

University Press Scholarship Online
Oxford Scholarship Online



Desire, Practical Reason, and the Good

Sergio Tenenbaum

Print publication date: 2010

Print ISBN-13: 9780195382440

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780195382440.001.0001

Good and Good For

Sergio Tenenbaum (Contributor Webpage)

DOI:10.1093/acprof:oso/9780195382440.003.0009

[-] Abstract and Keywords

This chapter tries to understand the notions of good and good for (and related evaluative notions) that seem to be in operation when, for instance, one says that by bringing about the greatest good, or by acting as morality requires us to act, we do something that is worse for Bill. The first section of the chapter tries to present more specific intuitions that seem to require that good and good for be capable of diverging, as well as presenting some intuitive reasons to think that these notions are also importantly related. This section ends by trying to express more precisely the difficulty in accounting for all these intuitions. The second section of the chapter argues that a number of seemingly promising ways of accounting for the relation between good and good for face serious problems. The third section suggests that the problems are actually more complex than they appear at first. When we consider personal relations, whatever reason we had to accept two different evaluative notions seems now to call for an indefinite number of such notions. For as we develop personal relationships, the same reasons that led us to drive a wedge between “good” and “good for” seem to work to generate a putatively infinite number of similar evaluative notions. In the last sections, I sketch a (hopefully) better account of “good” and “good for”, which I call “the appearance view”. According to the appearance view, the distinction between “good for” and “good” is not a distinction between two modes of evaluation, but a distinction between what appears to be good to someone and what is, in fact, good. However, the appearance view does not imply that any kind of evaluative appearance is constitutive of what is good for someone. The appearance view accounts for what is good or bad for someone in terms of how recalcitrant or persistent an evaluative appearance is. This account of “good for” allows us to show that a defense of “a guise of the good thesis” (or as I call it “a scholastic view”) can take the notion of good simpliciter as its central notion, and still be able to make room for a notion of well-being or good for.

Keywords: Good, welfare, well-being, practical reason, good for, Kant, Hume, Aristotle, Moore, sacrifice

1. Introduction

One of the most famous arguments in Rawls's *Theory of Justice* is the argument against utilitarianism. According to Rawls, by redistributing utility without any concern about who the "bearers" of utility will be, utilitarianism fails to respect the distinction among persons. As Rawls puts it:

This [classical utilitarianism's] view of social cooperation is the consequence of extending to society the principle of choice for one man, and then, to make this extension work, conflates all persons into one through the imaginative acts of the impartial sympathetic spectator. Utilitarianism does not take seriously the distinction between persons.¹

The principle of choice for individuals allows a person to sacrifice what is good for him at one time for the sake of what will be better for him overall; in Rawls's estimation, the utilitarian illicitly generalizes from that principle to the conclusion that one can also unproblematically sacrifice what is good for one person for the sake of what is overall better. This idea that there seems to be something wrong with the way that utilitarianism moves from intrapersonal to interpersonal trade-offs has been widely discussed and often endorsed. The idea that one cannot think about interpersonal redistribution of goods or utility simply on the model of intrapersonal redistribution of goods or utility without further argument now enjoys the status of a default position.² However, it is worth noting that Rawls's argument seems to depend on a duality between something like "good" and "good for," and despite the abiding influence of Rawls's argument, the relation between "good" and "good for" in the context of this argument has not received any adequate account, or so I will argue. This is particularly important for someone who is tempted by the guise of the good thesis. If the characteristic claim of the guise of the good thesis, or what I call "the scholastic view,"³ is that desires, and perhaps other practical attitudes, represent their objects as good, then the existence of these two evaluative notions make the scholastic view ambiguous. Are the objects of desire (p.203) conceived as good *simpliciter* or as good for the agent? In the absence of a proper justification for one of these options, it's not clear why we should accept *any* version of the scholastic view. In light of these concerns, I have a particular interest in the notions of "good" and "good for." Notions such as well-being and welfare are put to work for various philosophical jobs, and I doubt that there is one univocal sense of "well-being" and "good for" that is employed in all these debates. However, I am interested in these two notions only insofar as they are two candidates for being the formal end of practical reason and action,⁴ for being what an agent necessarily aims at when he or she acts or reasons practically, and thus the conditions of adequacy to which I'll hold an account of these notions are not the same that inform many theories of well-being.

Arguably, there are important historical precedents for the view that each of these notions is the one that plays this role in practical reasoning and action. According to some interpretations of Aristotle, *eudaimonia*, the end of all actions, is better understood as something that is good *for the agent*. On the other hand, Kant certainly holds that the good *simpliciter* is the object of rational volition; in fact, in Kant's view, the good is nothing but that which is necessarily the object of every rational agent's faculty of desire.⁵ No doubt that there are many other evaluative

notions, but no other seems to be a serious candidate to be the formal end of practical reason. No one would claim, for instance, that all action and all practical reasoning aims at the beautiful or the aesthetically good. However, the overall nature of both “good” and “good for” places each of them as a natural candidate for exactly this role.⁶

This chapter will argue that “good” is primary and “good for” should be understood as some way in which certain things persistently *appear* to be good to certain agents. In the course of arguing for my view, I try to show that the various explicit and implicit accounts of the distinction in the philosophical literature do a poor job of accounting for important intuitions about the distinction, but here again the importance of these intuitions should be gauged in accordance to whether they generate a notion that can be a plausible candidate for being the formal end of practical reason.

In order to get a better grip on the notions of “good” and “good for” as I understand them, it is worth examining how Rawls’s argument seems⁷ to be relying on such a distinction. This (putative) dependence can be better appreciated if we look at what happens when a distribution is fair. Suppose that a fair distribution would force Bill to sacrifice some of his income for the benefit of the worse off. Let us assume that there are no further complications; we accept that this is the just outcome, and that all things considered, this is what ought to happen. Now it would be natural to say that it is good that the money be redistributed, and insofar as Bill has a well-developed sense of justice, is fully informed, etc., he will agree with this judgment, and, one hopes, be willing to accept that the **(p.204)** income be redistributed. However, Bill must also experience the whole transaction as a sacrifice. After all, it was the whole point of the argument from the separateness of persons that Bill cannot, or at least need not, treat trade-offs with other agents in the same way that he would treat trade-offs among different time slices of his self. In particular, he cannot, or need not, think that his loss was fully compensated by the gain made by someone else in the same way that one is fully compensated when one makes a sacrifice at a certain point in one’s life so that one can benefit from a greater gain at a later point. Or, as one would ordinarily put it, the redistribution is good but not good for Bill; what is good and what is good for someone are capable of diverging. So it seems that there is a very natural path from Rawls’s classic argument to the conclusion that there are two independent overall evaluative perspectives captured, respectively, by the notions of good and good for. Of course, none of these remarks are supposed to help establish Rawls’s conclusion, or, for that matter, any other definitive conclusion. They only aim to show a quick route to the intuitive plausibility of the view that “good” and “good for,” suitably understood, mark two relatively independent evaluative perspectives. In fact, it is worth looking at what was said by someone who recently argued against the existence of these two independent evaluative notions:

Why should we not just promote all agents’ valuable activity, without worrying about anyone’s well-being? Of course, there’ll be trade-offs to be made between different people’s valuable activity, as there are trade-offs to be made between different valuable activities in our lives.... The same values are involved in these interpersonal trade-offs as in the intra-personal ones. We can make the trade-offs without thinking of well-being, so why not do so?⁸

This is exactly the kind of analogy that Rawls thought to be illicit in *A Theory of Justice*.

The first section of the chapter tries to present more specific intuitions that seem to require that “good” and “good for” be capable of constituting diverging overall evaluative perspectives, while also presenting some intuitive reasons to think that these notions are importantly related. This section ends by trying to express more precisely the difficulty in accounting for all these intuitions. The next section argues that a number of seemingly promising ways of accounting for the relation between “good” and “good for” face serious problems. In the last two sections, I offer an alternative suggestion of how to understand the relation between these notions, a suggestion that promises to do better than the accounts discussed in the third section while remaining compatible with various substantive views on the nature of well-being. I should point out that these issues deserve much more attention than I can give them here. This chapter should be seen only as a sketch of a problem and a solution; its cogency depends on the availability of a more detailed account of the view proposed here.⁹

(p.205) 2. The Intuitions and the Problem

The notion of “good for” has been identified in the literature with a notion of well-being. However, the notions that go under this heading are often trying to capture different concepts and often have different aims. Many philosophers developing the notion of well-being take it as a constraint that this notion will provide the material for moral or political theory construction. Although I’ll look at similar constraints, my aim is not to try to see what notion of well-being can support this larger theoretical purpose, but to try to account for much simpler intuitions such as, for instance, the intuition that certain choices that we ought to make involve personal sacrifices; that sometimes people who behave badly do well for themselves; that sometimes I experience my duties as a constraint that I reasonably wish I could somehow avoid; etc. The conjecture of this chapter is that these are the intuitions that suggest that we have two evaluative perspectives, each of them constituting a plausible candidate for being the formal end of practical reason, so that if we can account for these intuitions while privileging one of these two notions, we have thereby shown that the privileged notion is the only serious candidate for being the formal end of practical reason.¹⁰

