



According to one influential version of the derivation of Kant's Formula of Humanity, agents necessarily represent their ends as objectively good. In this paper I argue that there is good reason to think that this is false, that it is not the case that agents necessarily represent their ends as objectively good. The paper is divided into four sections. In the first, I explain what is at stake in the argument for FH. In the second, I explain the relevant terms and introduce some alternative positions that are available. In the third, I argue that there is good reason to think that it is not the case that agents necessarily represent their ends as objectively good and in the fourth, I consider a potential reply to this argument.

1 Introduction

According to one influential version of the derivation of Kant's Formula of Humanity (FH), the following claim is true:

GOG: Agents necessarily represent their ends as objectively good.¹

In this paper I argue that there is good reason to regard GOG as false. I argue there is good reason to think that it is not the case that agents necessarily represent their ends as objectively good. This, of course, does not go any way toward showing that Kant did not believe that agents necessarily represent their ends as objectively good. The influential version of Kant's derivation of FH referred to in the first sentence of this paragraph might be, for all that I say here, a good reconstruction of Kant's argument. I shall not be engaged in exegesis and so shall not be entering textual debates in this paper.

The paper is divided into four sections. In the first, I explain what is at stake in arguing that GOG is false. In the second, I explicate the terminology in this claim. I also contrast the claim with other possible claims one might make about how agents represent their ends. In the third, I argue that there is good reason to regard the claim as false and in the fourth, I consider a reply to the argument I make in the third section.

2 The stakes

The argument for FH that hinges on GOG is called the regress argument. In broad strokes, the argument runs as follows: (1) agents necessarily represent their ends as objectively good (GOG), (2) if agents necessarily represent their ends as objectively good, then agents necessarily represent their rational nature as unconditionally good, therefore (3) agents necessarily represent their rational nature as unconditionally good. But (4) agents necessarily represent their rational nature as unconditionally good *qua* rational nature rather than, say, *qua* their own and (5) if agents necessarily represent their rational nature as unconditionally good *qua* rational nature, then consistency demands that they represent rational nature in all its instances as unconditionally good or cease to represent their own rational nature as unconditionally good. But (6) agents cannot cease to represent their own rational nature as unconditionally good (this is based on (3)). Therefore (7) consistency demands that agents represent rational nature in all its instances as unconditionally good.

I am going to argue that there is good reason to regard GOG (the first premise of the regress argument) as false. But it might be the case that agents sometimes represent their ends as objectively good even if agents do not necessarily represent their ends as objectively good. And it might be the case that the regress argument can be used to show that if an agent represents its ends as objectively good in any given instance, then, in that instance, the agent ought to treat the rational capacities of others as unconditionally good. But this would weaken the conclusion of the regress argument considerably. Even in a best case scenario (the other premises of the regress argument are true and the inferences are valid), the regress argument no longer would show that it is always the case that agents ought to treat the capacity to set ends as unconditionally good.² It would show (at best) that *sometimes* agents ought to treat the capacity to set ends as unconditionally good. Since I suspect

that agents represent their ends as objectively good only if they *are* treating the capacity to set ends as unconditionally good in all of its instances, this is troubling, indeed; the regress argument would go precisely nowhere in showing what agents ought to do, for it would show only that agents ought to do that which they already do. That is, it would show only that agents ought to treat the capacity to set ends in all its instances as unconditionally good if (i.e., in those instances when) they already are treating the capacity to set ends in all its instances as unconditionally good.

But the point is, as noted above, that the goal of this paper is to show that there is good reason to think that GOG is false. If GOG is false, then the regress argument does not work, for it is based on a false premise. At a minimum, those are the stakes: the regress argument for FH. It should be noted that GOG also lies at the heart of a kind of neoPlatonic brand of moral psychology. If GOG falls, this school falls with it. However, in trying to keep a tight focus, I confine myself to the stakes just mentioned: the regress argument.

3 Terms and alternative positions

In this section, I discuss what it means to say agents necessarily represent their ends as objectively good. In order to understand GOG, the following four things need to be explained: first, the work being done by the term ‘necessarily’; second, the work being done by the term ‘represent’; third, the work being done by the term ‘end’; and fourth, the meaning of the term ‘objectively good’.

In saying that agents *necessarily* represent their ends as objectively good, GOG is making a strong claim. According to GOG, there is no possible world in which an agent has an end that it does not represent as objectively good. Regardless of whether the agent is behaving irrationally or rationally, passionately or incontinently, wantonly or consistently, deliberately or whimsically, the

agent is representing its ends as objectively good.

