FRED FELDMAN

BASIC INTRINSIC VALUE

1. SOME PUZZLES IN AXIOLOGY

When called upon to tell us what hedonism is, moral philosophers often start out by contradicting themselves. An extreme and blatant case would look like this:

Hedonism: the view that (i) pleasure is the only thing that is intrinsically good, and (ii) pain is the only thing that is intrinsically bad; furthermore, the view that (iii) a complex thing such as a life, a possible world, or a total consequence of an action is intrinsically good iff it contains more pleasure than pain.

The contradiction seems obvious, but since so many philosophers evidently don't see it, perhaps it will be useful to employ the sledge hammer. According to the first clause, pleasure is the only thing that is intrinsically good. This implies that nothing distinct from pleasure is intrinsically good. Since pleasure (whatever precisely it may be) cannot be identified with any life or with any possible world, it would follow that no life or possible world is intrinsically good. Yet, according to the third clause, complex things such as lives and possible worlds are also intrinsically good when they contain a favorable balance of pleasure over pain.

You cannot have it both ways: either (a) pleasure is the only thing that is intrinsically good (in which case no lives and no possible worlds are intrinsically good), or else (b) some lives and some worlds are intrinsically good (in which case pleasure is not the only thing that is intrinsically good).

I claimed that many philosophers have contradicted themselves in this way. I will cite two prominent examples.

In *Brentano and Intrinsic Value*, Roderick Chisholm seems to commit himself to this contradiction. In an attempt to sketch a

simple form of hedonism, Chisholm says, 'States of pleasure are the only things that are intrinsically good and states of displeasure or pain are the only things that are intrinsically bad.'[1] He goes on to specify the intrinsic values (according to the same theory of value) of various "mere sums" and other complex states of affairs including organic unities of various sorts. Some of these are compound states of affairs consisting of one person feeling some pleasure and another feeling some pain. Clearly, none of the complexes under consideration is a "state of pleasure" or a "state of pain" as Chisholm understands these concepts. Yet Chisholm says that some of these compounds have positive intrinsic value and so are intrinsically good. Thus, Chisholm states hedonism in a self-contradictory way, saying at the outset that pleasures are the only intrinsic goods and saying shortly thereafter that other things are also intrinsically good.

Richard Brandt seems equally guilty. In his article, 'Hedonism' in *The Encyclopedia of Philosophy*, he describes hedonism as the view that:

> . . . only pleasant states of mind are desirable in themselves; that only unpleasant states of mind are undesirable in themselves; and that one state of affairs is more desirable in itself than another state of affairs if and only if it contains more (in some sense) pleasant states of mind than the other . . . .[2]

Surely it is wrong to say both that pleasant states of mind are the *only* things that are desirable in themselves, and also that *other* states are desirable in themselves if they contain more pleasant states.

The first puzzle in axiology concerns this contradiction. How are we to avoid it? I will claim that an appropriate appeal to "basic intrinsic value states" would help to avoid this sort of contradiction, and thus to solve the first puzzle. Somewhat more substantially, I will suggest that the appeal to basic intrinsic value states will facilitate the formulation of any axiological theory.

The second puzzle concerns one sort of extrinsic value – value as a means. Since the time of Plato and Aristotle, moral philosophers have recognized a distinction between intrinsic value and this sort of extrinsic value.[3] Let us agree provisionally that something is intrinsically good if it is good "as an end", or "in itself", or "for its own sake", or "in virtue of its own nature". How then shall we explain value as a means (or "instrumental") value?

The standard answer is that something is good as a means if it causes something that is intrinsically good. Moore suggests something like this in *Principia Ethica*, where he says:

Whenever we judge that a thing is 'good as a means', we are making a judgment with regard to its causal relations; we judge *both* that it will have a particular kind of effect, *and* that that effect will be good in itself.[4]

C. I. Lewis says a similar thing in *An Analysis of Knowledge and Valuation*:

A thing *A* will never be said to have extrinsic value or instrumental value, unless it is meant to imply that there is some other thing, *B*, to which it is or may be instrumental, which has intrinsic value.[5]

In his description of what he takes to be "the standard account", Gilbert Harman introduces this formula:

S is instrumentally desirable to the extent that S does not entail but "leads to" something intrinsically desirable.[6]

The general idea here seems to be that a thing has instrumental value if it causes something with intrinsic value. This is typically conjoined with the further assumption (clearly suggested by Harman's formula) that the instrumental value of a thing is equal to the intrinsic value of something that it causes. Since instrumental values vary from world to world, depending upon what a thing happens to cause in a world, anyone adopting this approach should be careful to relativize instrumental values to worlds. Making use of the idea that the bearers of value are all states of affairs, we can put this point by saying:

EV1: The instrumental value of a state of affairs, P, at a world, w, is some number, n, iff P causes some state of affairs, Q, at w, such that the intrinsic value of Q is n.

This has strange implications. Suppose P causes an intrinsically good thing with value +10 and also causes an intrinsically bad thing with value –10 at w. Then EV1 implies that the instrumental value of P is both +10 and –10; that P is both instrumentally good and instrumentally bad. That seems wrong. The natural revision is this:

EV2: The instrumental value of a state of affairs, P, at a world, w, is the sum of the intrinsic values of all the things, Q, such that P causes Q at w.

EV2 gives each state of affairs a unique instrumental value at each world. Thus, nothing is both instrumentally good and instrumentally bad at the same world according to EV2. But EV2 has other problems. Consider a case in which Bob reads a delightful book. We may focus on this state of affairs:

B: Bob reads at 9:00PM on Monday evening.

Suppose B makes Bob happy to degree +10 while he reads, and sits in his easy chair, and wears his eyeglasses. Suppose B never makes Bob or anyone else happy or unhappy at any other time. In this case, B has many consequences, including the following:

C1: Bob being happy to degree +10 at 9:00PM on Monday evening.
C2: Bob being happy while reading.
C3: Someone being happy to degree +10.
C4: The man in Bob's easy chair being happy to degree +10.
C5: Bob being happy at least to degree +9.
C6: Bob being happy to degree +10 while sitting in his easy chair and wearing his eyeglasses at 9:00PM on Monday evening.
C7: Either Bob being happy to degree +10 or Babe being happy to degree +15.

If a certain simple and familiar form of eudaimonism is true, then each of these is "good in itself", or "good in virtue of its own nature". Some of these appear to have an intrinsic value of +10; others appear to have positive intrinsic value, though perhaps not +10. If so, and if B causes no bad things, then EV2 implies that the instrumental value of B is tremendously high. This is clearly wrong.