In all these cases, we seem to encounter two evaluative notions that are diverging; one that tells us what we ought to do, and the other that tells us that something else is desirable from a certain privileged perspective. Note, in contrast, for instance, Raz’s understanding of the notion of well-being. Raz says that the concept of well-being “captures one crucial evaluation of a person’s life; how good and successful is it from his point of view?”¹¹ This could be understood as compatible with the notion of “good for” I want to capture. However, when Raz considers the life of a person who “undergoes great deprivation in order to bring medical help to the victims of an epidemic,” he argues that this person is still doing well in terms of well-being, since “his life is no less successful, rewarding, or accomplished.”¹² It seems that Raz is right in saying that, in an important sense, the life of such a sacrificing individual is successful and accomplished, and I would venture to say that it is, in fact, a life of an agent who, as far as we can tell, chooses rightly. Moreover, such an agent must look back at his life and find that his life is or was a meaningful one. However, in many ways, not all is so great from the agent’s point of view, and, more important for our purposes, ordinarily one would not consider sacrificing oneself this way to be something that is *good for* the individual. In fact, it seems that this is exactly what constitutes this agent as a *sacrificing agent*;¹³ that is, we regard the choices of the agent as cases of sacrifice precisely because they are not good for the agent; they are the kind of losses

for the agent that on Rawls's view are not compensated by other people's gains. These are the intuitions that lead us in the direction of a wedge between "good" and "good for." All **(p.206)** these are pretty trivial intuitions, but, if I am right, they're enough to generate serious difficulties in understanding the relation between "good" and "good for," as I understand it.¹⁴

(a) Agents and Beneficiaries of Sacrifice

Teresa gives up her chance of having a stellar career in the music industry, a career that would generate a lot of money and fame for her, so as to help the poor in the third world. She has now just finished rescuing many children who were trapped in a cave. The rescue effort caused her to fracture many bones, but since she can't get any medical treatment she needs to endure a great deal of pain.

Now there would be something obviously inappropriate if Teresa's friend Karol would say, as he sees her at this moment, "Hi Teresa, you are doing really well. Once again you succeeded in promoting something great. Your life is just great; in fact, it's hard to see how things could be going much better for you!" Of course, there is a sense in which Teresa's life *is* great, but there seems to be an obvious sense in which her great pain and deprivation is *not* great. Now, if the children she saves go on to have lives much like the life that Teresa forsook, *they* will have very good lives, even though, in some sense, none of them has promoted as much good as Teresa did. In fact, it seems intuitive to say that this is exactly what Teresa did that was so great: by sacrificing her chance to lead a life that is good for her, Teresa allows each of the children to lead a life that is good for him or her.

(b) Getting Away with Murder and Being Framed for It

This is adapted from the movie *Body Heat*:

Matty murders her rich husband, but she is careful enough to frame Ned for it. Ned spends the rest of his life in a maximum-security jail. Matty moves to a tropical paradise, where she enjoys all the money she inherited from her husband as well as a successful career as a lawyer.

Here it would seem inappropriate, to say the least, for Matty's friend Dick to say: "Oh you poor Matty, things have gone awful for you. You made some poor choices in your life, and you failed to promote the good on various occasions," or to feel pity for Matty for being one of the constituents of a state of affairs that is intrinsically bad. Matty did something awful and, as we would say it, "she got away with it"; that is, she is now doing really well for herself despite having done something so awful.¹⁵ In fact, it seems that one would be tempted to say that it's not fair that things *go so well for her*¹⁶ after what she did. No doubt, for Ned life is awful, and things do go very badly for him. The fact that at least Matty is doing well can, if anything, only make things worse for him.

(p.207) (3) Wish and Duty

Let us look at the following case:

Marshall is about to talk to his grandmother on the phone. He hasn't had her over in a long time, so he thinks it is his duty to invite her to come for dinner sometime the

following week. He takes family duties very seriously, even though he finds his grandmother extremely dull, and he does not particularly enjoy cooking. He'll call her and invite her, and if she accepts the invitation, he'll go ahead and make dinner for her and entertain her for that evening. However, Marshall dreads the thought of cooking for his grandmother and spending an evening with her. As he calls her, he wishes she'll decline his invitation.

It seems perfectly reasonable that Marshall should have these attitudes, and the obvious explanation is that although he thinks it would be good to have his grandmother over for dinner, it would not be good for him to spend an evening in this manner. It would make very little sense for an impartial observer to have the same conflicting attitudes as Marshall. The impartial observer would simply wish for whatever she thought would be best to happen.

(d) Inheritance

Let us look at the following case:

Paris did not know that she had a very rich distant relative, a young man she had never met. She learned about his existence when he died without leaving a will and she turned out to be the nearest of kin, and the rightful inheritor of all his money.

Here it is perfectly appropriate for a friend to say (although it might be insensitive to put in these exact words): "I guess it's sad to hear that this guy died. But overall this is good news! You're so lucky; things have turned really well for you." Again the obvious explanation is that although the redistribution of the money could not possibly compensate for the badness of someone's death, it is still *good for Paris* that things turned out this way. Notice that it seems to make sense for Paris to wish that she would inherit money in this way, to see it as a desirable outcome, in a way that it would not make sense for an impartial observer who would regard this scenario, other things being equal, more or less as simply resulting in the net loss of one life. Similarly, it could not matter for an impartial observer, other things being equal, whether Paris or someone else inherits the money. In light of these cases, it seems that we can come up with the following desiderata for an account of the relation between "good" and "good for":

(1) Sacrifice—At least sometimes, when the agent chooses what is good *simpliciter*, she will be making a sacrifice. A sacrifice involves doing something that is in some way not good from the perspective of the agent.

(p.208) (2) Agent's Emotions and Attitudes—The appropriateness of certain agents' emotions such as pride or shame depends on whether what the agent chose was good for him or not (that is, even if paying myself a large bonus is what is best *simpliciter*, I should not feel the same pride in doing it as I would feel if I accepted a pay cut because it was what I thought to be best *simpliciter*).

(3) DIVERGENCE OF ATTITUDES—In some cases the attitudes of impartial observers and the recipient of a benefit or a harm will, or ought to, diverge. While it might make sense for me to be glad that I, rather than a stranger, inherit some money, it makes no sense for an impartial observer to have the same attitude.

(4) TEMPTATION AND NORMATIVE CONFLICT—Cases in which “good for” and “good” diverge are typically cases in which the agent faces temptation and in which there is at least an apparent normative conflict.

Moreover an adequate account of these notions should be able to answer a host of questions. What should a rational agent do when she could either promote the good or what is good for her? Should she always choose one over the other? Should she weigh the two? And how should she feel with respect to the fact that she failed to promote what is good or what is good for her? Should it matter (to her) in any special way that she failed to promote what is good or good for her? What should the attitude of a third party be when someone chooses what is good for him rather than what is good *simpliciter* or vice versa? These are the kinds of questions that, ideally, an account of the relation between what is good and what is good for would be able to answer.

Before we move on, I should first note that, as we saw even in Teresa’s case, although “good” and “good for” may potentially diverge, there are also strong intuitions that they are related. It seems that almost nothing, if anything at all, could count as good if it were not good for someone. Even if some people are convinced, for instance, that unobservable beauty is good, it would be hard to deny that many things are good because they are good for someone. As Mill suggests, a sacrifice is only good if it results in something that is good for someone.¹⁷

Now the most straightforward way of explaining the difference between “good” and “good for” would be to claim that there are two irreducible evaluative notions here, something like prudential value and moral value. Let us first distinguish between two ways of understanding these evaluative notions. The first one is purely descriptive: To say that something has a certain kind of value is just to attribute a certain kind of (natural or nonnatural) property to it. If the claim is that there are two such properties and one is not reducible to the other, this is all fine and good, except that it doesn’t seem to get us what we wanted. First, it tells us nothing yet about why agents should *care* about whether what they promote is good or good for them. Making the distinction this way does not capture our intuitions about these notions, since it seems that, if not in all cases at least in typical cases, it matters to us, and it should matter to us, that something (**p.209**) is good or good for us. Moreover, it does not tell us how to understand the seemingly incompatible demands that *the good* and *the good for someone* make on an agent. Finally, this way of making the distinction also tells us nothing about how the notions are *related*. No doubt a view that took these notions to be descriptive could try to add an account of the demands that these values make on us and of the relation between the values, but my point is only that, by itself, the strategy is incomplete. We’d need to add to this account a further account of the nature of these demands, and it is the plausibility of accounts of *that* kind that I am trying to examine.

In fact, the issues I want to examine can be brought into focus by considering a view that admirably answers a related question; namely, classical utilitarianism. Classical utilitarianism has a very elegant way of explaining “good” in terms of “good for.” According to classical utilitarianism, one’s greatest happiness is what is best for the agent (and the more happiness an agent enjoys, the better it is for her). One’s happiness is understood as total pleasure minus total pain. The general good is the sum of the happiness of all agents; the greater the sum, the better the state of affairs. However, classical utilitarianism by itself does not say whether one

should pursue one's own good or the general good, whether it should matter to an agent that she has to pursue one at the expense of the other, and if so, how it should matter, whether there are two competing notions of right corresponding to two competing evaluative notions, etc. But these are exactly the questions we will be interested in.¹⁸

More generally, we can put the issue in the following way. I will assume that evaluative notions are supposed to have normative implications. Now the question is how we are supposed to understand the normative implications. If they generate only *prima facie* or *pro tanto* reasons, then we can understand that sometimes one reason will override the other, and that, perhaps, the overridden reason might generate some regret. But we cannot understand, for instance, why it might be inappropriate to say to Teresa that her life is just great, or that she couldn't be doing much better. After all, the mere fact that there is *something* of *some* value that was not pursued can hardly give us a reason to think that someone's life is not going very well. And it's hard to see why we should not think that Matty does much worse than Teresa, since Matty presumably chose to act on a reason that was only a *prima facie* reason, whereas Teresa at least chose in accordance with what there was most reason to do. Although it is not clear how this suggestion would deal with cases (c) and (d), which do not center on choices, it is hard to see how it could serve as the basis of a satisfactory account. After all, it is not clear how these resources could explain the asymmetries in Marshall's and Paris's positions on the one hand, and the position of the impartial observers on the other hand, given that there's nothing in the distinction between *prima facie* and all-out reasons to distinguish their positions. For instance, one could say that it makes sense to regret that a lesser good was not promoted even **(p.210)** when a greater good was, and we can feel some satisfaction in the obtaining of a lesser good even when a greater good was not obtained. But this is very far from accounting what happens in *Inheritance*. After all, the impartial observer sees nothing good in what happens; rather than observing a rich person and a not-so-rich person, the impartial observer sees now a rich person and a dead person. But this is surely not how Paris sees things.