In saying that agents necessarily *represent* their ends as objectively good, GOG is not making a claim about phenomenology. Similarly, the idea behind GOG here is not that agents engaging in deliberation choose an end and then subsequently “decide” to represent their ends in some way (*viz.*, as objectively good, whatever that might be). That gets things exactly backward: it puts the cart before the horse, so to speak. And knowing that helps to explain what GOG is getting at with the word ‘represent’. The idea is that an agent will pursue an end only if that end is objectively good given the backdrop of that agent’s values, beliefs and desires. But the end does not have to *be* objectively good: for example, the agent can have mistaken beliefs about the world. Thus, an agent might believe that if she goes down to room X at time t, she will accomplish Y, and she therefore might regard going to room X at time t as objectively good. So she might set going to room X at time t as an end on those grounds. But we might know better: we might know that actually she needs to go down to room Z at time t (perhaps that is where M is, and it is only with M that she can accomplish Y. She just happens to believe, unfortunately, that M is in room X). The point is that given the agent’s values, beliefs and desires, going down to room X and *not* going down to room Z at time t will be objectively good, so according to proponents of GOG she will go down to room X at time t because that is what she represents as objectively good even though it really is not so.

In saying that agents necessarily represent their *ends* as objectively good, GOG is fencing off the realm of intentional actions. It is not the case that agents represent all of the actions that they perform as objectively good. Some actions are not intentional (Harding’s hand-movements; Bibbit’s stutter). Moreover, there might be a question as to whether any given action is something that the agent performs or something that merely happens to the agent (did Chief Bromden really raise his hand in the vote or was it the wires?). But that question does not get at the heart of GOG.

That question does not get at the heart of GOG because GOG deals with ends rather than actions. In saying that agents represent their ends as objectively good, GOG is saying that agents represent their purposes, intentions or goals as objectively good. An example will help to illustrate this point.

Consider an agent who tells a lying promise in order to get some ready money. In this case, the agent's end is to get some ready money. According to proponents of GOG, this agent represents getting some ready money as objectively good. Of course, the agent might want to get some ready money because of some further end that he has (e.g., to pay a debt). But GOG is not about "ultimate ends" or anything like that; GOG is about ends, plain and simple. Thus, in this example, both getting some ready money *and* paying a debt (or whatever the further end is) are represented as objectively good (and so on up and down the means-end ladder).

But we still do not have a rigorous understanding of what is meant by saying that agents necessarily represent their ends as *objectively good*. X is objectively good just in case "it is unconditionally good or . . . is conditionally good and the conditions under which it is good are met."³ For example, Kant thinks that the "highest good" is a world in which all agents are happy in proportion to their virtue, a world in which a conditioned good (happiness) is doled out in accordance with its condition (virtue) (5:110-114). For a Kantian, a world in which the highest good is realized would be objectively good.⁴

Objective goodness is a property that tracks reasons. If an agent's end is objectively good, then she has a sufficient (justifying) reason to pursue that end; if an agent represents her end as objectively good, then she takes herself to have a sufficient reason to pursue that end. For a Kantian, agents have a sufficient reason to promote a world in which the highest good is realized.⁵ Another example to which Kantians often appeal in explaining the meaning of objective goodness is happiness: "a conditionally good thing, like happiness, is objectively good when its condition is

met in the sense that it is fully justified and the reasons for it are sufficient.”⁶

In order to understand GOG more fully, it is useful to contrast it with other claims that philosophers have made about how agents represent their ends. Here are two:

GG: **Agents necessarily represent their ends as good.**⁷

NGG: **It is not the case that agents necessarily represent their ends as good.**⁸

GG (guise of the good) is a weaker version of GOG. According to GG, agents necessarily represent their ends as good. This would be true if GOG is true. However, it could be the case that agents necessarily represent their ends as good even if they do not necessarily represent their ends as objectively good (e.g., an agent might represent an end as conditionally good even if those conditions have not been met). NGG (negation of the guise of the good) is true only if GG is false (thus, also, only if GOG is false).

