judgment or a value appearance. Instead, desire itself should be construed as the attitude of taking-as-good, in analogy with belief, which is better construed as the attitude of taking-as-true rather than as involving any state with the content "x is true". In this way, in order to desire x, one need not entertain any evaluative content to the effect

that x is good. Evaluation is as it were built into the form (or as Franz Brentano put it, the "mode of consciousness") of desire itself.[9]

Like the appearance view, the attitude view might seem to afford the guise of the good a reply to Williams's objections. In general, the attitude view will agree with Williams (*contra* both the judgment view and the appearance view) that it takes an extra step, both conceptually and psychologically, to go from desiring x to thinking of x as good, *if* the latter means "entertaining a content to the effect that x is good". A judgment or even an appearance concerning value is not inscribed into the attitude of taking-x-as-good (notice the hyphens!). However, *contra* Williams, personal and immediate ends may still refer to the good as a matter of their very form or mode of consciousness simply *qua* ends, and this possibility is not obviously ruled out by Williams's arguments. Since the form of an attitude may remain opaque to us, we may sometimes get the feeling that what we pursue as a personal or as an immediate end is not evaluated by us as good in any way. At the level of representations, our feeling may well be right: we need not have any such evaluative representation. Nonetheless on the attitude view our personal or immediate ends would necessarily be evaluated, in the sense of being the object of a taking-as-good kind of attitude—a mode of evaluation that need not be transparent to ourselves.

As for the one-desire argument, it is useful to repeat what was said before. When Williams talks about an agent "putting all value into one desire", and concludes that, on the guise of the good, such an agent must possess only one desire, he must be assuming that any other desire of the agent would likewise be tantamount to an act of "putting value" into some other object,

[9] See Tenenbaum 2008, Tenenbaum 2018, Schafer 2013. For Brentano, see Brentano 1995: 186.

hence contradicting the starting hypothesis. But, like the appearance view, also the attitude view offers an alternative: the agent who puts all value into one desire is an agent who not only desires, i.e. takes-as-good, a certain object, but also and separately *judges* that only that one desire is worth satisfying. This is compatible with her possessing many other desires that she doesn't put value into, because their objects are merely taken-as-good without being further judged to be worth satisfying. Again, the guise of the good is not forced to ascribe only one desire to this type of agent.

5. Why these replies don't work

For all their ingenuity, the replies offered by the appearance view and the attitude view seem to miss the point of Williams's objections. I will start from their two parallel answers to the one-desire argument. The problem with these answers is that, upon closer inspection, they don't really respect the starting hypothesis—the idea of an agent who puts all value into the satisfaction of one desire.

Consider the appearance view. As explained above, this view can make sense of such an agent having other desires than the one she puts all value into, because there are other things (objects of desire) that appear good to her, but are not judged by her to be good. The question, however, is how to best describe possible conflicts arising between what she judges to be good and what merely appears to be good to her. If "appearing as good" is to be conceived as a modality of evaluation at all—and, on the guise of the good, it surely has to—then these cases of conflict are best described as *conflicts of value*, because the agent would have conflicting evaluations after all. If this is correct, then it's no longer obvious that she can put *all* value into the satisfaction of one desire only—some value seems to be put, if

reluctantly and without the assent of judgment, also into the satisfaction of other, conflicting, desires. So the appearance view violates the starting hypothesis of the one-desire argument. The agent who puts all value into one desire was, in a clear sense, an agent who could face no conflicts of value, because there is only one thing she puts value into. She of course can still face conflicting motivations—other desires may be incompatible with the "special one"—but there's no need to describe each conflict of motivations as a conflict of value.

The attitude view faces an exactly similar challenge. This view, as explained above, responded to the one-desire argument by saying that an agent who puts all value into the satisfaction of one desire can also have other desires, because there are other things (objects of desire) that are taken-as-good by her, but are not judged by her to be good. The question, however, is how to best describe possible conflicts arising between what she judges to be good and what she merely takes-as-good without judging it good. Again, if "taking-as-good" is to be conceived as a modality of evaluation at all—and, on the guise of the good, it surely has to—then these cases of conflict are best described as *conflicts of value*, because the agent would have conflicting evaluations after all. If this is correct, then it's no longer obvious that she can put *all* value into the satisfaction of one desire only—some value seems to be put, if reluctantly and without the assent of judgment, also into the satisfaction of other, conflicting, desires. Therefore also the attitude view violates the starting hypothesis of the one-desire argument.