Notice that among the consequences of B there are a lot of states of affairs that somehow or other relate to the fact that Bob enjoyed 10 units of happiness at 9:00PM on Monday night. Some entail it, some are entailed by it, and some seem, in a vague way, to coincide with it. C1 seems to contain precisely the information that we are interested in – provided that we endorse the assumed form of eudaimonism – that is, neither more (like C6) nor less (like C2) nor muddled nor misleadingly incomplete. Suppose we had some way to disregard all of these other states of affairs and to count just C1. Then we would have a way of explaining the instrumental value of B: it causes C1, and the intrinsic value of

C1 is +10. Therefore, we might say, the instrumental value of B is +10. However, what justification is there for taking account of the intrinsic value of C1 while ignoring the intrinsic values of all the other consequences of B? After all, each of them seems to be good in itself and each of them is just as much a consequence of B.

Thus, we have our second puzzle in axiology: how are we to explain the instrumental value of a state of affairs? I will claim again that an appropriate appeal to basic intrinsic value states will figure in the most plausible solution to the puzzle. Roughly, my suggestion will be that we understand basics in such a way that C1 is a basic intrinsic value state, while all the other consequences of B are not basics. Then we can say that the sum of the intrinsic values of the basics among the consequences of B is +10, and that for this reason B has an instrumental value of +10, which seems just right.[7]

The third puzzle concerns the intrinsic values of complex things, such as lives, total consequences, and possible worlds. For purposes of illustration, let us focus on lives. Since the time of Plato and Aristotle, it has been assumed that one of the main aims of axiology is to tell us something about "the good life". An axiology is called upon to provide principles specifying the intrinsic value of each life. How is this to be done?

The life of a person at a world might be taken to be the conjunction of all the facts about the person at the world. We might say that a state of affairs is "about" a person if any sentence of the most appropriate sort for expressing that state of affairs has, as its grammatical subject, a bare name of the person. Then, we can say that the intrinsic value of the life of a person at a world is equal to the sum of the intrinsic values of the states of affairs about him that are true at that world.

We can formulate this as follows.

LIV1: The intrinsic value of the life of s at world w = the sum of the intrinsic values of all states of affairs, q, such that q is about s and true at w.

I relativize to worlds since (as I see it) a given person might exist at several different possible worlds. His life at one world might be worth more or less than his life at another.

There are several apparent problems with LIV1. (a) Suppose S lives a tremendously boring life at w; he experiences just one instant of happiness, as indicated by the statement 'S is happy to degree +3 at time t'. Suppose no other resident of w has any happiness or unhappiness. Then LIV1 yields the result that S's life is infinitely valuable. For all of these are true and about S at w:

S1: S is happy to degree +3 at t.
S2: S is happy to degree +3 at t while 2 + 2 = 4.
S3: S is happy to degree +3 at t while 3 + 3 = 6.
S4: S is happy to degree +3 as a result of what happened to the chickens; etc.

(b) Suppose some other resident, Z, of this same world lives an even more boring life than S. He never experiences any happiness or unhappiness. Still LIV1 yields the result that Z's life is infinitely valuable. For each of these is true and about Z:

Z1: Z is such that S is happy to degree +3 at t.
Z2: Z is such that S is happy to degree +3 at t while 2 + 2 = 4.
Z3: Z is such that S is happy to degree +3 at t while 3 + 3 = 6.
Z4: Z is such that S is happy to degree +3 at t as a result of what happened to the chickens; etc.

Again, the natural solution seems to be that we should appeal to basic intrinsic value states and we should say that the value of S's life is equal to the sum of the intrinsic values of the basics that are true at w and appropriately about S. If we select our basics correctly, it will turn out that there is just one basic intrinsic value state that (a) is properly about S in the example cited above, and (b) that is true, and (c) has an intrinsic value other than zero, and that is S1. Then we can let the intrinsic value of S's life be equal to the sum of the intrinsic values of these basics. In other words,

LIV2: The intrinsic values of the life of s at world w = the sum of the intrinsic values of all states of affairs, q, such that q is a basic intrinsic value state; q is about s; and q is true at w.

Given my assumptions about basics, LIV2 correctly implies that the value of S's life is +3. In the case of Z, we will find that every basic intrinsic value state that is about him and true in this example has an intrinsic value of zero. It, therefore, turns out that LIV2 declares his life to be worthless, as our eudaimonistic axiology intends it to be.

Another traditional aim of axiology is the provision of principles about the evaluation of worlds. This is thought to be important for several reasons. For example, it plays a central role in many discussions of the problem of evil. It may also play a role in the formulation of certain views in the normative ethics of behavior.[8] Thus, there is some interest in finding a principle telling us the intrinsic value of a world. Again, we don't want to say:

> **WIV1: The intrinsic value of a possible world, w, = the sum of the intrinsic values of all states of affairs, q, such that q is true at w.**

but we might want to say:

> **WIV2: The intrinsic value of a possible world, w, = the sum of the intrinsic values of all states of affairs, q, such that q is a basic intrinsic value state and q is true at w.**

Again, if we have selected our basics in the way I am imagining, this will give positive intrinsic values to things like S1 and negative intrinsic values to things like

J1: Jones is happy to degree –10 at t.

If we select our basics correctly, and we assign intrinsic values to them correctly, it will be reasonable to say that the intrinsic value of a world is the sum of the intrinsic values of all the basics true there.

Similar considerations apply in the case of the total consequences of actions.

This gives us another reason to assume that there are basic intrinsic value states. It helps us to formulate principles that determine the intrinsic values of complex things such as consequences, lives, and worlds.

There are other reasons to introduce the concept of basic intrinsic value. Some of these have to do with the cluster of issues surrounding defeat and enhancement of intrinsic value and organic unities. They will be discussed later.

2. METAPHYSICAL DIGRESSION

A critic might raise the following criticism: 'You have assumed that the bearers of intrinsic value are states of affairs. States of affairs

are very "fine-grained" entities. (Consider C1–C6 above.) Once you make this assumption, you will need some way to avoid double counting, and maybe the appeal to basics will be unavoidable. However, if you had chosen "coarser" entities, such as concrete events, as the bearers of intrinsic value, the whole issue would have been avoided. There would have been no need for basics. So the puzzle is simply an artifact of your own metaphysics.'

I think it's important to see that this criticism is misguided. We will need to distinguish between basics and non-basics even if we assume that the bearers of value are concrete events.