Examples will help to disambiguate GOG, GG and NGG. Consider Milton’s Satan, who declares, “Evil be thou my good.” Or consider Augustine, who confesses that, in stealing pears, he was being “gratuitously wanton, having no inducement to evil but the evil itself.” Or consider Ovid’s Medea, who claims to be aware that what she is doing is wrong. On NGG, these agents are perfectly coherent. These agents might be behaving irrationally. But there is no logical impossibility in that. On GG and GOG, however, more must be said. For example, the proponent of GG or of GOG might say that Satan is embracing evil because Satan thinks that evil is the only means to liberty and to glory and Satan thinks that liberty and glory are (perhaps objectively) good. Similar stories would have to be told about Augustine and Medea. Whereas the proponent of NGG could accept Satan *et al* without further question, the proponent of GG or GOG would have to give a story about how these agents *really* do represent their ends as (objectively) good in some way.

In order to see where GG and GOG come apart, consider Leontius. Leontius seems to think

that looking at corpses would be good insofar as it would afford him pleasure. But he seems to think that this pleasure would be shameful. Leontius thus represents looking at the corpses as good but not objectively good. Looking at the corpses would afford him access to a conditioned good (pleasure) but not in such a way that its conditions would be met (the pleasure would be shameful). Nonetheless, Leontius looks at the corpses.⁹ Or consider Huck Finn. Huck's conscience tells him to return Jim to the slave-catchers. But his sympathies lie with Jim. Huck's conscience tells him to return Jim; his sympathies tell him to help Jim to escape. Huck decides to help Jim to escape. On GG, there is no incoherence in either of these stories. Leontius and Huck might be behaving irrationally. But there is no logical impossibility in that.

However, the proponent of GOG must say more. The proponent of GOG might say that Leontius *really* represents his ends as objectively good in looking at the corpses, for Leontius's end is not the pleasure (which is good but not objectively good); Leontius's end is satiating (and thus getting rid of) a troublesome desire that he is unable to tame at the moment. And something similar would have to be said about Huck.

Thus, there are three positions that one can take, namely: NGG, GG and GOG. My target is GOG. I shall not say anything about the debate between proponents of NGG and proponents of GG.

4 Against GOG

Now I shall argue that GOG precludes a distinction that I think many find intuitively plausible. I shall argue that GOG precludes the distinction between culpable and inculpable ignorance. It seems to me that the distinction between culpable and inculpable ignorance is basic enough and

important enough that its erasure or preclusion would constitute a *reductio* of a position. So if I am correct (if GOG precludes the distinction between culpable and inculpable ignorance), then the argument here constitutes a *reductio* of GOG.

Culpable ignorance is the lack of knowledge or understanding from the omission of ordinary care to acquire such knowledge or understanding. Inculpable ignorance is the lack of knowledge or understanding despite ordinary (and sometimes even exemplary) care to acquire such knowledge or understanding. The distinction can be illustrated by appeal to example. Consider Cecilia Jupe. Cecilia Jupe is not very bright. When it comes to tests of medical knowledge, for example, Cecilia is about as bad as they come. But that is not a problem; she is not a medical doctor, and she does not want to be one. However, Cecilia has a good heart. And when a doctor tells her to administer a certain treatment to a patient, she does so diligently and meticulously. Suppose that the treatment in this case is a course of penicillin and that the patient is allergic to penicillin. Cecilia does not realize that the patient is allergic to penicillin; she does not know even what penicillin is. And it is not her business to know: she is not in the medical line of work. She simply is a caring friend of the patient, someone who wants to help her friend in a time of need. But it is the doctor's business to know this. Now suppose that the doctor does not know that the patient is allergic to penicillin because the doctor failed to question the patient about medical history. Predictably, the patient has a bad reaction to the penicillin.

In this case (in the case of administering penicillin to a patient who is allergic to it), has anybody done anything wrong? Cecilia Jupe has not done anything wrong; her behavior was exemplary. The doctor, it can be supposed, did not intend to harm the patient; if the doctor had realized that the patient is allergic to penicillin, he would not have prescribed it. But as a medical professional, it is the doctor's business to know this. Moreover, it is not the case that the doctor asked the patient

and the patient lied or unintentionally uttered a falsehood. The doctor did not consult a previously taken medical history that was flawed. Rather, the doctor simply did not take a medical history.

It is true that there could be excusing conditions. Doctors are human; they get frazzled and distracted just as much as the rest of us. But (*ex hypothesi*) this is not a case in which excusing conditions apply. The doctor should have known that the patient is allergic to penicillin. The doctor's ignorance arises from negligence. Whereas Cecilia's ignorance is inculpable, the doctor's ignorance is culpable. There is a line in the sand and Cecilia is on one side of it; the doctor is on the other.