The advocate of the attitude view might reply that "taking-as-good" is not the kind of evaluation which could conflict with the evaluations issued by value judgments. Here is why. A conflict of evaluations requires the joint presence of contradictory value contents (no matter whether carried by a judgment or by a mere appearance): *x is good* vs. *x is not good*.

But the attitude view doesn't construe "taking-as-good" as reflecting or issuing any value content. So, when the agent in the one-desire argument has a desire (say, for sex) whose satisfaction is incompatible with the one desire she puts all value into (say, living a spiritual, including celibate, life), the agent is not faced with contradictory value contents. The conflict is between one taking-as-good attitude, directed at sex, with *no* such content as "some sex would be good", and a judgment *with* a content like "sex would be bad". Therefore there is no conflict of value, only conflicting motivations. As Williams wants it, the agent can after all put all value into one desire, while having other desires too.

This reply, however ingenious, won't rescue the attitude view from the one-desire argument. It might be correct that taking-x-as-good does not directly conflict with a judgment that x is bad. But the guise of the good cannot really buy the view that a conflict of evaluations requires the joint presence of contradictory value contents. Notice that the attitude view does recognize a conflict in the case described: there is a conflict between taking-sex-as-good and, effectively, taking-sex-as-bad; a desire for sex vs. a desire against sex. It won't do to claim that this is merely a conflict of desires but not a conflict of evaluations, because the attitude view, *qua* version of the guise of the good, must conceive of both "takings-as" as *bona fide* evaluations, or else the concept of good as in "taking-as-good" would be a mere embellishment rather than a distinctive feature of the view. Hence, since there is a conflict of evaluations at the level of "takings-as", then the agent did not after all put all value into one object of desire, and the attitude view, like the appearance view, violates the hypothesis made in Williams's one-desire argument.[10]

[10] As a referee noticed, it might be question-begging for Williams to simply assert that it must be possible for someone to have desires she does not put any value into. But I take it that the one-desire argument, while not conclusive, establishes at least a burden of proof on the guise of the good. I'll return to dialectical considerations in the Conclusion.

Let us now turn to the arguments from personal ends and from immediate ends. Throughout these arguments, Williams's concern is with the kind of *perspectives* that the concepts "good" and "good for someone" bring along or point towards, when these concepts are deployed to identify an agent's wants or ends. Now, it is immaterial whether these perspectives are such that the agent commits herself to them in judgment (as per the judgment view), or such that they merely appear to the agent to be valid (as per the appearance view), or again such as to characterize the form of desire, quite regardless of the agent's own evaluative representations (as per the attitude view). It doesn't matter which filler for "thinking as good" gets chosen—in each of these cases, the guise of the good holds that simply by desiring something, the agent places herself into a certain value perspective, whether that be the perspective of impersonal value or that of prudential value. It is the "good" that is problematic, not the "thinking". In this respect, the replies canvassed in the previous section simply miss the point of Williams's critique.

Williams regards placing oneself into a value perspective as inflating the notions of desire, wanting, and ultimately agency itself beyond their minimal requisites. While desires, as such, do reveal something about the agent's perspective, and even perhaps about her identity, there is no need to construe that as an evaluative perspective. The attitude view, whereby the value perspective that comes with desire does not require forcing the ascription to the agent of any particular value judgment or value appearance, might seem to have an advantage here. However, if any desire is a case of taking-as-good, then (again, assuming that "good" here is more than just an embellishment) ascribing a certain desire to an agent will still be a matter of interpreting where the agent stands with respect to value dimensions such as impersonal and prudential value. It's not as if "good" as in the hyphenated "taking-as-good" can seriously be

a different notion from "good" as in the unhyphenated "taking as good". If desiring means taking-as-good, then all desires must move within the range of the conceptually available permutations of the good (e.g. impersonal and prudential good), rather than present us with a *sui generis*, desire-specific dimension of good. In other words, to desire x is necessarily to have one out of an independently pre-set list of attitudes towards x: taking-x-as-impersonally-good, taking-x-as-prudentially-good, and so on. Therefore, also on the attitude view desire ends up placing the agent into some value perspective, hence inflating the notion of desire.

