



## Full Information Accounts of Well-Being

David Sobel

*Ethics*, Volume 104, Issue 4 (Jul., 1994), 784-810.

---

Your use of the JSTOR database indicates your acceptance of JSTOR's Terms and Conditions of Use. A copy of JSTOR's Terms and Conditions of Use is available at <http://www.jstor.org/about/terms.html>, by contacting JSTOR at [jstor-info@umich.edu](mailto:jstor-info@umich.edu), or by calling JSTOR at (888)388-3574, (734)998-9101 or (FAX) (734)998-9113. No part of a JSTOR transmission may be copied, downloaded, stored, further transmitted, transferred, distributed, altered, or otherwise used, in any form or by any means, except: (1) one stored electronic and one paper copy of any article solely for your personal, non-commercial use, or (2) with prior written permission of JSTOR and the publisher of the article or other text.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

*Ethics* is published by University of Chicago Press. Please contact the publisher for further permissions regarding the use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/ucpress.html>.

---

*Ethics*

©1994 University of Chicago Press

JSTOR and the JSTOR logo are trademarks of JSTOR, and are Registered in the U.S. Patent and Trademark Office. For more information on JSTOR contact [jstor-info@umich.edu](mailto:jstor-info@umich.edu).

©2001 JSTOR

# Full Information Accounts of Well-Being\*

*David Sobel*

Amartya Sen defines 'consequentialism' as the view that "an action [rule, motive, cooperative scheme, etc.] A is right if and only if the state of affairs x resulting from A is at least as good as each of the alternative states of affairs that would have resulted respectively from the alternative feasible acts [rules, motives, cooperative schemes, etc.]."<sup>1</sup> He goes on to define 'outcome utilitarianism' as the view that "any state of affairs x is at least as good as an alternative state of affairs y if and only if the sum total [or average] of individual utilities in x is at least as large as the sum total [or average] of individual utilities in y."<sup>2</sup>

If we combine consequentialism and outcome utilitarianism, we get the popular view that an act (rule, etc.) is right if and only if it produces at least as large a sum total (or average) of individual utilities as any other available act (rule, etc.).<sup>3</sup> Let us call the combination of consequentialism and outcome utilitarianism as defined above the standard consequentialist position or SCP. The SCP has it that an act (rule, etc.) derives its moral status from its effectiveness in furthering the average or total amount of individuals' well-being.

A defender of the SCP can be thought of as having a two-step task. First, she must isolate and commensurate an individual agent's well-

\* I gratefully thank Justin D'Arms, Stephen Darwall, Janice Dowell, Allan Gibbard, Daniel Jacobson, James Joyce, Peter Railton, Connie Rosati, David Velleman, and especially Elizabeth Anderson for their constructive comments on many of the issues raised in this article. I am also grateful to Richard Arneson and an anonymous reviewer for *Ethics* for their helpful criticisms of earlier drafts of this article.

1. Amartya Sen, "Utilitarianism and Welfarism," *Journal of Philosophy* 76 (1979): 464. Actually, Sen only offers there a characterization of act consequentialism. I intend to expand the characterization to include all of consequentialism by adding "rule, motive, cooperative schemes, etc."

2. *Ibid.*, p. 464.

3. It should be noted that fixing our terms in this way allows the possibility of an outcome utilitarian not being a consequentialist. One could accept the outcome utilitarian's account of goodness while defining the right independently from the good.

being. The defender of the SCP is committed to this commensurating first step because her aspiration is to be able to plug in the resulting account of an agent's well-being into the second step such that the instruction to maximize the sum or average of well-being provides enough content to serve as the basis for a theory of morality.<sup>4</sup> This aspiration requires the SCP to rely not only on the claim that interpersonal comparisons of utility are possible but also on the prior claim that the measure of the value to the agent of the various aspects of her good can be reduced to common units (perhaps with arbitrary origin and scale).

In the second step of the SCP theorist's argument she must defend the claim that it is morally appropriate to maximize the average or sum of individuals' well-being. It is this second step of argumentation that has received the lion's share of recent philosophical attention. Nozick has argued against maximizing individuals' well-being across persons in favor of a "side constraints" picture of morality. Rawls has charged that the SCP "fails to take seriously the distinction between persons" partially because it gives no intrinsic weight to issues of distribution. Williams has found that the SCP constitutes an attack on an agent's integrity because it requires an agent to regard even her most important projects as appropriately abandonable whenever the sum of individuals' well-being would be higher if she did so. Arrow has argued that all intelligible methods of moving from only the collection of individuals' utility rankings to a social welfare function must be, in some cases, either imposed (meaning that society will be said to prefer

4. In his candid and helpful article, "Pluralism, Determinacy, and Dilemma," *Ethics* 102 (1992): 720–42, Peter Railton argues that, "for the utilitarian, then, the amount of determinacy that is possible in evaluations of moral rightness will be a reflection of how things are in the domain of intrinsic value. Now it is unquestionably true that some utilitarians have held rather optimistic views about the determinacy of intrinsic value, its univalence, its measurability, and comeasurability with arbitrarily high degrees of accuracy, and so on. But these are questions in the theory of nonmoral value, not, so to speak, in utilitarian morality proper. . . . Nothing prevents the most orthodox utilitarian from discovering such pluralism—even 'fragmentation'—in the domain of nonmoral value" (pp. 735–36). While I agree that the SCP can allow that there are some incommensurabilities between values, it is nonetheless clear that the more pervasive the incommensurabilities, the less work the SCP framework is capable of doing. The SCP would seem to be committed to finding enough commensurability in well-being to allow the second aggregating stage to yield a substantial theory of morality. If well-being is found to contain systematic incommensurabilities, then the second stage of the SCP will be relatively mute. The instruction to maximize only offers direction after significant commensuration has been accomplished. Thus, it is a crucial and worrisome presupposition of the SCP that well-being can be significantly commensurated. Of course, the defender of the SCP who finds significant incommensurability might want to say that morality's demand that we maximize well-being actually gives us less instruction than we might have thought. But even if we can tolerate a fair amount of indeterminacy in our moral theory it would remain to be shown that the indeterminacy in well-being would lead to just that indeterminacy in the moral sphere that we could tolerate.

X to Y no matter what the individuals in that society prefer) or dictatorial. Arrow relied on a complaint that economists have made for years: that the notion of maximizing the sum or average of individuals' interests was unintelligible due to the impossibility of interpersonal comparisons of utility.<sup>5</sup>

Each of these famous attacks of the SCP, together with many others in the same vein, has been widely discussed. But no comparable group of debates which challenge the adequacy of the first step in the SCP exists. Often, even in works like the above which sharply criticize the SCP, it seems to be assumed with little discussion that a workable account of an agent's well-being adequate to the SCP's needs can be constructed. In addition, when the commensurability of an individual's well-being has been challenged there is too often no discussion of the most prominent and sophisticated strategies of value commensuration.<sup>6</sup> The battleground between the SCP and those who oppose it has been the second of the SCP's two-step strategy.

I believe that the first step of the SCP's project, where the attempt is made to determine and commensurate an agent's intrinsic well-being, is itself quite problematic. In this article I will criticize two quite popular strategies of commensurating well-being, namely, decision theoretic and full information accounts, which seem to offer much promise of producing an account of well-being which could play the role of the first step in the SCP's two-step proposal. I intend to offer arguments which shake the reader's confidence that these theories can be made adequate to that task.<sup>7</sup> My focus will be on full information

5. See Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic, 1974); John Rawls, *A Theory of Justice* (Cambridge, Mass.: Harvard University Press, Belknap, 1971); Bernard Williams, "A Critique of Utilitarianism," in *Utilitarianism: For and Against*, ed. J. J. C. Smart and Bernard Williams (Cambridge: Cambridge University Press, 1973); Kenneth Arrow, *Social Choice and Individual Values*, 2d ed. (New Haven, Conn.: Yale University Press, 1963); and Lionel Robbins, "Interpersonal Comparisons of Utility: A Comment," *Economic Journal* 43 (1938): 635-41. Actually, Robbins's position on the interpersonal comparability of utility is more complex than is standardly acknowledged.

