

THE NORMATIVE AUTHORITY OF OUR FULLY INFORMED JUDGMENTS

BY

GERALD DENIS BEAULIEU

**A Thesis
Submitted to the Faculty of Graduate Studies
in Partial Fulfillment of the Requirements
for the Degree of**

MASTER OF ARTS

**Department of Philosophy
University of Manitoba
Winnipeg, Manitoba**

(c) April, 1997



National Library
of Canada

Acquisitions and
Bibliographic Services

395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque nationale
du Canada

Acquisitions et
services bibliographiques

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-23218-2

**THE UNIVERSITY OF MANITOBA
FACULTY OF GRADUATE STUDIES

COPYRIGHT PERMISSION PAGE**

THE NORMATIVE AUTHORITY OF OUR FULLY INFORMED JUDGMENTS

BY

GERALD DENIS BEAULIEU

**A Thesis/Practicum submitted to the Faculty of Graduate Studies of The University
of Manitoba in partial fulfillment of the requirements of the degree
of
MASTER OF ARTS**

Gerald Denis Beaulieu 1997 (c)

Permission has been granted to the Library of The University of Manitoba to lend or sell copies of this thesis/practicum, to the National Library of Canada to microfilm this thesis and to lend or sell copies of the film, and to Dissertations Abstracts International to publish an abstract of this thesis/practicum.

The author reserves other publication rights, and neither this thesis/practicum nor extensive extracts from it may be printed or otherwise reproduced without the author's written permission.

Table of Contents

	Page
Abstract	ii
Acknowledgments	iii
Chapter 1: Introduction	1
Chapter 2: Personal Ideals and Full Information Accounts of Well-Being	11
Chapter 3: The Notion of Being Fully Informed	27
Chapter 4: Conclusion	46
References	51

Abstract

In my thesis, I consider some recent objections to the full information account of the good. I examine a particular line of argument advanced by Connie Rosati. Rosati argues that the judgments of our fully informed selves lack normative authority for us. She claims that the full information account does not allow us to address properly questions about what kind of person to be because the account identifies a person with her current motivational system. We may hold personal ideals that conflict with our current motivational system. But I argue that by treating personal ideals as a part of one's motivational system, the full information account can address properly questions about what kind of person to be. Rosati is also concerned with the changes that an agent would have to undergo in order to become fully informed. We might end up with a different motivational system once fully informed and so our fully informed judgments will lack normative authority for us. I argue that the changes to an agent's motivational system postulated by the full information account are beneficial changes. So I conclude that Rosati has not undermined the normative authority of our fully informed judgments.

Acknowledgments

I would like to thank the members of my thesis committee (Robert Shaver, Joyce Jenkins, and Edward Johnson) for their invaluable help. Special thanks to Michael Feld, Joyce Jenkins, and Robert Shaver for inspiring me in this area of philosophy. And finally, I would like to thank my parents for their support throughout this project.

Chapter 1: Introduction

I am interested in developing a naturalistic account of personal well-being. It is important, I believe, that the account be naturalistic. The normative concept 'good for a person' must be analyzable, without residue, in completely non-normative, or descriptive, terms. If we wish to maintain that value judgments are cognitive (i.e., that they are either true or false as opposed to, say, mere emotional expressions), naturalism seems like the only viable option. Non-naturalistic cognitivism leaves us puzzled about the epistemic status of value judgments and the ontological status of value. If value is, as G.E. Moore in Principia Ethica supposes it is, *sui generis* (metaphysically unique), one wonders how we could come to know about it. Do we have, for example, a special faculty by which we intuit the truth of value judgments? Naturalism, on the other hand, treats value as a natural fact accessible to us in much the same way that scientific facts are.

The problem for naturalists is to account for the normative content of value judgments. Natural facts, like the ones discovered by science, are not intrinsically motivating. Acknowledging the truth of scientific claims leaves open our attitude towards them. The truth of a value judgment, however, seems to guarantee that the person to whom it applies will be (at least in certain circumstances) moved by it. For example, if I say "Chocolate ice cream is good", I mean to recommend it to you. This is quite different from my saying "Chocolate ice cream is brown." Value, or the good, as J.L. Mackie in Ethics: Inventing Right and Wrong points out, has "to-be-pursuedness" built into it. The alleged facts alluded to by value, then, seem awfully queer.

But the account I am interested in developing identifies a person's good with her desires. A person's desires are, at the same time, a natural fact and

intrinsically motivating. There seems to be nothing particularly queer about these facts. We might, however, be suspicious of identifying our good with just any desires of ours. For example, I might have a desire to drink the clear liquid in the glass in front of me. I now believe that drinking the liquid will quench my thirst. But I could be mistaken about the contents of the glass. The clear liquid might actually be gasoline and not water as I had assumed. We must correct for false beliefs, then, if we are to identify our good with our desires. This is essentially the strategy of the full information account of well-being. We must now familiarize ourselves with this account.

Henry Sidgwick, in The Methods of Ethics, provides us with an early version of the full information account when he considers the Hobbesian view that "whatsoever is the object of any man's appetite or desire that is it which he for his part calleth *good*; and the object of his hate and aversion, *evil*." (Hobbes 28) He argues, contra-Hobbes, that we cannot simply mean by "X is good for A", "A desires X". Many things that A actually desires are desired only because A is mistaken about the facts or is using faulty reasoning. In these cases, if A knew all of the facts and was fully rational (i.e. if A was fully informed), A would no longer desire X. The suggestion, then, is that

a man's future good on the whole is what he would now desire and seek on the whole if all the consequences of all the different lines of conduct open to him were accurately foreseen and adequately realised in imagination at the present point of time. (Sidgwick 111-112)

Sidgwick points out that this seems like a plausible candidate for a naturalistic definition of 'good'. He writes that

[t]he notion of 'Good' thus attained has an ideal element: it is something that is not always actually desired and aimed at by human beings: but the ideal element is entirely interpretable in terms of fact, actual or hypothetical, and does not introduce any judgement of value, fundamentally distinct from judgements relating to existence; - still less any 'dictate of Reason.' (Sidgwick 112)

This account also satisfies a strong intuition that many philosophers share regarding a person's good. By linking the person's good to her desires, the account is internalist. In this context, internalism is the view that something is good for an agent only if it can, at least potentially, motivate her.¹ It seems implausible to suppose that something could be good for a person and yet completely fail, in appropriate circumstances, to motivate her. Peter Railton remarks that any account which did not satisfy this internalist constraint "would be an intolerably alienated conception of someone's good." (Railton 1986b: 9)²

But Sidgwick ultimately rejects this account of 'good' (qua naturalistic definition) despite the fact that it gives "philosophical precision to the vaguer meaning with which [it is] used in ordinary discourse." (Sidgwick 112) His reasons for doing so are not clear. Immediately following his favorable presentation of the full information account he states that

[i]t seems to me, however, more in accordance with common sense to recognise - as Butler does - that the calm desire for my 'good on the whole' is *authoritative* ; and therefore carries with it implicitly a rational dictate to aim at this end, if in any case a conflicting desire urges the will in an opposite direction. (Sidgwick 112)

Sidgwick is referring here to Butler's distinction between strength and authority. Butler maintains that what he calls "the principle of reflection" or "conscience" has *authority*. This is to be contrasted with the *mere strength* of our desires.

According to Butler,

that principle by which we survey and either approve or disapprove our own heart, temper, and actions, is not only to be considered as what is in its turn to have some influence, which may be said of every passion, of the lowest appetites, but likewise as being superior; as from its very nature manifestly claiming superiority over all others, insomuch that you

1 More specifically, we are dealing with what Stephen Darwall calls "constitutive existence internalism" where an agent's good *consists* in some subset of her desires (e.g., her fully informed desires). (Darwall 1992)

2 Internalism is rarely argued for but, instead, advanced as an intuitively obvious view. But see Rosati (1996) for five arguments for internalism.

cannot form a notion of this faculty, conscience, without taking in judgment, direction, superintendency. (Butler 41)

According to Sidgwick, the authoritative nature of what Butler calls "conscience" is not amenable to a naturalistic interpretation.

I believe that an important distinction between strength and authority is precisely what Connie Rosati has in mind when she argues against the full information account of well-being in "Naturalism, Normativity, and the Open Question Argument". Like Sidgwick, she maintains that Butler's distinction undermines naturalism. Unlike Sidgwick, she provides us with an argument for this view. This will be examined in chapter 2.

It is important to note that we will be concerned with what Rosati calls "the new naturalism."³ Full information theorists who fall under this heading include, among others, Richard Brandt and Peter Railton.⁴ The new naturalist offers a "reforming definition" of 'good'. It is not the business of reforming definitions to advance analytic or conceptual truths. Those who put forth a reforming definition propose, instead, an *a posteriori* identity claim. So, for example, when Brandt claims that 'good' means 'rationally desired', he is not claiming that there are no questions which can be asked using 'good' that cannot be asked using 'rationally desired'. Consequently, his account is immune to the kind of criticism advanced by G.E. Moore in Principia Ethica, namely the open question argument. "X is rational to desire, but is it good?" *is* a conceptually open question. But that, by itself, does not rule out the possibility that 'good' is to be correctly identified with 'rationally desired'. Compare this with the identification of water and H₂O. 'Water' and 'H₂O' are distinct concepts (i.e. their intensions differ) but they have, nonetheless, the same extension. They are, ontologically, one and the same thing. Similarly, we might maintain, with Brandt, that 'good' and 'rationally desired'

³ See Rosati 1995a: 47-52.