The distinction between culpable and inculpable ignorance sometimes comes up in legal settings. Sometimes punishment is mitigated or entirely avoided because of defendant's ignorance; sometimes the punishment is meted out precisely because of defendant's ignorance. But talking about the legal sphere brings in factors that are too difficult for me to grapple with here (e.g., the concepts of strict liability and absolute liability). I am interested only in the distinction between culpable and inculpable ignorance in the ethical sphere. The claim is that GOG precludes this distinction.

Now on NGG, there is a ready explanation of the difference between culpable and inculpable ignorance. Consider again the distinction between Cecilia Jupe and the doctor. On NGG, one could give the following explanation of this distinction. Cecilia Jupe always pursues ends that she represents as objectively good. Sometimes she makes mistakes. Cecilia's grasp of the mechanics of the world is sadly lacking. She might administer poison thinking that it is palliative. But she has a good heart and she never would hurt anyone willingly. She always pursues that which she believes is objectively good. The doctor, by way of contrast, has a far better grasp of the mechanics of the world. But the doctor does not have a good heart like Cecilia. The doctor might not be evil

through and through. He might be a very good parent or a very good friend. He even might be a good doctor in some respects. Perhaps he has a good bedside manner or is very solicitous of the well-being of his elderly patients. But he does not want to know much about his patients. He does not want to know much about his patients even when this might have adverse effects on his treatment of them. His end is ignorance about his patients, and although he is not so blatant as to do things like put in ear plugs when his patients begin to speak, he purposefully never asks for any information that is not volunteered. There is no silver lining; the doctor knows that he ought not to pursue ignorance in this way. His choice to do so anyway is not governed by some greater good; he simply does not want to know about any other medical conditions his patients might have, even if these other conditions might inform the present malady. “Have a bacterial infection according to the lab? Penicillin it is. Call me in the morning” (*sotto voce*: “now get out of my office and let somebody with a *real* problem come in”).

The doctor in this example chose something (namely, not to question the patient about medical history) knowing that this is a bad thing to do; he could not be bothered to do otherwise. The doctor’s ignorance is culpable even though Cecilia’s is not, and this difference is explained by reference to the ways in which these two agents represented their ends. Cecilia chose only ends that she represented as objectively good; the doctor did not. The doctor chose an end despite the fact that he did not represent it as good, and this choice resulted in the doctor’s (culpable) ignorance.

There is also a ready explanation of the difference between culpable and inculpable ignorance on GG. On GG, one could give the following explanation (which, perforce, one also could give on NGG). Cecilia Jupe is explained in much the same way. Cecilia has a good heart. She pursues only ends that she represents as objectively good. Sometimes she is incorrect in her beliefs about what

is objectively good. But that does not change matters. She has a good heart. The doctor, by way of contrast, does not have a good heart. At some point(s) while or before examining the patient, the doctor is faced with options. After weighing them, he chooses the lesser good. This need not be malicious; the doctor does not set out actively to hurt the patients that he sees. But he sees so many patients; he gets tired of the long days with patients tromping in and out every fifteen minutes. Sometimes it is clear what the patient needs without conducting a full medical history. He knows he ought to be more careful, but he cannot be bothered at the moment and he decides (consciously and deliberately) not to bother questioning the patient about medical history. The doctor wants to get home to dinner, and that is the end that governs his brusque rushing through with his patients. Or perhaps the doctor's pay is correlated with the number of patients he sees and he is somewhat avaricious (or simply stressed out about all those loans he took out to pay for medical school in the first place). For whatever reason, his end is not what it ought to be (giving his patients the best treatment he can); his end is to give his patients the fastest treatment he can. He chooses the good of going home early or of getting more money over the good of treating his patients as he ought. There is no good explanation for why he made this choice. Evil is not intelligible. There is simply the fact that the doctor made this choice, and this fact explains why the doctor's ignorance is culpable whereas Cecilia's is not.

I should point out that I am not wedded to either of these stories; there are other stories one could tell on NGG or on GG. My point is merely that there is a deep difference in kind between culpable and inculpable ignorance and that this difference cannot be explained on GOG. On both NGG and GG, the distinction between culpable and inculpable ignorance can be explained. In particular, the distinction between culpable and inculpable ignorance can be explained on NGG or on GG by showing that the agent's ignorance is a (direct or indirect) result of a bad choice.