More precisely, we can see the problem as arising from a conflict among the following tenets:

- (1) "X is good for A" and "X is good" express genuinely evaluative claims.
- (2) In some cases, it is true that all things considered, X is good for A, while it is true that all things considered Y is good, when X and Y are not jointly realizable states of affairs.¹⁹
- (3) Evaluative truths imply normative truths.

These three claims on their own are not mutually incompatible, and all the parties to the debate in section 2 accept (1) through (3). To generate a conflict we first need a much more specific version of (3):

- (3a) Evaluative truths of the type "all things considered, X is good for A" and also of the type "all things considered, X is good" imply rational requirements of the kind "All things considered, A ought to bring about X," when it is in A's power to bring about X.

And we need to add a clause with respect to the impossibility of true dilemmas of rationality:

- (4) It is never true that all things considered we ought to bring about X and Y when X and Y are not jointly realizable.

Obviously we cannot construct a parallel claim to (3a) for every single kind of evaluative claim. What would all things considered be most beautiful, for instance, is not necessarily what we should bring about. This is because aesthetic goods are what we can call a “merely contributory good”; these goods merely contribute, or may contribute in certain conditions, to our estimation of the value of a state of affairs. Claims about aesthetic goods are claims about a *kind of good*, not claims about *the good* or what is good *overall*. If we can show that either “good” *sans phrase* or “good for A” is a merely contributory good, or that one is just a kind of good of the other type, then we have an easy route to deny the truth of (3a) while staying clear of any contradiction. The most obvious way of doing this is by employing a reductive approach; it is to try to show that either to say that “X is good” *sans phrase* is just to say that X is a contributory good for A of a certain kind K, or to try to show that to say that “X is good for A” is just to say that X is a contributory good *sans phrase* of a certain kind K. Indeed these are the two reductive approaches we will consider below.²⁰

Claim (4) is reasonably intuitive but it has been denied exactly because of conflicts between impersonal and personal goods. Sidgwick has famously suggested that reason could not settle between the demand **(p.211)** to pursue the greatest good and the demand to pursue the individual good; or in our words, the demand to pursue what is good and what is good for oneself.²¹ Denying (4) would constitute what we might call “incommensurability approaches.” Ideally, we would investigate the plausibility of particular approaches of this kind, but I’ll leave them aside and simply assume that if we can come up with an alternative account, it would be better to come to the conclusion that there are not two incompatible, formal ends of practical reason.

As the astute reader has noticed, strictly speaking, (1), (2), (3a), and (4) do not generate a contradiction unless we also assume that some of the situations in which all-things-considered evaluative claims of the two kinds of conflict are cases in which the agent can bring about X or Y. This is, no doubt, a weak enough assumption, and it does us no harm to add it right here. However, even in cases in which the agent cannot bring about either X or Y, the conflict between the two all-things-considered evaluative judgments is still problematic. Whether things turn out for the best or not implies the appropriateness of various attitudes of rejoicing, regret, etc., or so I’ll assume. Some of these implications might conflict with our ordinary attitudes toward outcomes in which what is good and what is good for X turn out to be different. When examining the plausibility of various options, we need to see not only if they prescribe the right actions but also whether they explain why we think that certain attitudes are appropriate. Indeed, some of the cases above are cases in which we might be puzzled about exactly how to explain a certain attitude, and my contention is that often it’s harder for a theory to explain the appropriateness (or inappropriateness) of certain attitudes than to explain why a certain action is the one the agent ought to perform. Although I’ll not try to spell out more precisely the conditions of adequacy with regard to these attitudes, I’ll often examine whether a certain view matches our intuitive judgments about the appropriateness of certain attitudes. I would be remiss not to end the suspense here, and reveal from the outset that I’ll try to show that it is not possible to account for these intuitions and desiderata by rejecting either (4) or (3a) on its own. The better account, I will argue, rejects (3a) *by rejecting* (1). I want to argue that “good for” is

not a genuine evaluative notion; rather, it captures how the evaluative landscape *appears to be* to a certain agent.

3. Reductive Approaches

(A) Reduction to Good For

There are at least two kinds of approaches that are popular in the literature that can be broadly described as reduction of the “good” to “good for,” a Humean and an Aristotelian one.²² Let us start with the Humean approach. According to the Humean, our desires, projects, and, more **(p. 212)** generally, our “subjective motivational set” determine what is good for us, and we have only reason to pursue what is good for us. This easily explains what goes on in *Inheritance*. The outcome is obviously good for Paris, but not for her young rich relative. The Humean also allows, of course, that one’s desires or projects might involve commitments to a broadly impersonal pursuit such as morality. One can at first think of morality as a system that prescribes or gives some kind of positive valence that we can term “morally good.”²³ On its own, “morally good” is a purely descriptive term, and says nothing beyond the fact that a certain action is recommended by a certain system of rules. However, insofar as morality becomes an agent’s project, or is somehow incorporated into the agent’s motivational set, then such an agent has reason to pursue what morality recommends.²⁴ We can then describe scenarios (a) through (d) as cases in which what morality recommends conflicts with what would satisfy other members of the agent’s motivational set. The conflict between “good” and “good for,” on this view, is a conflict between different ends that an agent has.²⁵

Humean theories of morality have been widely criticized, but my concern is much narrower. Some of the well-known problems with Humean theories affect their plausibility of providing a good understanding of the relation between “good” and “good for.” So, for instance, many philosophers complain that Humeans make the reasons to be moral desire-dependent, and that this is a very counterintuitive understanding of the nature of moral reasons. So this view cannot capture the desire-independent nature of what is good. The Humean seems to thrive in explaining why Matty’s life is good for her; after all, she got all that she wanted, she succeeded in all her projects, and no lack of information seems to have misled her in her understanding of what was good for her. What many ethicists think is rendered incomprehensible by the Humean view is how there was any kind of normative demand that would have required Matty to have refrained from behaving the way she did.²⁶

However, even if we ignore these problems it is unclear that this way of conceiving the relation between *good* and *good for* will preserve the intuitions brought forth in (a) through (d). All that we need to see to make this point is that the theory cannot treat the difference between moral reasons and nonmoral reasons that are reasons for the agent, as anything but the difference between *prima facie* and all-out reasons. At best, the Humean can say that Teresa had to forego something that she had some reason to pursue in light of her stronger commitment to morality. But compare Teresa’s situation with the situation of someone who realizes that he has better reasons on this occasion to pursue his own good than the good of others. Suppose, for instance, that Milton, unlike Teresa, comes to the conclusion that he should not suffer great pain to save children who are undergoing some kind of deprivation. It’s not that Milton has no kind of commitment to morality; he is just like most of us. He is unwilling to sacrifice as much as Teresa

for the sake of unknown children. Assuming **(p.213)** that Teresa is not self-deceived, making any mistakes of deliberation, etc., the Humean has to regard Teresa's and Milton's situation as exactly the same with regard to the pursuit of their own good. Teresa is pursuing her own good to the same extent as Milton; it is just that the content of their own good is slightly different. And both are doing just as well for themselves. It is not just that the Humean view cannot explain why sacrifices are demanded from people who do not care enough about others to make them; the Humean view also fails to explain why it still counts as making a sacrifice when people *do* care. The Humean account makes Teresa's life, full of pain and deprivation, a life in which things are going well for her, and renders Teresa's friends' bizarre remarks mostly appropriate.²⁷

The Aristotelian view does not make one's reasons dependent on one's desires but rather on a conception of human flourishing. Now an Aristotelian view could claim that anyone's flourishing is equally important from anyone's point of view, but this would not be a reduction of "good" to "good for";²⁸ if anything, it would be a reduction that would go the other way around. The Aristotelian view I would like to consider takes it to be the case that each person aims, or should aim, at his or her own flourishing.²⁹ Now this general Aristotelian view is compatible with many conceptions of flourishing and many views about the relation between one's flourishing and actions that we ordinarily classify as morally good. In particular, we can distinguish between two kinds of view; one that allows that pursuing what is morally good would undermine one's flourishing at least in certain occasions, and a conception that claims that the pursuit of one's flourishing is never incompatible with doing what is morally required, or what is morally best.

I'll first consider views of the latter kind, which I consider the most successful of all the views we'll discard. So, according to this sort of view, a human being cannot flourish unless she leads a virtuous life. However, I take it that under any plausible version of this view being virtuous does not suffice for flourishing; in some cases, the absence of what Aristotle calls "external goods" will prevent the virtuous agent from flourishing. So the Aristotelian view can do well in explaining Teresa's predicament: Being virtuous has required that she deprive herself from external goods and thus from the possibility of having a truly flourishing life. Thus her life, despite being virtuous, is not going well. The Aristotelian can say this, without having to concede that Teresa chose the wrong thing. For things might have been even worse, with regard to her flourishing, had she not helped the children. The Aristotelian view can similarly explain our intuitions in *Inheritance*. For the external goods transferred in this manner from the young relative to Paris now contribute to *her* flourishing.