⁴ See Brandt 1979, Railton 1986a, and 1986b.

refer to the same thing despite their conceptual independence. Reforming definitions like Brandt's do not aim at conceptual closure. A reforming definition of 'good' should be, as Peter Railton supposes it can be, "tolerably revisionist", omitting only those aspects of our everyday usage that are superfluous to its essence.⁵

We are concerned, then, with the more sophisticated forms of naturalism. Railton, more than anyone else, is responsible for developing the full information account of well-being in the form of "sophisticated naturalism". He maintains that

an individual's good consists in what he would want himself to want, or to pursue, were he to contemplate his present situation from a standpoint fully and vividly informed about himself and his circumstances, and entirely free of cognitive error or lapses of instrumental rationality. (Railton 1986b: 16)

Again, since Railton puts this forward as a *reforming definition*, it will not do to simply object that we can still ask, regarding the things which our fully informed selves would want us to want, whether or not they really are good for us.

The "wanting to want" clause is also an important element in Railton's account. There are many things that our fully informed selves might not want for themselves but would want us to want were they about to step into our shoes. For example, once fully informed we would not want any more information since we already have it all. But we would want more information for our less than perfect selves were we about to assume their position.⁶

Finally, the information considered by the agent's epistemically preferred self, according to Railton, should include all of the possible experiences and lifestyles open to her. This is reminiscent of John Stuart Mill's suggestion, in Utilitarianism, that in order to determine which of two pleasures is the "higher pleasure" an agent must be "competently acquainted with both." (Mill 139) On

⁵ See Railton 1990: 158-159, 168-172.

⁶ See Railton 1986a: 174.

Railton's model, an agent's good is determined by the preferences of her fully informed self who is "competently acquainted" with every possible course her life might take. Accordingly, Railton's account seems to fare better than Brandt's regarding the internalist constraint. Brandt holds that an agent's good is determined by those desires the agent would have after undergoing a process which he calls "cognitive psychotherapy." This process involves "confronting desires with relevant information, by repeatedly representing it, in an ideally vivid way, and at an appropriate time." (Brandt 113) But, as J. David Velleman points out,

[w]e can alter our desires, not only by exposing ourselves to the facts, but also by exposing ourselves to other kinds of influence - to the influence of other people, of literature, of prayer, or of our own self-censure and self-praise. (Velleman 357)

So the problem seems to be as follows. If I now have a desire for X, and if that desire would survive cognitive psychotherapy, then the desire for X is a rational desire and, therefore, a part of my good. But if Y is something I could come to desire only through some noncognitive means, then Y cannot be a part of my good even if, once I acquired the desire for Y, the desire would not be extinguished by cognitive psychotherapy. Surely it makes sense to ask whether or not it would be good for me to desire Y. But Brandt's account seems to settle the issue too quickly. If I do not already have the desire for Y, and if the desire would not be acquired through purely cognitive means, then Y is simply not a part of my good. The internalism that underlies Brandt's account, then, is too narrow since it ignores potential desires which may figure into an agent's good. Railton's model, on the other hand, does not restrict a person's good in this way. The information considered by the agent's epistemically preferred self includes all of the possible experiences and lifestyles open to her.

The full information account of well-being, so stated, faces two serious worries. (1) It has been argued by many (e.g., Allan Gibbard, Connie Rosati, and J. David Velleman) that the judgments of our fully informed selves lack the normative authority that the full information account of well-being claims that they have. The full information account, then, does not provide us with what Railton calls a "tolerably revisionist" account of 'good'. There are some important features of 'that which is good for A' that are not captured by 'that which A would want herself to want were she fully informed'. Connie Rosati, for example, argues that there are some important questions that we can address using our everyday concept of 'good' that cannot be addressed using the definition of 'good' provided by the full information account.⁷ In particular, the full information account does not allow us to address properly questions regarding what kind of person to be. And these questions, she points out, are certainly relevant to discourse regarding a person's good.

Rosati believes that the judgments of our fully informed selves lack normative authority for us for another reason as well. She is concerned about the notion of being fully informed that the full information account relies on. The full information theorist must be able to make sense of fully informing someone while holding her motivational system fixed. That way we end up with the same person, only fully informed. But Rosati believes that in order to become fully informed, an agent must undergo various changes in her motivational system in order to appreciate certain information. This suggests that, after the process of becoming fully informed, the agent will no longer be herself. And even if she is herself, her judgments might still lack normative authority. After all, providing her with full information will result in **some changes** to her motivational system. She

⁷ We shall see, in chapter 2, that this is a sophisticated version of Moore's open question argument.

will, for example, lose desires that are based on false information. Finally, Rosati considers the possibility that the notion of being fully informed is incoherent. She maintains that an agent simply cannot (even in theory) be informed about every possible perspective open to her.

(2) It has also been argued that the concept 'that which a person would want if fully informed' is not an empirically determinate concept. J. David Velleman, in "Brandt's Definition of Good", argues that "[t]here is no single motivational impact associated with the facts in themselves. The facts would exert various impacts, when presented in various media, perspectives, and vocabularies." (Velleman 366) Together, these two worries present a formidable challenge to the full information account of well-being.

My goal, eventually, is to provide a complete defense of the full information account. It strikes me as the most plausible candidate for a naturalistic definition of 'good'. And I believe that it could serve well as the foundation for a teleological ethical system like utilitarianism. But I have a much more modest goal for this thesis. I will defend the full information account against the first type of worry (i.e., the type of worry raised by Connie Rosati). I want to show that the judgments of our fully informed selves **do** have normative authority for us.

In chapter 2, I attempt to meet Rosati's challenge that the full information account of well-being does not allow us to address properly questions regarding what kind of person to be. Rosati's point is that we may hold personal ideals that are at odds with our current motivations. Why, then, should we care about what our fully informed selves (which include these motivations) would want for us? So questions about what kind of person to be seem to be, on the full information account, inappropriately influenced by the person's current motivational system. But I argue that "our current motivational system" includes the personal ideals

that we now hold. And it is one's entire motivational system (which includes these personal ideals) that determines one's good under full information.

Still, one might be concerned that certain undesirable traits get too much weight in determining our good. In cases of weakness of will, more information often does nothing to prevent us from doing what we believe is wrong. To meet this worry, I argue that an agent's good is not to be identified with just any reactions that he might have to full information. Instead, I suggest that 'good for a person' refers to what that person would "endorse" under full information. An agent endorses a desire when she wants that desire, and not other conflicting desires that she may have, to be effective in bringing about action. This basically is what Harry Frankfurt, in "Freedom of the Will and the Concept of a Person", has called "second-order volitions".

In chapter 3, I attempt to meet Rosati's worries concerning the notion of being fully informed that the full information account relies on. First I deal with her concern that we are not the same person after undergoing the process of becoming fully informed. I argue that memory can serve to appreciate information that initially required a different motivational system than the one we now have. With the help of an extended memory, it is possible for an agent to appreciate all of the information which requires different motivational systems from the perspective of her current motivational system. She will thus be the same person (a person with the same motivational system), only fully informed. I then argue that the changes in her motivational system that will result from **being fully informed**, and not merely from the **process of becoming fully informed** (e.g., the elimination of desires based on false beliefs), will be welcome. The judgments of her fully informed self will then have normative authority for her. Finally, I will try to show that the examples which Rosati provides to suggest that the notion of being fully informed is incoherent are ineffective.

I believe that if Rosati's worries can be met we will have good reason to suppose that the judgments of our fully informed selves do have normative authority for us. These judgments seem to speak to the concerns we normally have when considering our good under ordinary optimal conditions.⁸ And this suggests that the full information account of well-being provides us with an account of 'good for a person' which is, at worst, tolerably revisionist.

⁸ This will be addressed in chapter 3.

Chapter 2: Personal Ideals and Full Information Accounts of Well-Being

We now turn to Rosati's objection in "Naturalism, Normativity, and the Open Question Argument". We have seen that the mere fact that it is possible to ask the question "This is something I would want myself to want were I fully informed, but is it good?" is no objection to Railton's account. Moore's open question argument is no challenge to the new naturalist who advances reforming definitions instead of conceptual truths. However, Rosati makes use of what we might call a more sophisticated version of the open question argument. Indeed, as she points out, Railton himself recognises the possibility of revising Moore's argument to render it effective against the new naturalist. (Rosati 1995a: 48)

Railton writes that

[e]ven if Moore's "open question" argument cannot be deployed directly against an interpretation of discourse about a person's good that does not purport to express analytic truths, a significant critical function may still be served by pressing Moorean questions against such interpretations. For it would be a challenge to any theoretical identification or reforming definition of P in terms of Q to argue that there is something central to the notion of P that does not appear to be captured by Q; this would make the question 'I can see that this is Q, but is it P?' genuinely compelling, not just barely possible. (Railton 1990: 158)

Now, if Rosati is to avail herself of this strategy she must show that there are in fact some important, or "genuinely compelling", questions that can be asked using our everyday conception of 'good' that cannot be asked using Railton's reforming definition. And that is exactly what she tries to do by introducing what she calls "ideals of the person." (Rosati 1995a: 60)

According to Rosati, any theory of personal well-being must take into account the fact that we are autonomous agents. What is fundamental to humans (i.e. what separates us from, say, animals) is our capacity for self-

invention. We need not identify with our current motivations and desires. We have higher-order desires which allow us to step back from our current desires and perhaps dissociate ourselves from them. Specifically, we have personal ideals that we identify with and we may choose to change ourselves in various ways in light of them.