Suppose that the bearers of value are concrete events. Suppose that these expressions all serve to pick out the same concrete event:

C1: Bob being happy to degree +10 at 9:00PM on Monday evening.
C2: Bob being happy while reading.
C3: Someone being happy to degree +10.
C4: The man in Bob's easy chair being happy to degree +10.
C5: Bob being happy at least to degree +9.
C6: Bob being happy to degree +10 while sitting in his easy chair and wearing his eyeglasses at 9:00PM on Monday evening.

We can say that there is just one concrete event involved here and that it has an intrinsic value of +10. The occurrence of this event at a world increases the value of the world by just 10 points. There may seem, therefore, to be no need (or even *possibility*) to distinguish the basics from the non-basics.

However, the appeal to concrete events does not solve the problem. Consider the concrete event, e1, that consists in Bob's being happy at 9:00PM on Monday evening. Perhaps each of C1–C6 somehow indicates this single event. Now consider the event, e2, that is Bob's whole life. In light of the fact that e2 begins long before e1, and ends long after e1, and contains many parts that e1 does not contain, we surely cannot identify e2 with e1. Yet if e1 and e2 are distinct and each has a positive intrinsic value, then it may seem that each contributes to the intrinsic value of the world. But to count them both is to engage in double counting, since the intrinsic value of e1 is already included in the intrinsic value of e2.

There are deeper problems about the move to concrete events. One of these turns on the fact that we may want to make a number of

claims about the modal features of intrinsic value. For example, we may want to say that each thing has its intrinsic value of necessity. But, there is no consensus about the modal features of concrete events. Consider the concrete event that consists in Bob's being happy to degree +10 at 9:00PM. Could it have happened a bit earlier in the evening? Could it have involved a slightly smaller amount of happiness? Could it have happened to Babe instead of to Bob? If the item in question is truly "concrete", I see no way to assure negative answers to these questions. Thus, it appears that a concrete event that has an intrinsic value of +10 might have had an intrinsic value of +9; a concrete event that serves to enhance the intrinsic value of Bob's life might have served to enhance the intrinsic value of Babe's life. It would be better to start with metaphysical assumptions that rule out such possibilities.

Thus, the move to concrete events (a) will not obviate the need for basics and (b) will introduce puzzles and problems of its own. So I prefer to stick with the assumption that the fundamental bearers of intrinsic value are very finely individuated states of affairs. Given that assumption, it seems reasonable to try to distinguish the basics from the non-basics.

3. GENERAL CHARACTERIZATION OF BASICS

Let us agree, then, that there might be some justification for the assumption that there are some selected states of affairs that are basic intrinsic value states ("basics"). One natural question at this point is this: which states of affairs are the basics?

But of course this question cannot be answered without begging many of the central substantive questions of axiology. Different axiological views entail different views about which states of affairs are the basics. What is a basic on one axiology might have no value, basic or otherwise, in another axiology.

Since I believe that the recognition of basics is useful no matter what our axiology, I want to give a general characterization of basics, one that will hold no matter what axiological theory we ultimately decide to accept.

I think there are six main features of basics.

1. *Every basic is a pure attribution of a core intrinsically valuable property or relation.* Ideally, an axiological theory should precisely identify the states of affairs that it takes to be the ultimate sources of intrinsic value. These are the basics on that axiology. Thus, on the simple form of eudaimonism I mentioned earlier, the basics would be states of affairs such as the one expressed by this:

J12: Jones being happy to degree +12 at noon, March 25, 2000.

One crucial feature of J12 is that it is a "pure attribution" of *happiness*. The state of affairs attributes it directly to a person and not via any of his properties. The phrase used to express this state of affairs picks out Jones by the use of a meaningless "tag" and not by any meaningful description. The state of affairs attributes a precise degree of happiness (+12) to that person. It attributes it to him at a specific, named, time.[9] It does nothing else.

I assume that some states of affairs are "directly about" certain people, times, and numbers. Thus, I assume that there is a state of affairs that says, with respect to Jones, now, and +12, that he is happy at it (the time) and to it (the number). Some philosophers of language might think that this state of affairs is composed of Jones, now, +12, and happiness. If we take this view (and I am not sure that I want to) we can say that happiness is the only "conceptual component" of the state of affairs. The other components of the state of affairs are "objects" such as Jones, now, and +12. For this reason, I will say that it is a pure attribution of happiness.

We can say in general that a state of affairs is a pure attribution of F iff there is something, x, such that the state of affairs is the state of affairs of x's having F (where x appears directly in the state of affairs). That covers the one-place case. For relations: p is a pure attribution of x,y|Rxy if there are individuals, x and y such that p is the state of affairs of x's bearing R to y (again where x and y appear directly in the state of affairs); and so on for relations involving larger numbers of terms.

Thus, the basics of the simple form of eudaimonism are all pure attributions of *x,t,n|x is happy at t to n*. Since each such basic contains no other information beyond the mere identity of the recipient of the happiness, the time at which the recipient receives it, and the amount received, I will say that it contains no "extraneous informa-

tion." It contains just the core information that a simple eudaimonist will think is relevant to intrinsic value.

Perhaps this helps to explain why the ancients persisted in saying that eudaimonism is the view that *happiness* is the sole intrinsic good. Such talk certainly appears to attribute intrinsic value to a property. But, what they surely meant was that the intrinsically good basics are all pure attributions of happiness. Those who said that hedonism is the view that *pleasure* is the sole intrinsic good must similarly have meant that, according to hedonism, every intrinsically good basic was a pure attribution of pleasure.

It also gives us a way of explaining what's distinctive about "pluralist" axiologies. It is not sufficient merely to say that pluralism is the view that there are several different sorts of intrinsically good states of affairs. This would make even the most hardcore monistic hedonism into a form of pluralism, since even in such an axiology the good states of affairs differ in a number of respects, e.g. some are about you, some are about me, some are about large amounts of happiness, and some are about small amounts of happiness. But the appeal to basics may help. We can say that an axiology is pluralist just in case it implies that some good basics are pure attributions of one property, and other good basics are pure attributions of another property. Thus, for example, if someone says that some good basics are pure attributions of pleasure and others are pure attributions of knowledge, then that person is a pluralist in axiology.

