

2015

The Guise of the Bad

Joseph Raz
Columbia Law School, jr159@columbia.edu

Follow this and additional works at: https://scholarship.law.columbia.edu/faculty_scholarship



Part of the [Law and Philosophy Commons](#)

Recommended Citation

Joseph Raz, *The Guise of the Bad*, JOURNAL OF ETHICS & SOCIAL PHILOSOPHY, VOL. 10, No. 3, p. 1, 2016, KING'S COLLEGE LONDON LAW SCHOOL RESEARCH PAPER NO. 2015-24, OXFORD LEGAL STUDIES RESEARCH PAPER NO. 21/2015; COLUMBIA PUBLIC LAW RESEARCH PAPER NO. 14-457 (2015).

Available at: https://scholarship.law.columbia.edu/faculty_scholarship/1909

This Working Paper is brought to you for free and open access by the Faculty Publications at Scholarship Archive. It has been accepted for inclusion in Faculty Scholarship by an authorized administrator of Scholarship Archive. For more information, please contact cls2184@columbia.edu.

The Guise of the Bad

By Joseph Raz

*'Evil is unspectacular and always human,
And shares our bed and eats at our own table.'*

W.H. Auden: *Herman Melville*

My remarks will focus primarily on the connection between what I shall call the Thesis, meaning the thesis of the Guise of the Good, and actions under the Guise of the Bad. I will argue that to the extent that action under the Guise of the Bad is possible it does not contradict the Thesis.

1) The two versions: Reason and Motive

The discussion will proceed on the assumption that actions (and I use the term to refer to actions, activities and omissions) can be bad in some regards, as well as bad all things told. Furthermore the discussion assumes that the fact that an action is bad, or bad in some regard, is not a reason to perform it. So what would constitute an action under the Guise of the Bad? The expression 'the Guise of the Bad' is a stipulative one, whose meaning is not well entrenched in the philosophical lexicon. Various theses can reasonably claim the name. But its meaning cannot be so wide as to include action taken for a reason that is believed to be defeated, meaning taken to be weaker than conflicting reasons that apply in the circumstances.¹

Ordinary akratic action falls under that description. It involves action done

¹ I will follow the custom of often referring to all the reasons supporting one option as a single reason, and to all the reasons conflicting with them as one reason, relying on context to disambiguate the meaning.

for what the agent believes to be a defeated reason, defeated by considerations that establish that the action should not be done, but it is done, and it is done for what the agent believes to be a reason that shows that it has some merit, that there is something good about doing it. The view that I will examine says that **it is possible to perform an action that one believes to be bad (to have bad-making features) and for the reason that it is, as the agent believes, bad.**² I will call that version of the Guise of the Bad Thesis the normative version.

As is obvious, by ‘reason’ I refer to a normative reason. Reasons generally are facts that explain. Normative reasons are somewhat different. They may explain nothing. There may be a reason for an agent to perform an action (which is not also a reason for any other action, nor for any belief, emotion, intention etc.) and because that action is never performed the reason for it does not explain anything.³ Normative reasons, however, can explain (or be central parts of explanations of), for example, actions that are taken for those reasons.

Normative reasons can explain various objects: beliefs, emotions, intentions, or actions. Given our topic, I will ignore reasons other than reasons for action. The view that we examine is not about whether bad properties can in themselves provide or constitute a reason for an action, but about the possibility of actions taken in the belief that they are bad and that that is a reason to take them. One way of expressing the difference is that it is not about whether the badness of actions can be a reason for them (I proceed on the assumption that it cannot), but about whether it is possible to

² The discussion to come will clarify one ambiguity in this formulation of the thesis.

³ Though it can figure in the explanation of its normative and logical implications – in that sense everything is a reason, i.e. figures in the explanation of what it entails.

believe that the bad features of an action are reasons to take it. And that is close to asking whether actions taken in the belief that they are bad are susceptible to normative explanations. This oracular statement itself requires clarification.