6. Michael Stocker, in *Plural and Conflicting Values* (Oxford: Oxford University Press, 1990), focuses on hedonistic strategies of value commensuration and thus offers little discussion of what I take to be the most plausible options available. David Wiggins, in "Weakness of Will, Commensurability and the Objects of Desire," in his *Needs, Values, Truth*, 2d ed. (Oxford: Basil Blackwell, 1991), argues against the prospects for significant value commensuration by assuming that commensuration can only take place if it can be established that there is but one genuine value. However, the commensurator only needs to suggest that there is an accurate way of measuring and reducing to common units the respective merits to oneself of different values, not that they are all, at bottom, one.

7. There are obviously many theories of intrinsic value commensuration that I will not be discussing here, most notably the many versions of hedonism and objective list theories. I have no intention of making a full case in this article that no theory of well-being could adequately fulfill the role of the first step of the SCP. I find James Griffin's

accounts as I am to some extent using the shortcomings of decision theory in filling this role to motivate the full information account.

## DECISION THEORY

I will discuss, under the heading of decision theory, a family of theories which purport to offer a way of commensurating an agent's well-being without significantly idealizing the information available to the agent in question. Typically decision theorists will present several axioms, alleged to be constraints on rational preference orderings, and prove that any preference ordering that obeys these axioms will be able to be represented as a utility function. However, even if we do not take issue with the decision theorist's debatable axioms, we are given little reason to think that the utility function which can be generated from the agent's axiom obeying preferences corresponds to the agent's well-being.

Perhaps many who use and further decision theory do not conceive of it as an attempt to determine an agent's well-being. I do not here criticize understandings of decision theory other than as a theory of well-being. Decision theory is sometimes thought of as analogous with logic. Ensuring that one's beliefs are logically consistent does not ensure that one's beliefs are true. Nonetheless, being internally logically consistent is normative for beliefs. Similarly, a decision theorist might argue, adjusting one's preferences such that they obey the standard axioms does not ensure that one's preferences correspond with one's well-being, but nonetheless obeying the axioms is necessary for preferences to be fully coherent. Money pump arguments and the like can be seen as an attempt to demonstrate an internal contradiction in preferences that fail to obey the axioms.

However, absence of internal contradictions in one's preferences is not sufficient, as we will see, to ensure that we prefer the option that is actually better for us even if we accept a deeply subjective account of an agent's well-being. It is too frequently unclear exactly what interpretation an author has of the utility function which can be generated from an agent's axiom-obeying preferences. Nevertheless, it is clear that many who use decision theory, especially social choice theorists, implicitly or explicitly interpret decision theory as a theory of well-being.<sup>8</sup>

---

criticisms of objective list accounts persuasive (see his *Well-Being* [Oxford: Clarendon, 1986], pp. 40–55). I am working on my own criticisms of the various forms of hedonism.

8. For example, Kenneth Arrow's "Extended Sympathy and the Possibility of Social Choice," in *Collected Papers of Kenneth J. Arrow* (Cambridge, Mass.: Harvard University Press, 1983), vol. 1, p. 148. See also Amartya Sen's "Description as Choice," and "Rational Fools," both in his *Choice, Welfare, and Measurement* (Cambridge, Mass.: MIT Press, 1982). John Broome's "Utility," *Economics and Philosophy* 7 (1991): 1–12, is useful in sorting out the importantly different uses of the word 'utility' in economics and ethics.

There are several different ways in which one's unrestricted actual preferences<sup>9</sup> which are constructed without idealized information, even if fixed up to obey the axioms, are likely to stray from one's true interests. I will insist here only on the most obvious way this can happen. Our actual preferences can be based on false beliefs or lack of propositional or experiential knowledge in the sense that were we to be disabused of our false belief, come to know some fact, or have some relevant experience firsthand we would not value the option as we initially did. Typically in such cases it is clear that our initial estimation of the value of the option to us was distorted. For example, Taniquill might believe that the life of a rock star (the adoring fans, the spotlight, the volume) would be so exciting and artistically challenging that she vastly prefers such a life to the life of a violinist in a symphony. However, it could (and often does) turn out that that which seemed so exciting and challenging to Taniquill from a distance is seen to actually be, upon closer scrutiny and with firsthand familiarity with the option, empty and boring. Similarly the symphony, which seemed so stuffy and lifeless, could, with firsthand familiarity, come to be seen as challenging and creative. This is not a case in which at first being a rock star was better for Taniquill and now being a classical musician is better for her. Unbeknownst to Taniquill, she would actually have preferred the life of the orchestra member at all stages of her life if she had had firsthand experience with the two options.

This simple story of the way in which one's actual preferences which are formed with less than full information can, with increased information, be revealed to fail to track one's well-being is so familiar from everyday experience that it is surprising that some decision theorists overlook it. John Broome has suggested that perhaps economists inappropriately extend the convincing liberal thought that there is a sphere of self-regarding acts in which a person should, as a matter of public policy, be free to live in accordance with her own decisions, into the clearly false doctrine that whatever a person voluntarily

9. Scanlon adopts this useful categorization of the typical decision theoretic approach in "Value, Desire, and Quality of Life," in *The Quality of Life*, ed. Martha Nussbaum and Amartya Sen (Oxford: Clarendon, 1993). The term 'unrestricted' implies that all one's preferences are included in the utility ranking. In other accounts the set of preferences which make up the utility function are restricted to, e.g., nonmoral or nonantisocial preferences. These two restrictions respond to importantly different pressures. The former is an attempt to accommodate the fact that one can intrinsically want to promote the well-being of others even when doing so does not promote the intrinsic value of one's life to oneself. The latter restriction attempts to respond to the thought that not all of one's preferences, even if they are correlated with one's well-being, make claims on others in proportion to their strength even at some initial stage of deliberation.

chooses must therefore be best for her.<sup>10</sup> But, as Mill's writings demonstrate, one can fully respect the liberal's claim above while admitting that people can make mistakes about where their own good lies.

I believe that decision theory is best interpreted not as a theory of well-being but rather as a theory of choice worthiness given one's information. Of course such a choice-worthy option might still disappoint. But the best one can do in making choices is make the best use of the information that one has (where this might entail gathering more information). That the option was disappointing need not overturn the claim that given one's information the option was the most choice-worthy. However, it would tend to overturn the claim that the option best served one's interests. Recognizing the distinction between the virtues appropriate to a theory of choice worthiness and those appropriate to a theory of well-being could prevent confusion and help ethicists (who are typically searching for theories of well-being) recognize the genuine merits and problems of decision theory as a theory of choice worthiness. However, in this article I am assessing theories for their adequacy as an account of what intrinsically makes one's life go well, and it is clear, as the example with Taniquill shows, that decision theory is not fit for this role.