2. *Every basic has a determinate intrinsic value.* Eudaimonism of the sort imagined presupposes that every basic intrinsic value state has a determinate amount of intrinsic value. Our eudaimonist would say that the basics are all pure attributions of happiness. He could go on to say that the "H-value" of a basic is equal to the precise amount of happiness (positive or negative) attributed by that basic. Then he could say that the intrinsic value of each basic is equal to its H-value. Thus, given this eudaimonistic axiology and the imagined view about which states of affairs are the basics, it follows that every basic has a fully determinate, absolutely precise, intrinsic value.

What's true of these basics on this eudaimonistic axiology is the universal feature of basic intrinsic value states: each of them has a perfectly determinate intrinsic value. But again, different axiologies identify different sets of basics. As a result, we have to be careful

to state this thesis in a slightly more guarded way: *each axiology selects a set of basic intrinsic value states; for each such axiology, the basics it selects must have perfectly determinate intrinsic values*.

States of affairs that are not basic intrinsic value states may fail to have fully determinate intrinsic values. A number of philosophers have tried to formulate principles that determine the intrinsic value of each state of affairs (including the non-basics). They have struggled to identify the intrinsic values (given specified axiologies) of various states of affairs.[10] Philosophers in this tradition evidently assumed that, since eudaimonism implies that *Jones being happy about the circus* is good in itself, it must be good in itself to some particular degree.

But how intrinsically good is it?

In his 1967 paper, "Toward a Theory of Intrinsic Value",[11] Gilbert Harman introduced a general formula that seems to be relevant here. He proceeded by first considering a formula that does not work. It is this: 'S is intrinsically desirable to the extent that S entails something that is intrinsically desirable'[12] Harman's remarks are not entirely clear to me, but in light of what he later says, he might have been thinking of this:

IV1: IV(p) = the sum, for all q, such that p entails q, of IV(q)

Of course this is not acceptable, since if S is intrinsically good, it entails infinitely many other intrinsic goods. Then its intrinsic value goes through the roof. Harman makes some remarks that suggest that he would reject IV1 for just this reason.[13]

Then Harman says:

IV2: IV(p) = the sum, for all q, such that p entails q, of BIV(q)

We may assume that BIV is a function that takes us to the intrinsic value of P if P is a basic intrinsic value state; otherwise, it takes us to zero. In this way, IV2 can be understood as the principle that the intrinsic value of any state of affairs is equal to the sum of the intrinsic values of the basics it entails; or, equivalently, equal to the sum of the basic intrinsic values of all the states of affairs it entails. This generates acceptable results in some cases. For example, (still assuming the form of simple eudaimonism that we have been using) consider:

N: S being happy to degree +12 at t1 & Z being happy to degree
     +10 at t2.

N entails exactly two basics and the sum of their intrinsic values
is +22. IV2 implies that the intrinsic value of N is +22. That seems
correct.

However the proposal is wrong in just about every other sort of
case. Consider *someone is happy to degree +12*. It entails no basics,
and so gets an intrinsic value of zero according to IV2. Yet according
to the traditional way of thinking of intrinsic value, it is viewed as
being intrinsically good. Consider *the man in the blue hat is happy
to degree +12*. It also gets an intrinsic value of zero on IV2 and it
too would traditionally be taken as good. Consider:

M: S being happy to degree +12 at t1 or Z being happy to degree
     +10 at t2.

M seems pretty good, it guarantees at least 10 units of intrinsic
value, yet IV2 implies that it has an intrinsic value of zero. A wide
variety of other examples fail to receive the desired intrinsic values
on IV2.

Another proposal makes the intrinsic value of S equal the sum
of the intrinsic values of basics that "make it true" at w. Consider
the disjunction M. If the first disjunct makes it true at w, then the
intrinsic value of the disjunction is +12. If the second disjunct makes
it true, then its intrinsic value is +10. If both disjuncts make it true,
then its intrinsic value is +22. In the case of the man in the blue hat,
if it's true, then some basic must make it true. Maybe it is *Smith
being happy to degree +12 at t*. If Smith is the man in the blue hat,
this might be it. In that case, the intrinsic value of *the man in the
blue hat is happy to +12* is +12.

However, there is reason to be uncomfortable with this proposal.
It conflicts with what we may call the "necessity principle"

**NIV: For any state of affairs, p, and number, n, if IV(p) = n,
then it's necessary that IV(p) = n.**[14]

Since a given state of affairs might be made true by basics with
different intrinsic values at different worlds, the present proposal
implies that a given state of affairs might have different intrinsic
values at different worlds. It also seems obscure. Consider *someone
is happy*. Suppose several different people are in fact happy. Which

of these is such that a basic about him "makes true" the generalization that *someone is happy*?[15]

Michael Zimmerman has suggested[16] that the search for a precise number indicating intrinsic value in every case is a wild goose chase. Maybe some things are not intrinsically good and not intrinsically neutral either. For example, disjunctions of good and bad. Maybe some other things, such as "vague and indeterminate intrinsic goods" are intrinsically good, but not to any specific degree. Maybe they are indeterminately intrinsically good. This means that basics bear their intrinsic values in a manner different from non-basics. Every basic has a precise intrinsic value; but some non-basics have indeterminate intrinsic values. It is not clear to me that this is a problem. Suppose we have determinate intrinsic values for all basics. Suppose we have determinate intrinsic values for lives, worlds, and total consequences. Why should we be dismayed when we discover that we do not have precise intrinsic values for certain non-basics?[17]

So my point here is this: on any axiology, each basic must get a precise intrinsic value. But states of affairs that are not basics may be left with vague, indeterminate, or undefined intrinsic values.

3. *Completeness*. Following Harman, let us assume that there are two functions, IV and BIV. IV takes us to the regular intrinsic value of a state of affairs; BIV takes us to the basic intrinsic value.

Where p is a basic intrinsic value state, $BIV(p) = IV(p)$. This is in every case a determinate number. I assume that, for any state of affairs, p, if p is a basic intrinsic value state, then there is some number, n, such that n = the basic intrinsic value of p, or $BIV(p)$. But what shall we say about states of affairs that are not basics? What does BIV do in those cases?

It will be convenient to say that if a state of affairs is not a basic intrinsic value state, then its basic intrinsic value is zero. This is nothing more than the claim that such states of affairs do not have any intrinsic value *in the most basic way*. It is consistent with the idea that such states have plenty of intrinsic value in some derived way.

Given these assumptions, the functions BIV and IV differ in an important respect: BIV is "complete" – for every state of affairs (whether basic or not), BIV takes us to some precise number (of

course, the number is often zero). But IV is not "complete" – for many states of affairs; IV is undefined. In the case of many non-basics, p, there is no number, n, such that IV(p) = n.