That is, in cases of culpable ignorance, the agent's ignorance is *a result* of the choice of an end that the agent does not represent as objectively good. In cases of inculpable ignorance, the agent's ignorance arises *despite* the fact that the agent chooses only those ends that s/he represents as objectively good.

Proponents of GOG, in contrast to proponents of GG or NGG, do not seem to have the conceptual resources to make the distinction between culpable and inculpable ignorance. Proponents of GOG cannot explain the distinction between culpable and inculpable ignorance by appeal to whether an agent's ignorance is a (direct or indirect) result of an end that is not represented as objectively good. Proponents of GOG cannot explain the distinction between culpable and inculpable ignorance in this way because, according to proponents of GOG, agents necessarily represent their ends as objectively good.

This should come as no surprise: the flip-side of GOG is to say that moral evil arises when (and only when) that which an agent represents as objectively good is not actually so, and this occurs when (and only when) an agent is ignorant about that which is objectively good. That is, the point is not that the only way to understand culpability is by appeal to the adoption of an end that was not represented as objectively good. This is certainly how proponents of NGG and GG make sense of culpability. But it does not follow from the fact that this is unavailable to proponents of GOG that proponents of GOG can make no sense of culpability in general: proponents of GOG make sense of culpability in general by saying that an agent is culpable for a wrong if it arises from *ignorance*, from *misrepresenting* an end as objectively good. And *that* is the point: to say that some ignorance is culpable but some not undermines GOG. To flip this back over to push the thesis I am arguing for here: GOG cuts above the distinction between culpable and inculpable ignorance.

Now the proponent of GOG might agree with much of what I have said here about GG, NGG,

Cecilia Jupe and the doctor. But the proponent of GOG might say that I am going too far when I say that this shows that GOG precludes the distinction between culpable and inculpable ignorance. The argument for this objection would be as follows. Whether a given agent's ignorance is culpable depends on whether the agent has acquired a certain body of knowledge, and (here is the crucial step) that body of knowledge is determined by the agent's role, so to speak. To make this more concrete, the idea is that Cecilia is not a medical professional; she is "just" a friend trying to help. So the standards that determine that which she is responsible for knowing are more relaxed than those for the doctor. In particular, Cecilia is not responsible for knowing whether her friend is allergic to penicillin, whence it may be inferred that Cecilia's ignorance is inculpable. The doctor, as a medical professional, must meet higher standards of knowledge. The doctor does not meet those standards, whence it may be concluded that the doctor's ignorance is culpable.

The trouble with this response is that it misses its target. This can be seen by the fact that (as already pointed out) there are excusing conditions under which the doctor's ignorance is not considered culpable. So it does not suffice merely to appeal to the fact that there are different standards for Cecilia and for the doctor. These different standards are granted; there is still something more to explain. In other words, if appealing to the fact that Cecilia is not a medical professional and the doctor is sufficed to explain why the doctor's ignorance is culpable whereas Cecilia's is not, then there are no excusing conditions, conditions in which the doctor's ignorance of his patient's allergy is inculpable. But there are excusing conditions, so appealing to the doctor's role as a medical professional does not suffice.

I suppose the proponent of GOG might argue that there are no excusing conditions for culpable ignorance. I think this objection is so counterintuitive that nobody seriously would make it. But even if somebody did, it still misses the mark. The appeal to excusing conditions was made merely

to make obvious what proponents of GOG are lacking: *viz.*, an *explanation* of the distinction between culpable and inculpable ignorance. The denial of excusing conditions notwithstanding, what the proponent of GOG is missing here in this attempted rebuttal is that the distinction between culpable and inculpable ignorance has become (at best) *ad hoc*. Denying excusing conditions makes maintaining the distinction between culpable and inculpable ignorance like accounting for the precession of Mars by adding a handful of epicycles when one has laid bare before one the possibility to switch to a heliocentric model with elliptical orbits governed by a universal attractive force. So let us overlook this objection and assume that there are not only role-based expectations but also excusing conditions.