I must say that even at this point the Aristotelian account already encounters some difficulties. Even if the Aristotelian understanding of *Inheritance* is fully adequate, it's not clear that this understanding of Teresa's predicament is so compelling. After all, if the Aristotelian view wants to claim that Teresa did what she had most reason to do (or that she chose **(p.214)** her good, or some similar claim), he must say that Teresa did relatively well in terms of flourishing, better than if she had just ignored the plight of the children and returned home to safety, or just led a life that most of us lead. The Aristotelian can say that it would have been even better if Teresa had ended up not suffering pain and deprivation, but cannot, at this point, recover the intuition, that had Teresa chose not to undertake such enormous sacrifices, she would have chosen what would be better for her.³⁰

The Aristotelian view seems also to have difficulties in accounting for *Wish and Duty*. It seems that either having dinner with his grandmother contributes to Marshall's flourishing, in which case, he should make the promise and wish she could come, or it does not contribute to his flourishing, in which case, Marshall should not invite her. Aristotelians have tried to explain similar phenomena by trying to make an argument roughly along the following lines.³¹ Certain dispositions, such as loyalty, are essential for the agent's flourishing. But having this disposition requires that the agent act in accordance with the disposition even when the disposition leads us to perform actions that do not contribute to our flourishing. One cannot flourish unless one is a loyal friend, but once one is a loyal friend, one will be ready to sacrifice oneself even when such a sacrifice would curtail one's flourishing. And perhaps the same strategy will help us account for Teresa's plight; her commitment to the cause of humanity might have been essential for her flourishing, but now living up to this commitment will be harmful to her flourishing.

I cannot give a full examination of the plausibility of this view. I will here mostly register my reasons for being skeptical that this move will help the Aristotelian much. In particular, I can't see how this move can be made without making the Aristotelian view into what Parfit calls a "self-effacing" view. The loyal friend cannot recognize both that being loyal to one's friend on this occasion is in all things considered detrimental to her flourishing, and that her flourishing is her "ultimate end," and yet reasonably conclude that she should be loyal to her friend in this occasion.³² But to explain our intuitions in terms of a self-effacing theory is basically to postulate an error theory. It's hard to see how it is true on this view that Teresa ought to help the children (or that Teresa should judge that helping the children is overall good on this occasion), or that Marshall ought to make this promise to his grandmother. These seem to be just beneficial illusions; one would need to deny that these very plausible claims are, strictly speaking, true.

(B) Good For to Good

What about trying the other way around? What about saying that there is only "good" and that "good for" should be reduced to good? The central idea of this view is that various things are good, and that all agents have reasons to pursue these goods, but that "good for" should be understood **(p.215)** in terms of this more general notion of "good." Now there are various ways one could try to reduce "good for" to "good." As we'll see, the view I favor can also be classified as a reductive view of this sort. But in this section, I'll be concerned mostly with attempts to reduce "good for" to "good" in which "good for" is a kind of good. In other words, the reduction that I now have in mind always takes the form:

(Red) X is good for A if and only if X is good and *p*.

or

(Red*) X is good for A if and only if X is good and X is F.

I'll only look at one possible substitute for the second conjunct proposed originally by Moore, and more recently by Regan.³³ According to this view, "X is good for A" should be understood as "X is good and X occurs in the life of A"; I'll call this view the "Moorean view."

The Moorean view can only be plausible if combined with a view that all the relevant goods are good experiences. Otherwise, it's hard to see how it could match what we ordinarily take to be

good for someone. Let us assume that instances of beautiful singing, rather than experiences of listening to beautiful singing, are themselves good. I have a beautiful voice, but I am also deaf and it hurts my throat when I try to sing. Now, if for any reason, I end up singing, the instance of good singing would occur in my life, but it would be rather counterintuitive to think that my singing is something that is good for me. However, even if one thought that other things were good, one could still claim that only experiential goods that occur within the agent's life constitute what is good for the agent. There is a lot that this view can account for at a certain level of explanation. We can say that what's wrong with Matty, from the point of view of an impartial spectator, is that she gets undeserved goods in her life, and undeserved goods call for resentment rather than pity. We can say that Teresa sacrifices herself because she brings about the most good by bringing a lot of pain to her life. The money will result in more goods occurring in Paris's life, and Marshall's life will be spared the pain of boredom if his grandmother releases him from his obligations.

But there is immediately something dissatisfying about this account. After all, given that all we add by saying that the good occurs in the life of the agent it is that we give it a certain location in the agent's mental life, the fact that something is good for the agent makes no extra normative claim on the agent. So there is no reason why, on this view, it should matter more to the agent that a good occurs in his life than that it occurs in the southwest corner of San Antonio. Now the Moorean might object that whether or not it *should* matter to the agent, it *does* matter to the agent. Even if "good for" makes no *normative* claim on us, the fact that X is good for A should lead us to expect that A cares for it. Although this move may explain, for instance, why Paris's friend might be happy for her, it still leaves quite a bit unexplained. For instance, why, on this view, are (p.216) Teresa's actions a sacrifice? It is not true that Teresa cares more for the goods in her life than for the goods in the lives of the children; the fact that she chose to sacrifice herself for their good is clear evidence that she cares more for their good. But in this case there is nothing that the Moorean can offer to restore our sense that things are not going well for her in order to explain the impression that Karol's remarks are wholly inappropriate.

Of course, the Moorean can say that as much as Teresa is interested in sparing the children from pain, she also cares about the good in her life. But this, first, does not explain why the fact that she is in pain should be any different for her than the fact that many other children cannot be saved by her actions (especially given that, again, Teresa is the sort of person who obviously does care for other children). Moreover, not all can be perfect. Although her actions brought about some bad things, they obviously brought about much more good. Karol's remark that Teresa's life is great might be off the mark on this view, but overall when someone thinks that things are going badly for Teresa they cannot be using any overall evaluative claim. There is no reason for Teresa to, say, turn to God and complain "Why me?" Overall, she must think that things have turned out really well from the point of view of all that matters and should matter to her, and thus from all the relevant points of view.

It is similarly difficult for the Moorean to explain Marshall's position. After all, Marshall does care enough about doing nice things for his grandmother that this is what he'll do if she doesn't stop him. In assuming that it is good that Marshall entertains his grandmother this way, we also assume that this is what Marshall should care about overall. So why does it make sense for him

to wish that she would turn down the invitation? Again the Moorean, one can insist that although Marshall *has no reason* to wish for his grandmother to turn down the invitation, he does so wish. But here the Moorean appeals to an irrational and unexplained psychological fact to account for something that seems quite reasonable. Perhaps Marshall could be a more loving grandson, such that he would fully enjoy cooking for his grandmother and entertaining her. It might even be better if he were a more loving grandson. But given that he isn't, his attitudes seem perfectly reasonable.

I find the Moorean view most problematic when we try to understand questions of desert. We want to say that Matty doesn't deserve to end up doing so well; if anyone deserves this kind of life it is Teresa, not Matty. One way to make sense of the idea that she does not deserve her life is to say that things should not be so good for someone who is so evil. You only deserve things to go well for you if you are a good person. But if "good for" just marks the location at which a certain good occurs, why should it matter whether it occurs in the life of someone who did evil or if it occurs in the life of someone who did good? If, as Regan suggests, we should "just promote all agents' valuable activity, without worrying about anyone's well-being," why should it matter that valuable activity happened to be in **(p.217)** the life of an evil person? Why should it be any different from promoting valuable activities on Sunset Boulevard, given that Sunset Boulevard was the location of so much that was evil? To be sure, the Moorean can say that it is just a fact that there is value in good things occurring in the lives of good people and disvalue in them occurring in the lives of bad people. I have nothing in general against having such primitives in one's theory. However, it is a strike against one's view when one must make into a primitive something that seemed to have an obvious explanation.

We can put matters in an admittedly unfair and oversimplifying, but I think helpful, manner. Reducing the *good* to the *good for* makes moral actions a matter of self-indulgence, and thus fails to leave room for understanding what one gives up when one makes a sacrifice, as well as the possibility of normative *demands* to make genuine sacrifices. Reducing the *good for* to the *good* alienates the agent from his life in such a way that one cannot understand how there could even be any such thing as genuine sacrifice and genuine rewards, let alone genuine demands to make sacrifices or to offer rewards to those (and only those) who deserve them.

4. The Appearance View

I want to suggest a different way of understanding the distinction between *good* and *good for*. It is a reductive account in the sense that we end up with only one kind of evaluative dimension. However, it is not an account that reduces "*good*" to "*good for* and *p*" or vice versa. In fact, I think that the problem with the approaches surveyed is that it took each case in which we want to say that "X is good" or "X is good for A" as a genuine instance of value, and the question then was whether one kind of value could be reduced to the other, or if they were two disparate, incommensurable kinds of values.³⁴ Rather, my view is that the difference is not between kinds of value, but a difference in perspectives. "Good for" marks the things that will seem good from the perspective of the agent, and "good" marks what is, in fact, good.