4. *The Supervenience Thesis*. Each axiology must identify the basics in such a way that all the intrinsic value at any world is determined by the basics true there. In the case of a simple axiology such as our sample eudaimonism, the basics are all things like *Jones is happy to degree +3 at noon today*. Each such thing that's true at a world contributes its intrinsic value to the world. The value of the world is the sum of the values of these true basics. On this axiology, nothing else can affect the value of the world. As a result, any two worlds that are alike with respect to basics will have to be alike with respect to intrinsic value. In other words, the value of a world supervenes upon the values of the basics that are true there.

5. *Indefeasibility*. If we have chosen our basics correctly, then we have chosen our basics in such a way that their value is indefeasible, in a certain specifiable sense. This is a complex matter and a number of possible views have been proposed.[18] According to one conception of defeat, the intrinsic value of a given state of affairs may vary from circumstance to circumstance. Those who view things this way may start out by supposing, for example, that *Jones being happy to degree +12 at noon today* (or "J12n") usually has an intrinsic value of +12. But they may go on to suppose that the intrinsic value of J12n is somehow diminished when Jones gets this happiness undeservedly; or when he gets his happiness as a result of seeing that Smith is suffering. They might think that under those circumstances IV(J12n) is less than +12; maybe it is zero. They might also suppose that J12n sometimes has greater intrinsic value, as for example when Jones has been very good and deserves a lot of happiness but so far has not received any.[19]

However, there is another way to view these phenomena of defeat and enhancement. We may be moved by the notion that a thing's intrinsic value is supposed to be the value it has "in itself", or "in virtue of its own nature". We may think, then, that since J12n has remained precisely the same "in itself", its intrinsic value cannot change no matter what circumstances it occurs in. We therefore

think that J12n (as well as everything else) has its intrinsic value of necessity.

If we want to view things in this way, we may say instead that J12n has precisely the same intrinsic value (+12) in every possible circumstance. Yet, when J12n occurs as part of certain larger circumstances, the value of J12n may be "obliterated". By this I mean to suggest that the "larger" state of affairs somehow fails to reflect the intrinsic value of J12n. Consider, for example, the state of affairs

> D: *Jones being happy to degree +12 at noon today when Jones does not deserve to be happy at all.*

We may suppose that D has an intrinsic value of 0. In this case, D contains a good part (J12n); it contains no bad part, but its intrinsic value is significantly less than the intrinsic value of this good part. J12n continues to have its intrinsic value, but it seems to have been made irrelevant.

When I say that the intrinsic value of J12n is "made irrelevant", part of what I mean is that when J12n occurs in the context of D, then although J12n continues to have positive intrinsic value, that value does not help to make the world better. Because it has been defeated by its context, the intrinsic value of J12n fails to make any contribution to the intrinsic value of the world.

Let us say that the intrinsic value of a state of affairs, p, is *obliterated* by p's occurrence in a larger state of affairs, q, iff p is part of q and the intrinsic value of p, q worlds directly reflects the value of q, but not of p.

This second view about intrinsic value gives us yet another way to distinguish between regular, old-fashioned intrinsic value and basic intrinsic value. For we can say that regular intrinsic values can be obliterated, but basic intrinsic values cannot be obliterated. In other words, what happened to the intrinsic value of J12n in the example just discussed could never happen to the basic intrinsic value of a basic intrinsic value state.

Since the basics on any axiology include all the things with non-zero basic intrinsic value, this thesis puts a constraint on what we may identify as the basics. We have to be sure to choose them in such a way that their values will never be obliterated. The point is that if BIV(p) = +n, then any world with p true in it must as a result be n units better than the world just like it except that p is false.[20]

Those who believe that the possibility of obliteration is pervasive may, as a result, be forced into a sort of axiological holism. Suppose we start out thinking that something like J12n has an intrinsic value of +12. But suppose we come to think this value may be obliterated when J12n occurs in the context of D; and suppose we also think that the value of D may be obliterated when D occurs in some wider context. Suppose, more generally, we think that every apparently intrinsically valuable state of affairs runs the risk of having its value obliterated. Whenever we are inclined to assign an intrinsic value to a state of affairs, we stop ourselves and request further information about the circumstances of its occurrence. This fear of obliteration does not relax until we reach whole possible worlds. Then we rest assured that we have found something whose value cannot be obliterated by its occurrence in some larger defeating situation. (We reach this conclusion, of course, entirely because we think that worlds are so large that they cannot occur in any larger situations.)

If we are attracted to this way of viewing intrinsic value, and we agree that basic intrinsic value cannot be obliterated, then we will have to conclude that the only states of affairs that are basics are whole possible worlds. If we reach this conclusion, we will be committed to "world holism". This would be unfortunate, for it would mean in effect that worlds are organic unities whose basic intrinsic values are not functionally dependent upon the basic intrinsic values of the things that happen within them. But the view seems to me to be coherent, as does a similar view we might call "life holism".

6. *Familiality; Falling into families.* On the eudaimonistic axiology I have been imagining, each basic is a pure attribution of some amount of happiness to some person at some time. If we hold the person and the time constant, but allow the amount of happiness to vary, we get a certain set of basics. This list suggests such a set:

J+3n: Jones is happy to degree +3 at noon today.
J+2n: Jones is happy to degree +2 at noon today.
J+1n: Jones is happy to degree +1 at noon today.
  J0n: Jones is happy to degree 0 at noon today.
J–1n: Jones is happy to degree –1 at noon today.
J–2n: Jones is happy to degree –2 at noon today.
J–3n: Jones is happy to degree –3 at noon today.

Consider the set containing these and other states of affairs like them. Each of these is directly about Jones. Each is also directly about noon today. Each is directly about some degree of happiness, but they differ with respect to the precise amount of happiness attributed. Thus, if Jones exists at noon today, exactly one member of this set must be true. Let's use the term "Jones noon family" for the set containing all of these plus the state of affairs *Jones does not exist at noon today* (which I am not counting as a basic intrinsic value state, even though I am assuming that it is a member of the family). Now we can say: exactly one member of the Jones noon family must be true.[21]

Since the basics on each axiology are pure attributions of whatever is asserted by that axiology to be the fundamentally intrinsically good property, it follows that the basics selected by any axiology will have to fall into families such as this one. In general, suppose some axiology selects some relation, $x,n,t|Rx,n,t$, as the source of all intrinsic value; suppose this relation relates a person, x, to a number, n, and a time, t. Then, for any selected pair of person, S, and time, t, there will be an S,t-family. The members of that family will be all the basics about S and t, differing only in the precise amount of R that they attribute to S at t (plus *S does not exist at t*). Exactly one member of this family must be true.