So the proponent of GOG needs to explain why excusing conditions do not apply to the doctor in this case. But perhaps that can be done in the following manner: Cecilia's actions sometimes do not turn out well. Her actions have bad effects. But she always has good intentions. That is, Cecilia not only represents her ends as objectively good. Her ends *are* objectively good. The fact that the world does not cooperate is contingent (and, for present purposes, irrelevant). The doctor, by way of contrast, chooses ends that he represents as objectively good but he is mistaken in doing so. This, of course, involves a slight shift: in this attempted reply, the proponent of GOG is saying, *in esse*, that Cecilia is not ignorant of the objectively good whereas the doctor is.

The trouble with this attempted explanation on the part of the proponent of GOG is that it breaks down almost as soon as one tries to spell it out. The section of the means-end chain here has three nodes (as I am imagining it). Cecilia follows the doctor's orders because she has the end of treating her friend; she treats her friend with penicillin because she has the end of following the doctor's orders; she gets the medication because she has the end of treating her friend with penicillin. This chain could be extended in either direction (she treats her friend because she has

the end of helping others in need; she goes out to the pharmacy once a week because she has the end of getting any medication prescribed by the doctor). But the point is that if Cecilia is given bad medical advice about how to treat one of her friends and carries it into effect, then she has an end that is *not* objectively good despite her best efforts. To claim otherwise is not to do justice to Cecilia, who presumably would be shocked and appalled to learn that the treatment she had been administering actually was harming her friend. Cecilia's end of treating her friends in the way prescribed by a medical professional is represented correctly as objectively good. But her subordinate end of administering penicillin is represented *incorrectly* as objectively good. Thus, this attempted explanation of the distinction between Cecilia and the doctor does not work. It does not work because both Cecilia and the doctor really do have ends that they incorrectly represent as objectively good: it is just because of Cecilia's ignorance that she mistakenly adopts an end (*viz.*, administering penicillin to her friend) that is not objectively good.

Now the proponent of GOG might want to back up a bit here. The proponent of GOG might complain that I did not do justice to the attempt to explain why excusing conditions do not apply to the doctor in this case. Clearly Cecilia has a good heart and never would adopt deliberately an end that she does not represent as objectively good. The difference between Cecilia and the doctor is that Cecilia's end really is objectively good *given what she knows about the world* whereas the doctor's end is not. That is, both Cecilia and the doctor incorrectly represent their ends as objectively good. But the difference between them (Cecilia's ignorance is inculpable whereas the doctor's is culpable) lies in the fact that the doctor knew better whereas Cecilia did not. The doctor knew better than not to take a medical history; Cecilia did not know better than not to administer penicillin.

This, it seems to me, is a step in the right direction. But it overlooks a crucial point. Agents

cannot decide on a whim to represent something as objectively good. It is not that easy. Whether something is represented as objectively good by an agent depends on the backdrop of the agent's various beliefs. It is precisely because of this that an agent can come to change its mind about whether something is objectively good. But representing X as objectively good does not occur in isolation. In other words, the problem with this response is that it misconstrues the word 'represent' in GOG: it commits the cart before the horse misunderstanding noted above (in section 2 of this paper). If the doctor's beliefs, values and desires antecedently commit him to representing X as objectively good, then he will continue to represent X as objectively good even if he chooses to pursue some other end. Of course, proponents of GOG think that he cannot pursue some other end. And that is precisely the point: the proponent of GOG has no way of accounting for how the doctor got to where he is (*viz.*, in a case of culpable ignorance) that does not apply also to Cecilia. So the proponent of GOG has no way of accounting for why excusing conditions do not apply to the doctor.

The basic problem with GOG at which I am pointing is quite simple. GOG embeds a necessity; it is trivially true that any necessity claim rules out the possibility of something. In claiming that agents necessarily represent their ends as objectively good, GOG limits the range of possible kinds of moral error to one, namely: being mistaken about what is objectively good. There is not even a way of being mistaken about what is objectively good that does not involve representing an end as objectively good. Thus, GOG rules out the possibility of willfully pursuing an end that one does not represent as objectively good. In so doing, it relegates moral error to the domain of ignorance, thereby precluding the distinction between culpable and inculpable ignorance. Indeed, that is the point: the intuitive pull behind GOG just is the idea that agents would do the "right" thing if only (not only if!) they knew better. But limiting the range of possible kinds of moral error in this way

is not necessary. As I pointed out in the previous section, there are alternative theses about how agents represent their ends (*viz.*, GG and NGG). In addition, as I have tried to show in this section, limiting the range of possible kinds of moral error in this way is not plausible. If I am correct about this, then GOG should be dropped.