Explanations are of various types. For example, Aristotle famously distinguished four types of causes and four types of (causal) explanations depending on which kind of cause features in them. Whether an explanation is a good or successful one often does not affect the type it belongs to. If it fails because it asserts the existence of facts that do not exist it nevertheless belongs to the same type of explanation it would have belonged to had they obtained. Similarly, if it claims that certain facts explain in a certain way (e.g. are an efficient cause of the explanandum) while in fact they do not explain in that way, it nevertheless belongs to the kind it would have belonged to had its claim been true (i.e. efficient cause explanations in my example). To give an example, an epidemiological explanation is an epidemiological explanation even if the statistical connection it relies on does not obtain or the theory of statistical explanation it relies on is mistaken, with the result that the explanation fails.

Matters are a little more complex with what I will call reason explanations. Successful reason explanations are explanations whose core is expressed in statements commonly made using sentences of the form: 'X (an agent) φ ed because of F', where 'because of F' means because F shows the action to be worth doing. For example: Jane ate the apple because it was tasty, watched "Away from Her" (a film) because it is insightful about the way advancing dementia affects couples, etc. In other words, reason explanations connect a reaction of the person (in the cases

we are examining – actions) to features of the world that make the reaction appropriate, when that is why the agent reacts as he does.

But of course people may act because they are mistaken about how things are, or about what is an appropriate reaction to the way things are. In such cases the action taken is not an appropriate response to the facts that prompt it either because those facts are not a reason for the action or because the belief that they exist is false. The agent is not connected to the world in the way he thinks he is. But in acting as he did he attempted to react as one does to normative reasons. That is why the explanation of such actions is similar to successful reason explanations. Of course, it is not a successful reason explanation because there is no reason that can explain the response. It is a failed reason explanation. But there is a successful explanation nearby. It contains a segment of a successful reason explanation: it explains the action by the agent's attempt to conduct himself in a way that is appropriate to how things are. Agents, we may say, take themselves to be normatively guided, guided by a reason, and that is what led them to act as they did. And even though they are not guided by a reason, they tried to be. The explanation, the successful explanation, of their action is therefore an explanation (of at least one kind of case) of attempting to be guided by a reason. Therefore it is an explanation of normative guidance. As a terminological abbreviation (which roughly conforms to the way the terms are often used) let me call explanations of conduct in which agents attempt (successfully or not) to be guided by reasons 'normative explanations'. Successful reason explanations, because they embed in them normative explanations, can also be said to be normative explanations.

So far I have been explaining the normative version of the Guise of the Bad. It says that there can be normative explanations of people's actions in which they take the badness of some actions to be reasons for their performance. Another version, which I will call the motive version, concerns the possibility of acting out of bad motives. That version asserts that **an agent can, without having any relevant false beliefs, perform actions motivated by the badness of those actions, namely by features of the actions that make them bad.**

The condition that the agent is free of relevant false beliefs excludes from the scope of the thesis those cases in which the agent believes that the features that motivate him are good, are features that make the action good in some respect. The condition may be too strong, as it excludes from the scope of the thesis cases in which the agent is ambivalent or self-deceived, cases in which he knows that the features are bad, but deceives himself into believing that they are good, and other more complex psychological ambivalences. Later other kinds of motive explanations are tacitly introduced.

Motive explanations are productive explanations. They explain what brought about the performance of actions. As is the custom I will sometimes refer to them as causal explanations, not meaning by that more than that they explain what produced or brought about the explanandum. So do normative explanations, but unlike the latter, motive explanations do not imply that the agents knew, or believed, that they had any particular motive, let alone that they knew or believed that they acted out of the motives that explain their actions. Motives can be guided or triggered by reasons and they can bring people to act for certain (believed) reasons, but

they can also bring people to act without reason, as when they induce an accidental act, or a false and masking belief about one's reasons.

Let me explain: First, accidental actions, e.g. accidentally knocking over and breaking a wine glass, are not done for a reason (though they may happen in the performance of an intentional action that is taken for a reason, as when we break the glass while passing the salt, as requested). Needless to say, such accidental actions are caused, and the agent's motives can be among the causes of the accident, often – perhaps normally – without the agent being aware of the fact.

More complex are cases in which a motive causes what I called a false masking belief. In these cases agents deceive themselves into a belief. Such beliefs are motivated, e.g., by wishful thinking, or by anxieties and fears about facing what they know to be the case, etc. They think that they act because of their beliefs, but in fact their actions are explained by their motives and not their beliefs, whose role is limited to masking from the agents the true nature of their actions. I will return to masking beliefs shortly.