In order to lay out this proposal more clearly, it is worth starting with a traditional view about the relation between desiring and the good that I call "the scholastic view." According to the scholastic view, to desire X is to conceive X to be good, and to be averse to X is to conceive X to

be bad. Now many philosophers think that not all desires are for the good,³⁵ and, although this is less discussed, I take it that the same philosophers would think that not all aversions are for the bad. I do think that the scholastic view is true in full generality,³⁶ but my argument will not depend on this strong claim. All that we need is that in many cases desiring goes together (causally or conceptually) with conceiving something as good, and that in many cases being averse goes together (causally or conceptually) with conceiving something as bad. But perhaps more important to the view I will defend is that “conceiving to be good” or “conceiving to be bad” does **(p.218)** not imply “judging it to be good” or “judging it to be bad.” “Conceiving” is to be understood in terms of appearance: something that inclines us or tempts us to judge in a certain manner, something that is a *prima facie* (though not necessarily *pro tanto*) reason to judge, but that we sometimes can recognize as being illusory.

We can distinguish three kinds of illusory appearances. Some kind of illusory appearances just go away once we realize they’re just an illusion. Suppose as I am hiking I stop and notice something that seems to me to be a sleeping wild animal. I focus my attention and I realize that it’s just a rock formation. Once I can see what is in front of me as a rock formation, it is likely that it no longer appears to me to be a sleeping animal. I’ll call such illusions non-recalcitrant. Not all perceptual illusions are non-recalcitrant. The well-known Müller-Lyer illusion is not like that. If I see lines drawn in this manner, they’ll continue to appear to me to be of different sizes even when I know that they are not. These kinds of illusion are recalcitrant. However, even though the illusion is recalcitrant, it often does not affect belief formation. Even if the lines still appear to be of different sizes, I have no problem in these situations sticking to the belief that I form once I measure the lines with a ruler; my knowledge is after this point completely stable. I’ll call this kind of illusion a benign recalcitrant illusion.

One might think that in the theoretical realm at least, all recalcitrant illusions are benign. However, it is not so clear that this is true. Take, for instance, a non-perceptual appearance. Many people are susceptible to the illusion that some kinds of “jinxing” are possible. So one might think that boasting about how one will win the next race will have some influence on the outcome of the race, or that being overconfident about one’s chances that one’s poem will be selected by the prize committee will diminish the probability of this happening. Often one is fully convinced that this is just a superstition, but one still might find it hard to shake off its influence in belief formation. I have no doubts that jinxing is not causally efficacious, and yet, when someone tells me about how my favorite team is likely to win the World Cup, I still feel compelled to issue cautionary remarks, and point out things that may prevent this desirable outcome.³⁷ Illusions that still influence one’s beliefs even when we know that they are illusions can be called “malignant recalcitrant illusions.” Now if we accept that desires and aversions are at least sometimes appearances in the practical realm, we can think that they also may constitute illusions of all three kinds. My general view is that something counts as good for the agent or bad for the agent if it appears to the agent that a certain object is good or bad in a certain stable manner. One kind of stability is the kind that is generated by practical knowledge; in this case, X appears good to the agent, because the agent correctly judges it to be good in a stable manner.³⁸ Cases of practical knowledge of this kind are cases in which what is good and what is good for the agent do not come apart.

(p.219) But another kind of stability is generated by the existence of a recalcitrant illusion, especially the existence of a malignant recalcitrant illusion. These are cases in which what appears good to the agent is not good, or cases in which what appears to be bad for the agent is not bad, but the illusion does not go away even if one knows that the desires and aversions in question are not accurate conceptions of value. These are cases of divergence between what is good and what is good for the agent. The divergence is not between two kinds of values, but between an objective and a subjective perspective on value.

Let us start by looking at a simple comparison. Let us look at two cases in which I think that something bad must happen for a greater good.

(i) *Staying Fit*

Paula has made a New Year's resolution to run every day. Paula thinks that staying fit is not only instrumentally good but also intrinsically good. However, Paula now realizes that she has an unexpected opportunity to climb the corporate ladder. But to take advantage of this opportunity, she needs to work long hours, and given her other commitments, she won't be able to keep her running schedule. Paula is very ambitious, and she knows that this is a unique opportunity. Given these considerations, she decides it's better to give up on her New Year's resolution.

(ii) *Anesthesia*

Jerzy needs to have a certain dental procedure, and he lives in a country where one has to pay a very high price for anesthesia. Jerzy comes to the conclusion that the best thing is to endure the pain so that he can use the money to help his ailing wife.

When Paula wakes up in the morning to go to her job, she might regret, to some extent, not being able to run to stay fit. She recognizes this as a valuable activity. But it would certainly not be far-fetched to think that this regret in no way threatens the stability of her judgment that, all things considered, or overall,³⁹ it's better to advance in her job. There's no serious temptation for her of going running, or any real chance that Paula will wake up one day, and just go running instead of showing up for her job.⁴⁰ We can say in this case that there is no recalcitrant illusion, or at least no malignant illusion, associated with Paula's desire to run every morning. The same is not true in Jerzy's case. Certainly the pain gives rise to temptation, and it's hard to rule out the possibility that if he can change his mind during the procedure, he would. Even if Jerzy is fully convinced that the money is better spent in helping his wife, it'll take a great deal of effort for him to stick to his decision because *it constantly appears to him that suffering the pain is worse*—or so I'll contend. It is common to think that a victim of torture breaks down by saying, or thinking, something **(p.220)** like: "I'll do anything if the pain will just stop." My contention is that the victim of torture is at this point under an inescapable illusion (or is nearly so) that nothing is worse than the pain.

But what is true in the torture case is true more generally; when we are focused on pain, especially intense pain, other evils will pale in comparison, and the good of pain relief will appear to be particularly great. Now even if the torture victim does not confess, it is still true that given the nature of pain, it'll keep appearing to him that very little is worse than the pain, even if he can keep coming back to his view that it is better not to cooperate with the torturer.

The fate of the person being tortured is similar to someone who knows she is under massive perceptual illusions. As, say, illusions of her friends coming in and saying hurtful things to her would assail her, she would have to make an immense effort not to be taken by the illusion. She might, by dint of this effort, never act badly, or feel hurt by what her illusory friends say, but it would constantly appear to her that they were in front of her, and it would be plausible to describe her even as momentarily forming false beliefs that she would keep correcting. Similar things, *mutatis mutandis*, can be said about pleasure.

If one accepts this view about how pain affects our evaluative point of view, we can put forward the following hypothesis, which I'll call "the appearance view" with regard to what makes the case that X is good for A:

(GOOD FOR) X is good for A if and only if X contributes in a persistent way to making it appear to A that the actual state of affairs is better (more good) than it would appear to A if X did not obtain.

(BAD FOR) X is bad for A if and only if X contributes in a persistent way to making it appear to A that the actual state of affairs is worse than it would appear to A if X did not obtain.

There are obvious questions about how to understand the counterfactuals in these definitions, but I will assume that just as we have an intuitive understanding about whether it would have been better or worse if certain events had or had not occurred, we can also have an intuitive sense of whether it would make things appear better or worse for the agent if certain things had or had not happened.⁴¹ And as we'll see, we might want to modify each definition so that we restrict the range of appearances to genuinely practical appearances.

It seems that the appearance view has no trouble explaining the fact that the notions of good and good for are related. After all, one is just a subjective conception of the other. However, although this shows that the notions are related, it does not seem to show that they are related in the right way. After all, this does not explain how what is good for each agent is a *constituent of* what is good. Any reasonable theory of the good *sans phrase* would take what is good for an agent to be of tremendous importance; for some theories, it is all that could possibly be good. It seems that, if anything, the fact that something merely appears to someone to be good would be of no importance whatsoever.

(p.221) But this is incorrect. First, and most obviously, sometimes things appear just as they are, so the fact that something appears to be good is obviously not incompatible with its being good. Of course, this is unlikely to lay to rest those who think that what is good for the agent is a constituent of what is good, but it is important to notice that this account does not claim that "good for" refers to a set of nonveridical illusions in the life of an agent. However, the appearance view also does not preclude that what is good for an agent is a constituent of what is good. Let us call a state in which it appears to the subject of the state that **p** an "appearance state"; we will be mostly concerned with states in which it appears to the agent that something is good or bad, or as I will call them, "evaluative appearance states." Now describing a state as an appearance state does not necessarily describe all that it is relevant to evaluating this state. For instance, a state in which it appears to me that I am going to fall to my death could also be a

thrilling experience. Pain, if I am correct, is a state in which it appears to the subject of pain that overall the state of affairs is bad, and it can also constitute a malignant recalcitrant illusion; even when one judges that it is worth undergoing some pain for a greater good, the pain might threaten the stability of the overall judgment that one ought to endure the pain. But, of course, one can accept this understanding of pain and yet think that the existence of pain typically makes a state of affairs worse than it would otherwise be. In particular, the following popular theories are immediately committed to the relevance of states in which things appear for the agent in a certain way:

(i) Endorsement Theories

Theories that claim that it is a necessary condition of certain states being valuable that a certain agent judges it to be valuable will be theories that accept that appearance states are, if not valuable in themselves, at least conditions of value. The same goes for a theory that takes an appearance state to be a constituent, or a condition of something being a constituent of the good if it accepts any of the following: The value of certain states of affairs depends on whether an agent approves of it; or the value of the state of affairs depends on the agent endorsing its desirability; or the value depends on the agent having the appropriate desire for the state of affairs; or an end is only valuable or worthy of being pursued if it is reflectively endorsed, or properly incorporated into a maxim, etc. After all, approving something involves judging it to be good, and thus, approving something is an evaluative appearance state. The same goes for endorsing. Not everyone accepts that every case of desiring is one in which something appears to be good for the agent; however, it is hard to deny that many of them are, and, in particular, it is plausible to assume that only desires in which their objects appear to be good in some way to the agent can be determinants of whether something is or is not a constituent of the good.