Where the full information account fails, according to Rosati, is in its attempt to tailor the good for humans as self-inventors. She writes that

[a]n account of a person's good will fail to be suited to persons as self-inventors insofar as it treats a person as identical with certain motivations or traits, for this is to accord her current features a normative authority that they lack. (Rosati 1995a: 61-62)

Naturalists in general, and full information theorists in particular, treat the person as identical with her current motivations and character traits. On Railton's model, the idealisation process simply involves providing the individual with full information while her motivational system remains fixed. That way it is the same individual, only fully informed, who acts as the ideal advisor.⁹ But the problem is that we are often concerned precisely with questions about what kind of person we should be - about which motivational system we should have. We may hold personal ideals which conflict with our current motivational system. The personal ideals that we hold have normative force for us even though certain other desires that conflict with them might be stronger in the sense that we end up acting on them (much to our dismay). To use Butler's terminology, our personal ideals have *authority* and not just *mere strength*. One might not care, then, about what one's fully informed self would want oneself to want. Let us consider an example of Rosati's in order to make clear this last point.

⁹ Rosati uses this label because of the structural similarities between the full information account of well-being and Roderick Firth's "ideal observer" theory. (Rosati 1995b: 298)

Sandy is thinking of becoming more like her friend Madelyn. Sandy feels that she has been living life too cautiously. Madelyn lives her life in a carefree way. According to the full information account, it is good for Sandy to become less cautious if she, in an epistemically preferred state, would want herself now to be less cautious. But will Sandy want her own motivational system, which includes her cautiousness, to be informing her decision regarding whether to be less cautious or not? According to Rosati, "[g]iven her concern about the sort of person she is, she may reasonably wonder whether what someone like herself would want for her if fully informed and rational is indeed good for her." (Rosati 1995a: 54) Indeed, Sandy might care more about what Madelyn's fully informed self would want for Sandy. After all, she would be fully informed about everything, including facts about Sandy, and she has the motivational system and character traits with which Sandy identifies.

So it seems that the full information account does not allow us to address properly questions about what kind of person to be - questions which certainly seem relevant to discourse regarding a person's good. Rosati believes that the answers to such questions are, on the full information account, inappropriately influenced by the motivations and character traits up for evaluation (e.g., cautiousness in Sandy's case). The personal ideals that we hold have normative force for us. And discourse about our good is often tied up with these personal ideals. This, according to Rosati, is the important feature of 'good' that does not get picked up by the full information account such that the Moorean question "I see that this is something that I would want myself to want were I fully informed, but is it good?" becomes "genuinely compelling, not just barely possible."

But I believe that Rosati's worry can be met. Although the naturalist's account of 'good for a person' does identify a person with her current motivations and character traits, it is important to note that those motivations and traits may

(and often do) include holding an ideal of the person. So the account will "accord her current features a normative authority that they lack" only if we leave out important features of the person like her holding a particular personal ideal. The full information account of well-being seems perfectly suited to us as self-inventors. It is good for me to act on my desire to change myself (in light of the personal ideals that I now hold) if that desire would still be had by my ideal advisor (i.e., my fully-informed self about to step into my shoes). Again, 'me' here refers to my entire motivational system which includes any ideal of the person that I might now hold.

However, Rosati objects to the idea of treating ideals of the person as brute facts about people in the way that I have just outlined. In a footnote, she comments that

it is a fact about a person that she identifies with a particular ideal. If the fact that she identifies with a particular ideal favored that ideal, then the facts about a person would favor a certain standpoint for making determinations about her good. But to allow this fact to determine a person's good would be to allow a person's choice to determine her good in a way unintended by the new naturalists. (Rosati 1995a: 68)

By treating ideals of the person in the manner I have suggested, Rosati believes that substantive evaluative content has been introduced into the definition of 'good' and that the naturalist's program has therefore been abandoned.

But I do not think that any evaluative content has been imported into the definition of 'good.' My holding a personal ideal, again, is simply a part of my motivational system. It involves the desire to lose some of my other traits and perhaps acquire some new ones. But the fact that I hold some ideal of the person does not alone "favor a certain standpoint for making determinations about my good." What is good for me, according to the account under consideration, is what my fully-informed self would want myself now to want. My

fully-informed self would of course take into consideration the ideal that I now hold and how important it is to me. But it is my *whole* self and not just *some* motivations, like my holding a particular ideal of the person, that determines my good under full information.

To recapitulate, Rosati's worry is that either (1) the full information account of well-being accords our current motivations, those which conflict with the personal ideals we hold, undue normative authority, or (2) the naturalism which underlies the full information account of well-being is compromised by allowing the fact that we hold a particular personal ideal to favor a certain standpoint for determining our good. But we seem to have settled both of these worries by including as a part of "our current motivations" the personal ideals that we hold.

It seems to me that we have no other option but to treat personal ideals this way. What else could holding a personal ideal involve but the desire to lose and / or acquire certain motivations or character traits? And there is no reason to suppose that this type of desire is any less a part of one's motivational system than any other motivation or character trait. For example, cautiousness and the desire to loosen up (the desire to be like Madelyn) are among Sandy's motivations.

One might be tempted to resist this and maintain that although holding a personal ideal involves certain desires, the personal ideal itself is not a desire. We might want to explain why we have certain desires (e.g., Sandy's desire to be less cautious) by appealing to an ideal of the person. And it seems that we lose the ability to give such an explanation by treating personal ideals as I have suggested. This is an instance of the general strategy employed by objective list theorists against desire-satisfaction theories of the good (like the full information account). David O. Brink, in Moral Realism and the Foundations of Ethics, writes that

we desire certain things *because we think these things valuable*. This is true of our preferences for many activities and relations as well as states of the world. It is not that these things are valuable because we desire them; rather, we desire them because we think them valuable. (Brink 1989: 225)

He concludes that "[t]he need for value-laden explanations of many desires shows that we think that desire-satisfaction theories are false." (Brink 1989: 226)

The problem with this is that the objective list theorist has given us no positive account of what value is. Brink merely claims that we often appeal to value in order to explain certain desires and that this suggests that the desire-satisfaction theory gets things backwards. But what exactly is it that is supposed to be explaining these desires? The suggestion is that what explains Sandy's desire to be like Madelyn is her **perception** that being like Madelyn is valuable. Sandy **recognises** that becoming like Madelyn is a worthy personal ideal and that explains why she has the desire to be like her. But what is it that Sandy is perceiving? Brink does not help us here. In fact, if we want to maintain that what is explaining Sandy's desire to be like Madelyn is her perceiving the value of being like Madelyn, the full information account **can** help us here. In this case, what Sandy is perceiving is a particular disposition of hers, namely, to desire being like Madelyn under full information. I believe that an explanation of desire that ends with an appeal to value is unenlightening unless we are given some positive account of what value is.

Moreover, it is not clear that we do actually appeal to value this way when explaining our desires. Sandy might initially claim that the reason she desires to be like Madelyn is that being like Madelyn is valuable. But if we pressed her on this she would likely point to certain features of Madelyn that she finds appealing. She might, for example, point out that Madelyn seems happy almost all of the time. And Sandy **desires** to be happy. We would find it puzzling if, instead of

giving this kind of response, she simply maintained that being like Madelyn is valuable. To be sure, such explanations often do end this way. We do not usually keep pushing until we get a response in terms of desire. But I believe that if we do keep pushing we will, in the end, get this kind of response. And this is a good thing since, as we have just seen, a response that ends with an appeal to value will be unenlightening.

Since Brink gives us no positive account of the value being appealed to in the explanation of certain desires, such explanations remain unilluminating. And so the charge that desire-satisfaction theories lose the ability to provide such an explanation is hardly a knock against them. Therefore, we can maintain that personal ideals are simply a part of one's motivational system without any explanatory loss.

It should be noted that even if these considerations do not satisfy the objective list theorist, Rosati does not, and cannot make use of this strategy. That is, she cannot maintain, with Brink, that certain activities, relations, and states of the world are objectively valuable. She cannot maintain that X is good for A even if A could not, under any circumstances, be moved by X. Rosati is an internalist. As we shall see in the next chapter, her concern with the full information account is that it does not satisfy a strong enough version of internalism (what she calls two-tier internalism). She is worried that, on the full information account, we are alienated from our good. Even though A might be motivated to pursue X under certain circumstances, A might not care about those circumstances. In this case, X is something which can motivate A but is nonetheless alien to her. On the objective list theory, we have an even greater concern regarding alienation from our good. For on that view, X might be good for A even if X fails to move A under any circumstances.

So although the objective list theorist might maintain that certain personal ideals are objectively valuable, Rosati cannot. And it seems that her only other alternative is to treat personal ideals as I have, namely, as a part of our motivational system. She does have, however, one final alternative that we might look at briefly. If Rosati cannot be an externalist because of her concern with alienation, perhaps there is another form of internalism that she might endorse. Recall that the internalism we are concerned with is what Darwall calls "constitutive existence internalism". Our good is such that it must be capable of motivating us (existence internalism), and our motivations **constitute** our good. But Rosati might endorse, instead, a form of non-constitutive existence internalism. On this view, our good is such that it must be capable of motivating us, but our good is not constituted by our motivations. This is the view that J. L. Mackie, in Ethics: Inventing Right and Wrong, attributes to Plato. He writes that

[t]he Form of the Good is such that knowledge of it provides the knower with both a direction and an overriding motive; something's being good both tells the person who knows this to pursue it and makes him pursue it. (Mackie 1990: 40)

So personal ideals might be something akin to Platonic forms - things which are not desires but which do motivate us once we recognize or perceive them.