All this was to show that acting for a bad or worthless motive is a distinctive phenomenon, different from taking the fact that an act is bad or worthless to be a reason to perform it. We thus have at least two different phenomena that could reasonably be taken to be action under the Guise of the Bad, a normative version and a motive version. I remarked earlier that in some respects the motive version as I defined it is too strong. Some would say that in another respect it is too weak. It does not require that the agents either know or believe that their motive or its object is bad. There is of course no objection to additional versions of the Guise of the

Bad. I will not consider this stronger view for two reasons. First, in so far as judgments of people (rather than their actions) go, people are bad if significant actions of theirs are motivated by bad motives, regardless of whether they believed them or their objects to be bad. But that of course is a topic for another occasion. Second, to the extent that the stronger version is philosophically problematic in ways that the weaker version is not, the difficulty or problem is with the possibility of believing that a feature that makes an action bad is a reason for it because it makes it bad. That problem will be examined when considering the normative version, which is the main focus of this paper.

2) The theoretical case for the Guise of the Good

One difficulty in explaining the possibility of acting for the reason that the action would be bad arises out of the case for The Guise of the Good Thesis. Here too there are diverse versions of the Thesis. The version closest to the truth (simplified to avoid various qualifications and complexities) is that when people act with an independent intention their action is done in, and because of, a belief that the action has some feature that makes it good, at least in some respect.⁴ Not all intentional actions are done with independent intentions. However, to abbreviate and avoid awkward formulations, and as we will not be concerned with other kinds of intentional action, I will use ‘intentional action’ to refer to actions done with an independent intention.

Why accept the Thesis? Because intentional human actions are performed by agents who identify them, and take them to have, among other features,

⁴ For the more complete statement and defense of the thesis see *From Normativity to Responsibility* ch. 4.

some that constitute a reason for doing them, that is making them worth doing (at least in some respect). And the features of actions that are reasons for performing these actions are those that make the actions good. That is what makes them worth doing.⁵

This assumes not only that agents believe the action they are about to take will have some feature that appeals to them. It assumes that the agent is aware of the feature and its appeal.⁶ That assumption reflects the fact that people who act intentionally recognise that they may be mistaken, that it is possible that the action does not possess the feature they think it does or that they are wrong to think that the feature makes the action worth doing. That kind of mistake is different from dissatisfaction or disappointment with the action once it is performed because it is not as enjoyable or is not as admired by others, etc. as one wishes that it were. That disappointment can be experienced without the thought that one made a mistake, and agents can tell the difference between these two reactions.

I should underline that this view of human intentional action does not derive from the nature of intentions or of choice. There are animals of species incapable of judgements of that kind yet capable of intentional action. There are humans of whom that is true. However, *Homo sapiens* is among the species whose members can possess rational powers of a kind that enables not only choice of action, but also recognition of the value of things, and these powers once possessed are automatically used (unless

⁵ If one chooses an option, and not being able to identify it, succeeds by luck in performing it rather than some alternative, then the action, while intentional, is not performed with the intention that motivated it. To be done with that intention the intention must guide the action, including identifying what one is doing as that action.

⁶ This could be that it is more important to act quickly than to do what would otherwise be the best action to perform. In that case the chosen action is one that satisfies that condition: it can be chosen and performed quickly. The feature can also be one that is believed to be manifested by several actions among which the agent thinks that there is nothing to choose.

disabled by sleep, drugs, etc.). Hence, human choices and intentional actions depend on beliefs that humans have and are able to reassess. The point relevant to our purpose is that those beliefs attribute to one or more of the options that agents take to be available to them features in virtue of which it merits choosing. That is what the Thesis asserts. And if that is true, then how can one act for the bad, i.e. how can one choose an action for the reason that, as the agent believes, it is bad?

3) Clarifications

Some possible misunderstandings of the Thesis can be easily clarified.

Others pose greater difficulties.

(a) It does not maintain that agents always have a view as to which feature of the action makes it worthwhile, only that they believe that it has some such feature.

(b) It does not assume that agents have adequate competence to identify verbally the features that lead them to think that the action is worthwhile. Their knowledge may be implicit and exceed their ability to articulate it.