(p.222) (ii) Theories in Which Pleasure and the Absence of Pain Are Important Goods

Given the above understanding of pain, and a similar understanding of pleasure, both pains and pleasures are evaluative appearance states. So any view that accepts that these are important constituents of the good will accept that evaluative appearance states are important constituents of the good.

(iii) Kantian Theory

Here is a coarse sketch of such a view: Suppose one thinks that the objects of one's inclinations are putatively good, but in order to be, in fact, good, it is necessary and sufficient that they also conform to certain conditions. Such a view would also take evaluative appearance states to be constitutive of the good, albeit in a more complex way.

It is worth mentioning that in not all of these cases does it turn out that something is good *because* it is good for the agent. Especially in (ii), it would be more precise to say that what makes something good is also what makes it good for the agent. However, this does not seem to be a problematic implication; our intuitions do not so clearly favor one option over the other. But it is also important to note that the fact that the appearance view is compatible with all these views enables it to have, or at least to borrow, an advantage over the Moorean view. The Moorean view has troubles explaining why so much of what is good is also good for someone; the appearance view, on the other hand, can simply piggyback on whatever explanations such other theories offer for their account of the good.

(A) Teresa and Matty Reconsidered

The best way to test further the plausibility of the appearance view is to see how it deals with each example. Let us start with Teresa. Teresa does bring about much good to the world. But she does by causing herself to suffer a great deal of pain. The consequence of this fact is that she can do it only by having this state of affairs constantly appearing to her to be bad. Things are not great for her in the sense that no matter how convinced she is that she did the best possible thing, it keeps appearing to her that what she did was bad. One might compare Teresa's predicament with, say, a mother who can only save her son by then suffering from constant illusions that he is dead. No matter how much the mother cares for her son, this is obviously not an enviable position, since it keeps *appearing* to the mother that the son is dead. Given that her pain constitutes a malignant recalcitrant illusion, from Teresa's point of view things are not as good as one might have thought they would be simply because from her point of view things keep appearing to be bad.

(p.223) It is true that on this view, the fact that Teresa judges that her choice was correct will make it the case that things are better for her than if she were a selfish person whose same suffering accidentally helped the children in the same way. After all, the fact that she judges it to be good must contribute to things not appearing as bad to her, or at least not in the same way, as if she did not think it was good to help the children. But I do not find this consequence counterintuitive; the fact that Teresa accomplished something that she finds important must indeed contribute to things not being as bad for her as they would otherwise be.

Similar things can be said about Matty. Given Matty's evaluative views, she is now in the position that things appear good to her; she is in a special kind of fool's paradise. The state of affairs is one that stably appears to Matty to be good. As for Ned, who was framed for the murder, the state of affairs will appear to him even worse than it actually is. It is worth considering in more detail Matty's situation. Our attitudes to someone in more traditional versions of fool's paradise tend to be ambiguous. Whether it is better to know or not that one's spouse has had an affair is a primary example of an unsettled bar controversy. On the other hand, almost no one doubts that it would be particularly cruel to point out to a dying mathematician a major fallacy in the "proof" she thought had been the greatest achievement of her career. We might pity the betrayed spouse or the dying mathematician, but we certainly pity them mostly for having a disloyal spouse and for not having succeeded in finding the proof. It is difficult in the case of the betrayed spouse to say whether we should pity him more or less for not knowing that the affair happened. But in the case of the dying mathematician we would pity her even more if she were to learn that her proof was fallacious. At any rate, we must note that the kind of mistake that leaves you in a fool's paradise is not always necessarily the subject of pity. And although we do not want to exchange our position with the person who is in a fool's paradise (after all, we know that that state of affairs would be, in fact, worse), there is something obviously enviable about a fool's paradise. And this is similar to our attitude toward Matty. Although there is something clearly enviable about Matty's situation, the morally good agent would not want to be in her place. But it seems hard to accept that there is no significant difference between Matty's case and the case of the betrayed spouse or the dying mathematician, and it would be hard to accept that our attitudes need to be the same in all these cases.

The source of the difference must be that Matty's illusions are not theoretical but evaluative. Matty's views of the nonevaluative world are perfectly accurate. The illusion in Matty's case (just as in Teresa's case for this matter) lies in her evaluative attitudes, not in her theoretical attitudes. And since practical mistakes are more blameworthy than theoretical mistakes,⁴² it might be inappropriate to feel pity for Matty even if it turned out to be true that it was appropriate to feel pity for the betrayed spouse or the dying mathematician solely on account of their mistake.

(p.224) But Matty's case might also be different from the betrayed spouse and the dying mathematician for more profound reasons. For we can think that evaluative appearances and theoretical appearances make different contributions to what is good for an agent. So far we have presented the appearance view as not discriminating between things appearing good or bad in virtue of the evaluations of the agent and things appearing to be good and bad in virtue of the agent's beliefs. But an appearance view that makes *good for* and *bad for* only relative to evaluative appearances will be able to distinguish between Matty's case and the case of any fool in more traditional versions of a fool's paradise.

One could argue that what is good for an agent is determined solely by evaluative appearances; we judge how good or bad various states of affairs are for the agent by examining how good or bad they would appear to the agent if she were to know how the world actually is. We could justify such a view by arguing that what we're interested in is how the agent responds to the actual facts; how what's the case would affect how things appear to him as good or bad. Such a view, however, could not do justice to the case of the dying mathematician, and it would judge that it is indifferent to someone's good whether they know or fail to know that, for instance, their beloved daughter is still alive. More plausibly, one could argue that evaluative appearances and theoretical appearances contribute in different ways to what counts as good for the agent; but of course the plausibility of such a view depends on whether we can explain more precisely how each appearance contributes to an agent's good.

We can think about our judgment of how things are going for the agent as an imaginative exercise. As we imagine how things are going for her we form two sets of appearances: how good or bad things appear to the agent, and how good or bad things would appear to the agent if she were to learn what we know. Suppose someone erroneously thinks that he has won the Nobel Prize. If we imagine the truth revealed to the false Nobel Prize winner, he'll be no doubt saddened by it. But if he's like most of us, he'll also not just think "Oh, at least it was fun while it lasted." The revelation of the truth would "spoil" how he perceives the past; the celebrations will now appear completely inappropriate and silly, the gloating shameful, and the sense of superiority misplaced. The memories of this time are more likely to be experienced with shame than with joy. It will be very different from, say, someone who had an extremely fun birthday party on the wrong date. Finding out the mistake later would probably do very little to retrospectively spoil the fun.⁴³ But things are even more complicated if we think about a father who erroneously think that his daughter is dead; this situation is not parallel to the case of our fake Nobel Prize winner. If the father were to find out that his daughter was alive, he would not look back at the way he felt as inappropriate or would feel any kind of retrospective joy for it; grief and sadness are appropriate not only when one's loved ones are in fact in distress but also if we believe that they are in this situation, or even when we just suspect that they are.

(p.225) A more satisfactory view would say the following: We count not only how good or bad a state of affairs appears to the agent now but also how it would appear if the truth were to be revealed to the agent, including how the past would retrospectively appear to the agent once the truth was revealed. Although this is the view I favor, I must confess our conception of what's good for someone becomes somewhat messy and ambivalent at this point. We, or many of us, seem sometimes tempted by the thought that ignorance is bliss, especially in situations in which we want to spare someone the truth. A fully satisfactory account of "good for" would try to accommodate or explain away these intuitions more precisely. But for our purposes what matters is that any such refinement would be a matter of refining how different appearances make a contribution to what is good for the agent; but my aim here is just to argue for the claim that the referents of "good for" and "good" stand to each other as appearance to reality.

Before we move on, it should be noted that the appearance view gives a more satisfying explanation of moral desert than the Moorean view does. It seems quite compelling to think that it is unfair that the price of bringing about some good is that things will appear bad to you. And it seems in general unfair that things will appear bad exactly for someone who brought about much good. On the hand that everything will appear to be good from the point of view of someone who brought about so much that is bad seems equally unfair. Of course, the latter is true only if one was culpable in bringing about bad things. And one might think that the appearance view makes Matty innocent, since it makes Matty just mistaken about an evaluative fact. It is worth noting that this kind of ignorance is what Aristotle calls "ignorance of the universal." And although I cannot argue here the plausibility of the view that ignorance of the universal is blameworthy, it is at least a respectable view in the thorny landscape of philosophical views about moral responsibility.

(B) Marshall and Paris Reconsidered

Obviously Paris's receiving the inheritance will bring many changes to Paris's evaluative appearance states. Not only will it change much that is recalcitrant in making things appear good or bad for Paris, but also Paris, with the money, can promote what *she* takes to be good. Now Marshall's case is a bit more complicated. Marshall recognizes that it is best to invite his grandmother, but he faces the prospect of boredom and painful cooking, which will recalcitrantly make things appear bad to him (or it will at least make having his grandmother over for dinner appear worse than not having her over). But it also seems safe to assume that the prospect of pain is itself something that is inherently unpleasant; that is, the prospect of pain is not just a representation of a future state in which it will appear to the agent that things are bad, but a representation of a future state that also *now* appears to be bad—in many cases, a representation of a future **(p.226)** state that appears even worse than it actually is. But if this is the case, we can understand that Marshall's wish is the offspring of this recalcitrant illusion; from his point of view, he cannot but experience the prospect of his grandmother not coming with relief, since given the nature of his evaluative appearance states, it can't but stop appearing to him that it would be better if it were not to happen.

5. Conclusion: Some Complications

I have said at the beginning that I am here only sketching a view that needs to be further developed to be fully persuasive. I want to briefly point out a couple of important complications

that a more detailed treatment of the issue would have to address; I do not claim that these are the only ones.