6. Replies: The varieties of goodness

I now come to the second kind of reply that the guise of the good advocate might offer in response to Williams's arguments. This is, in particular, a reply to the arguments from personal ends and from immediate ends. As far as I can see, the one-desire argument is not at all touched by this reply, so the guise of the good still owes a response to it even if the second kind of reply were successful against the other arguments.

The second kind of reply is that Williams has failed to consider other possible dimensions of value than just impersonal or prudential value, i.e. other "goodnesses" under the guise of which we can desire something. Of course there are other dimensions of value that Williams doesn't mention in his brief discussion, like instrumental, symbolic, or non-instrumental (or final) value. But these are not relevant, because they are, as it were, merely structural distinctions—we can safely assume that Williams focuses throughout on things we desire as final ends, and so, as per the guise of the good, as final values.

In this connection, it might be thought that Williams's arguments can be easily met by saying that Williams has wrongly focused on overall or all-in goodness (be that impersonal or prudential), whereas the guise of the good only requires that one's objects of pursuit be seen or thought of as good *pro tanto* or good in some respect. In the case of immediate ends, one may reply to Williams that the agent at least sees her action under the guise of *pro tanto* prudential goodness, even if not under the guise of overall prudential goodness. However, the notion of *pro tanto* prudential value would still include a view to the agent's long term or whole life—to see something, e.g. a fleeting sensual pleasure, as *pro tanto* good for me is to see that pleasure as making *my life* better, as far as it (the pleasure) goes, while leaving it open whether, overall, my life would be better without that pleasure. By contrast, an immediate want or end, as Williams understands it, imports no such view of one's own long term. The *pro tanto*/overall distinction is another merely structural distinction in value which cannot be used to address Williams's concerns.[11]

The reply from varieties of goodness I explore here is rather that Williams has failed to consider other possible dimensions in the *subject-matter* of value beyond the impersonal and the prudential. If so, then the objects of personal ends and immediate ends could be seen by an agent under the guise of these other dimensions of value. In other words, his argument from elimination hasn't eliminated all plausible candidates.

Before moving on, there is an interpretive obstacle to face: does Williams take the argument from personal ends to be superseded, once he recognizes that the guise of the good can allow prudential value to be one of the many guises of the good? It might seem so, given the dialectical flow of the text:

[11] Thanks to a referee for inviting me to consider this potential answer to Williams.

Williams's Objection (from personal ends): the "idea of something's being good imports an idea, however minimal or hazy, of a perspective in which it can be acknowledged by more than one agent as good" (ibid.: 58). Some of our ends imply no such idea.

Reply (by the guise of the good): No such idea is necessary to desire; what is necessary is that at least the agent think what he desires "would be good for him".

Williams's New Objection (from immediate ends): "Even this modest claim implies a perspective that goes somewhere beyond the agent's immediate wants, to his longer-term interests or well-being" (ibid.: 59).

If this were the correct reading, we should really only focus on the argument from immediate ends here. However, another possible reading is that personal ends—the ends that the agent does not see under the guise of impersonal value or to be acknowledged as values by other agents—need not consist *only* of those of my ends which, unless they are immediate wants, could be accounted for by the guise of prudential value (long-term interest). As Williams elsewhere writes, the answer to the Socratic question "how should I live?" need not be either morality or prudence, but for example this: "the best way for me to live is to do at any given time what I most want to do at that time" (ibid.: 20). Williams also refers to "personal commitments that are not necessarily egoistic but are narrower than those imposed by a universal concern or respect for rights" (ibid.: 70), such as commitments to friends, family, or to a sports team. Lastly, he is ready to allow the existence of both counterethical motivations (and hence desires), such as malevolence and "counterjustice", and "counterprudential" ones, e.g. what he calls the "masochism of shame" (ibid.: 13-4). All these motivations have little in

common with each other. The point, though, is that each of these is a plausible instance of a personal end the agent may, but *need not*, think of as good under the concepts of either impersonal or prudential value. Nor are they necessarily classified as immediate wants (though, at times, they might occur as immediate wants). Therefore, it appears plausible to claim that, for Williams, the argument from personal ends is not sufficiently answered by resorting to the guise of prudential value, and is thus not superseded by the argument from immediate ends.