As a last ditch effort, the proponent of GOG might point to the role of representation: the proponent of GOG might say that excusing conditions do not apply to the doctor because something went wrong with the way in which the doctor came about representing his end as objectively good. Thus, the proponent of GOG might grant on the one hand that both the doctor and Cecilia have ends that are not in fact objectively good, and the proponent of GOG might grant on the other hand that both the doctor and Cecilia have ends that they represent as objectively good. But the proponent of GOG might say that the doctor brought himself to represent his end as objectively good in a faulty way. The idea would be as follows.

Cecilia sometimes incorrectly represents her ends as objectively good. However, when this happens, it is through no fault of her own. It is because she is deceived by someone else (e.g., a malicious doctor) or something of this nature. The doctor in the example above also incorrectly represents his ends as objectively good. But there is an important difference between the way in which the doctor came to do this and the way in which Cecilia came to do this. Whereas Cecilia's mistake can be traced back to the activity of another, the doctor's mistake cannot be traced back to the activity of another; the doctor's mistake is the product of self-deception. Against the backdrop of the doctor's beliefs, the end that the doctor adopted is not objectively good. The doctor must have recognized this at some level. Therefore, in order to represent the end as objectively good (despite the backdrop of his beliefs) the doctor must have decided to exercise self-deception.

Now it might be the case that the only way to understand some bad agents is to say that they

subjected themselves to self-deception. (I am thinking, here, of the perpetrators of genocide.) But the trouble with this attempted explanation here is that self-deception does not seem to help proponents of GOG. Self-deception does not seem to help proponents of GOG because, assuming that the doctor in the example above *was* subject to self-deception, on GOG the doctor would have represented the end of engaging in this self-deception as objectively good. But self-deception is not an end in itself: the doctor must have represented the end of engaging in this self-deception as objectively good because he must have represented the belief that he would form as a result of this self-deception as objectively good, which in turn could be the case only if he represented this belief as a necessary (or at least sufficient and perhaps efficient) means to the realization of some end that he represented as objectively good. In other words, follow the chain down and the result is that the doctor's *original* beliefs and values commit him to representing the end as objectively good: no self-deception is necessary!

5 The infection theory

One final attempt might be made on the part of the proponent of GOG. The proponent of GOG might advance an infection theory. According to an infection theory, the difference between the doctor and Cecilia does not lie in how they represent their ends. Both the doctor and Cecilia incorrectly represent treating the patient in the example with penicillin as objectively good. According to an infection theory, the difference between the doctor and Cecilia consists in how far up in the means-end ladder their respective mistakes lie. That is, the difference consists in how many of their respective actions and projects these mistakes “infect.”

In the case of Cecilia, the mistake lies close to the surface; Cecilia adopts the prescribed treat-

ment on the basis of the fact that she wants to help her friend and she believes that the doctor can be trusted as a medical authority. Cecilia's mistake "infects" only her treatment of this one person. In the case of the doctor, the mistake lies somewhat higher up; the doctor, recall, decided not to take a medical history on various grounds. This will come up again with various patients (whenever the grounds are satisfied).

This idea needs some working out. (What if Cecilia was a nurse? Or what if the doctor had not slept the night before and it was the last patient of the day?) But hopefully the idea is clear enough for now. Indeed, it seems plausible that something like an infection theory is the rationale for saying that Cecilia has a good heart. Some kinds of mistake are more serious than others, for some kinds of mistake are more fundamental. Erroneously judging that the members of a certain class of people are non-rational is much more serious and much more fundamental than erroneously judging that it is permissible not to take a medical history when one is tired and the patient's problem seems straightforward.

However, this attempt on the part of the proponent of GOG seems to me to be a distractor ("wolf!"). The infection theory might be true after a fashion. It might be true that certain mistakes are more fundamental than others. But if one wants to maintain the distinction between culpable and inculpable ignorance, then the infection theory will not suffice. And this should be clear, for the infection theory was brought in as a way of meliorating the fact that this distinction is precluded by GOG. On GOG, there is no distinction in kind between a shopkeeper who has not mastered the multiplication tables and an Eichmann or a Himmler.

Again, the infection theory might be true after a fashion; it might be true that certain errors are more fundamental than others. But it also might be true that there is a distinction between culpable and inculpable ignorance. And given the prevalence of this distinction in our ordinary

moral reasoning, I think that the argument in this paper constitutes a strong presumption against GOG.