As many philosophers have noted, the relationship between my good and the good of others is similar to the relationship between my immediate good and my future good. My immediate good can be the source of temptation; when I go to the dentist I experience it as a sacrifice for my later good. And yet, here the Moorean view might seem perfectly adequate: Something is an immediate good if and only if it is a good and it would be brought about in the immediate future. Is this a problem for my view? I think that, if anything, the opposite is true. Even if we grant this modified “Moorean” that this is a correct description of the immediate good, it does not explain why we are tempted to pursue what is immediately good for us even when a greater good lies ahead in the future. In order to explain this fact, we need to invoke a similar structure of recalcitrant appearances; the immediate good is the source of temptation exactly because it keeps appearing as better to us than what lies ahead. In fact, George Ainslie’s groundbreaking work on hyperbolic discounting of future goods can be understood as explaining the nature of the illusion in question, an illusion structurally similar to familiar perceptual illusions.⁴⁴ Of course, we would still need to sort out how different appearances contribute only to what counts as what is immediately good and to what counts as good for the agent in a more general manner, and, again, this issue deserves more attention than I can spare here. But the basic idea is that we would want to exclude from our more general notion of the *good for* any recalcitrant appearances that simply reflect a temporal bias.

We might also suspect that the distinction between *good* and *good for* obscures the different ways in which we respond to the goods of different people. Let us take one aspect of how we relate differently to what is good and what is good for someone, an aspect that I have emphasized at different points here: We experience the promotion of the good as a sacrifice and the promotion of what is good for us as a temptation. We can now ask how the good of a loved one figures in our experience in relation to what is good and what is good for us. No matter how much I love (p.227) someone, there will be (at least possibly) cases in which I ought to forego promoting my own good so as to promote his good. Moreover there’ll also be cases in which I ought to forego promoting the good of my loved ones in order to promote the good. When we look at these cases, we do find that we can easily assimilate my concern for the good of my loved ones to the concern I have for my own good or for the good in general. A dramatic case of the former would be a father who is being tortured by terrorists who want to know the whereabouts of his daughter. I take it that in this case we would think that the father must be *tempted* to give in, and if he doesn’t give in, we’ll admire the kind of sacrifice he performed for his daughter’s sake. So this would lead us to assimilate the good of someone’s loved ones to something that contributes to what is good, but not to what is good for him. On the other hand, suppose Francis can somehow illicitly influence a job search for a job that his daughter Sofia very much wants, a job search that would otherwise certainly favor Martin, who is much more qualified for the job. In both cases, we can assume Francis knows that what is good, and what he ought to do overall, is not to help his daughter. In these cases, it seems that for Francis the relation between Sofia’s good and the good in general is similar to the relation between his own good and the good. Choosing rightly will be something that he’ll experience as a sacrifice, and he’ll be tempted to pursue Sofia’s good, rather than what is simply good. Here Sofia’s good seems to play the same role in Francis’s life as his own good.⁴⁵ These two cases seem to suggest that the good of our

loved ones does not function exactly as a part of our own good, or as just one kind of a good *sans phrase*, but it seems to somehow squeeze between the two.

Of course, so far all that these two cases seem to show is that sometimes the good of a loved one is a constituent of our own good, and sometimes it isn't, a rather trivial result. However, matters are not so simple, when we look at cases of multiple conflicts. Suppose Frank is a repentant mafioso now trying to be an honest and law-abiding citizen. However, his former buddies think that he owes them a lot of money, much more money than Frank has right now. If Frank does not pay the money to his buddies, they will do horrible things to him. Frank suddenly remembers where he hid some money stolen from Bing that had not been previously found. The money is enough to pay off the Mafia, but could also be used to send his drug-addicted daughter Nancy to a very expensive, but also very effective, rehab facility. Frank judges, let us assume correctly, that what he should do is to return the money to Bing. But when he thinks about his beloved Nancy, and how this money could turn her life around, Frank fears that he'll succumb to the temptation of using the money for this end. Frank knows that if the money was rightfully his, he should spend it this way, even if this would lead him to great suffering. He certainly thinks that if he's going to fail to give the money back to Bing, it would be better to use the money for Nancy's sake than to use it to save his skin. But he is really scared of his old buddies, and he's afraid that if **(p.228)** they'll show up at the door, he'll be too weak to resist the temptation to give them the money.

It seems here that in the *very same scenario*, in relation to Francis's own good, the pursuit of Nancy's good is a sacrifice, and its effective pursuit requires that Francis overcome the temptation of the pursuit of his own good, but, in relation to what is good, the pursuit of Nancy's good is a temptation, a temptation Frank thinks he should resist in order to pursue what he thinks is good. In this case of multiple conflicts, it does not seem that the good for one's loved ones can be assimilated completely to a constituent of either one's own good or the good *sans phrase*. It seems instead to also occupy an intermediary position between the two. And to make things worse, it is not hard to imagine how the good of a good friend will stand as an intermediary between the good of one's daughter and one's own good, and the good of one's spouse, one's close relatives, one's distant friends, etc. might occupy any place in between one's own good and the good. Needless to say, there is no reason to stop here at three distinctive evaluative rankings; we can easily imagine the goods of various friends and relations occupying various intermediate positions.

The problem of how to understand our relation to the good of our loved ones is a difficult one, and insofar as it presents a problem for, or a gap in, my view, it also presents a problem for (or a gap in) all other views. However, I think the view put forward here is in a good position to allow for the existence of intermediaries between what is good and what is good for someone. After all, how recalcitrant an appearance is can be not only a matter of degree but can also vary from situation to situation. Thus the good of my loved one might give rise to recalcitrant appearances only in the presence or absence of other evaluative appearances. These facts about evaluative appearances can at least provide us with the basic materials to carve up an intermediate notion; however, a final assessment of the suitability of these materials will require a more detailed account than I can provide here.

I have claimed that my view is compatible with various substantive accounts of well-being, but the reader might have noted that at least one major account seems to be left aside; namely, objective lists. Although, strictly speaking, it is compatible with my view that the items in the list would be such as to generate recalcitrant appearances, this wouldn't be a particularly plausible position. However, objective list views do not aim to answer the desiderata I presented above, so I think it is fair to suggest that they are not trying to understand the same concept. In my understanding of the conceptual landscape, these are actually theories about what is, in fact, good *simpliciter*, but theories that claim that what is, in fact, good *simpliciter* is agent-relative. Here again, more remains to be said, but I hope to have taken at least one important step in giving an account of the notion of good for that does not leave us with competing candidates to be the formal end of practical reason.

Notes

References

Bibliography references:

Ainslie, G. (2001). *Breakdown of will*. New York: Cambridge University Press.

Cocking, D., & Kennett, J. (2000). Friendship and moral danger. *The Journal of Philosophy*, 97, 278–296.

Davidson, D. (1980). How is Weakness of the Will Possible? In D. Davidson (Ed.), *Essays on actions and events* (pp. 21–42). Oxford: Clarendon Press.

Foot, P. (1985). Utilitarianism and the virtues. *Mind*, 94, 196–209.

Foot, P. (2002a). "Moral Belief" in her *Virtues and vices and other essays in moral philosophy*: Oxford University Press.

Foot, P. (2002b) "Morality as a System of Hypothetical Imperatives" in her *Virtues and Vices and Other Essays in Moral Philosophy* (New York: Oxford University Press) pp. 157–174

Helm, B. (2005). Friendship. *Stanford Encyclopedia of Philosophy*.

Hursthouse, R. (1999). *On virtue ethics*: Oxford University Press.

Kant, I. (1997). *Critique of practical reason* (M. J. Gregor, Trans.). New York: Cambridge University Press.

Korsgaard, C. M. (1997). The Normativity of Instrumental Reason. In G. Cullity & B. Gaut (Eds.), *Ethics and Practical Reason* (pp. 215–254). New York: Clarendon Press.

Kozhevnikov, M., & Hegarty, M. (2001). Impetus beliefs as default heuristics: Dissociation between explicit and implicit knowledge about motion. *Psychonomic Bulletin and Review*, 8(3), 439–453.

- Kraut, *Aristotle on the Human Good*, (Princeton: Princeton University Press, 1989).
- McDowell, J. (1996). Two sorts of naturalism, *Mind, value, and reality* (pp. 167–197). Cambridge, Mass.: Harvard University Press.
- Mill, J. (1979). *Utilitarianism* (G. Sher, Ed.). Indianapolis: Hackett Publishing Co.
- Moore, G. E. (1993). *Principia Ethica (Revised Edition)*. New York: Cambridge Univ Pr.
- Overvold, M. (1980). Self-Interest and Self-Sacrifice. *Canadian Journal of Philosophy*, X(1), 105–118.
- Rawls, J. (1971). *A theory of justice*. Cambridge, Mass.: Belknap Press of Harvard University Press.
- Raz, J. (1986). *The morality of freedom*: Oxford Univ. Press.
- Regan, D. H. (2004). Why Am I My Brother's Keeper? In R. J. Wallace (Ed.), *Reason and Value: Themes from the Moral Philosophy of Joseph Raz* (202–230). Oxford: Clarendon Press.
- Sidgwick, H. (1981). *The Methods of Ethics*. Indianapolis: Hackett.
- Stocker, M. (1979). Desiring the Bad: an Essay in Moral Psychology. *Journal of Philosophy*, 76, 738–753.
- Tenenbaum, S. (2007). *Appearances of the good: an essay on the nature of practical reason*. New York, Cambridge: Cambridge University Press.
- Velleman, J. D. (1992). The Guise of the Good. *Nous*, 26(1), 3–26.
- Williams, B. (1982) "Internal and External Reasons" in his *Moral Luck* (New York: Cambridge University Press).