However, as we saw in chapter 1, this view is not plausible. For now we must suppose that there are intrinsically motivating features of the world, where these features are not simply the desires or motivations of agents. As Mackie points out, things (other than desires) that have to-be-pursuedness built into them would have to be awfully queer. An account would have to be given explaining how it is we come to know about these things, and how exactly knowing about them motivates us. Certainly Rosati provides us with no such account. Indeed, she gives us no reason to suppose that she endorses anything like non-constitutive existence internalism. I simply point out that it is a **possible** (not

plausible) option for treating personal ideals which does not alienate us from our good as externalist accounts do.

So Rosati is left with treating personal ideals as a part of our motivational system. As a result, her second worry is settled straightforwardly. A personal ideal that one holds does not alone favor a certain standpoint for determining that person's good since it constitutes only a part of her current motivational system. In order to determine her good, it must compete with other motivations and traits when confronted with full information. The naturalism underlying the full information account, then, is not compromised.

Treating personal ideals as particular character traits, as we have already seen, also seems to settle Rosati's first worry. The full information account accords one's current motivational system undue normative authority only if we leave out important aspects of that motivational system like holding a particular personal ideal. Again, *all* of the motivations and character traits which make up the individual's current motivational system compete with one another under full information in order to determine her good. But perhaps it is precisely this competition among motivations and traits, ***even if those motivations and traits include the holding of a particular personal ideal***, that is the root of Rosati's first worry. Let us now pursue this issue.

Consider once again Sandy's case. Sandy *really* wants to be carefree like her friend Madelyn. Why, then, should Sandy care about the effects that full information would have on her current motivational system which includes her cautiousness? On the full information account, all of Sandy's motivations and character traits, including her cautiousness and her desire to be less cautious, get a say in determining her good under full information. But Sandy might not want her cautiousness to get *any* say in the matter.

Sandy's desire to be like Madelyn is what Railton calls "goal-setting." According to Railton, "one embraces a desire, or accepts it as goal-setting, when one desires that it be effective in regulating one's life." (Railton 1986b: 14) These desires are to be distinguished from our more basic first-order desires.

Embracing a desire as goal-setting involves not only certain desires but also certain beliefs. We feel compelled to prove to ourselves, and to others, that our goals, our values, are worthwhile, and we do this by showing that they are somehow rooted in reality. We do not treat our more basic desires this way.

Railton comments that

it does not do much to explain to myself or others the worthwhileness or point of what I have done with my life to say that I have simply acted upon whatever desire happened to be most urgent at the moment. Higher-order desires of the sort that are involved in embracing a desire are more responsive to changes in belief, and so not only do they become more closely tied to our identity, they become the basis of the idea of value. (Railton 1986b: 16)

We believe that our goals are supported by facts and this support, in turn, proves to us, and to others, that our goals are worthwhile. Consequently, "[i]f we discover that our values are psychologically dependent upon ignorance or error, we lose this source of support." (Railton 1986b: 16)

So Sandy believes that becoming more like Madelyn is a worthwhile goal. As a result, she does not identify with her cautious nature. Perhaps she believes that she would be a happier person if she were less cautious. Madelyn seems so happy all of the time and Sandy believes that this is a result of her carefree nature. Or perhaps she believes that she will be a better person morally without her cautiousness.¹⁰ If more and more information confirmed her belief that being less cautious would make her, say, a happier person, then her goal (to be more

¹⁰ I do not wish to give happiness or moral uprightness a privileged theoretical role here. These are only possible explanations of Sandy's case.

like Madelyn) would be further supported. On the other hand, if she discovered that many of her beliefs bearing on the matter were false, she would doubtless reconsider her goal. Sandy now believes that the facts support her goal. This gives her life meaning since the values she identifies with appear to her to be rooted in reality. It might cause her a great deal of distress to think that she might be mistaken about her beliefs on this matter. But this, by itself, is no reason to exclude any of her motivations or character traits (e.g., her cautiousness) in determining her good under full information. After all, Sandy, when contemplating her good, wants to be sure that she is getting the facts straight and dealing with them rationally. It would be irrational for her to rule out certain possibilities regarding her good simply because she does not now identify with some of her current motivations or character traits. She might be mistaken about the facts that led to her dissociation with these motivations or traits.

Perhaps Sandy does not want her cautiousness to have any say in the matter not because she is worried that the facts might not support her goal, but because of her concern with weakness of will.¹¹ Sandy often has the opportunity to be more carefree. And each time such an opportunity comes up she believes that doing so would be the right thing to do. But unfortunately her cautious nature always prevails and she ends up playing it safe. Regardless of how much information she gathers to confirm her beliefs which support her goal, she still ends up doing what she believes is wrong (i.e., acting cautiously). It seems likely, then, that Sandy's fully informed self, who differs from her only in that she has more information, would react the same way.

We must remember, however, that we are not identifying Sandy's good with what her ideal advisor wants for herself but with what she wants for Sandy

¹¹ Robert Shaver has recently suggested (in Shaver 1997) that the problem of weakness of will is what drove Sidgwick to reject the full information account of well-being.

now. If full information would support Sandy's goal, then her ideal advisor would identify with that goal; she would strongly embrace the desire to be like Madelyn.¹² And despite the fact that she (the ideal advisor) might not herself act appropriately, she would want Sandy now to act appropriately were she about to step into her shoes. Sandy's ideal advisor endorses the desire to be like Madelyn. Being more like Madelyn, in this case, is good for Sandy.

Still, we might be concerned that our ideal selves would endorse things that just seem wrong to us. Why should we think that full information will affect us the "right" way? Perhaps Sandy's cautiousness is so strong that it affects virtually every facet of her life. Full information might reveal to her even more opportunities to act cautiously. Sandy's ideal advisor might be motivated, in the end, to do nothing more than stay in bed all day (after all, it's safe there). Ultimately, her cautiousness might affect what she will endorse. Recognizing what kind of person she is (a person with an extremely cautious nature), fully informed Sandy might no longer endorse the desire to be like Madelyn. Instead, she might endorse desires that, to most of us, just seem crazy (e.g., the desire to do absolutely nothing).

We might suppose that one would in fact **never** endorse such things once fully informed. In other words, we might maintain that full information will always affect us the "right" way. But I do not think that this is the right move to make. I agree with Rosati when she claims that "[w]e lack sufficient reason to think that the process of becoming fully informed works uniformly in the direction of acquiring desirable traits." (Rosati 1995a: 57) We must allow for cases like the one just given. But such cases can be handled by the full information account of well-being. If fully informed Sandy endorses the desire to always play it safe (no

¹²Of course she might identify equally with several other goals. This would simply mean that several options are equally best for Sandy. See Railton 1986a: 176.

matter how small the risk), then playing it safe is good for Sandy. And this goes for any other "undesirable trait" that an agent might endorse under full information. In these cases, we must simply revise our judgment regarding the "desirability" of the trait in question. Of course we will want to maintain steadfastly that certain immoral traits are undesirable. For example, one's fully informed self might endorse the desire to hurt people. We want to say that this is wrong. But our judgment here is a moral one and we are considering an account of a person's non-moral good.

But just how effective a response is this? We might not be happy with the idea of holding our good hostage to the psychological outcome of becoming fully informed. For one thing we might be concerned that the process of becoming fully informed will radically change us so that we are no longer the same person. That is, we might end up with a completely different motivational system. This will be the focus of the next chapter. But even if we grant that we are the same person after becoming fully informed, we might still be uneasy about identifying our good with that which we would endorse under full information given that we might endorse things that we now think are "undesirable". We might be surprised by what is in fact good for us.

Surprises, though, are not necessarily a bad thing. Indeed, one of the two criteria that any realist conception of value (or, for that matter, any realist conception of any domain of discourse) must meet is what Railton calls "independence". (The other criterion, which will be discussed shortly, is what Railton calls "feedback".) If 'personal good' is real then "it exists and has certain determinate features independent of whether we think it exists or has those features, independent, even, of whether we have good reason to think this". (Railton 1986a: 172) The full information account of well-being satisfies this criterion. Although it is a subjective account in the sense that our good depends

upon our psychological features, it is not subjective in the sense that our good is determined by whatever we might happen to think about it. We have identified a person's good with a dispositional feature of that person, namely her disposition to endorse certain things under full information. As Brandt points out, we have an account in which we "frame our questions clearly and then go out to find answers, letting the chips fall where they may." (Brandt 1979: 3) The results might be surprising.

Now, clearly if everyone, once fully informed, endorsed "undesirable things" we would have the right to be suspicious of the view that what is good for us is what we endorse under full information. But we have no reason to suppose that this is the case. After all, on the full information account, we are intimately connected with our good. The full information account satisfies Railton's second criterion for realism, namely feedback. If 'personal good' is real then "it is such - and we are such - that we are able to interact with it, and this interaction exerts the relevant sort of shaping influence or control upon our perceptions, thought, and action." (Railton 1986a: 172) Surely Sandy's case (where she endorses her extreme cautiousness under full information) is not the norm. Given that we are intimately connected with our good, and are therefore constantly learning about it, it is more likely that most of us now endorse things that more or less approximate those things which our better selves would endorse. An example of Railton's in "Moral Realism" might help to make this clear.

Lonnie is a homesick traveler who is suffering from a sore stomach due to dehydration (an affliction, Railton tells us, that is common to tourists). Usually when Lonnie has a stomach ache he takes a glass of milk. But milk, being hard to digest, would only worsen Lonnie's condition. If Lonnie knew that his sore stomach was caused by dehydration he would no longer want his desire for milk to be effective. Instead, he would want himself to want clear liquids like water.