One might well argue that well-being is not the best candidate for playing the role of the first step in the SCP. This would open the possibility that the decision theoretic framework provides the notion of utility that is appropriately maximized, at least in social policy decisions. One might recognize that persons can be in significant error about what is good for them yet argue that it is nonetheless appropriate for the state to respond to the actual preferences of individuals in shaping policy. (Although the state surely may not ignore the presence of toxins in the public water supply, e.g., simply because most people are unaware of this health hazard.) Some decision theorists appeal to antipaternalist considerations to support this contention. Although the decision theoretic notion of utility has its attractions (at least when concerning competent adults) in the context of social policy, it is surely not this notion of utility which typically should be attended to in other contexts. In parenting or advising a friend we often should not consult the actual preferences of the person whose good we are trying to promote. The decision theoretic notion of utility is clearly inadequate as the fundamental object of moral concern in all contexts. Well-being

10. John Broome, "Choice and Value in Economics," *Oxford Economic Papers* 30 (1978): 313–33. See also Allan Gibbard, "Interpersonal Comparisons: Preference, Good, and the Intrinsic Rewards of a Life," in *Foundations of Social Choice Theory*, ed. J. Elster and A. Hylland (Cambridge: Cambridge University Press, 1986), for a helpful discussion.

has a better claim to be fit for this role. I will limit my attention to attempts to commensurate it.

### FULL INFORMATION ACCOUNTS

The full information accounts that I will consider share the decision theorist's aspiration of determining and commensurating an agent's well-being.<sup>11</sup> The full information account is an attempt to overcome the kind of objection raised against decision theory in the previous section by idealizing the information and reasoning ability available to the agent. By counterfactually granting the agent unconstrained access to information and freeing her from all errors of instrumental rationality, the worry offered above against decision theory might be quieted. But as we will see, the full information theorist's attempts to rectify the inadequacies of decision theory are highly problematic.

The full information account involves specifying a privileged epistemic standpoint that is sufficiently knowledgeable about all the different goods or lives an agent could have such that the preferences of the knowledgeable agent (or some subset of them) definitively determines the value of those goods or lives to the nonidealized agent. Mill sparked this project by rejecting the Benthamite identification of an agent's good with the sensation of pleasure (or any other substantive end) and replacing this account with his "competent judges' test." Mill writes, "Of two pleasures, if there be one to which all or almost all who have experience of both give a decided preference, irrespective of any feeling of moral obligation to prefer it, that is the more desirable pleasure."<sup>12</sup> The root idea behind this test is that if a person intrinsi-

11. One must not confuse such an approach with those theories which derive the appropriate moral judgment, rather than the agent's own well-being, directly from the determinations of the informed agent. Roderick Firth's "Ethical Absolutism and the Ideal Observer," *Philosophy and Phenomenological Research* 12 (1952): 317-45, is the modern classic of the former approach.

12. J. S. Mill, *Utilitarianism* (Indianapolis: Hackett, 1979), chap. 2. David Hume, "Of the Standard of Taste," in his *Essays, Moral, Political and Literary*, ed. E. Miller, Liberty Classics (Indianapolis: Liberty, 1985), provides an earlier and similar attempt to understand the criticizability of evaluative responses in cases in which the response does not seem straightforwardly apt for truth or falsity. He writes, "If in the sound state of the organ, there be an entire or a considerable uniformity of sentiment among men, we may thence derive an idea of the perfect beauty; in like manner as the appearance of objects in day-light, to the eye of a man in health, is demonstrated their true and real colour, even while colour is allowed to be a mere phantom of the senses" (p. 234). A few pages later he argues that the ability to notice subtle variations in a thing is a sure sign of the soundness of the organ used. In developing this delicacy of taste, Hume finds it to be of the utmost importance that one engage in "practice in a particular art, and the frequent survey or contemplation of a particular species of beauty"—in short, that one acquire "experience in those objects" under consideration. Notably both Mill and Hume insist that agreement between competent judges is crucial to substantiate the claim that one good is higher than another. They are not in these instances focusing



cally prefers X to Y, independent of moral considerations, while fully acquainted with both options and in the “cool hour,” then X is more conducive to the agent’s well-being than Y no matter what other properties X and Y have.

This account has a tremendous amount going for it as an account of the good for a person. It makes sense of the common observation that experimenting with alternatives frequently increases the agent’s knowledge of her good. The account explains why we often regard our authority concerning the whereabouts of our own good as limited primarily by our lack of information and experience. It also respects the seemingly powerful thought that if a person cannot be brought to value some option intrinsically then it is not intrinsically valuable to him because it would be “an intolerably alienated conception of someone’s good to imagine that it might fail in any way to engage him.”<sup>13</sup>

(However, it is easy to overestimate the force with which this last consideration tells for the full information account if we continue to ignore an important distinction in kinds of internalism. Consider the difference between what I will call “who” and “how” internalism. The former asserts a constraint on what could be in an agent’s intrinsic interests by demanding that it be the agent whose good we are investigating that gets idealized. This internalism places no constraints at all on the idealization process so long as the process is done to the agent whose good is in question. This yields an uninterestingly weak version of internalism since it amounts only to the claim that something is not part of my good if I am constitutionally incapable of caring about it. Any otherwise plausible account of well-being will be able to accommodate this claim. In contrast, “how” internalism adds to the above by imposing constraints on acceptable methods of idealization. Here the idealization process itself is subject to the constraint that the agent in question find the proposed idealization process to enhance greatly their ability to see their own good. Here, how the agent is to be idealized is subject to the internalist constraint of having to hook up with the agent’s actual motivational set. But “how” internalism is implausibly strong. Some people are unreasonable or misguided precisely in failing to see that certain alterations to themselves would place them in better circumstances to see their own good. Thus the interesting and true version of internalism has not yet been identified, let alone shown to do work particularly for the full information account.)

---

on determining an individual’s well-being but, rather, making sense of a kind of objectivity that judgments of taste or value can have across persons. Thus what I call the “root idea” of Mill’s competent judges’ test in the text is what seems to have been made of Mill’s test by those interested in commensurating an individual’s well-being.

13. Peter Railton, “Facts and Values,” *Philosophical Topics* 14 (1986): 5–29, p. 9.

Sidgwick modified Mill's account, expanding both the knowledge given to the perspective from which the agent is to assess the options and the options to be simultaneously chosen between. He argued 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 realized in imagination at the present point in time."<sup>14</sup>

Sidgwick's proposal implies that there is a single vantage point from which we can accurately assess the value of all our possible futures as well as a single pro-attitude that measures the value of all the diverse kinds of goods. Later full information theorists have followed him in these implications. However, Mill's competent judges' test does not seem committed, considered by itself, to this claim since, for example, one need not accept that any single vantage point could be competent with the values of all of one's possible futures.

Richard Brandt greatly elaborated the framework that Sidgwick had put forward, offering the account as a reforming definition of rationality. "I shall pre-empt the term 'rational' to refer to actions, desires, or moral systems which survive maximal criticism and correction by facts and logic." "This whole process of confronting desires with relevant information, by repeatedly representing it, in an ideally vivid way, and at the appropriate time, I call *cognitive psychotherapy*." Brandt goes on to claim that what is good for an agent is determined by what she would "self-interestedly" desire after cognitive psychotherapy.<sup>15</sup>

But the Sidgwick/Brandt version of the full information account quickly runs into difficulties. Consider that our fully informed and rational self would never want more information for herself but we are firmly convinced that it can sometimes be intrinsically in our interests to gain information.<sup>16</sup> Further, our better selves no doubt have a

14. Henry Sidgwick, *The Methods of Ethics*, 7th ed. (Indianapolis: Hackett, 1981), pp. 111–12. But see Sidgwick's masterful development of this position from pp. 105–15.

15. Richard Brandt, *A Theory of the Good and the Right* (Oxford: Clarendon, 1979), pp. 10, 113, 329. A truly impressive and diverse list of contemporary ethicists have found the full information account of well-being congenial. In addition to those mentioned in the text above and below, there is R. M. Hare's explicit agreement with Brandt's position (see *Moral Thinking* [Oxford: Clarendon, 1981], pp. 101–5 and 214–16). See also D. Senior and N. Fotion, eds., *Hare and Critics* (Oxford: Clarendon, 1990), pp. 217–18. Griffin, pp. 11–17, takes exception to Brandt's view but ultimately endorses a version of the full information account. See also Rawls, pp. 407–24; David Gauthier, *Morals by Agreement* (Oxford: Clarendon, 1986), chap. 2; Stephen Darwall, *Impartial Reason* (Ithaca, N.Y.: Cornell University Press, 1983), pt. 2; and John Harsanyi, "Morality and the Theory of Rational Behavior," in *Utilitarianism and Beyond*, ed. Amartya Sen and Bernard Williams (Cambridge: Cambridge University Press, 1982), p. 55.