Notes

¹See Henry Allison, *Kant's theory of Freedom* (Cambridge: Cambridge University Press, 1990), p. 91p2 and 97p2; Stephen Engstrom, "The Concept of the Highest Good in Kant's Moral Theory", *Philosophy and Phenomenological Research*, 52 (1992): 747-780, p. 750, 760 and 764; Barbara Herman, *The Practice of Moral Judgment* (Cambridge, Massachusetts: Harvard University Press, 1993), p. 217 and 221-223; Christine Korsgaard, *The Sources of Normativity* (Cambridge: Cambridge University Press, 1996), p. 122; Christine Korsgaard, *Creating the Kingdom of Ends* (Cambridge: Cambridge University Press, 1996), chapters 4, 8 and 9; Allen Wood, "Humanity as an End in Itself", in *Proceedings of the Eighth International Kant Congress* (Milwaukee: Marquette University Press, 1995), I.1:301-319; Allen Wood, "Humanity as an End in Itself", in Paul Guyer (ed.), *Kant's Groundwork of the Metaphysics of Morals: Critical Essays* (Lanham, Boulder, New York and Oxford: Rowman and Littlefield Publishers, Inc., 1998), 165-189; Allen Wood, *Kant's Ethical Thought* (Cambridge: Cambridge University Press, 1999), chapter 4; and Allen Wood, *Kantian Ethics* (Cambridge: Cambridge University Press, 2008), chapter 4.

²For critiques of the regress argument that do not question GOG, see Hannah Ginsborg, "Korsgaard on Choosing Nonmoral Ends", *Ethics* 109.1 (1998), 5-21 and Donald H Regan, "The Value of Rational Nature", *Ethics* 112 (2002): 267-91.

³Christine Korsgaard, *Creating the Kingdom of Ends* (Cambridge: Cambridge University Press, 1996), p. 118p3.

⁴*Ibid.*, p. 118p3-119p1, 258p1, 260p1 and 262p2.

⁵See Allen Wood, *Kant's Moral Religion* (Ithaca, NY: Cornell University Press, 1970).

⁶Christine Korsgaard, *Creating the Kingdom of Ends* (Cambridge: Cambridge University Press, 1996), p. 118p3-119p1; see also Allen Wood, *Kant's Ethical Thought* (Cambridge: Cambridge University Press, 1999), p. 128p3: "... I have a (justifying) *reason* to seek my own happiness (as opposed to a mere *desire* for my happiness, providing me with a *motive* to pursue it) only if I bring my desire under a universal principle ascribing objective value to my happiness; for example, by seeing it as the desire of a being whose objective worth is such that the satisfaction of all its desires is

objectively good.”

⁷See Alan Gewirth, *Reason and Morality* (Chicago: University Press, 1978), p. 49; David Pears, *Motivated Irrationality* (Oxford: Clarendon Press, 1984), chapters 9-10; and Joseph Raz, “On the Guise of the Good”, Oxford Legal Studies Research Paper, Working Paper Series 43 (2008). Although Davidson adopts GOG explicitly, in practice he actually seems to vacillate between GG and GOG (Donald Davidson, *Essays on Actions and Events* (Oxford: Clarendon Press, 1980), p. 21-40).

⁸See Michael Stocker, “Desiring the Bad: An Essay in Moral Psychology”, *Journal of Philosophy* 76 (1979): 738-753; J. David Velleman, “The Guise of the Good”, *NOÛS* 26.1 (1992): 3-26; and David Velleman, *The Possibility of Practical Reason* (Oxford: Oxford University Press, 2000), chapter 5. Hill argues against GOG but it is not clear where he stands with regard to GG and NGG (Thomas E. Hill, Jr, *Human Welfare and Moral Worth: Kantian Perspectives* (Oxford: Clarendon Press, 2002), chapter 8). Kerstein argues against taking GOG as Kant’s view, but it is unclear where Kerstein stands philosophically, and it is unclear whether Kerstein thinks that Kant subscribes to GG or NGG (Samuel J. Kerstein, “Deriving the Formula of Humanity (GMS II, 427-437)”, in Christoph Horn and Dieter Schönecker (Ed.’s), *Groundwork for the Metaphysics of Morals* (Berlin, Germany: Walter de Gruyter, 2006), 200-224).

⁹See Raz, *op. cit.*, 22n23: “we must always remember... [that an agent] might have acted for what he took to be the lesser reason...”