Notes:

- (1) . John Rawls. 1971. *A Theory of Justice*. Cambridge, Mass.: Belknap Press of Harvard University Press.
- (2) . Some accounts of personal identity (nearly) collapse the distinction between interpersonal and intrapersonal redistribution, and obviously not everyone is convinced by Rawls's view that it is illicit to generalize from the intrapersonal to the interpersonal case as it will be clear below. However, Rawls is clearly tapping on an important intuition about redistribution.
- (3) . See Sergio Tenenbaum. 2007. *Appearances of the Good: An Essay on the Nature of Practical Reason*. Cambridge: Cambridge University Press.
- (4) . For a more detailed explanation of what I mean by the "formal end of practical reason and action," see *Appearances of the Good*. For brevity's sake, from now on, I'll simply refer to the "formal end of practical reason."

- (5) . Immanuel Kant. 1997. *Critique of Practical Reason*. New York: Cambridge University Press.
- (6) . I come back to this point later.
- (7) . “Seems” is appropriate here, since Rawls’s argument is certainly suggestive, but vague enough that one would not want to commit oneself to be advancing the only possible interpretation of the argument.
- (8) . Donald H. Regan. 2004. “Why Am I My Brother’s Keeper?” In Donald Wallace, ed. 2004. *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, xxx-xx. Oxford: Clarendon. Pp. 202-30.
- (9) . The last section discusses some of the complications that need to be addressed.
- (10) . Assuming, of course, there are no further candidates for such a role.
- (11) . Joseph Raz, *Morality of Freedom*, (Oxford: Oxford University Press, 1986), p. 289.
- (12) . *Ibid.*, p. 296.
- (13) . See M. Overvold. 1980. “Self-Interest and Self-Sacrifice,” *Canadian Journal of Philosophy* 10:105-18, on the claim that an appropriate notion of what is good for the agent needs to make sense of cases of sacrifices as cases in which the agent chooses something that is not good for her.
- (14) . For obvious reasons I’ll from now on omit the qualification “as I understand it.”
- (15) . Stephen Darwall defines “well-being” in terms of what someone who cares for the agent would rationally want for the agent for the agent’s own sake. Darwall’s account comes closest in the literature to the one I am putting forward here, but it is not clear what we should say about what we would desire for the agent’s own sake in cases of immoral actions. Darwall explicitly uses cases of sacrifice to argue against certain views of well-being. He rightly points out that many desire accounts are unable to account for the fact that an agent might desire, and even rationally desire, to sacrifice herself for the sake of an other and should not lead us to conclude that the sacrifice is good for the agent. But by the same token, it also would seem wrong to say that sacrificing herself would be good for the agent in cases in which the agent is morally *required* to sacrifice herself. At the same time I find it hard to say that it would be rational to want for someone’s sake that they perform an immoral action; at least, a quite plausible view about the rationality of morality would rule out this possibility. Of course, one could insist that even if it would be *irrational* to want that she does not sacrifice herself in these cases, it is rational to want that she does not sacrifice herself *for her sake*. But I find it hard to make sense of this idea, unless “well-being” is explicating “for her sake” rather than the other way around.
- (16) . At this point I am only claiming that this is at a first glance an intuitive view. Many philosophers have denied that in such situations things are going well for her; Plato, for instance, famously claims that the tyrant harms himself.

- (17) . It is not exactly what he says, but it is obviously implied by J. S. Mill. 1979. *Utilitarianism* (G. Sher, ed.). Indianapolis: Hackett.
- (18) . It is worth noting that I am not saying that the utilitarian cannot add to his view a specific account of the relation between good and good for of the kind that I am suggesting. In fact, the account I propose in section 4 is compatible with utilitarianism.
- (19) . I am ignoring grammatical issues of whether we should say “X is good” or “X would be good” for not-yet-realized states of affairs for the sake of simplicity.
- (20) . I should point out that when considering these views, I am not claiming that they were put forth with the aim of answering the questions I am posing here.
- (21) . Henry Sidgwick. 1981. *The Methods of Ethics*. Indianapolis: Hackett.
- (22) . I am of course not claiming that these views are explicitly offered as attempts to address the problems I have raised.
- (23) . This approach is compatible with thinking that the subject matter of morality cannot be described simply in terms of the pursuit of the greatest good. This is the view defended by Foot, within a broad Humean framework, in “Utilitarianism and the Virtues,” *Mind*, 94, 196–209. However, since this does not make much difference to my argument, I will use it as a simplifying assumption that every action that morality recommends gets classified as “morally good.”
- (24) . For classic formulations of views along these lines, see B. Williams, “Internal and External Reasons” in his *Moral Luck* (New York: Cambridge University Press, 1982) and Philippa Foot, “Morality as a System of Hypothetical Imperatives” in her *Virtues and Vices and Other Essays in Moral Philosophy* (New York: Oxford University Press) pp. 157–74.
- (25) . Or perhaps in case (b) between an end that Matty has and an end that *other* agents have, since it’s not clear that Matty is in any way committed to morality.
- (26) . See, for instance, Christine M. Korsgaard. 1997. “The Normativity of Instrumental Reason.” In Garrett Cullity and Berys Gaut, eds. *Ethics and Practical Reason*. Oxford: Oxford University Press.
- (27) . Overvold makes a similar point. See Overvold. “Self-Interest and Self-Sacrifice.”
- (28) . I don’t really know of anyone who defends exactly such a view. Richard Kraut, however, does think that “in saying that our ultimate end should be happiness, Aristotle must be taken to mean that ultimately we are and should be aiming at *someone’s* happiness, whether our own’ or another’s (Richard Kraut, *Aristotle on the Human Good*, (Princeton: Princeton University Press, 1989), p. 145).
- (29) . See, for instance, Rosalind Hursthouse. 1999. *On Virtue Ethics*. Oxford: Oxford University Press.

- (30) . We'll see below, when discussing *Wish and Duty*, a possible way around the objection for the Aristotelian.
- (31) . See P. Foot. 2002. "Moral Beliefs." In P. Foot. 2002. *Virtues and Vices and Other Essays in Moral Philosophy*. New York: Oxford University Press.
- (32) . McDowell makes a similar point in J. McDowell. 1996. "Two Sorts of Naturalism." In *Mind, Value, and Reality*. Cambridge, Mass.: Harvard University Press.
- (33) . G. E. Moore. 1993. *Principia Ethica*, rev. ed. New York: Cambridge University Press; Regan. "Why Am I My Brother's Keeper?" 202-30.
- (34) . This might seem not quite true for the Humean approach.
- (35) . Michael Stocker. 1979. "Desiring the Bad: An Essay in Moral Psychology," *Journal of Philosophy* 76:738-53; J. David Velleman. 1992. "The Guise of the Good," *Nous* 26:3-26.
- (36) . See Sergio Tenenbaum. *Appearances of the Good*.
- (37) . For fascinating cases of recalcitrant illusions in the realm of theoretical reason regarding our beliefs about motion, see M. Kozhevnikov, and M. Hegarty. 2001. "Impetus Beliefs as Default Heuristics: Dissociation between Explicit and Implicit Knowledge about Motion," *Psychonomic Bulletin and Review* 8:439-53.
- (38) . I am using "judging to be good" and "judging it to be bad" in a way that should be neutral between thinking that the attitude in question is an (evaluative) belief, or thinking that it is a different kind of all-out attitude peculiar to the practical realm. See, on this issue, Tenenbaum. *Appearances of the Good*.
- (39) . Davidson famously distinguishes between "all things considered" and "all-out" judgments in Donald Davidson. 1980. "How Is Weakness of the Will Possible?" In Donald Davidson, ed. 1980. *Essays on Actions and Events*. Oxford: Clarendon.
- (40) . I'll assume that at least some cases of *akrasia* involve a temporary reversal of one's all-out evaluative judgment. Although my view can be presented and argued for independently of this assumption, it does simplify matters a great deal to make it.
- (41) . One must also make sure that one does not run afoul of the conditional fallacy, etc. But I'll not get into these complications here.
- (42) . Of course, in light of difficulties justifying treating these mistakes differently, one might argue that we should hold agents that make these different kinds of mistakes to be equally blameworthy. But this position would also require a revision of our attitudes; all that I argue is that I can explain why these attitudes are warranted, *assuming* that our ordinary views about blameworthiness and related matters are also justified.

(43) . Of course someone could have the same attitude of finding out that he didn't win the Nobel Prize that most people would have if they found out that they celebrated their birthdays on the wrong date. But I think in this case our views about how he's doing when he falsely believes would also change; after all, the Nobel Prize is, for this person, just an excuse to have a party.

(44) . G. Ainslie. 2001. *Breakdown of Will*. Cambridge: Cambridge University Press. Ainslie does not endorse this understanding of hyperbolic discounting as a form of illusion, though he does explicitly compare it with perceptual illusions.

(45) . I take it that some such phenomenon is what leads Cocking and Kennett to suggest that friendships can be morally dangerous. See D. Cocking, and J. Kennett. 2000. "Friendship and Moral Danger," *The Journal of Philosophy* 97: 278-96. Even if they are wrong that a good friend would help her friend to hide a body (see B. Helm. 2009. *Stanford Encyclopedia of Philosophy*. S.v. "Friendship."), a good friend would be *tempted* to help.



Access brought to you by: University of Toronto Libraries