16. Peter Railton offers this critique in "Moral Realism," *Philosophical Review* 95 (1986): 163–207, p. 174.

refined palate and will hence highly value expensive wines that taste just like the cheaper stuff to us. It is implausible that one wine is much better for me than another when I cannot tell the difference (assuming that it is only the taste of the expensive wine which causes our idealized self to prefer it over the cheaper stuff).<sup>17</sup> The idealization process turns us into such different creatures that it would be surprising if the well-being of the two of us, my informed self and my ordinary self, consisted in the same things.

In response to problems such as these Railton has revised the account, proposing 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.”<sup>18</sup> The adoption of a “wanting to want” framework neatly eschews the implausible identification of interests between our informed and our ordinary self while retaining the insight that “the advice of someone who has this fuller information, and also has the deepest sort of identification with one’s fate, is bound to have some commending force.”<sup>19</sup> The “wanting to want” framework also seems to work nicely in cases of weakness of the will. In such cases we think that unless we can capture some sense in which the agent advises herself that the shunned option is the more valuable, we do not have a potential case of weakness of the will. Clearly in cases of weakness of the will it is the advice the agent gives herself about what to go for rather than what the agent actually goes for that has the better prospect of corresponding with what is in the agent’s interests. More recently David Lewis has taken up a somewhat similar account of value to an agent.<sup>20</sup>

17. Griffin offers this critique in *Well-Being*, p. 11. Surprisingly Griffin, in his otherwise exemplary book, confidently proceeds with the claim that value is measured by the strength of informed desire after having noted serious objections of this sort to the standard full information accounts. He seems to feel little pressure to develop the details of a version of the account which is not subject to such simple and telling critiques.

18. Railton, “Facts and Values,” p. 16.

19. *Ibid.*, p. 14. But consider that our idealized self could want our ordinary self to want X because the idealized agent knows that our ordinary self’s doing so will be instrumentally effective in bringing-about, albeit unintentionally, Y which is what the idealized agent finds to be best for our ordinary self. If we say that what is good for our ordinary self is what our idealized self wants our ordinary self to want we seem to misdescribe these cases of indirection. Perhaps it would be better to focus on the kind of life the idealized agent wants the ordinary self to have.

20. David Lewis, “Dispositional Theories of Values,” *Proceedings of the Aristotelian Society*, suppl. ser., 63 (1989): 113–37. A departure of Lewis’s view from Railton’s is Lewis’s claim that the idealized self requires only full imaginative acquaintance with the life, not facts and logic. One might suspect that this signals Lewis’s conviction that only states of mind and not states of the world can affect one’s intrinsic interests. Mark

There are at least four general strategies one could take in arguing that such an informed viewpoint is inadequate in capturing and commensurating what is in an agent's interests. First, one could argue that the notion of a fully informed self is a chimera. This would likely involve the worry that from the fact that any of the lives that one is to assess the value of must be in some sense available to one (otherwise it could not be a valuable life for one to live) it does not follow that all of them together must be available to one's consciousness. To make good this suggestion against the full information account one would have to provide reasons to think there are substantive worries about uniting the experience of all lives one could lead into a single consciousness.

Second, one could argue that even in cases in which an agent is adequately informed of the different life paths she is choosing between, there is no single pro-attitude, such as preferring, which appropriately measures the value of the diverse kinds of goods available to an agent. Rather, one might claim, different goods are appropriately valued in irreducibly different ways. Elizabeth Anderson has given a robust defense of this view. She argues that "we need a plurality of standards to make sense of the plurality of emotional responses and attitudes we have to things. The things that sensibly elicit delight are not generally the same things that merit respect or admiration. Our capacity for articulating our attitudes depends upon our understandings of our attitudes, which are informed by norms for valuation. To attempt to reduce the plurality of standards to a single standard, ground, or good-constituting property threatens to obliterate the self-understandings in terms of which we make sense of and differentiate our emotions, attitudes and concerns. To adopt a monistic theory of value as our self-understanding is to hopelessly impoverish our responsive capacities to a monolithic 'pro' or 'con' attitude or to mere desire and aversion."<sup>21</sup>

Third, one could argue that a vivid presentation of some experiences which could be part of one's life could prove so disturbing or alluring as to skew any further reflection about what option to choose. Allan Gibbard has suggested the example of "a more vivid realization of what peoples' innards are like" causing a "debilitating neurosis" which prevents me from eating in public. One could also imagine that vivid presentations of ways one could be tortured could lead to an

---

Johnston effectively critiques Lewis's account and finds that the incommensurability of values renders "conceptually amiss the consequentialist idea of maximizing value *simpliciter*" ("Dispositional Theories of Value," *Proceedings of the Aristotelian Society*, suppl. ser., 63 [1989]: 139–74, p. 165).

21. Elizabeth Anderson, *Value in Ethics and Economics* (Cambridge, Mass.: Harvard University Press, 1993), p. 5.

excessive concern for personal safety which could not be extinguished by reflection on the extremely small chance that such an event would occur even if one were less obsessed with personal safety. Gibbard summarizes that there are cases in which a person may reasonably worry that he is “an unreliable transformer of vivid realizations into rational desires, and so avoids dwelling on the facts he thinks would lead him astray.”<sup>22</sup>

Fourth, one could worry against naturalistic versions of the full information account that the purportedly naturalistically described informed viewpoint essentially invokes unreduced normative notions. This would undermine this full information theorist’s attempt to discover, after the SCP’s second step of aggregation, moral facts lurking within less suspicious ontological entities. Worse, this would undermine the project of offering an account of an agent’s well-being which does not presuppose substantive values beyond those imposed by instrumental rationality.<sup>23</sup> In this vein one could argue that no acceptably naturalistic account of “ideally vivid information” or “one’s non-moral good” can be produced.<sup>24</sup>

22. Allan Gibbard, *Wise Choices, Apt Feelings* (Cambridge, Mass.: Harvard University Press, 1990), pp. 20–21. Gibbard’s criticism is leveled specifically against Brandt’s version of the full information account. He discusses full information accounts on pp. 18–22 and 183–88. Gibbard’s first presentation of this type of criticism is in “A Non-Cognitivist Analysis of Rationality in Action,” *Social Theory and Practice* 9 (1983): 199–221.

23. Sidgwick writes that “the notion of ‘Good’ thus attained has an ideal element: it is something that is not always desired and aimed at by human beings: but the ideal element is entirely interpretable in terms of *fact*, either actual or hypothetical, and does not introduce any judgment of value” (p. 112). Brandt, in *A Theory of the Good and the Right*, argues that his “proposed definition of ‘rational’, it is important to see, does not import any substantive value judgments into the concept of ‘rational’” (p. 13). Railton’s and (less clearly) Lewis’s use of this framework to support their substantively naturalistic reforming versions of value realism suggests that they stand in this tradition.