During the course of his travels Lonnie might come across some drinking water and, in the absence of milk, he might drink some. After a few such occasions, he might come to recognize that his condition has improved (he no longer has a stomach ache) and he might eventually acquire a genuine taste for water in similar situations. Our good tends to shape our desires this way. And we can explain, using the full information account's conception of 'good for a person', why some people are more or less content, or satisfied with their lives. Consider another tourist who, unlike Lonnie, continues to drink milk. She will remain miserable while Lonnie will perk up. This is because she, unlike Lonnie, has not (yet) discovered and acted on what is good for her.¹³

So we can interact with and learn about our good (sometimes through trial and error as in Lonnie's case). The criteria for realism (independence and feedback) seem to be met by the full information account. And this is a good thing since realism seems to be suggested in discourse about our good. When we argue with someone about whether or not something is good for us we take it that one of us is right and the other is wrong. That the account will sometimes yield surprising results (as in Sandy's case) is not a cause for concern. Indeed, it is precisely what we should expect from a realist account of value. But we have no reason to suppose that it will yield only surprising results. Most of us have desires that have evolved by interacting with our good (as in Lonnie's case) and it is likely that we now endorse things that roughly approximate our good.

I submit, then, that Rosati's two worries can be met by treating personal ideals as a part of one's motivational system. By treating personal ideals this way, the full information account of well-being neither accords one's current motivational system undue normative authority nor does it compromise the

¹³ Railton suggests that this meets Gilbert Harman's challenge, in The Nature of Morality, that moral facts are explanatorily impotent.

naturalism which underlies it by favoring a certain standpoint for determining the agent's good. The agent's motivational system includes all of her motivations and character traits, including the personal ideals that she holds, and each motivation and trait is considered under full information in determining her good. So it seems that the full information account of well-being does allow us to address properly questions regarding what kind of person to be after all and thus seems perfectly suited to us as self-inventors. Those interested in pursuing what I have called the sophisticated version of Moore's open question argument against the full information account of well-being, then, must search for other important features of 'X is good for A' that are not captured by 'X is something A would want herself to want if A were fully informed.'

Chapter 3: The Notion of Being Fully Informed

In "Persons, Perspectives, and Full Information Accounts of the Good", Rosati explores the notion of being fully informed. We have, up until now, assumed that this notion is unproblematic. We have taken it for granted that it is (at least theoretically) possible to fully inform someone while holding her motivational system fixed. The idea here, recall, is to ensure that it is not just anybody who acts as an agent's ideal advisor - it is the agent herself. Sandy's ideal advisor is Sandy, only fully informed. There are three related worries in "Persons, Perspectives, and Full Information Accounts of the Good" all of which center around the notion of being fully informed. They are as follows:

- (1) The person who emerges after the process of becoming fully informed may not be you. Her judgments will therefore lack normative authority for you.
- (2) Your fully informed self (even if it is you) might not have the qualities suitable for an ideal advisor. Her judgments will therefore lack normative authority for you.
- (3) The notion of being fully informed is incoherent. It is theoretically impossible to be informed about all perspectives from one's own perspective.

I will take up each of these concerns in turn. I want to show that Rosati's arguments are unconvincing and that the notion of being fully informed while holding one's motivational system fixed is quite plausible.

Rosati's first and second worries stem from her concern that the notion of being fully informed which the full information account relies on is not compatible with any respectable form of internalism. We will begin by looking at what Rosati has in mind by a "respectable form of internalism." We will then consider whether or not this form of internalism is compatible with the notion of being fully informed.

"Simple internalism" provides us only with a necessary condition for a person's good. (Rosati 1996: 300) For X to be a part of an agent's good, X must be capable of motivating her. But it seems that there are many things which are capable of motivating an agent under certain circumstances that we would not be inclined to consider as a part of that agent's good. For example, I might come to desire X only after being lobotomized. In this case, even though X is something which is capable of motivating me, it is not something that can plausibly be considered as a part of my good. Recall that Railton's reason for adopting internalism is that accounts which are not internalist yield "intolerably alienated" conceptions of personal well-being. But, as the lobotomy case suggests, simple internalism does not seem to fare any better. Rosati's concern with simple internalism is that "it counts as possible goods for a person things that would violate the intuition that inspires internalism - the intuition that a person's good must suit her, that it cannot be something alien to her." (Rosati 1996: 301)

Rosati believes that the intuition which drives internalism suggests "a double motivational link." (Rosati 1995b: 300, and 1996: 307) For X to be good for A, X must not only be capable of motivating A under conditions C, but A must also care about the fact that she would desire X under C. Obviously we cannot maintain that the appropriate conditions, under which A's desire for X determines her good, are to be determined by the agent herself "whatever her present state might be." (Rosati 1996: 305) Appropriate counterfactual conditions, according to Rosati, are determined by the agent's reflecting on them under "ordinary optimal conditions." For example, the person must be attentive, sober, calm, and under "whatever normally attainable conditions are optimal for reflecting on questions about what to care about self-interestedly." (Rosati 1996: 305) X is good for A, then, only if the two following conditions are met:

(1) Were A under conditions C and contemplating the circumstances of her actual self as someone about to assume her actual self's position, A would care about X for her actual self;

(2) conditions C are such that the facts about what A would care about for her actual self while under C are something A would care about when under ordinary optimal conditions. (Rosati 1996: 307)

Rosati calls this form of internalism "two-tier internalism."

Now, it is not at all clear that the full information account fails to satisfy these conditions. After all, as Rosati points out, Railton's view seems plausible to us precisely because "the counterfactual conditions in [his] account speak to certain concerns we ordinarily have in contemplating whether something is good for us - Do I know enough about the thing I desire? Am I being sufficiently rational?" (Rosati 1996: 308-309) And again she writes that

[b]ecause Ideal Advisor views hold a person's personality constant and constrain counterfactual conditions to meet those epistemic concerns we ordinarily have when we wonder whether something is good for us, their proponents believe that these views capture both the critical character and the recommending force of judgments about our good. The information that you would desire X for yourself under the specified conditions will have pull for you, even if X itself still fails to move you. (Rosati 1995b: 303)

Where, then, do Ideal Advisor views fail? According to Rosati, the second link to motivation is not secured by the full information account because of the process that a person must undergo to become fully informed. A person's personality cannot be held constant while being fully informed. The person that emerges after this process will be radically different from the person prior to the process and so her judgments will lack normative authority. We must now examine Rosati's argument for this claim.

It seems that, in many circumstances, informing someone about something involves nothing more than presenting them with the information. For example, merely telling someone that the glass in front of her is filled with gasoline and not

water will probably be sufficient in getting that information to "sink in". Chances are that she will no longer desire to drink the contents of the glass.¹⁴ But getting information to register is not always this easy. Certain information will not register with certain people due to their motivational system. In these cases, the agent must "undergo education or experiences of a kind sufficient to render [the] information fully vivid." (Rosati 1995b: 308) Someone who is lazy, for example, will not really be informed about the fact that hard work towards one's goals can be rewarding merely by being told this. Thus, if we wish to fully inform this person we must first change her motivational system (by having her undergo the appropriate education or experiences) so that the otherwise impotent information can be appreciated. But if we have to change a person's motivational system in order to provide her with certain information then it seems that her idealized self will not really be her after all and her judgments will lack normative force. This is Rosati's concern. She writes that "[i]f an account of what it is to be fully informed is to have normative force, it must overcome what I shall call the problem of appreciation: the problem of the gap between merely having information and appreciating it." (Rosati 1995b: 304)

But although certain information may require a change in one's motivational system in order for it to register, I believe that we can still make sense of the notion of being fully informed required by the full information account of well-being. Suppose that I have a motivational set which consists of A, B, C & D. Having this particular motivational set allows me to appreciate certain kinds of information but not others. In order for me to appreciate a certain bit of information (call it I1), I will need to add E to ABCD. And in order to appreciate

¹⁴ Of course, we do not want to identify "being informed" with particular motivational responses for we are concerned with what effects full information has on one's motivations. I only mean to suggest that this response seems to indicate that the information has registered with the agent.

I2, I will need to lose A and B, and perhaps add F. And so it goes, altering my motivational system until "I" am fully informed. At the end of this process "my" motivational system might consist of something like XYZ. But if I am ABCD (recall that the naturalist treats a person as identical with his current motivational system), then why should I care about what someone having motivational set XYZ would want for me?

However, we are not yet done. We are not concerned, in this example, with what is good for a person having a motivational set XYZ. We want to know what is good for me - for someone having a motivational set ABCD. Now, it is possible that after undergoing all of the changes necessary for becoming fully informed that I end up with motivational set ABCD. In this case there seems to be no problem. The difference between me then and me now is that now I have experienced all of the information. I can **remember** what it was like to appreciate I1 and I2 and all of the other information the appreciation of which required a change in my motivational system. And although a change in my motivational system was necessary for getting this information to "sink in" the first time, it seems that I no longer need those motivational traits in order to appreciate it. An atheist who was at one time a devout Christian can still recall how uplifting it can be to think that one's prayers are being answered even though she is no longer motivated to pray. On the other hand, it is likely that someone who has always been an atheist will not be able to fully appreciate this information.

It seems that memory is a reliable means of considering information that one was able to appreciate initially only in virtue of having a different motivational system. This suggests a solution to the problem of appreciation. We have just considered a case where the agent's motivational system ends up being the same as it was before he underwent the process of becoming fully informed. But we have no reason to suppose that this will always, or even ever, be the case.

How can we deal with the possibility that the process of fully informing me might cause my motivational system to change from ABCD to XYZ? Quite simply, we need to change my motivational system back to ABCD. As long as this is done after I have received and appreciated all of the information, I can rely on memory to appreciate those bits of information that required a change in my motivational system. We can now say that I (a person with motivational set ABCD) am fully informed.