As anticipated, the guise of the good advocate will reply that Williams hasn't eliminated all the available conceptual options, and propose value concepts that could meet the challenge of personal and immediate ends. I consider two such proposals, and argue that both fail.

The first proposal is in terms of *radically relative value*. Personal ends might be seen by the agent under the guise of the *proprietary good*. Roughly, what is proprietarily good is what is good for someone or for someone's life, when first-personally regarded by the subject as "my life", rather than as "the life of anyone exactly like me".[12] It is obvious why this notion goes beyond that of impersonal value, and thus does not import any perspective of interpersonal acknowledgment. But proprietary goodness also goes beyond prudential value, because, unlike proprietary evaluations, prudential evaluations are universalizable: if a certain action is good for me (i.e. promotes my long-term interest), the same kind of action will be good for anyone who is exactly like me in the relevant respects.[13] We could say that, under prudence, I

[12] See Sayre-McCord (2001: 347ff). The notion of proprietary goodness also bears affinity with the concept of "goodness for" famously rejected by G. E. Moore in *Principia Ethica* (1993: 150-1).

[13] For the same reason, proprietary goodness is not what is usually called *agent-relative value*, because the latter is universalizable value: my friendships have agent-relative value for me only if your friendships have similar value for you, etc. Since agent-relative values can, in this sense, be acknowledged as values by other agents, and the agent acting under the guise of

am *normatively* special, because reasons of prudence for me ultimately have to do with my life, but I am not *evaluatively* special, because the fact that "it's my life" doesn't make an evaluative difference per se to what is prudentially good for me. Proprietary goodness, instead, makes me evaluatively special.

Likewise, immediate ends could be seen by the agent under the guise of the *good here-and-now*. Roughly, what is good-here-and-now is what is good (for someone) for, or at, a certain time and place, when indexically regarded by the agent as "now" and "here", rather than as a time or place of a certain kind. Evaluations in terms of here-and-now goodness, like those in terms of proprietary goodness, are not amenable to universalization, and make the "now and here" evaluatively special, hence not importing any perspective of the agent's long-term interest.

Can we expand the guises of the good in these radically relative ways to meet Williams's challenge? Williams is likely to answer "no". One of his assumptions is that we should only consider *ordinary* understandings of good in order to test the theory: "In any ordinary understanding of *good*, surely, an extra step is taken if you go from saying that you want something or have decided to pursue it to saying that it is good" (ibid.: 58). Proprietary goodness and here-and-now goodness appear to be theoretical constructs created ad hoc to account for recalcitrant types of desire, rather than notions we need to accept for independent reasons. Guise of the good advocates can certainly question Williams's assumption, but it seems fair to say the burden of proof is on them rather than on Williams.

---

agent-relative value knows or at least is committed to this, then the guise of the good cannot account for personal ends in terms of the agent's seeing her personal ends under the guise of agent-relative value.

The second proposal is to understand the evaluation present in personal and immediate ends as referring to the very *fact* that the agent finds x valuable (here and now). The guise of the good would resort to a concept of what may be called *subjective value*: x having value insofar as the agent finds x valuable. When I pursue a personal end, I see what I desire under the guise of subjective value: valuable *qua* valued by me. No perspective of interpersonal acknowledgment is imported here. When I pursue an immediate end, I see what I desire under the guise of time-indexed subjective value: valuable *qua* valued by me here and now. No perspective of my own long-term interest is imported.[14]

The main problem with this proposal is not that it resorts to a purely theoretical construct in order to save the theory. Some ordinary people might find the notion of subjective value perfectly intelligible. The question is rather whether it is coherent with the guise of the good to appeal to this notion of subjective value. Asked the question "what's the good in it?", where "it" stands for an instance of what I've called a personal or an immediate end (already realized or to be realized), the agent's answer would be: "the good in it is the fact that I find it good (here and now)". But this is no answer—or at best, a very unstable answer—because it invites the further question: "what is it that you find good in it?", which in effect asks for the very same thing as the initial question.[15] The inability of subjective value to furnish an acceptable answer to "what's the good in it?" strongly suggests that subjective value cannot

[14] It is tempting to say that to desire something under the guise of subjective value is to desire something *under the guise of the guise of the good*. However, this phrase would be misleading, because "the guise of the good" as appears in this phrase would not refer to a variety of *value*—it would only refer to the psychological fact that something appears good to the agent.