We can then make a general statement about basic intrinsic value states: they fall into families. In some cases, admittedly, the families might be fairly degenerate.

Facts about families may shed some light on the concept of "isolation". In a number of places, Moore claimed that the intrinsic value of a thing is the amount of value it would have "in isolation".[22] Moore's remark is suggestive, but problematic. It becomes even more problematic if we assume (as I do) that the fundamental bearers of intrinsic value are states of affairs. What could be meant by saying that some state of affairs exists "in isolation"?[23]

Many traditional axiologies yield families that are logically independent. In the case of any such axiology, if p is a member of one family, and q is a member of another family, then p and q are logically independent. There are worlds where both are true, worlds with neither true; and yet others with one but not the other.[24]

This fact about simple axiologies suggests a possible interpretation of Moore's isolation thesis. Let us confine our attention to axiologies according to which every basic is such that there is some possible world in which it is the sole true basic. On these axiologies, if p is a basic and true at w, then no other member of its family is true at w. Furthermore, (in the case of simple axiologies) each basic is such that it's possible for it to be true even though no extrafamilial basic is true. Therefore, in these cases, there is a world at which no other basic is true.

Let us say that if p is a basic, and w is a world at which p is the only true basic, then p is *basically isolated* at w. I have argued that for many basics, there is a world at which it is basically isolated. A restricted version of the isolation thesis may now be stated: if p is a basic, and there is a world at which it is basically isolated, then IV(p) is equal to the intrinsic value of any world at which p is basically isolated.

Of course this won't work for axiologies that permit basics to entail basics in other families.

4. SOLUTIONS TO PUZZLES

At the outset I mentioned some puzzles and I claimed that the appeal to basics might be useful for their solution. Let us briefly review the proposed solutions.

i. The first puzzle concerned a contradiction into which we may fall when we try to state an axiological theory. If we try to formulate the simplest sort of hedonism, we may be inclined to say, for example, that (i) hedonic (doloric) states are the only intrinsic goods (bads); and (ii) a world or other complex thing is intrinsically good (bad) iff it contains more (less) pleasure than pain. As I pointed out in Section 1, this is a contradiction.

I propose that we make use of basics in our statement of any axiological theory. Thus, for example, if we want to formulate a simple form of eudaimonism, we should start out by identifying the things that are basic intrinsic value states on that axiology. We might say that a *basic eudaimonistic state* is any pure attribution of happiness, something like:

J12n: Jones is happy to degree +12 at noon today.

And then we could say that on our eudaimonistic axiology, the basic intrinsic value states are all and only these basic eudaimonistic states. We could go on to say that the intrinsic value of each such thing is equal to the precise amount of happiness attributed. Thus, IV(J12n) = +12.

Then we could say that the intrinsic value of a world, life, or total consequence is equal to the sum of the intrinsic values of the basics that are true within it.[25] Thus, on this axiology, J12n will inevitably contribute exactly +12 units of intrinsic value to any world at which it occurs. Given this way of calculating the intrinsic value of a world, we can conclude by saying that a world, life, or total consequence is intrinsically good (bad) if its intrinsic value is greater (less) than zero. If we proceed in this way, our formulation will not be internally inconsistent.[26]

Suppose we prefer a much more complex axiology, such as the one Moore presented in the final chapter of *Principia Ethica*. Moore did not endorse anything simple, like hedonism or eudaimonism. Rather, he said that some of the "great intrinsic goods" are things such as "the love of beauty" and "the hatred of evil". Moore's discussion makes it clear that each of these is in fact a very complicated type of state of affairs involving "proper appreciation", knowledge, and the actual existence of the appreciated object. What then are the basics?

As I see it, one part of Moore's view is that something such as this is a basic intrinsic value state:

LB: Jones taking aesthetic pleasure of intensity +10 in the beautiful qualities, F1, F2, F3, etc. of object A at noon today while knowing that object A in fact exists and has F1, F2, F3, etc., and that these are beautiful qualities.

LB is an instance of "the love of beauty". It is intended to be a pure attribution of a complex relation: *x,n,K,y,t|x taking aesthetic pleasure of intensity n in the beautiful qualities in set K of object y at time t while knowing that y in fact exists and has the members of K and that the members of K are beautiful qualities*. The intrinsic value of LB might be thought to depend upon the intensity of the aesthetic pleasure taken. If we wanted to capture some other thoughts Moore suggests, we could consider a slightly more complex vari-

ant in which there are places for numbers indicating the extent to which the object has the beautiful properties, and perhaps even more.

Note that LB has many "proper parts", states of affairs that it entails but which do not entail it. Some examples are:

> JP: Jones taking pleasure of intensity +10.
> JKA: Jones knowing that object A exists.

None of these is a basic intrinsic value state on the Moorean axiology. None of these has any basic intrinsic value and it is not clear that any of them has any determinate (regular) intrinsic value. Similar comments apply to all the other proper parts of LB. Thus, the intrinsic value of LB is not determined by performing some mathematical operation such as addition on the intrinsic values of its parts. When those basically worthless parts come together, something of great value emerges. Perhaps this gives some sense to Moore's insistence that the great intrinsic goods are all organic unities.[27]

As an example of a case of "hatred of evil", we could consider:

> HE: Jones experiencing hatred of intensity +10 in the vicious qualities, V1, V2, V3, etc. of person Z at noon today while knowing that person Z in fact exists and has V1, V2, V3, etc., and that these are vicious qualities.

To complete the statement of a Moorean axiology, we would have to identify all the main sorts of basic intrinsic value states, and we would have to provide principles specifying, for each type, the basic intrinsic value of states of that type. Then we could say that the intrinsic value of a life, a total consequence, or a world is a simple function of the intrinsic values of the basics true therein. Again, if this were done carefully, it would be internally consistent and true to the spirit of Moorean axiology.

Thomas Hurka has recently discussed a type of axiological theory that involves a sort of "embedding", or recursion. On one theory of that sort, we could start by saying that pure attributions of pleasure are intrinsically good basics. Then we could say that knowledge of something intrinsically good is also intrinsically good. Now consider this sequence:

1. S1 feeling pleasure of intensity +10 at noon.

2. S2 knowing that S1 feels pleasure of intensity +10 at noon.

3. S3 knowing that S2 knows that S1 feels pleasure of intensity +10 at noon.

4. S3 knowing that S1 feels pleasure of intensity +10 at noon.

On the imagined axiology, each of these would be a basic intrinsic value state and each would independently contribute some intrinsic value to any world in which it is true. Assuming that each has a basic intrinsic value of +10, we might want to say that if these four are the only basics true at a world, then the value of that world is +40.