(c) In taking features that make an action worth performing as properties that make the action good, at least in some respect, 'good' is stipulatively assigned a wider meaning than it has in English. English is naturally more specific: 'moral', 'attractive', 'enjoyable', 'rewarding', 'helpful' and many others are the sort of concepts used, and while it is always true that an action is good because it is moral or because it is attractive, etc. it would often be awkward or misleading just to describe the action as good. Moreover, many features are contextually good: being funny is good at the right place and the right time, and can be anything but good otherwise. While this makes the meaning with which 'good' is used in stating the

Thesis somewhat stipulative, it does not undermine the truth of the Thesis.

(d) The good-making features of an action can be relational, as when it is instrumentally good, relationally good (good for the friendship, etc.). As the example of instrumental goods illustrates they can be contingent features of the action.

(e) The feature that makes an action good may be that it is the action the agent has a duty to perform. It is therefore good because it is good to do what you have a duty to do. In other words, the Thesis does not assume what is sometimes asserted or denied as “the primacy of the good over the right”. It is neutral regarding that matter. The controversy, as I understand it, is whether there could be actions that are the right thing to do, actions that are one’s duty, that one morally must do, etc. which have these properties not because they are good in some way. The Thesis is silent on that issue.

(f) Moreover, the Thesis does not assume that all people capable of intentional actions believe that when they or others act intentionally they do so in the belief that there is some good in their actions. That is a philosophical thesis that can be and is denied by many people capable of intentional action. The Thesis is that regarding any intentional human action, it is performed by the person who performs it in the belief that there is some good in it. People who have that belief regarding each of their actions need not believe that the same is true of others, nor that it must be true of themselves or of others.

4) The Difficulty

The Guise of the Good Thesis implies that if an act is intentional it is done because of a belief that the action is good in some respect. The Guise of

the Bad says that one can act intentionally because of a belief that the action is bad in some respect. While an action can be both good in some respect and bad in some respect, and it can be believed to be both good in some respect and bad in some respect, its being done because the agent believes it to be bad in some respect cannot make it intentional, at least not if the Thesis is true. For according to it what makes a human action intentional is that it is done in and because of belief that there is some good in it.

5) Conciliation?

But perhaps, appearances to the contrary notwithstanding, the two guises are not in conflict; perhaps the Guise of the Good is a true thesis, while acting under the Guise of the Bad is possible.

The Guise of the Good is a philosophical not a linguistic thesis. It does not imply that if one says or thinks that one's or someone else's action is intentional when that person took it because of a belief that the action's being bad is a reason to take it, one is using English incorrectly, deviating from a linguistic norm etc. Could not that fact be relied upon to show that people can act under the Guise of the Bad, even though the Thesis is correct? Could it not be claimed that the Thesis simply does not apply to all linguistically correct uses of intentional action?

A related approach to conciliation, overlapping the preceding one, would deny that the Thesis is a strictly universal one. Rather, it will be said, it is to be understood as an ordinary generalisation, which like all other generalisations allows for exceptions.

But neither conciliation succeeds. The Guise of the Good Thesis connects intentional action with (a) action that is taken by agents in light of their

view of their situation; and therefore with (b) actions for a reason, namely it holds that agents acting intentionally act because they think that their action is worth doing, that there is a reason for it; and therefore with (c) that to be worth doing the action must have something good about it. Thereby it sees intentional actions as intentional because they can be given normative explanations, showing that the agent is trying to do something good.

Actions under the Guise of the Bad too are subject to normative explanation: they are actions taken because of, and guided by, the action being – as the agents see things – bad. That is why they are problematic. They are problematic because the Thesis makes action under the Guise of the Bad, if it is possible, not merely exceptional but incomprehensible.

The Thesis as presented ties up with reason explanations. Successful reason explanations have to explain why features of an action are, in certain contexts, reasons for it, and why some features can be thought to be such features. The brief explanation, though crude and requiring elaboration, is that features of the action can be reasons for it if they show it to be good, in some respect, and therefore worth doing. Normative explanations need not be successful reason explanations. But they depend on understanding the agents as taking themselves to be acting for reasons. Action under the guise of the bad is also susceptible to normative explanations: it is action for the (believed) reason that the action is bad. But for action under the guise of the bad to be possible we need to understand how something bad in an action can be thought to make it worth doing. Failure to do so – it can be claimed – establishes that there cannot be actions under the Guise of the Bad. Success in providing such an explanation will refute the Thesis of

the Guise of the Good. Either way, conciliation fails. Either the Thesis is correct or there can be action under the Guise of the Bad; but not both. Or at least that is the claim.