[15] Of course, the fact that people find certain things good might be a value in itself, i.e. it might be a bearer of intrinsic value. But if *this* were the answer to "what's the good in it?"—in terms of the agent pointing out the impersonal value of people finding something good—Williams would correctly object that the guise of the good offers an implausibly detached picture of how an agent pursues personal and immediate ends.

do any service for the guise of the good. One of the alleged advantages of the guise of the good is that it can explain, better than other theories, what it means to act for a reason: you act for a reason when you act in the light of something that appears good to you.[16] So, when asked for the reason for which you acted or will act, the answer "because something in it appears good to me" (or "because I find it good") can only be either the beginning of an answer, or the admission that you are not quite sure just what is the reason for which you acted or will act. What matters, for the purposes of understanding the reasons for which you act, is *what* appears good to you, not *the fact that* something appears good to you. *Bonum est multiplex* (good is multiform), as Elizabeth Anscombe stressed (1963: 75), but subjective value cannot be one of the forms of the good to which the theory can turn for addressing Williams's objections.[17]

To sum up the last three sections: Williams's arguments for rejecting the guise of the good – the argument from personal ends, from immediate ends, and the one-desire argument— cannot be successfully met either by coming up with better candidates for "thinking as good", or by positing alternative concepts of the good (varieties of goodness) that can be tailored to fit the nature of personal and immediate ends.

[16] See Raz (2010), Tenenbaum (2007). Also critics (such as Setiya 2010) agree that this is what is initially attractive about the guise of the good.

[17] The guise of the good is compatible with the substantive thesis that the subject's desire, endorsement, value judgment, or value appearance may be a constituent of some of what's good (impersonally or prudentially) (see Tenenbaum 2007: ch. 3). But subjective value, as I understand it here, is a kind of value predicate (a variety of *goodness*), not a substantive value. Also, the guise of the good could allow "I care for it, I love it" as acceptable answers to "what's the good in it?", because, whether or not care or love contain a value judgment, they include *more* than that. But Williams's personal and immediate ends cannot be accommodated in this way by the guise of the good, because Williams overall makes it clear that he is talking about ends we pursue but don't *necessarily* matter to the agent in the way typically expressed by "I care for it, I love it".

7. Normative reasons to the rescue?

As a final reply, it could be said that I have been mistakenly fixating on the notion of *value* as an essential feature of the guise of the good. Recent literature has brought out an attractive alternative in the vicinity: the guise of normative reasons, understood as the view that in order to desire or intentionally do something, the agent needs to see it under the guise of normative reasons—reasons which are taken by the agent to count in favor of the desire or in favor of the action.[18]

The guise of normative reasons might seem to answer Williams's objections better than any of the variations of the guise of the good presented so far. Against the argument from personal ends, one could claim that seeing one's end or action under the guise of a normative reason, quite in general, does not import "an idea, however minimal or hazy, of a perspective in which it [i.e. the end] can be acknowledged by more than one agent as good". Against the argument from immediate ends, one could likewise claim that acting under the guise of a normative reason, even when this reason tracks something good for the agent, does not imply "a perspective that goes somewhere beyond the agent's immediate wants, to his longer-term interests or well-being".

As for the one-desire argument, the guise of normative reasons does not imply that the agent who puts all value into one desire will have only that one desire. She may put all value into

[18] See Bedke 2008, Gregory 2013, Singh 2019. My formulation doesn't exactly follow any of these authors, though. Bedke claims: "end-setting attitudes are attitudes through which we see reasons, or more precisely attitudes through which we see various considerations as favoring our actions and attitudes" (2008: 90). Gregory identifies desire (to act in a certain way) with a belief about normative reasons (to act in that way). Singh claims: "an agent Φ-s on the basis of some (motivating) reason R just in case: (a) she represents R as a normative reason to Φ, and (b) her representation explains, in the right way, her Φ-ing" (2019: 419).