24. In different ways David Velleman, “Brandt’s Definition of ‘Good,’” *Philosophical Review* 97 (1988): 353–71; and Connie Rosati, “Persons, Perspectives, and Full Information Accounts of the Good,” *Ethics* 105 (1995), in press, persuasively make this sort of case against naturalistic conceptions of “ideally vivid information.” In this vein, I would press the naturalist’s criterion for distinguishing between an option’s being desired because it is part of the agent’s nonmoral well-being and being desired for other, perhaps moral, reasons. Mark Overvold’s “Self-Interest and Getting What You Want,” in *The Limits of Utilitarianism*, ed. Harlan B. Miller and William H. Williams (Minneapolis: Minnesota University Press, 1982), provides an often cited, but I think ultimately unsuccessful, attempt to draw this divide. Decision theorists, insofar as they are attempting to determine an agent’s well-being, should also be worried about this latter issue (see Sen’s “Rational Fools”). See also Allan Gibbard’s “Ordinal Utilitarianism,” together with Kenneth Arrow’s reply, “Reflections on the Essays”—both in *Arrow and the Foundations of the Theory of Economic Policy*, ed. George Feiwel (New York: New York University Press, 1987).

In the rest of this article, I will develop the first line of argumentation above. In doing so I will point out different interpretations of the full information framework which have not been adequately disambiguated. The places where I want to press the full information account have not received much attention and I will thus sometimes be unable to report the precise positions which proponents of the view hold. Rather my main task will be to present what seem to me to be the options available to the account and suggest problems for each.

#### INTERPRETATIONS AND CRITICISMS OF FULL INFORMATION ACCOUNTS

In Lewis's theory of value we can see a possible ambiguity in the full information account. Lewis recognizes that the evaluative perspective ideal for appreciating some values places one in positions ill equipped to appreciate others fully. For example, the evaluative perspective ideal for appreciating the value of independent world travel unconstrained by strong attachments to others differs sharply from an evaluative perspective ideal for appreciating the value of deep, long-term commitments to a loved one and a family. Thus Lewis conceives of the ideal value appreciator as a specialist lawyer, ideally situated to experience the full worth possible to the agent only of a particular type of value. Lewis's analogy seems to invite us to suppose that the hoard of value lawyers (one for each purported value) each presents their case to an arbiter or judge who ultimately determines the value of that way of life to the agent in question. In this way commensuration of the disparate values might be thought to be accomplished.<sup>25</sup> At the same time, Lewis recognizes that full appreciation of a possibly valuable aspect of one's life typically requires "getting into the skin of the part," something the judge mentioned above might well be unable to do since she is "dispassionate and impartial."

Two different full information models of commensurating value are suggested by the above considerations. I will call them the "report" model and the "experiential" model. The crucial difference between the two is that in the report model the idealized agent typically lacks firsthand experience with the lives she is choosing between while in the experiential model the ultimate decision maker is held to lack no such information. The report model is suggested by the passionate and partial lawyers presenting a case to the impartial judge and is

25. Lewis, p. 126. Despite the naturalness of this extension of Lewis's analogy, he writes that it is not his intent to have his theory speak to the issue of commensurating values (p. 124, n. 17). Rather, Lewis simply intends his theory to detect different values without offering a method of commensuration. I do not attribute the ambiguity mentioned in the text above to Lewis's view but rather intend to show how one could use his insights to justify different commensurating strategies.

motivated by the thought that appreciating all the values one could possibly appreciate requires adopting incompatible perspectives. The experiential model is suggested by the claim that full appreciation of a possibly valuable aspect of one's life requires "getting into the skin of the part" and is motivated by the thought that often one cannot fully appreciate, or be sure that one does fully appreciate, the value of an experience to oneself without having that experience.

I will consider the report model first. We often need firsthand information if we are to be ideally situated, and if we are to be assured that we are ideally situated, to measure the value of a life to ourselves. Hence the report model does not offer a reliable method of accurately commensurating one's well-being. Some experiences are revelatory in the sense that they alter our responses to facts and descriptions. Revelatory experiences enliven our appreciation of facts and descriptions such that although we were vividly aware of the facts and descriptions of the case, we had previously been dead to the import that we now find in them.

Railton seems to recognize this general problem with the report model. He requires of the idealized agent that she have had "whatever experience or education would be necessary" to give the "bit of information . . . life." But he adds, "Fortunately, as fiction and drama show, not every fact need be directly experienced in order to make a profound impression. A well-told or well-acted or well-filmed tale, perhaps one that connects with whatever kinds of experience one already has had, may do the job." Lewis concurs that novels are often the sort of thing needed to produce the requisite effect.<sup>26</sup> This suggestion is an attempt to make use of the report model at least for part of the job of conveying full information. Railton's idea is to require first-hand information when it is necessary but to make do without it when it is possible to convey sufficiently a sense of what the life would be like by other means.

Two issues arise when we consider this suggestion of a method of attempting to give an agent an accurate understanding of what it would be like for them to live a particular kind of life. First, in which cases is it possible to convey accurately such an understanding through such means? Second, in the cases in which it is possible, how do we determine which possible book, play, movie, etc. would do the trick? Railton, in the passage above, does not seem sanguine about the scope of the lives which novels and such can give us an understanding adequate for the purposes of the full information account.

I will here stress troubles for the account stemming from the second question. I do not see how we are to know when the novel

26. Railton, "Facts and Values," pp. 30–31, n. 16; and Lewis, pp. 123–24.

gives us the experience necessary to give a piece of information the appropriate life. The full information theorist's project should not be simply to convey a profound impression in these instances but rather to convey an accurate impression of what it would be like to live the life. The bit of information must be given life in a particular way. We can be moved, even profoundly so, and still widely miss the experience as it would have felt had it been part of our life. When we have cause to worry that this is so we have cause to worry that we have only a distorted measure of the value of that life to ourselves.

It may seem a bit unfair to raise what appear to be epistemic worries against a view which is sometimes held to do little more than demonstrate the coherence of the conception of well-being employed by the SCP by instructing us how, in principle, to determine the referent of the term. However it is not sufficient merely to claim that a vantage point is correct about the value of a life to ourselves. We must be able to see why there will always be a very tight connection between the preferences of the vantage point proposed and our well-being if we are to be assured that the preferences of that vantage point definitely determine our well-being. We should therefore find that the report model does not offer a reliable method of determining our well-being because we have had explained to us no trustworthy method of consistently and across a wide range of contexts conveying to our idealized self an adequate understanding of what the lives would be like that she is choosing between.

Undoubtedly in many cases a good piece of fiction or just a talk with a friend can go some way toward conveying a sense of what lives rather unlike our own would be like. Yet, it can also happen that the differences between living a life and reading about it cause us to distort, or worry that we may have distorted, the way such a life would feel if it were our own. Attempting from one's armchair to determine what it would be like to face grave danger or an easy opportunity to steal a great amount of money is hardly a surefire strategy for success. Even if we thought that for an hour and a half a show gave us a genuine taste of a life that we might live, this still might fail to convey an accurate sense of the relative worth of that life being ours. Some condiments taste great when sampled but they still might make a lousy dinner. The full information theorist could claim that each of these potential sources of error would be known to the informed self, but knowing of sources of error and knowing how to correct for them are different things.

Additionally, we should keep in mind that in everyday life we must make our way with a good deal of uncertainty and lack of knowledge. Because of this we inevitably come to rely upon and make use of quite fallible methods because they are the best available. Thus part of our willingness to heed the direction of our fully informed self could be



due to our generally sorry state of knowledge in which we must cling to whatever floats. We all heed the advice of experts, friends, and loved ones but the trust the full information theorist asks us to have in our informed self is of a different order. A critic of the full information account need not suggest that we should pay no heed to the direction of our idealized self. Rather, such a critic need only argue that we should not take the direction of our informed self to be the be-all and end-all concerning our good.