6) Narrowing or Revising the Thesis?

Given that there is a strong case for the Thesis it is unlikely that the normative version of the Guise of the Bad can refute it. More plausibly either no action under the Guise of the Bad can take place, or the scope of the Thesis has to be narrowed to accommodate such actions. I will examine one type of objection to the Thesis, an objection that the case that is presented for it justifies a weaker or narrower thesis only.

First, why assume that regarding each of their intentional actions agents believe that there is some good in them rather than that they believe that each possesses some property that does in fact makes the action good in some respect, but that they may be unaware that it makes the action good in some respect? I assume that understanding any of the value-, i.e. good-making, properties, involves knowing that there are others like them, that is that they are instances of a more general category whose other instances would, among other things, also make the action worth doing. Therefore, if agents believe that the action has some specific feature that is in fact good-making, and they broadly know what that feature is like at least to the extent that they believe it makes the action worth doing, they at least implicitly understand that it must be of the kind of properties that can have other instantiations that would also make other actions worth doing, namely that it is a good-making property.⁷ That is what is meant by the

⁷ That is not a terminological stipulation. It attributes to agents a substantive belief that what makes actions worth doing is something that makes them (in the extended sense of the word used here, and in philosophical discussions generally) good or of value in some respect.

claim that if they believe that the action has a feature that is a good-making one they also believe that it has some good.

Second, these observations assume that in deploying concepts we are, normally vaguely and implicitly only, familiar with their rough inter-relations. We know that if we run to the station because the bus will leave in 2 minutes then the bus's impending departure is only part of the reason to run to catch it, and therefore that when the other facts that are part of the reason are absent its departure will not be a reason to get to it. Etc.

But circumstances do not always determine what beliefs people have, or do not have. Consider:

I told him it wasn't long till morning, and how in the morning somebody would find them, and then all of it, me and Dick and all, would seem like something they dreamed. I wasn't kidding him. I didn't want to harm the man. I thought he was a very nice gentleman. Soft-spoken. I thought so right up to the moment I cut his throat. I didn't realize what I'd done till I heard the sound. (Perry Smith as reported in Capote's *In Cold Blood*)

Did he act intentionally? Did he believe he had a reason to kill? Given that description it is difficult to say, and one possibility is that it was a marginal case: the act was intentional and is subject to a normative explanation, but abnormally so: Smith's control of the initiation and course of his action being limited. He knew not what he was doing in a somewhat literal sense. Does that call for a modification of the Thesis? I doubt it. Psychological theses are always subject to marginal and anomalous cases. That is their nature. It is not part of the content of these theses that they are.

Third, let's return to the case of masking beliefs. Recall failed attempts to

be guided by reasons: Agents act – as they see matters – for a reason, but their belief that the reason is there is mistaken. Their actions can be given normative explanations, even though they cannot be given successful reason explanations. While they are not brought about by a normative reason, they are brought about by belief in a normative reason. In this they differ from cases in which agents’ motives induce a false belief in a reason, as when a desire to revenge an injury induces belief that the person who accidentally and blamelessly caused it is guilty of deliberately causing it, thus leading the agent to believe that he has reason to retaliate. These are cases of self-deception, of motivated irrationality. They are not susceptible to normative explanations. In them while the belief has a role in explaining the action, its role is limited to disguising from the agent what “really” brings it about, namely that it is the motive not the belief in the reason, which, being a self-deceived belief, is not a full or regular belief. It does not fulfill the role of ordinary belief in the life of the person.