one desire while possessing also other desires—other things she regards as favored by some normative reason—to which she doesn't accord the kind of importance that comes with "putting value" into them. While, on this view, all desires imply taking something as a normative reason, taking something as a normative reason might look like a much lighter affair than taking something as good. In general, we can take ourselves to have normative reasons which need not be, even by our lights, underwritten by *any* conception of value. Williams himself seems to stress this much in a footnote to this passage: "The truth is that not every reason for action is grounded in an evaluation" (ibid.: 232). In sum, Williams's contention that it takes an "extra step" from desiring or pursuing something to thinking it as good seems to have somewhat less bite if we replace "thinking as good" with "thinking one has a normative reason for pursuing it".

Interestingly, the guise of normative reasons appears also to address the passing remark made in the course of the one-desire argument: "Merely to have one desire might well be to have no value in my life at all" (ibid.: 59). The guise of the good seemed forced to assign at least some apparent value to such a life just as long as the agent has that one desire, because any desire implies at least apparent value. The guise of normative reasons instead says: it is indeed possible, like Williams says, that there is neither apparent nor real *value* in such a life at all, even though the person has at least some apparent normative reason for action as long as she has that one desire, because any desire implies taking something as a normative reason.

8. Why the guise of normative reasons doesn't work

As just seen, the guise of normative reasons appears at first glance to avoid Williams's objections against the guise of the good. However, I will suggest here that this may be nothing more than an effect generated by the lexical shift from value to the more plastic notion of a normative reason (and cognates such as the relation of favoring an action). Once one starts pressing against the guise of normative reasons some of the very same concerns I above pressed against the various solutions offered by the guise of the good, it is no longer clear that the guise of normative reasons is on any safer ground against Williams's challenges. For reasons of space, I'll only focus on the arguments from personal and immediate ends.[19]

First, in section 5 I claimed that it does not ultimately matter whether the guise of the good is carried by a judgment, an appearance, or a taking-as-good attitude, because Williams's concern is with the kind of value perspectives brought along with concepts such as "good" and "good for someone". Now, while I have just pointed out that thinking in terms of normative reasons need not import either a perspective in which one's action can be acknowledged by more than one agent as good, or a perspective that includes the agent's long-term well-being, it is by no means obvious that thinking in terms of normative reasons imports no analogous *normative* perspective. In particular, it seems plausible that, when I regard my actions as favored by a normative reason, I do expect others to acknowledge that much, even when the reason is agent-relative (e.g. grounded in my personal relationships) and thus not shared by others who are not similarly placed. Even reasons in competitive settings (e.g. my reason to win the tennis match) are expected to be acknowledged by our competitors. The perspective of normative reasons need not be as impersonal as (one kind of) value perspective, but it may still make a claim to intersubjective acknowledgment in a way

[19] Thanks to a referee for prompting me to rethink this section of the paper.

that does not sit well with Williams' remark about personal ends and, *a fortiori*, about immediate ends (at least to the extent that these are a subset of personal ones).

Second, and relatedly, in section 6 I criticized the proposal to appeal to notions such as proprietary value or subjective value to make sense of the guise of the good under which one pursues personal and immediate ends. In particular, I argued that a notion like proprietary value (and the related notion of good here-and-now) seems to be created ad hoc, especially due to the fact that evaluations in such terms are, by hypothesis, not amenable to universalization. Now, thinking in terms of normative reasons appears to commit to universalizability: if A has a reason to do x under circumstances C at time t1, then anyone under exactly the same circumstances has a reason to do x at time t*n*. Like for the case of intersubjective acknowledgment, the commitment to universalizability marks normative reasons in a way that does not sit well with Williams's insistence on the particular character of both personal and immediate ends. Again, it is thus by no means obvious that turning to the guise of normative reasons is the right move to make in response to Williams: pursuing something would mean taking oneself to have a normative reason to do so, and this would burden the agent with commitments which are not, after all, too different from the commitments incurred when thinking in terms of goodness.