This idea of changing an agent's motivational system back to what it was prior to the process of becoming fully informed might seem a little unclear. We have spoken of changing one's motivational system in order to accommodate certain information by providing the agent with the appropriate education or experience. So we provide the life-long atheist with religious experience in order for her to appreciate that having one's prayers answered is uplifting. But we might wonder what kind of education or experience is required in order to restore the agent's original motivational system since, in this case, we are not trying to accommodate any particular information.

But perhaps the same point can be made another way. Instead of changing the agent's motivational system to accommodate certain information and changing it again to restore the original motivational system, we might simply provide the agent directly with the experience of remembering what it was like to appreciate the information. This way we can present all of the information to the agent without changing her motivational system at all. Of course, talk of memory implants for the purpose of fully informing someone may seem far fetched. In practice it may indeed be impossible to fully inform someone this way, or any other way. But we must remember that what we are concerned with is not the practical applicability but the theoretical possibility of being fully informed while holding one's motivational system fixed.

Whether we think about my answer to the problem of appreciation in terms of memory implants or whether we think about it in the way I originally presented it, where the agent's motivational system is changed to accommodate information and changed again to restore her original motivational system, I hope that the general strategy is clear. We can now appreciate certain information via memory even if our motivational system is such that we could not appreciate this information without these memories. An atheist who was once a Christian can appreciate that having one's prayers answered is an uplifting experience even though she is no longer motivated to pray. In the idealized case we simply suppose that we have all of the relevant memories that allow us to appreciate all of the information. We have the same motivational system (i.e., we are the same person) only now we are fully informed. And since we are concerned, when contemplating our good under ordinary optimal conditions, about whether we have got all of the information straight, the judgments of our fully informed selves will have normative authority for us. So it seems that the changes due to the process of becoming fully informed pose no barrier for the second link to motivation.

But perhaps the second link will be threatened not by the changes that an agent's motivational set undergoes **due to the process of becoming fully informed**, but by the changes that will result **due to being fully informed**. My motivational system (ABCD) will not be (permanently) altered just because it requires certain (temporary) changes in order to appreciate certain information. We have seen that I may simply rely on memory to appreciate this information. But we must allow for changes in my motivational set that are the result of the information itself. After all, we are identifying a person's good with what he desires once fully informed. One will probably lose certain desires (those based on false beliefs) and acquire new ones (those based on facts he was not

previously aware of). So I may end up with a different motivational set once fully informed. But now I might once again wonder if the judgments of my fully informed self have normative authority for me. He might want things for me that I do not now care for, and he might not want things for me that I now do care for. Is my fully informed self really my **ideal advisor**?

Richard Brandt, in A Theory of the Good and the Right, suggests an affirmative answer to this question when he considers "the force of knowing what it is rational to want." (Brandt 1979: 149-162) And I believe Brandt's suggestion is worth looking at. In what follows, I shall, for convenience, treat Brandt's expression "rationally desired" as identical to "that which the agent endorses for herself now once fully informed". For although the two expressions differ (as we saw in the last chapter), I believe that here Brandt is addressing the very worry that concerns us now. He is concerned with why rational desires recommend themselves to someone who does not now have them (and why irrational desires do not recommend themselves to someone who does now have them).

Brandt's solution is simple and it has the rigorously naturalistic tone that we are seeking. He claims that as a matter of fact people do prefer rational desires over irrational ones. "Since they do have this preference, a showing that a desire is rational will be a recommendation: a showing that it is irrational a discommendation." (Brandt 1979: 157) I say that the solution is rigorously naturalistic since not only is the concept 'rationally desired' completely interpretable in terms of fact, but so is its recommendation. There are no unreduced normative terms in Brandt's attempt to bridge the gap between recognizing that X is something one would desire if fully informed (i.e., recognizing that X is a rational desire) and being moved by X. He writes that "I am not offering any reason why this **should** be the case. I am asserting that, as a fact, people - including the reader - **do** dislike having to think that their desires

are irrational in my sense."¹⁵ (Brandt 1979: 157) We must now take a closer look at this strategy.

Brandt suggests that our motivation to acquire rational desires (and lose irrational ones) is similar to our motivation to acquire consistent beliefs (and lose inconsistent ones). He appeals here to psychological literature under the heading "cognitive dissonance theory". According to cognitive dissonance theory a person feels psychological tension, or dissonance, when she discovers that some of her beliefs are inconsistent, and she is motivated to relieve this tension. Dissonance theorists (e.g., Leon Festinger) maintain that this kind of psychological dissonance is not restricted to incoherent beliefs. Brandt considers one of Festinger's examples. It seems that one can be made uncomfortable "from awareness that one is afraid although there are only friends in the vicinity - apparently awareness that this is conditioned fear in a situation where there is no objective threat." (Brandt 1979: 156) We might call this a kind of "affective dissonance". Brandt concludes, by analogy, that dissonance also occurs when one recognizes that her desires are irrational.

The proposal here is that awareness of the fact that one has irrational desires works in a way similar to awareness that one has incoherent beliefs or unjustified fears. One is made uncomfortable by the awareness, and is motivated to remove its source. (Brandt 1979: 157)

This seems like a plausible strategy. But we need to take a closer look at the analogy that Brandt draws here. Consider what goes on with cognitive dissonance. We feel uncomfortable when we discover that a belief (or several beliefs) of ours is inconsistent with other beliefs that we hold. This moves us to either reject the belief(s) in question or revise our belief set in order to accommodate the new belief(s). In other words, the actual truth or falsity of the belief(s) in question does not necessarily determine whether we accept or reject

¹⁵ Emphasis is mine.

the belief(s). Consider an example.¹⁶ The founder of a radical religious sect announced that she had received a message from the "Guardians" of outer space. There was to be an enormous flood on a particular day and all but the true believers (the members of the sect) would perish. When the day came and no flood occurred, the true believers were faced with a dilemma. They either had to admit that their belief was false or else modify their belief set in order to accommodate it. They chose the latter. Instead of believing that they were mistaken about the messenger's prediction they believed that the world had been spared thanks to the faithfulness of the sect. This case seems to show that in cases of cognitive dissonance we are motivated to acquire **consistent** beliefs and not necessarily **true** ones since most of us would agree that the belief adopted by the sect was false. Of course the members of the sect believe that it is true (and not just merely consistent with other beliefs). But the point is that, in this case, the truth of the belief plays no role in resolving the psychological dissonance.

Truth seems to play more of a role in what I have called affective dissonance. In Festinger's example, what causes the psychological tension is the awareness that one's fear is not based on reality. The person recognizes that there are only friends in the room, that there is no reason to be afraid, and yet she is still afraid. It seems that in this case one would naturally try to relieve the tension by trying to eliminate the fear. However, we should not be too quick in accepting the idea that truth plays this kind of role in all cases of affective dissonance. For example, one might strive to keep an optimistic demeanor despite bad news from the doctor that one has only a short time to live. This person might be moved to relieve the psychological tension (the tension that

¹⁶ The example is taken from Henry Gleitman's textbook Psychology (p.458).

might come from awareness that her optimism is not based on reality) not by accepting her plight (i.e., adopting a less optimistic attitude) but by ignoring the doctor's information, perhaps choosing to believe that he is a quack. In this case (unlike the Festinger case), the affect is not tailored to fit the belief. Rather, the belief is modified to fit the affect. So, as with cognitive dissonance, affective dissonance does not seem to depend on the truth of one's belief(s) for resolving the psychological tension.

It seems, though, that truth must play a role in resolving the dissonance that comes from awareness that one's desires are irrational (which, of course, is a kind of affective dissonance). At least it must play a role if we are to maintain, with Brandt, that rational desires recommend themselves to us. After all, rational desires just are desires that a person would have if all of her beliefs were true. But in these cases, are we always moved to resolve the dissonance by changing (or wanting to change) our desires in accordance with the truth? Certainly we are moved to resolve the dissonance one way or another (as we are with all cases of psychological dissonance). But is it always in favor of rational desires?

There are two cases to consider here. First, there are those desires that we do not now have, but would have if we were fully informed. Suppose that I would desire X after being fully informed. If I somehow learn that this is the case, will the desire for X move me now? It probably will. To be sure, I might not desire X immediately upon learning that X is rational for me to desire. But it seems that I will want to take the necessary means to acquire the desire for X. However, it is not hard to imagine a situation where this is not the case. Suppose, for example, that X would conflict with another desire (Y) that I now believe to be more important. Since X is the rational desire and it conflicts with Y, it follows that Y is an irrational desire. But if I now believe that Y is more important, I will likely resolve the dissonance (which was the result of my

recognizing that X is a desire that I would have if fully informed) by ignoring X (perhaps choosing to believe that it is not really a desire that I would have if fully informed).

The second case involves desires that we now have but would not have were we fully informed (i.e., irrational desires). Here again it seems that we would likely resolve the dissonance that would result from recognition that one's desires are irrational by getting rid of the desire (or at least taking the appropriate steps to rid ourselves of the desire). We often change our desires very quickly upon learning that they are based on false information. For example, one would quickly lose the desire to drink the contents of the glass in front of him upon learning that it is filled with gasoline and not water. But once again there are many counterexamples. There are many things that we now believe we should desire. In other words, we now endorse certain things. If we somehow learn that some of these desires are irrational we will surely experience psychological dissonance. But in these cases it is not at all clear that we will resolve the dissonance by trying to lose the desire. These desires matter to us and we may therefore choose to resolve the dissonance by modifying our belief that they are based on false information rather than by giving them up.

So it seems that we need something more in order to show that rational desires recommend themselves to us. It will not do to simply maintain that people, as a matter of fact, always do seek rational desires (and avoid irrational ones) when they are faced with this type of psychological dissonance. We have seen that this is not always the case. Although most people are likely to resolve the dissonance in favor of rational desires most of the time, we have no reason to suppose that everyone will do so all of the time.