Now it might well be claimed in defense of the report model that we do not need to actually live through enslavement in a concentration camp to know that this would be very bad for us. Surely, it might be suggested, supplying us with true propositions about the conditions in the camp, one's health there, and so on would suffice to give us a reasonable feel for how bad this life would be for us.<sup>27</sup> No doubt the extent to which this is true will depend on the extent to which one has had relevantly similar kinds of experiences and one's ability to empathize and imagine. Perhaps there are lives such that almost no matter which of our possible evaluative perspectives we have, we still see that those lives are wonderful or horrible for us to live. But there are surely many lives that we could lead in which our evaluative perspective in that life plays a large role in how we would experience that life from the inside. In many of these cases the way the life feels from the inside will be an important factor in determining the relative value of that life. I therefore see little reason to be optimistic about the attempt to capture the worth of these lives when we are provided only with the facts and experiences accessible to a perspective outside the evaluative perspectives which would be ours in the life in question.

Another way of getting at some of the concerns above about the report model is to open the question of who is to play the role of judge.<sup>28</sup> Ourselves as we are? But particularities of our current evaluative perspective could presumably cause us to discount the value of unfamiliar experiences which do not resonate with positive experiences that we have had even though, if we had the unfamiliar experiences, we would enjoy them greatly. Further, to return to the version of the report model that we can construct from Lewis's metaphor of the judge and the lawyers, the judge will surely have to discount the testimony of some of her value scouts as being beguiled by love or shallowly taking all of her delight in frivolous things. Which value scouts the judge discounts the testimony of will depend on the judge's evaluative perspective. But if the judge's evaluative perspective is not

27. I thank Richard Arneson for this objection and example.

28. Rosati effectively pursues this line of thought. I have benefited greatly from this interesting paper.

to be our own as it is now then what perspective should the judge have which is able to unmask those blinded value scouts and resonate with the reports of the value scouts who discover genuine values? To suppose that we know the proper evaluative perspective for the judge to have seems to presuppose that we already know where our good lies. Additionally, we cannot assume that each value scout is highly articulate or persuasive about her genuinely valuable experiences. Some valuable lives we could live are incompatible with such articulateness and force of argumentative powers. The best person to present a case for the value of a kind of life need not be the type of person best suited to enjoy fully that life.

I take these sorts of problems with the report model to suggest that the idealized agent typically needs firsthand acquaintance with the purported good in question if she is to be in, and if we are to be rationally persuaded that she is in, an ideal position to measure the relative value of that good to her. While the report model may be sufficient for the task of fully informing an agent about the value of goods substantially familiar to us in slightly altered settings, it leaves the vast bulk of the job of fully informing the agent to be done by other means. I will argue below that those other means are not sufficient to take on this job left over from the report model. Thus I claim that even Railton's sophisticated attempt to make use of different models of conveying information in different contexts will be inadequate in constructing a single perspective fully informed about the different paths a person's life could take.

Before considering those other means of conveying full information we should note the significance of the distinction between what I will call isolatable and relativistic value-commensurating strategies. If the amount of well-being an option would give an agent were generally determinable in isolation from the amount of well-being of all other options available to them let us say that well-being is isolatable. If the amount of well-being an option would give an agent is only determinable relative to the value of other options (as is the case with decision theory and full information accounts) let us say that value is relative. If well-being were isolatable then the above problems with the report model could be significantly silenced. Thus if well-being was, for example, as hedonists who conceive of pleasure as a sensation think it is, then something like the report model would seem workable.<sup>29</sup> There would then be no need to force any single consciousness to experience what all the different lives available to the agent would be like. The report of each of Lewis's value scouts could then be summed up in a number rather than an attempt to convey a much more full sense of

29. Daniel Jacobson helped me see this point.

what it was like to experience the distinctive values of that life. But theories which suggest that well-being is isolatable are less plausible than those which admit the relativity of value. However, I am not trying to convince the reader of that here. Rather I just need to note that the theories of commensuration I am considering are committed to the relativity of well-being.

The second model of conveying information about what our possible lives would be like I will call the “experiential” model. This model would have it that it is the same agent who experiences firsthand each type of life they might live. I will distinguish between two different versions of the experiential model which I will call the “serial” and “amnesia” versions. The rest of this article will consist of an examination and criticism of these two versions of the experiential model of the commensurating full information account.

In the serial version, our idealized self is expected to achieve full information by acquiring firsthand knowledge of what one of the lives we could live would be like, retaining this knowledge, and moving on to experience the next life we could lead. I will call this acquiring full information incrementally. A reason this will not work is that to experience adequately the value of many types of lives one must be unaware of the multitude of alternative ways one could live.

Consider the case of an Amish person who does not know what other options society holds for her. The experience of such a person could differ significantly from the experience of the same person who did have knowledge of many other options that society offers. I do not suggest that the appropriate vantage point to assess the value of such lives is always from inside them. We can become seduced by a way of life such that from inside the life the experience of the value of that life is out of proportion with the genuine value of the life. Rather, the point is that to be able to claim that one knows what it is like to lead such a life one must experience what it would be like to be in those shoes (explicitly not what it would be like to be in those shoes with the accumulated knowledge of what it would be like to have lived a multitude of alternative sorts of lives). Attempting to give the idealized agent direct experience with what it would be like to be such an Amish person, while this agent has the knowledge of what it would be like to live many significantly different sorts of lives, will in many cases be impossible.

Consider, as a second example, trying to experience incrementally what each of one’s possible first kisses would have been like. What we are expected to be able to do on the serial model is to experience firsthand one way our first kiss might have gone, remember this experience such that we remain fully appreciative of all the valuable aspects of the experience, and then try to experience fully and accurately another way our first kiss would have felt had it been the first kiss

that we experienced, and so on for all the ways our first kiss might have gone. But the fact that one has experienced and clearly remembered several such kisses would seem, at least in many people's cases, to cause the experience of some possible first kisses to differ importantly in ways that distort the experience from what it would have been like to have it actually be one's first kiss. Similarly, we might not be able to experience what it would be like to be hopelessly depressed such that we thought nothing could be of value to us (not that none of the things that would be valuable to us will occur) if we had already experienced (and retained knowledge of) lives we might lead in which our life is filled with intrinsic value.

Providing us with experiential information can alter our ability to experience certain types of lives. The introduction of propositional knowledge can have a similar effect. Suppose it is a fact that God does not exist, and we are given this knowledge as part of the process of becoming fully informed. Insofar as we retain this knowledge and accept it, as we are expected to, unhesitatingly, we may well find ourselves unable to experience adequately what it would be like for us to be born-again Christians.

This latter point about God can, I think, be handled by full information views, and the method for handling such cases might be thought to provide us a model for dealing with the Amish, first kiss, and hopelessly depressed cases. If we were first granted all and only phenomenological facts and only later granted facts about physics, psychology, and such we could, it seems, avoid the worry that some phenomenological facts would be made inappropriately inaccessible to us by our having been fed ontological facts inconsistent with some worldviews that we might have held. But if carefully ordering the information to the idealized agent works in this case why not in the others?

If we could order our experience of the lives we could lead from most naive to most sophisticated, or most hopelessly depressed to least hopeless, it might seem that we could finesse the problem of some phenomenological facts interfering with our ability to experience others. However, this strategy is hopeless. Any young person we select could lead many different lives which would be in several relevant respects "naive." What order could we give the agent's experience of the different "naive" lives such that they could experience what all of them would be like from the inside? Obviously after experience with (and retention of) a few of such naive lives our agent would realize that she could play many different social roles and will hence often be unable to experience what some of her other naive lives would have been like.