Thus, though (3) "guarantees" forty units of intrinsic value, I have assigned it a basic intrinsic value of just +10. (3)'s guarantee is of course fulfilled; if (3) is true, then so are (1) and (2) and (4). Since each of these others is also a basic on this axiology, each of them will contribute ten more points to the value of the world.

I prefer to give (3) a basic value of just +10 rather than +40. One reason for this is that I want to appeal to (3)'s basic intrinsic value when I calculate the value of the life of S3. As I see it, anyone endorsing this axiology would want to say that the value of S3's life should be increased by just 10 points as a result of (3)'s truth. Even though the truth of (3) guarantees 40 units of value, to increase the value of S3's life by 40 points would be excessive. S3's life is only 10 points better as a result of the truth of (3). So I say that BIV(3) = +10. Then the value of the world comes out right, as does the value of each of the lives.

ii. The second puzzle concerns the concept of instrumental value. Clearly, we cannot say that the instrumental value of a state of affairs is equal to the sum of the regular intrinsic values of all the things it causes. To do so would invite double counting in cases such as the case involving Bob who reads a delightful book. Recall this state of affairs:

    B: Bob reads at 9:00PM on Monday evening.

I imagined a case in which B makes Bob happy to degree +10 while he reads, and sits in his easy chair, and wears his eyeglasses. In this case, B has many consequences, including the following:

> C1: Bob being happy to degree +10 at 9:00PM on Monday evening.
>
> C2: Bob being happy while reading.
>
> C3: Someone being happy to degree +10.
>
> C4: The man in Bob's easy chair being happy to degree +10.
>
> C5: Bob being happy at least to degree +9.
>
> C6: Bob being happy to degree +10 while sitting in his easy chair and wearing his eyeglasses at 9:00PM on Monday evening.
>
> C7: Either Bob being happy to degree +10 or Babe being happy to degree +15.

Since each of these is a consequence of B, and each seems to have some positive (regular) intrinsic value, the simple aggregative principle yields an incorrectly high estimate of the instrumental value of B.

As a first step toward a solution, I propose:

EV3: The instrumental value of a state of affairs, P, at a world, w, is the sum of the intrinsic values of all the basic intrinsic value states, Q, such that P causes Q at w.

Again assuming a simple sort of eudaimonism, EV3 implies (much more plausibly) that the instrumental value of B is +10.

Some philosophers think that a state of affairs gets a boost in instrumental value if it *prevents evil*. Thus, we might say that it's instrumentally good to get your cavities filled since this prevents the unhappiness you would experience if you left them unfilled. EV3 makes no provision for such preventive instrumental value. It is easy enough to incorporate this idea. We do it in two steps as follows:

First we introduce the notion of "preventive value". To find the preventive value of a state of affairs, P, first consider all the states of affairs it prevents; then locate the basic intrinsic value states among these; then find the sum of their intrinsic values. Since it's good to prevent evil, but bad to prevent good, the preventive value is equal in amount but opposite in sign to the sum of the intrinsic values of the basics prevented. So if some state of affairs prevents basics whose intrinsic values sum to –10, then that state of affairs has a preventive value of +10.

Now we can define instrumental value:

EV4: The instrumental value of a state of affairs, P, at a world, w, is the sum of the intrinsic values of all the basics, Q, such that P causes Q at w plus the preventive value of P at w.[28]

iii. Our third puzzle concerned the calculation of the intrinsic values of complex things such as lives, total consequences, and possible worlds. Since I have incorporated aggregative principles into the formulation of the sample axiologies, it is not necessary to devote much attention to this puzzle. The solution is the same in all cases: the intrinsic value of a complex thing is in each of these cases equal to the sum of the intrinsic values of the basic intrinsic value states contained within.

I acknowledge that this leaves the intrinsic values of certain states of affairs undefined. For example, consider disjunctions such as:

Jv: Jones being happy to degree +12 at t1 or Jones being happy to degree +10 at t2.

Even on the simplest sort of eudaimonism, Jv is neither a basic nor a world nor a life. Nothing I have said implies anything about its intrinsic value. So far as I can see, this is not a problem. Let its intrinsic value be undefined. If it is true and a part of the life of Jones at some world, then either its first disjunct or its second disjunct or both will also be included in Jones' life at that world. One or both of those basics will, therefore contribute its full value to the value of the life and the world. Thus, the axiological significance of the disjunction is entirely accounted for by the basic intrinsic value states that make it true.

Roderick Chisholm, Gilbert Harman, Thomas Hurka, and Warren Quinn.

I acknowledge that I have not incorporated every suggestion that I have received. Nevertheless, I am grateful for them all.

I have presented various bits of this paper in seminars and talks over a period of many years. Many participants in those discussions provided useful insight and challenge. Although I cannot now cite them individually, I am grateful to them all.

NOTES

[1] Roderick Chisholm, *Brentano and Intrinsic Value*, 74.

[2] Richard Brandt, 'Hedonism', 432.

[3] See, for example, *Republic* 357; *Nicomachean Ethics*, 1096b13.

[4] G. E. Moore, *Principia Ethica*, 22.

[5] C. I. Lewis, *An Analysis of Knowledge and Valuation*, 385.

[6] Gilbert Harman, 'Toward a Theory of Intrinsic Value', 796.

[7] In Section 4 below, I discuss some further difficulties with this account of instrumental value and I sketch a way of dealing with some of them.

[8] In my *Doing the Best We Can* I present a normative theory according to which what we ought to do depends upon what happens in the "best accessible worlds". Clearly, this view makes essential use of the idea that worlds can be evaluated in terms of intrinsic value.

[9] Some problems remain. If we assume that the temporal indicators refer to *instants* of time, and that there are infinitely many of these, each of which is unextended in time, we will run into trouble. Since even a short period of happiness will presumably involve happiness at infinitely many of these instants, it will be hard to find the sum of the values of these things. Perhaps a better approach involves the notion that time can be exhaustively broken down into a finite number of non-overlapping brief segments. The temporal indicators may then be taken to refer to these tiny intervals.

[10] See, for example, Chisholm's 'The Intrinsic Value in Disjunctive States of Affairs' and Quinn's 'Theories of Intrinsic Value'.