Do masking beliefs present a case for revising the Thesis? The answer depends on the way we resolve an ambiguity in the Thesis. It sets two conditions: when acting with an intention people act in the belief that there is some good in the action – that condition is met. The self-deceived believe that there is some good in the action. The second condition is that they act because of this belief. As I remarked earlier, the self-deceived belief is part of the explanation of the action: they need to deceive themselves to allow their motives full reign. But this explanatory role of the belief is not one that renders the action subject to a normative explanation. Should we say therefore, that when understood as intended the Thesis has to be narrowed down to allow that if agents’ belief in the value of the actions is

self-deceived then they can act intentionally without their actions being explained normatively by that belief? That is not clear. The self-deceived believe that their actions are motivated by (their belief in) reasons. And that belief, though incorrect, need not be self-deceiving.⁸ That shows that they are trying to be so guided. That, as you will recall, was what the concept of normative explanation tried to capture. It turns out that it does not cover all cases of attempts to be guided by reasons. But that may tilt one towards an understanding of the Guise of the Good that applies to their cases as well. The case is moot, and ultimately not much depends on which way we go. So, let us assume for the time being that the Thesis does not need to be narrowed.⁹

7) The Luciferian option

It is time to discuss the Guise of the Bad. But first – an analogy and a contrast with weakness of the will. The analogy is limited to one point: weak-willed action both is and is not susceptible to normative explanation. It is, because it is taken for the reason that, as the agent sees matters, the action is good in some respect. The food is tasty, or having it will assuage the pangs of hunger, etc. At the same time it is not (altogether) susceptible to a normative explanation for the agent does not believe that there is sufficient reason for the action. In his view the reason for the action is defeated by reasons against it. That is what makes the action weak-willed and while the fact that the agent acted akratically may be explained, the explanation will not be a normative explanation. It will be an explanation of

⁸ Though in some cases it may be best to narrow it. We need not consider all the possible complexities of such cases.

⁹ That reminds us that motives as well as beliefs can identify actions that are to be the object of intentions, enabling those intentions to guide the performance of those actions. What they cannot do is replace the Thesis and establish another route to forming intentions – that is why we need to deceive ourselves into believing that there is some good in the action.

why normative explanations fail to apply in this case.¹⁰

By its very identification, action taken for the reason that it is bad (or has some bad feature) is subject to normative explanation: it is an action for a (believed) reason. But the explanation cannot be complete, or completely successful, at least not if the Thesis is correct. The agent attempts to act for a reason, but fails. He fails because that the action is bad is no reason to take it. But there is more to his failure. It differs from that of a person who mistakes a poison for wholesome food. The agent acting for the bad makes no mistake about the character of his action, except for those that are entailed by his main mistake, which is about the possibility that a bad feature of the action would make it worth doing.

Is it possible to make such a mistake? Do those who act under the guise of the bad make it? It is possible to be mistaken or confused about concepts. But not all mistakes about a concept are consistent with being able to deploy the concept. Those, if there are such, who act under the Guise of the Bad deploy the concept of a reason, whether or not they would express themselves in these words. They take the badness of an action as favouring the action. That is, they act for a reason, and to do that they must have some mastery of that concept, or of closely related ones. They have the concept to the extent that their thoughts and words can be described using it. Given that they have the concept of a reason, is it possible that they believe that the badness of an action favours it?

Suppose that one performs an action, believing it to be bad, and saying or thinking that one is taking the action for that reason, yet, while employing

¹⁰ Though it is worth remembering that the fact that the action (or any other condition) cannot be normatively explained, or the fact that it is irrational, does not entail that it is necessarily bad all things considered.

the concept of a normative reason, having no thought about that concept. Perhaps in that case it is impossible to understand the agent to be taking the badness of the action as his reason. We are forced to understand his reason in a way that avoids attributing to him this incomprehensible thought: He may take the fact that it is bad for someone as his reason, for he may enjoy harming that person, or he may take the fact that it is thought by some or all to be bad as his reason, for he aims to shock them. But in such and any other case we may imagine there is, in the agent's eyes, some good in the action, and that is his reason.