The seemingly forced move on behalf of the serial version of pushing back the point at which the idealized agent is to be granted

ontological and psychological facts merely highlights another problem for the serial version. If the idealized agent is granted first all phenomenological facts and only later ontological and psychological facts, and if the idealized agent is supposed to be gaining full information incrementally, then familiar psychological effects will again disrupt our ability to experience what we would have experienced in life A after experiencing what we would have experienced in life B. For example, Tversky and Kahneman report that a commonly used but fallacious method of determining the probability of an outcome is based on the ease with which one can remember or imagine similar instances.<sup>30</sup> But the ease with which one will be able to remember or imagine similar instances will be partially determined by one's past experiences. Thus it seems that one's estimation of the probability of events will fluctuate depending on what one has experienced. But then on the road to incrementally collecting full information we will sometimes be more or less likely to expect a certain event than we would have been had we not experienced a variety of alternative ways our life might have gone. Thus as we attempt to acquire full information incrementally, our experience of frustrations and happy surprises will sometimes differ from what they would have actually felt like to us, and those opportunities deemed worth the risk will shift. Similarly, since one's satisfaction with an outcome can depend upon one's expectations, which are themselves largely determined by past experience, one's satisfaction with outcomes could similarly shift.

The full information theorist might be tempted to grant the agent attempting to incrementally experience all of her possible futures the knowledge that persons often adopt such useful but potentially misleading heuristics and biases prior to the attempt to experience all phenomenological facts. This might be thought to prevent past experience that the agent attempting incrementally to achieve full information has (but which the possible life-path that she is investigating does not) from distorting the experience of the possible life-path in the ways broached in the paragraph above. But again this move will not solve the problem since this knowledge itself (i.e., the psychological knowledge of human heuristics and biases) could alter the agent's ability to experience what uninformed possible selves would have experienced.

Thus we have seen two different problems with the experiential strategy of granting a person full information. First, having some

30. Amos Tversky and Daniel Kahneman, "Judgment under Uncertainty: Heuristics and Biases," *Science* 185 (1974): 1124–31. Their works are rich in ways that people can be misled by past experience. See especially Daniel Kahneman and Carol Varey, "Notes on the Psychology of Utility," in *Interpersonal Comparisons of Utility*, ed. J. Elster and E. Roemer (Cambridge: Cambridge University Press, 1991).

experiences can preclude or significantly effect the having of others even when either of the two could have been part of the agent's actual life. Second, granting the idealized agent ontological and psychological facts prior to her attempt to experience a life that the actual agent could lead can cause the idealized agent's experience of that life to deviate substantially from the experience that the actual agent would have of that life were she to live it. But attempting to push back the point at which ontological and psychological facts are fed to the incrementally idealized agent until after all phenomenological facts have been experienced will not work because of the psychological effects of past experience in determining one's estimation of the probability of an event, the likelihood of one's success in a project, and one's satisfaction with an outcome.

One of the fundamental inadequacies of the serial version of the full information account is its failure to appreciate adequately the ways in which past experiences help shape future experiences. This happens in many ways. One which is relevant here is the way that some valuable experiences can occupy unique places in a life's narrative. Some decisions take on the phenomenology that they do because they radically affect the structure of one's life narrative. The scarcity of opportunities to have certain experiences can affect the phenomenology of an experience. The decision to attempt to become a philosopher or a doctor (and indeed what it would like to be either) likely would be experienced differently than it would be by us if the decision were made by someone who suspected they would eventually get around to trying both. Many rituals and initiations are designed to mark transitions from one stage of one's life to another. Such transitions can indelibly affect one's self-perception. Some crucial decisions or acts express aspects of one's character which can demand a place in constructing one's self-image. Someone who exhibited the heroism of Henry V and later cowered in fear at slight challenges will experience a different form of diminishment from the coward who never knew what it was to be brave. The same cowardly acts likely will be experienced differently because of the different narrative that the acts fit into, a disorienting loss of nerve rather than the continuing expression of a timid nature.

In demanding that the idealized agent be able to experience fully these aspects of nonidealized agents like us, we demand that she do incompatible things. The idealized agent is to be affected as we can by fundamental life choices which close off some life paths and shape who we are and what we can make of our lives, but at the same time she is to try everything and be practically immortal. She is to experience a life transition as it would be experienced by the nonidealized agent but she is to not let the experience affect her future experiences with other possible lives in which the transition does not take place. The idealized agent is to know what it is to be virtuous and base but not

fit these experiences into a life narrative which allows the character exhibited in these acts to stem from a single self.

The second version of the experiential model is the amnesia version. Here the hope of being able to experience adequately what it would be like to lead all the different lives we could lead and unify that knowledge within a single consciousness rests on being able to hide experiential knowledge from consciousness while storing it. That is, this full information theorist will have to posit controllable amnesia. The agent must have an experience of what some life would be like, then forget this and be ready to learn what some other life would be like without the latter process being affected by its position in the series. Then at the end of the learning and forgetting process we would have to remove (serially or all at once) each instance of amnesia while, on some views, adding factual information and immunity to blunders of instrumental rationality somewhere along the way. No one I am familiar with has ever endorsed, or indeed even mentioned, this rather desperate understanding of how full information is attained.<sup>31</sup>

This position is an attempt to finesse the above problems of capturing in one consciousness the firsthand experience of what it would be like to lead each of the lives one could live. The idea is to insist that, because it was accurate firsthand information about what each life that one might lead would be like which was forgotten, it must be accurate firsthand information about each life that is remembered when the forgetting is undone.

But we need more than just firsthand experience with all the lives we could lead. For the purposes of the full information account we need a perspective whose preferences between the lives in question accurately determine their value to the agent. But now a problem arises. Our actual evaluative perspective changes over time. We can therefore expect that we would respond differently to factual and experiential knowledge at different times in our future. Thus we do not have a single informed perspective to deal with, but several. And each will offer occasionally conflicting assessments of where the agent's well-being lies.<sup>32</sup> How are we to render univocal this discordant chorus such that value commensuration can be accomplished?

31. Likewise for a possible alternative which would abandon the attempt to construct a single perspective whose preferences were definitive of the agent's well-being. Instead, the idea might be to construct a complete ranking of the value of possible futures through a series of informed pair-wise comparisons of such lives. This would represent a significant break from the traditional full information account and perhaps deserves a different name. This version likely would rely heavily upon, among other things, familiar decision theoretic axioms. I find this an insecure foundation for the type of view broached here but I regrettably cannot discuss such issues here.

32. Richard Brandt's "Two Concepts of Utility," in his *Morality, Utilitarianism, and Rights* (Cambridge: Cambridge University Press, 1992), makes this case more fully.

Three replies seem possible. First, the full information theorist could attempt to construct a method of weighing the preferences of the different idealized selves which does not seem ad hoc. Second, the full information theorist could attempt to construct an allegedly temporally privileged vantage point on one's life such that the preferences from that temporal vantage point decisively determine one's good. Finally, the full information theorist might claim that the differences between one's actual evaluative perspective at time  $t$  and then later at time  $t_1$  cannot cause the preferences of the two to diverge when each is fully informed. Such differences, it might be claimed, wash out when the agent is properly idealized.

Unfortunately I cannot do justice to these possible replies in this article. But I do want to follow out a bit two ways one might flesh out the above replies. In the spirit of the first reply one might suggest that one's well-being at a time is determined solely by the idealized agent who started from the agent at that time. This would solve the problem of too many voices in determining one's well-being at a time. But we still lack an understanding of the agent's well-being across time, which is what is needed to play the role of the first step of the SCP. We need to understand when trading off one valuable option at a time for another valuable option later improves one's overall well-being. But I cannot see what agent is to make such a determination on this scheme.

This problem strikes me as worse than the traditional worries about the interpersonal comparability of utility since some version of existence internalism looks more irresistible in the case of an agent's own well-being than in the case of moral value.<sup>33</sup> That is, we might think one need not find motivating, even after idealization, the fact that an act promotes the aggregate of individuals' interests. But this maneuver seems unavailable in the case of an agent's own good. Hence the constraints on resolving intrapersonal comparisons of utility appear even more severe than those in the case of interpersonal value.