[11] Throughout this paper I often have in mind things that Gilbert Harman said in his 1967 paper 'Toward a Theory of Intrinsic Value'. I have to admit that I don't fully understand every detail of his argument. Nevertheless, he comes to the conclusion that we should recognize basics, and he seems to make use of some considerations that are at least similar to things I will say.

[12] Harman, 'Toward a Theory of Intrinsic Value', 798.

[13] Harman, op. cit., 799.

[14] In Section 16 of *Principia Ethica*, Moore seems to be discussing this principle, though he speaks of the "universal truth" of judgments of intrinsic value. Later in 'The Conception of Intrinsic Value' he again seems to endorse a version of the

necessity principle when he says (260–261) that it is *impossible* for one and the same thing to have a certain value at one time or in one set of circumstances, but not at another time or in another set of circumstances. This is complicated by the fact that he seems also to be assuming that nothing can change its intrinsic nature.

[15]   A variety of suggestions have been made. Here are three: (a) the intrinsic value of a state of affairs such as someone is happy is equal to the intrinsic value of least good basic that could make it true; (b) such things have vague intrinsic values; they are good, but they do not have any precise amount of goodness – perhaps the value of such a thing can be represented as a range between worst and best basic that could make it true; (c) such things have, as their intrinsic values, the average intrinsic value of basics that make it true. Some of these have been discussed in the literature. None of them seems particularly plausible to me.

[16]   Zimmerman presented the outline of this idea in his 'Evaluatively Incomplete States of Affairs' back in 1983. He discusses it further in his forthcoming 'Virtual Intrinsic Value'.

[17]   For an argument against this idea, see Chisholm, *Brentano and Intrinsic Value*, 81–82.

[18]   Thomas Hurka discusses two conceptions of defeat in his very insightful paper, 'Two Kinds of Organic Unity'.

[19]   Chisholm seems to see things in this way in his Presidential Address, 'The Defeat of Good and Evil'. I inadvertently suggested that I endorse it in my 'Adjusting Utility for Justice'.

[20]   Here, I am speaking loosely. In many cases there is no world that differs only with respect to one selected state of affairs.

[21]   I first encountered the idea that basic intrinsic value states fall into families in Warren Quinn's, 'Theories of Intrinsic Value'. Quite a bit of what I say in this paper is influenced by Quinn.

[22]   Moore, *Principia Ethica*, 91.

[23]   For a perceptive discussion of problems about the isolation test, see Eva Bodanszky and Earl Conee, 'Isolating Intrinsic Value'.

[24]   On more complex axiologies, this may sometimes fail. Sometimes a member of one family will stand in a logical relation to a member of another. Thus, for example, we can consider an axiology according to which pleasure and knowledge are both intrinsically good. Then consider this pair:

J10: Jones is pleased to degree +10 at t.

SK10: Smith knows that Jones is pleased to degree +10 at t.

We might want to say that each of these is intrinsically good; they are members of distinct families, and SK10 entails J10. Thomas Hurka discusses this sort of axiology in detail in his 'Virtue as Loving the Good' as well as in his forthcoming *Virtue and Vice: A Perfectionist Account*.

[25]   Note that we do not say that the value of a world depends upon what basics *exist* within it; that would yield an incorrect number. Since basics are abstract entities, each of them exists at each world. Rather, we say that it depends upon what basics are *true* within it.

[26] Although it may have other problems involving the need to find the sum of an infinite number of addends. It may also have some problems due to the fact that the substance of the axiology is somewhat naive.

[27] I do not offer this remark as a general characterization of organic unities.

[28] I have benefitted from discussions on this topic with Ben Bradley. See his doctoral dissertation, *Species of Goodness*, University of Massachusetts, Fall, 1998.

REFERENCES

Aristotle, *Nicomachean Ethics*.

Bodanszky, E. and Conee, E. (1981), 'Isolating Intrinsic Value', *Analysis* 41, 51–53.

Bradley, B. (1998), *Species of Goodness*, University of Massachusetts.

Brandt, R. (1967), 'Hedonism', in P. Edwards (ed.), *The Encyclopedia of Philosophy*, vol. 4, New York: Macmillan Publishing Co. & The Free Press), pp. 432–435.

Chisholm, R. (1986), *Brentano and Intrinsic Value*, Cambridge: Cambridge University Press.

Chisholm, R. (1968), 'The Defeat of Good and Evil', Presidential Address delivered before the Sixty-fifth Annual Meeting of the Eastern Division of the American Philosophical Association in Washington, D.C. December 27–30, 1968. *Preceedings and Addresses of The American Philosophical Association* Volume XLII – 1968–69, pp. 21–38.

Chisholm, R. (1975), 'The Intrinsic Value in Disjunctive States of Affairs', *Nous* 9, 295–308.

Feldman, F. (1978), *Introductory Ethics*, Englewood Cliffs: Prentice-Hall, Inc.

Feldman, F. (1986), *Doing the Best We Can: An Essay in Informal Deontic Logic*, Dordrecht: Reidel.

Harman, G. (1967), 'Toward a Theory of Intrinsic Value', it The Journal of Philosophy 64, 792–804.

Hurka, T., *Virtue and Vice: A Perfectionist Account*.

Hurka, T. (1998), 'Two Kinds of Organic Unity', presented at Twenty-first Annual Symposium in Philosophy at the University of North Carolina at Greensboro, April 3–5, 1998. Forthcoming in *The Journal of Ethics*.

Lewis, C.I. (1946), *An Analysis of Knowlege and Valuation*, LaSalle, Illinois: Open Court Publishing Co.

Moore, G.E. (1993), *Principia Ethica*, revised edition edited and with an introduction by Thomas Baldwin, Cambridge: Cambridge University Press.

Moore, G.E. (1922), 'The Conception of Intrinsic Value', in *Philosophical Studies* London: Routledge & Kegan Paul, pp. 253–275.

Plato, *Republic*.

Quinn, W. (1974), 'Theories of Intrinsic Value', *American Philosophical Quarterly* 11(2), 123–132.

Zimmerman, M.J. (1983), 'Evaluatively Incomplete States of Affairs', *Philosophical Studies* 43, 211–224.

Zimmerman, M.J. (in press), 'Virtual Intrinsic Value', forthcoming in *Philosophy and Phenomenological Research.*

Department of Philosophy
University of Massachusetts
352 Bartlett Hall
Box 30525
Amherst, MA 01003-0525, USA
(E-mail: ffeldman@philos.umass.edu)