One might reply that, in the case of personal and immediate ends, the agent might still think in terms of normative reasons to pursue those ends, while rejecting or ignoring commitments to intersubjective acknowledgment or universalizability. However, this is not a happy move. The guise of normative reasons would be able to accommodate Williams's personal and immediate ends at one or the other of these significant costs: either personal and immediate ends require the agent's thinking in terms of an ad hoc notion of normative reasons—with

which we are back to the same problems faced by ad hoc notions of goodness—or else agents are systematically mistaken whenever they pursue personal and immediate ends, because in these cases their desires would imply mistaken thinking about normative reasons. If Williams is right that personal and immediate ends characterize large part of our normal life as agents, the guise of normative reasons would entail a worrying error theory of sorts about agency.

Before finishing, I should point out that these considerations are not meant to be conclusive. Advocates of the guise of normative reasons may defend the view that taking oneself to have a normative reason does not come with any commitment which would be at odds with the nature of personal or immediate ends. In that case, turning to normative reasons might be a more promising move to make in response to Williams's challenges to the guise of the good. But I do think the burden is on such advocates to prove that much. After all, the guise of normative reasons is still a fairly young proposal in this area of debate and, though certain advantages over the guise of the good have been highlighted (e.g. Gregory 2013), its broader implications have not yet been sufficiently explored.[20]

9. Conclusion

[20] What if one understands normative reasons for A to do x simply as considerations that would move A to do x if she deliberated soundly from her existing motivational set, as proposed in Williams's own theory of internal reasons (Williams 1981)? While this notion of normative reasons is relatively free of commitments, surely no one would suggest that, in simply wanting or pursuing an end, I need to think that there are considerations such that I would be moved by them if I deliberated soundly from my existing motivational set. That would be an implausibly intellectualistic (and perhaps implausibly self-referring, too) view of wanting. In this connection, I should add that Williams's critical stance on the guise of the good probably dovetails with his positive remarks about practical deliberation being essentially first-personal, involving "an *I* that must be more intimately the *I* of my desires than" certain accounts of agency appear to allow (Williams 2006: 67). But whether the guise of normative reasons can accommodate this point is a difficult question better left for another work.

How to move forward in the dispute Williams vs. the guise of the good/normative reasons ("guise-of" for short)? I believe that there is an important dialectical difference among the three arguments I have reconstructed and discussed. The arguments from personal ends and from immediate ends do place some burden of proof on advocates of guise-of theories. But it would be wrong to stake the success of these theories on those arguments alone, because, after all, those appear to be arguments from counterexample. The intuitions Williams marshals here—people can sometimes have, or at least be coherently described as having, personal and immediate ends with no thought to values or even to normative reasons—may be worth jettisoning if one of the guise-of approaches proves to be an otherwise powerful account of desire and agency. Those arguments will need to be weighed within an overall comparison of guise-of approaches against rival accounts.

The one-desire argument, however, seems to be a more bothersome thorn in the side for the guise-of. It trades on an intuition that cannot be easily dismissed: it is possible that an agent puts *all* that is evaluatively or normatively relevant to her into one desire only, while at the same time possessing other (and, as I added, possibly conflicting) desires she doesn't put *anything* evaluatively or normatively relevant into. Guise-of theories have significant trouble allowing for this possibility, despite the various solutions canvassed in this paper. Of course, again, it might be that the benefits of a certain guise-of theory are so significant that we should lay aside also the intuition behind the one-desire argument. But another reaction might be to try out a looser kind of guise-of account, whereby the guise of the good/normative reasons is a feature that *must* characterize any agent's motivational set without necessarily

"trickling down" to *each and every* (final) desire, end, or action.[21] On this account, the intuition behind the one-desire argument can be accommodated, as this agent can have other desires while not needing to associate them with any evaluative or normative guise. Because and as long as she does put value or normative reasons into one desire, she is still perfectly intelligible as an agent. On the other hand, it remains to be seen whether any such looser guise-of can still deliver the same goods promised by stricter versions.[22]

Bibliography

Altham, J. E. J. and Harrison, Ross (eds.) (1995), World, Mind, and Ethics: Essays on the Ethical Philosophy of Bernard Williams. New York: Cambridge University Press.

Anscombe, G. E. M. (1963 [1957]), Intention. 2nd ed. Cambridge, MA: Harvard University Press.

Bedke, M. S. (2008), "Practical Reasons, Practical Rationality, Practical Wisdom", Ethical Theory and Moral Practice, 11 (1): 85-111.