In the spirit of the third reply one might suggest, as Railton and Lewis have, that perhaps the fixation of beliefs imposed by the idealization process is sufficient to substantially fix desires across persons.<sup>34</sup> That is, perhaps we only value things differently from others and ourselves at different times (aside from differences in circumstances and position) because we have not been idealized in the way the full information account specifies. If we had, perhaps we would find that we agreed in what we each wanted for ourselves. Thus,

33. Darwall introduces the distinction between judgment and existence internalism in *Impartial Reason*, p. 54.

34. Railton, "Moral Realism," pp. 176–77, n. 19; and Lewis, pp. 128–29.



perhaps the idealization has made it unlikely that our selves at different times would, when idealized, disagree about where our well-being lies. Railton and Lewis offer this possibility very hesitantly and avoid placing much weight on this claim, and with good reason. The evidence to support it is surely indeterminate and one would be wise not to tie one's philosophical project to such a seemingly indeterminable and unlikely possibility.

The problem of too many voices is not a difficulty peculiarly for the amnesia version. I want now to mention briefly two problems which are specific to the amnesia version and a final problem which is not. First, it would surely require much very complicated research to have anything to say about the question of the similarity of experience between the original having of an experience and a later sudden recollection of it in cases in which the evaluative perspectives held at the time of the original experience differs significantly from that held by the person who is suddenly remembering. Second, even if we can convince ourselves that some system of controlled amnesia would enable us to avoid the problems I have mentioned above, there is no way of estimating the psychological shock of experiencing such a large number of instances of amnesia and loss of amnesia to our idealized selves. The full information theorist cannot simply stipulate that the idealized agent remain sane through this process, and I would have concerns about this issue. Finally, another worry that the nonidealized agent might have about what the idealized agent wants her to want is that the idealized agent might, after having experienced what it is like to be so nearly perfect in understanding, come to think of her actual self as we might think of ourselves after a serious brain injury (i.e., better off dead). Of course there might well be cases in which a person really is better off dead, but the worry here is that the idealized self might want the actual self to want this due to the only relatively degraded condition of the actual self compared to the idealized self.

## CONCLUSION

The plausibility of full information accounts is seemingly derived from the thought that when we can get our experiences with goods, as it were, side by side for comparison, our preferences between them are decisive concerning our well-being. A paradigm case might be the decision between two candy bars that are equally unhealthy for one that one has eaten frequently and recently sampled. In cases such as this there are surely strong grounds for taking one's preferences between the two candy bars as decisive concerning one's well-being.

But this does not support the inference that preferences between options that one knows of only secondhand should be taken to be so authoritative, and we have seen reasons to doubt this claim. Further, when we try to place firsthand experiences with the multitude of lives

we could lead side by side we consistently run into daunting difficulties which do not arise in the simpler cases which provided the motivation for the full information account. The important aspects of one's life remain fixed when we choose between candy bars, and this provides an ample context to make a decision from and an ample notion of the self left to make the decision. However, when we must choose between more important aspects of our life (how to carry on without a loved one, which projects to mingle one's identity with, what aspects of one's character to repudiate), it can become less clear where we are to deliberate from and who is left to do the deliberation. When we are asked to live all of our possible futures and express preferences over them, such worries overwhelm. The full information theorist must confront these issues of extending the scope of the account from the local cases in which the justification of the full information account seems most secure to the far grander aspirations of commensurating an agent's well-being.

The narrative unity of a life can provide the context to make sense of choosing one option over another, but this context is significantly dropped when we are choosing between lives rather than from within them. It is not just the different kinds of life paths that we could lead that are to be chosen between, it is also who we are to be; what kind of person we want to be who is having these different lives. Without anything like the context provided by our actual lives, the chooser becomes disorientingly "unencumbered."<sup>35</sup> Of course we can meaningfully and helpfully consider the question of which kinds of lives would be best for us to lead and what kind of person to be and, of course, we can reasonably decide that living lives rather unlike our own and being persons rather unlike ourselves would be best for us. But if such questions make sense it is only because they are considered by creatures who can bring more of their selves to the question than could a self which was capable of being fully informed of all of its possible futures.

My arguments rely upon contingent facts of psychology and, perhaps, culturally specific modes of valuing. I do not find the very idea of being fully informed of all one's possible futures incoherent. Neither have I suggested that this could not be an attractive method of commensurating the well-being of some creatures. Rather, I have argued that given what we are like and how we value, the hope of commensurating our well-being by constructing a vantage point fully informed about our possible futures is misguided. I have argued that

35. The term 'unencumbered' is meant to underline parallels with Michael Sandel's discussion of the "unencumbered self" in *Liberalism and the Limits of Justice* (Cambridge: Cambridge University Press, 1982).

the full information account is not adequate for us because some of the limitations which are idealized away by the full information account play a fundamental role in shaping our capacity to value in the ways that we do. In order to have many experiences one must be a particular kind of person. The idealized self which the full information theorist recommends is not the kind of person who could have some of the experiences which could be ours.

It is generally admitted by proponents of the SCP that interpersonal comparisons of utility pose special problems. One way of thinking of my criticism of full information accounts is that I am finding similar complications with that account of intrapersonal comparisons of utility. Typically defenders of the SCP have sought to reduce interpersonal comparisons to the case of intrapersonal comparisons which were thought to be conceptually unproblematic. But I suggest that the far-fetched method of making interpersonal comparisons by somehow taking the preferences of all others into oneself and then determining what one wants most is not much more conceptually problematic than the full information account of well-being. I find problematic in both cases the assumed capacity of the deliberating agent to appreciate fully what it would be like to be an impossibly wide diversity of personalities and sort them out into a coherent set of preferences while maintaining an account of the normative authority of such a deliberating agent.

This article does not address some important pressures to accept a commensurating account of well-being. The seldom spoken but powerfully influential thought that reasoning between incommensurable values is impossible can seem to force us to choose between either admitting that most moral (and prudential) choices are necessarily made arbitrarily or finding a method of commensurating value. If it were true that there can be no reasoning between incommensurable values we would be wise to accept a theory which commensurated well-being even if it came with grave difficulties.<sup>36</sup> Therefore a complete argument against the above commensurating accounts of well-being must address the issue of reasoning our way through incommensurable values or point toward a more promising method of commensurating well-being. I believe this pressure to adopt a commensurating theory of well-being should be rejected, but it would require at least a separate article to make this claim persuasive.<sup>37</sup>

36. Griffin writes, "Nor can we first fix on the best account of 'well-being' and independently ask about its measurement. One proper ground for choosing between conceptions of well-being would be that one lends itself to the deliberation that we must do and the other does not" (p. 1).

37. Anderson's *Value in Ethics and Economics*, chaps. 1–5; and Joseph Raz's *The Morality of Freedom* (Oxford: Clarendon, 1986), chaps. 12 and 13, forcefully present this case. See also Henry Richardson's "Specifying Norms as a Way to Resolve Concrete Ethical Problems," *Philosophy and Public Affairs* (Fall 1990), pp. 279–310.

Too often the SCP is simply granted, without discussion, that a theory of an agent's well-being adequate to its needs can be constructed. In part, the effectiveness with which proponents of the SCP have been able to avert or accommodate criticisms directed against it as a theory of morality is due to the large concessions overwhelmingly made to it in the theory of value. The full information account enjoys such popularity today partially, no doubt, because, unlike other value commensurating strategies such as the many versions of hedonism and objective list accounts, it is so infrequently challenged. I believe my arguments above demonstrate that the full information account is not so clear or obvious as to merit this silence. If we find that we cannot ultimately accept the full information account, it becomes far less clear that an agent's well-being is commensurable to the extent needed to make the SCP a plausible foundation for a theory of morality.