But there is another consideration, equally naturalistic, which recommends rational desires. Brandt calls it a 'pragmatic' consideration. He writes that "both

irrational desires and irrational aversions are apt to be costly, or to stand in the way of benefits, in one way or another." (Brandt 1979: 157) The examples that we have just considered help to illustrate this point. In the Festinger example we have a person who recognizes that she is afraid despite her belief that there are only friends in the room. Now in this case we saw that she would likely resolve the psychological dissonance by trying to rid herself of the fear. But again, this is not necessary. She may choose to modify her belief that only friends are present. However, if it is true that there are only friends in the room (and hence no reason to be afraid), remaining fearful will prove to be costly. She could be enjoying herself at a party with friends. This consideration will recommend to her losing the fear.

The example involving the patient who receives bad news from her doctor might seem to present a problem here. Suppose that the doctor's diagnosis is correct and the patient has only a short time left to live. In this case, not only is it plausible that she might resolve the psychological dissonance by choosing to maintain her optimistic demeanor, but doing so might actually benefit her. And losing her optimism might prove to be costly. Remaining optimistic might improve her quality of life for the time she has left. It might even increase the time she has left. So we seem to have a case where an "irrational desire" pays off and a "rational desire" is costly.

But I believe that it is a mistake to label the patient's desire to remain optimistic "irrational", and the opposite "rational". Recall that something is rational for an agent to desire if she, after being fully informed, would desire it for her less than perfect self. The fact that the patient's epistemically preferred self would know that her prognosis is bad does not entail that she will want herself now to give up hope. Indeed, if it is true that keeping an optimistic attitude would benefit her (whether it would improve or extend her life), then surely that is what she will

want herself now to want. In this case it is the desire to remain optimistic (despite the doctor's news) which is the rational desire.

So rational desires recommend themselves to us not because we always do resolve psychological dissonance in favor of them, but because acting on them pays off (while acting on irrational desires is costly). Of course certain desires that are rational for us will be costly if we act on them right away. One might discover that one's fully informed self would want him to quit smoking. For this person, the desire to quit smoking is a rational desire. But it is possible that immediately acting on this desire would be too much to bear for this person. Perhaps he would suffer from severe depression if he just quit cold turkey. What recommends itself to him is taking the necessary steps to rid himself of the desire to smoke. This will have to be done over a period of time if it is to be cost-effective.

I submit that the notion of being fully informed poses no threat to securing the second link to motivation. The changes that result from the **process** of becoming fully informed pose no threat since they are only temporary. Someone with a motivational set ABCD will require different motivational sets to appreciate various information. He might, after the process, end up with motivational set XYZ. But we are not concerned with what someone who has motivational set XYZ wants even if he is fully informed. So we ask instead what the fully informed person with motivational set ABCD wants. This person will rely on memory to appreciate information that initially required a different motivational set. We have, then, the same person only fully informed. Of course, this person's motivational set will likely undergo some changes after being fully informed (he might even end up with motivational set XYZ). But these changes, the changes that result from **being** fully informed (as opposed to the process of becoming fully informed),

will be welcome. Taking the necessary steps to acquire these new motivations will be beneficial while not doing so will be costly.

But Rosati believes that the problem of appreciation goes deeper than simply posing a threat to the second link to motivation. Her third worry in "Persons, Perspectives, and Full Information Accounts of the Good", recall, is that the notion of being fully informed is incoherent. She writes that

proponents of Ideal Advisor views attempt to resolve the problem of appreciation by requiring that a person undergo whatever education and experiences are necessary to render information fully vivid. The problem of appreciation does not arise, however, simply because we need experiences or education for information to register. It has its root in what it is like to be a particular person. (Rosati 1995b: 307)

Recall from the last chapter that the full information account of well-being must be able to address properly questions regarding what kind of person to be. It seems, then, that our fully informed selves must be able to survey the various lives open to us and the various kinds of people we might become. We have been concerned, in this chapter, with whether or not it is possible for an agent with a particular motivational system to appreciate all possible bits of information. We have seen that memory provides us with a means for appreciating information that otherwise would require a different motivational system. But can we extend this strategy so that it allows an agent to appreciate not only every possible bit of information, but also every possible life that she might live - every possible perspective?

A perspective, though, seems like nothing more than a disposition to be able to appreciate certain bits of information (and not others). If we can make sense of appreciating individual bits of information via memory (and I think that we have seen that we can) then there should be no problem in making sense of appreciating particular bundles of information via memory. If the atheist can now

appreciate that having one's prayers answered is uplifting, it seems she can also appreciate that the belief that heaven awaits us can get one through tough times. And it so it goes with all of the other individual bits of information that an atheist with no past religious experience might fail to appreciate. The atheist can now appreciate, via memory, what it was like to be a theist. Appreciating different perspectives does not seem to be a problem.

Rosati asks us to consider what it is like to compare thirty-one different flavors of ice cream. (Rosati 1995b: 315) In order for us to make an informed decision regarding our preference in this case, it seems that we must be familiar with all thirty-one flavors. Having tasted each flavor in the past will suffice. As Rosati remarks, "remembering how a certain flavor of ice cream tasted is curiously like tasting it now." (Rosati 1995b: 315)

But comparing different lives, according to Rosati, is radically unlike comparing different flavors of ice cream. Consider a person who is obtuse. This person will react to a friend in distress not with sympathy but by pointing out that she is behaving irrationally. Can such a person know what it is like to be a sympathetic person? Rosati writes that

[i]f she cannot now see another's distress as calling for a sympathetic response, it is unclear how she could accurately remember what it was like to so view it. In order to recall what it is like to see another's distress as calling for a sympathetic response, it would seem that she must now be capable of seeing another's distress in that way. (Rosati 1995b: 320)

So comparing different lives is unlike comparing different flavors of ice cream in that the latter does not depend upon one's perspective. "The obtuse person's obtuseness will not influence whether she enjoys chocolate ice cream, but it will influence whether she responds to a crying friend with sympathy or irritation." (Rosati 1995b: 316) Perspectives that are at odds with each other cannot be

compared. Some perspectives cannot be appreciated from within others. Thus the notion of being fully informed about all perspectives is incoherent.

But, as Rosati points out, "[o]btuseness may seem a special case, since it involves an incapacity to grasp certain information." (Rosati 1995b: 320) The obtuse person cannot appreciate the fact that her friend's condition calls for a sympathetic response because not being able to appreciate such a fact is precisely what we mean by being obtuse. Someone who is obtuse suffers from a cognitive deficiency and the full information account corrects for such deficiencies. To be sure, someone who is fully informed might still desire to respond to her friend's situation by pointing out to her that she is acting irrationally. But this fully informed person will appreciate that the situation calls for sympathy and so she will not be acting from an obtuse perspective.

What Rosati needs, in order to show that the notion of being fully informed is incoherent, is a clear case where two perspectives cannot be simultaneously appreciated where this is not simply due to a lack of knowledge. Unfortunately, she provides us with no such example. She provides us with only one other example. She asks us to consider a person who loses her sense of humor. Can such a person remember what it was like to have a sense of humor? According to Rosati, "[i]f she can no longer find things funny, it is hard to see how she could remember what it is like to find things funny. When we remember what it was like when we found something funny, we are likely to find it funny again." (Rosati 1995b: 320)

Perhaps in this case it is not as tempting to suggest that someone who lacks a sense of humor is missing certain information (i.e., lacking knowledge). It is controversial to suppose that finding something funny, or not finding it funny, constitutes a cognitive shortcoming. But what is not clear is Rosati's claim that we cannot remember what it was like to find something funny which we do not

now find funny. Most of us have a sense of humor which develops over time. Many things that we once found funny do not seem funny to us now. Certain things that once seemed funny might even offend us now (e.g., certain tasteless jokes). But I see no reason to suppose that we cannot accurately remember what it was like to find these things funny. We might even be able to explain exactly why we did find these things funny by citing various circumstances in which we found ourselves (e.g., the social position we held, our character traits, etc.). One might, for example, try to explain to you why jokes which degrade certain ethnic groups were funny to her as a young teenager. She might explain that as a teenager she was driven by a need to fit in with her circle of friends. She strongly identified with this group whose members were all from similar ethnic backgrounds. As a result of this strong identification she took differences in other people (i.e., people from different ethnic backgrounds) as something to poke fun at. Of course, if she gives you such an explanation you might not appreciate what it is like to find these things funny (you may have never found such things humorous). But she has had first hand experience with the circumstances that she cites and so it would seem that she can appreciate what it is like to find these things funny even if she no longer finds them funny.

This type of consideration also applies to Rosati's first case. Suppose that instead of an obtuse person we imagine a person who is unsympathetic. This person will respond to her friend with irritation not because she does not recognize that the situation calls for sympathy (like the obtuse person), but because she simply does not have the desire to respond sympathetically. There is no lack of knowledge in this case. But once again it is not at all clear that the unsympathetic person cannot appreciate what it was like to be sympathetic. We can imagine a sympathetic person who has become cold and uncaring through bitter experience. There seems to be no conceptual difficulty in picturing this

person accurately remembering what her life was like prior to this experience. Indeed, much like the case involving the ethnic jokes, she might even be able to explain to someone what it was like to be sympathetic, citing various circumstances in which she found herself back then. Of course, given only this explanation, this person might not appreciate what it is like to be sympathetic (perhaps he has always been unsympathetic). But the one providing the explanation will appreciate what it is like to be sympathetic since she has had first hand experience with the circumstances that she cites.