[21] Such a looser account is suggested, but not developed, in different ways by different authors: "Although the claim that we must act with an eye to some good turns out not to be true of any particular action, this condition will still hold for any concerns that could serve as organizing aims of our lives" (Sussman 2009: 618-9); "Human beings want what seems good to them; and they act for reasons they regard as good…these propositions could be necessary truths of human nature even though some of us want to own saucers of mud and others act from spite or vanity or despair" (Setiya 2010: 103).

[22] Thanks to two anonymous reviewers of this journal for their very helpful feedback on a previous version. This research has been supported by the European Union through the European Regional Development Fund (the Centre of Excellence in Estonian Studies, TK 145), and the University of Tartu, grant PHVFI21914.

Brentano, F. (1995). Psychology from an Empirical Standpoint. Edited by Linda L. McAlister. Translated by A. C. Rancurello, D. B. Terrell, and L. L. McAlister. London and New York: Routledge.

Chappell, S. G. and van Ackeren, M. (2019). Ethics beyond the limits: New essays on Bernard Williams' Ethics and the Limits of Philosophy. London and New York: Routledge.

Gregory, A. (2013), "The Guise of Reasons", American Philosophical Quarterly, 50 (1): 63–72.

Heuer, U. and Lang, G. R. (eds.) (2012), Luck, Value, and Commitment: Themes from the Ethics of Bernard Williams. Oxford Oxford University Press.

Moore, G. E. (1993), Principia Ethica. Cambridge: Cambridge University Press.

Oddie, G. (2005), Value, Reality, and Desire. Oxford: Clarendon Press.

Orsi, F. (2015), "The Guise of the Good", Philosophy Compass, 10 (10): 714-724.

Orsi, F. (2023), The Guise of the Good: A Philosophical History. London and New York: Routledge.

Raz, J. (2010), "On the Guise of the Good", in S. Tenenbaum (ed.), Desire, Practical Reason, and the Good. Oxford: Oxford University Press. 112–136.

Sayre-McCord, G. (2001), "Mill's 'Proof' of the Principle of Utility: A More than Half-Hearted Defense", Social Philosophy and Policy, 18 (2): 330–360.

Schafer, K. (2013), "Perception and the Rational Force of Desire", Journal of Philosophy, 110 (5): 258–281.

Setiya, K. (2007), Reasons Without Rationalism. Princeton: Princeton University Press.

Setiya, K. (2010), "Sympathy for the Devil", in S. Tenenbaum (ed.), Desire, Practical Reason, and the Good. Oxford: Oxford University Press. 82–110.

Singh, K. (2019), "Acting and Believing Under the Guise of Normative Reasons", Philosophy and Phenomenological Research, 99 (2): 409–430.

Stampe, D. W. (1987), "The Authority of Desire", Philosophical Review, 96 (July): 335–381.

Stocker, M. (1979), "Desiring the Bad: An Essay in Moral Psychology", Journal of Philosophy, 76 (12): 738–753.

Sussman, D. (2009), "For Badness' Sake", Journal of Philosophy, 106 (11): 613–628.

Tenenbaum, S. (2007), Appearances of the Good: An Essay on the Nature of Practical Reason. Cambridge: Cambridge University Press.

Tenenbaum, S. (2008), "Appearing Good: A Reply to Schroeder", Social Theory and Practice, 34 (1): 131–138.

Tenenbaum, S. (2013), "Guise of the Good", in H. LaFollette (ed.), The International Encyclopedia of Ethics. Oxford: Wiley-Blackwell.

Tenenbaum, S. (2018), "The Guise of the Guise of the Bad", Ethical Theory and Moral Practice, 21 (1): 5–20.

Tenenbaum, S. (2020), "The Guise of the Good", in R. Chang and K. Sylvan (eds.), The Routledge Handbook of Practical Reason. London: Routledge. 226–236.

Thomas, A. (ed.) (2007), Bernard Williams. Cambridge: Cambridge University Press.

Velleman, J. D. (1992), "The Guise of the Good", Noûs, 26 (1): 3–26.

Williams, B. (1981), "Internal and External Reasons", in his Moral Luck. Cambridge: Cambridge University Press. 101-113.

Williams, B. (2006), Ethics and the Limits of Philosophy. London: Routledge.