Are matters different if his thoughts that lead to the action are, in part, about the concept of a normative reason, or about the Thesis? There is a traditional explanation of an affirmative answer, a secular version of the Luciferian motive. One may have an urge to defy the limits of thought, the limits of what we can do or think, limits that are expressed in our concepts. Ordinarily, the drive to extend boundaries is different. It is the drive to create things of new kinds, requiring new concepts for their description. One may wish to write a narrative text that is neither fact nor fiction, and people who had that desire created a new genre of writing that is sometimes called 'non-fiction novel' or 'faction'. Things are different when the desire is to defy basic concepts: to produce an object that will not be an object, to add one to one without their sum being two. Or, to take the badness of an action as something that makes it worth doing, that is a reason for doing it. Alternatively, and that is another way to try to defy the Thesis, one may want to do something for a reason but not by taking anything to show that it is worth doing. To do that one could perform the action because it possesses a feature that does not make it good. It need

not be one that makes it bad. One chooses a bad feature as one's reason just to make abundantly clear that one is acting for a reason but without anything assumed to be good about the action being one's reason, and therefore without taking the reason to establish that the action is worth doing in any respect at all. There are probably other variants of Luciferian motives.

Some people may deny that there can be Luciferian motives or Luciferian goals. Stating them does not make sense. To be sure, they will say, one would be using grammatical sentences, but they have no content. One is merely verbalising, and perhaps getting some feeling of satisfaction or comfort from one's own thoughts or words. But there is no content to those thoughts or sentences. But it is generally recognised that in some sense we do understand some kinds of nonsensical motives and goals. They are stated in sentences that do not give rise to the incomprehension of the famous 'Green ideas sleep furiously' or 'Saturday is in bed'. We do not need an account of what sense they make and how. We can assume that people can have these motives and goals. The question is to what extent one can succeed in them.

We do not expect complete success. Lucifer cannot win. But he can have something that can be thought of as a partial success. Like the weak willed, the Luciferian aspires to both eat his cake and have it. The weak willed aims, metaphorically speaking, to satisfy reason while flouting it. He succeeds in acting for a reason, and to the extent that he does that his conduct is subject to normative explanation. But he knowingly follows a reason while acting against reason, and that cannot be given a normative explanation. All we can do is explain (non-normatively) the motivation that

leads him to attempt the impossible.

The Luciferian aspires to act for a reason by taking something that cannot be a reason for a reason. He knows what he is doing. By definition the Luciferian aspires to break the mould, so he knows what it is. In effect he is trying by his action to **make** the bad a reason for action, knowing that it is not, or has not been so far, or cannot be for beings lesser than he is.

I said that he couldn't succeed. But it may appear that he succeeds in something. He does not make the bad into a reason. But he performs the action that is bad because it is bad. Does not that show that he is guided by it and that his action is susceptible to a normative explanation, namely that he showed that it is possible to think that the bad is a reason?

Assume that he succeeds in that. I mean: he really acts out of a belief that the badness of the action is a reason for it. In itself that would not be remarkable. Many act out of a false belief that they have a reason when they do not. In the case of the Luciferian the belief itself would be irrational: he irrationally believes – on our supposition – that by his choice he makes the badness into a reason, and that is irrational. But even so his action is susceptible to a normative explanation, though his having that irrational belief is not. It can only be explained non-normatively, by reference to his motivation. If this is how to understand the Luciferian then action under the Guise of the Bad is possible and the Thesis has to be scaled down to allow for the Luciferian exception. This exception itself may be unexceptional. It is but one way in which theses about what can and cannot be thought are subject to exceptions, including exceptions through defiance.

But is this really how the Luciferian has to be understood? We have already encountered, in another context, the alternative. In discussing masking

beliefs, we saw that motives can lead to intentional action in a way independent of a belief in there being something good in the action, provided they can induce an irrational belief that there is some good in the action. Perhaps the Luciferian action is similar: the bad action is to be explained non-normatively by the motive, but the rationalisation is the irrational belief that the bad can be a reason. That belief does not explain the action, but it enables the Luciferian to believe that he achieved his goal. He is self-deceived, and the Thesis is intact. The Luciferian is acting intentionally but not for a reason, let alone the reason that the action is bad.

But does the Luciferian act under the Guise of the Bad? He does in the motivation version. Does he in the normative version? I suppose that the answer is Yes and No. As he sees matters he does. But that is due to his self-deceived belief that the bad can be a reason. So, in the objective sense, from an objective perspective, he does not. Which is the more important of these perspectives? Clearly, the subjective: Being the perspective of the agent it affects his actions and impact in the world. The objective perspective is only conceptually important.