Rosati's examples, then, do not show that the notion of being fully informed about all perspectives is incoherent. Perspectives that are at odds with each other because of a lack of information (e.g., sympathetic vs. obtuse) pose no threat since the full information account corrects for such cognitive deficiencies. And she provides us with no reason to believe that perspectives which otherwise seem at odds with each other (e.g., sympathetic vs. unsympathetic) cannot both be appreciated from one perspective provided that the agent has experienced both. Many of us, it seems, can now remember what it was like to view the world from an entirely different perspective.

So it seems that we can make sense of the notion of fully informing someone while holding her motivational system fixed. An agent could (theoretically) be informed about all of the possible lives open to her, relying on memory to appreciate that information which otherwise would not register given her current motivational system. Since we ourselves are the ones who occupy the ideal standpoint, and any changes to our motivational system that are the result of being fully informed are welcome, the judgments of our fully informed selves will have normative authority for us and so the second link to motivation will be effected.

Chapter 4: Conclusion

We have examined several criticisms of the full information account of well-being which purport to show that the judgments of our fully informed selves lack normative authority. According to Rosati, our fully informed selves do not make ideal advisors. Let us now review Rosati's arguments for this view and the reasons we found for rejecting them.

We began, in chapter 2, by looking at Rosati's concern that naturalism in general, and the full information account of well-being in particular, does not allow us to address properly questions regarding what kind of person to be. Her argument for this claim went something like this:

- (1) Humans are, fundamentally, self-inventors.
- (2) The good for A must be suited to A *qua* A.
- (3) Therefore, the good for humans must be suited to us as self-inventors.
- (4) The naturalist's account of 'good for a person' treats the person as identical with her current motivational system and so is not suited to us as self-inventors.
- (5) Therefore, the naturalist's account of 'good for a person' fails.

We saw that the weak link in this argument is premise (4). For although it is true that the naturalist's account of 'good for a person' does identify a person with her current motivational system, we must take into account that one's motivational system includes the personal ideals that one now holds. We may choose to change ourselves in light of these personal ideals. Indeed, on the full information account, changing ourselves in light of these ideals (if that is what we desire for our current selves once we are fully informed) is what is good for us. So the naturalist's account of 'good for a person, at least on the full information account, seems perfectly suited to us as self-inventors.

And we have seen that Rosati really has no other alternative but to treat personal ideals as a part of one's motivational system. She cannot treat them as the externalist does. The externalist claims that a personal ideal might have value for a person, where 'value' does not simply consist of the person's desires. That way, we might appeal to things like personal ideals in order to explain certain desires. For example, Sandy's desire to be like Madelyn might be explained by her perception, or recognition that being like Madelyn is valuable. I have suggested that such explanations are unilluminating unless some positive account of value is given. At any rate, Rosati cannot endorse such a strategy. She believes that an account of the good must satisfy a double link to motivation to ensure that we are not alienated from our good. It is not enough that X is something that is capable of motivating us. We have to care about the conditions under which we desire X. On externalist accounts, we might be radically alienated from our good. X might be good for a person even if X completely fails to move that person under any circumstances.

Rosati could, instead, endorse a form of internalism which claims that the good for a person is such that it must be capable of motivating that person but is not simply constituted by her desires (Darwall calls this non-constitutive existence internalism). But we found this view to be implausible. If we are to maintain that certain things (other than desires) are intrinsically motivating, or, as Mackie puts it, have to-be-pursuedness built into them, we will need an account of how it is that we come to know about these things, and how exactly knowing about them entails being motivated by them. Rosati provides us with no such account.

We concluded, in chapter 2, that Rosati's attempt at a sophisticated version of Moore's open question argument fails. Her strategy is to find a question that was genuinely compelling (and not just barely possible) which could

be addressed using our everyday concept of 'good for a person', but could not be addressed using the analysis provided by the full information account. In other words, it is not enough merely to point out that we can ask "This is something that I would want if fully informed, but is it good for me?". We can ask the same kind of question with respect to any *a posteriori* identity claim. For example, one might ask "This is H₂O, but is it water?". Such questions are possible to ask, but they are not genuinely compelling. According to Rosati, the genuinely compelling questions that cannot be addressed on the full information account are questions regarding what kind of person to be. But we have seen that by treating personal ideals as a part of our motivational system, the full information account can address such questions.

In chapter 3, we examined Rosati's concern with the notion of being fully informed that the full information account relies on. There were three related worries:

- (1) The person who emerges after the process of becoming fully informed may not be you.
- (2) Your fully informed self (even if it is you) might not have the qualities suitable for an ideal advisor.
- (3) The notion of being fully informed about all possible perspectives is incoherent.

The first two worries stem from Rosati's concern that the full information account does not satisfy a strong enough form of internalism, namely, what she calls "two-tier internalism". If X is good for me then X must be something that is capable of motivating me (the first link to motivation) and I must care, under ordinary optimal conditions, about the circumstances under which I would desire X (the second link to motivation). But we will have to change markedly in order to become fully informed and so, on the full information account, the second link

to motivation will not be effected. The judgments of our fully informed selves will lack normative authority for us.

There are two types of changes that an agent's motivational system might undergo. The two types of changes correspond to Rosati's first two worries. First, there are those changes that an agent must undergo in order to appreciate certain information. These are the changes that are **due to the process of becoming fully informed**. Second, there are those changes that are **due to being fully informed**.

We saw that the first type of change poses no real threat. A person can rely on memory to appreciate those bits of information that initially required a different motivational system. With an extended memory, she will be able to appreciate all information from within her current motivational system. She will then be the same person, only fully informed. And so the process of becoming fully informed does not threaten the second link to motivation.

Next, we dealt with the second type of change. We found that the changes that a person's motivational system might undergo due to being fully informed will be welcome. Acting on (or taking the necessary steps to acquire) desires that one would have if fully informed will be beneficial, while not acting on these desires (or acting on desires that one's fully informed self would not have) will be costly. So these changes do not threaten the second link to motivation either.

Finally, we examined Rosati's claim that the notion of being fully informed about all possible perspectives is incoherent. We found that the examples she provides in support of this claim are unconvincing. She provides us with no clear case where two perspectives cannot be simultaneously appreciated where this is not simply due to a lack of knowledge.

So it seems that the judgments of our fully informed selves do have normative authority for us. And this suggests that our fully informed selves make ideal advisors. These considerations provide strong support for the full information account of well-being. There is, of course, much more that needs to be done in order to provide a complete defense of the full information account. For instance, we have not considered the type of worry raised by Velleman. We have simply assumed that the concept "that which one's fully informed self would endorse" is an empirically determinate concept. Recall, from chapter 1, that this is what I take to be the second serious worry that the full information account faces. It seems, though, that we have made considerable progress in meeting the first serious worry. I believe that dealing with Rosati's worries has helped to demonstrate the normative authority of our fully informed judgments.

References

- Brandt, Richard B. A Theory of the Good and the Right. Clarendon Press of Oxford University Press, 1979.
- Brink, David O. Moral Realism and the Foundations of Ethics. Cambridge University Press, 1989.
- Butler, Joseph. Five Sermons Preached at the Rolls Chapel. The Liberal Arts Press, 1950.
- Darwall, Stephen L. "Internalism and Agency," Philosophical Perspectives, 6. Ethics, 1992: 155-174.
- Frankfurt, Harry G. "Freedom of the Will and the Concept of a Person," The Journal of Philosophy, Vol. 68, 1971: 5-20.
- Gibbard, Allan. "A Noncognitivist Analysis of Rationality in Action," Social Theory and Practice, Vol. 9, Nos 2-3, Summer-Fall 1983: 199-221.
- Gibbard, Allan. Wise Choices, Apt Feelings: A Theory of Normative Judgment. Harvard University Press, 1990.
- Gleitman, Henry. Psychology: Third Edition. W.W. Norton & Company, Inc., 1991.

- Harman, Gilbert The Nature of Morality: an Introduction to Ethics. Oxford University Press, 1977.
- Hobbes, Thomas. Leviathan. Hackett, 1994.
- Loeb, Don. "Full Information Theories of Individual Good," Social Theory and Practice, Vol. 21, No. 1, Spring 1995: 1-30.
- Mackie, J.L. Ethics: Inventing Right and Wrong. Penguin Books, 1977.
- Mill, John Stuart. On Liberty and Other Essays. Oxford University Press, 1991.
- Moore, G.E. Principia Ethica. Cambridge University Press, 1903.
- Railton, Peter. "Moral Realism," The Philosophical Review, XCV, No. 2, April 1986a: 163- 207.
- Railton, Peter. "Facts and Values," Philosophical Topics, Vol. XIV, No. 2, Fall 1986b: 5-31.
- Railton, Peter, "Naturalism and Prescriptivity," Social Philosophy and Policy, Vol. 7, Issue 1, 1990: 151-174.
- Rosati, Connie S. "Naturalism, Normativity, and the Open Question Argument," Noûs, 29, 1995a: 46-70.

Rosati, Connie S. "Persons, Perspectives, and Full Information Accounts of the Good," Ethics, 105, January 1995b: 296-325.

Rosati, Connie S. "Internalism and the Good for a Person," Ethics, 106, January 1996: 297-326.

Shaver, Robert. "Sidgwick's False Friends," Ethics, January 1997.

Sidgwick, Henry. The Methods of Ethics (seventh edition). Hackett, 1981.

Sobel, David. "Full Information Accounts of Well-Being," Ethics, 104, July 1994: 784-810.

Velleman, J. David. "Brandt's Definition of 'Good'," The Philosophical Review, 97, 1988: 353-371.